Non-cooperative Facility Location and Covering Games

Jean Cardinal* Martin Hoefer†

Abstract

We consider a general class of non-cooperative games related to combinatorial covering and facility location problems. A game is based on an integer programming formulation of the corresponding optimization problem, and each of the k players wants to satisfy a subset of the constraints. For that purpose, resources available in integer units must be bought, and their cost can be shared arbitrarily between players. We consider the existence and cost of exact and approximate pure-strategy Nash equilibria. In general, prices of anarchy and stability are in $\Theta(k)$ and deciding the existence of a pure Nash equilibrium is NP-hard. Under certain conditions, however, cheap Nash equilibria exist, in particular if the integrality gap of the underlying integer program is 1, or in the case of single constraint players. We also present algorithms that compute simultaneously near-stable and near-optimal approximate Nash equilibria in polynomial time.

*Computer Science Department, Université Libre de Bruxelles, Belgium, jcardin@ulb.ac.be.

†Corresponding author. Department of Computer Science, RWTH Aachen, Germany, mhoefer@cs.rwth-aachen.de.

1 Introduction

Over the last decade the area of algorithmic game theory has evolved, in which computational environments are analyzed using game-theoretic models. Motivated in large parts by the Internet, the resulting effects and dynamics of introducing selfish behavior of independent agents into a computational environment are studied. In this paper we follow this line of research by considering a general class of non-cooperative games based on general covering integer problems. Problems concerning service installation or clustering, which play an important role in large networks like the Internet, formally are modeled as some variant of covering, partition, or facility location problems. The Internet is composed, built, and maintained by a multitude of different parties - companies, governmental institutions, as well as private users - which all have their personal interests and goals. In particular, investments into the infrastructure of the network like establishing connections or placement and installation of service technology are done by telecommunication companies that strive to sell products and maximize profits. This requires to model these scenarios using economical, and in particular, game-theoretic approaches. Our games can serve as a basis to analyze service installation and investment problems in the presence of independent non-cooperative selfish agents.

The formulation of our games generalizes an approach by Anshelevich et al. [6], who proposed games in the setting of Steiner forest design. In particular, we consider a covering optimization problem given as an integer linear program and turn this into a non-cooperative game as follows. Each of the k non-cooperative players considers a subset of the constraints and strives to satisfy them. Players offer money for the purchase of units of resources, which are modeled by the variables. A variable is raised in *bought* integer units if the resulting cost in the objective function is paid for by the players. Bought units can be used by *all* players simultaneously to satisfy their constraints – no matter whether they contribute to the cost or not. A player strives to minimize her cost, but insists on satisfaction of her constraints. A variety of covering problems, most prominently variants of set cover and facility location, can be turned into a game with the help of this approach. In these games, the problem of finding the social optimum solution is represented by the underlying integer program. The problem of finding the best response strategy for a player i represents an adjusted subproblem. It is derived by dropping all constraints of players other than i and reducing the cost coefficients in the objective function by the share that is contributed by other players.

We investigate our non-cooperative games in terms of stable solutions, which are the pure strategy Nash equilibria of the game. We do not consider mixed strategy equilibria, because our environment requires a concrete investment rather than a randomized action, which would be the result of a mixed strategy. Hence, when using the term "Nash equilibrium" we mean pure strategy exact Nash equilibria unless mentioned otherwise. First, we study the *price of anarchy* [35, 42], which is the ratio of the cost of the worst Nash equilibrium over the cost of a minimum cost cover satisfying all constraints of all players for a game. In addition, we consider the *price of stability* [5], which measures the best Nash equilibrium in terms of the optimum cost instead of the worst equilibrium. Both of these measures, however, can be prohibitively high, and pure Nash equilibria might not exist in our games. Hence, we focus on a trade-off between stability and optimality by considering a two-parameter optimization problem to find (α, β)-approximate Nash equilibria (denoted as (α, β)-NE). These are states in which each player can reduce her cost by unilateral deviation by at most a factor of α and that represent a β-approximation to the socially optimum cost. We refer to α as the *stability ratio* and β as the *approximation ratio*. Intuitively, we strive to obtain a state with small values for α and β. This guarantees that it is near-stable and near-optimal.

Note that, in contrast to the additive concept of ε-Nash equilibrium [23] used in the game theory literature, our parameters are multiplicative ratios. The reason that we chose to use multiplicative ratios is that in our games each player is required to optimize a combinatorial optimization problem to find a best response strategy. Most of these problems turn out to be NP-complete. In this case a multiplicative stability ratio has the advantage that results directly relate to approximation algorithms with multiplicative approximation factors, which might be used by players to find good strategies. Such algorithms have been proposed for many of the underlying optimization problems we consider in this paper.

1.1 Outline and contributions

We study our games with respect to quality of pure strategy exact and approximate Nash equilibria. Throughout the paper we denote a feasible solution by $\mathcal{C}$ indicating that it forms a cover, and denote the social optimum solution by $\mathcal{C}^*$. Our contributions are as follows.

Section 2 introduces covering games using the special case of vertex cover and provides some initial observations. In Section 3 we show lower bounds for the vertex cover game. The price of anarchy is k even if the underlying graph is a tree. There exist simple unweighted and weighted games for two players without exact Nash equilibria. They can be used to prove that the price of stability can be arbitrarily close to $k-1$. Determining existence of exact Nash equilibria for a given game is NP-hard, even for unweighted games or two players.

In Section 4 we specify classes of set cover games for which exact Nash equilibria exist and the price of stability is 1. For games, in which the underlying integer program has integrality gap 1, our proof is based on linear programming duality. This yields a polynomial time algorithm to calculate cheap exact Nash equilibria. For the class of *singleton games*, in which each player owns exactly one element, the proof is based on a local improvement method. This can be turned into an FPTAS that yields $(1+\varepsilon, O(\log n))$-NE in polynomial time, for any constant $\varepsilon > 0$. Thus, every player can improve her payments through selfish defection by a factor of at most $1+\varepsilon$, and the bought cover is an $O(\log n)$-approximation of the minimum cost cover for all players.

In Section 5 we further consider the problem of finding (α, β)-NE, i. e., cost sharings of covers that are cheap and allow low incentives for players to deviate. We propose a simple algorithm for set cover games that finds (f, f)-NE, in which f is the maximum frequency of any item in the sets. For vertex cover games with $\mathsf{f} = 2$ we show that any game has a $(2, 1)$-NE. This is tight, because there are general vertex cover games without an (α, β)-NE for any $\alpha < 2$. Recent progress on the complexity status of the minimum vertex cover problem can be used to reasonably conjecture that there can be no polynomial time algorithm with a better guarantee for the approximation ratio β as well. For planar games our argument extends to a lower bound of 1.5 on α. It can be increased close to 2 by forcing β to be close to 1, indicating a trade-off between the ratios.

In Section 6 we discuss facility location games. The results and proofs mainly translate previous arguments given for covering games. Even in the most simple variant, the metric uncapacitated facility location (UFL) game, the price of anarchy is exactly k, the price of stability as high as $k-2$, and it is NP-hard to determine whether a game has an exact Nash equilibrium. However, if every player has only a single client, it follows from [6] that the price of stability is 1 and that $(1+\varepsilon, 1.5)$-NE can be found in polynomial time, for any constant $\varepsilon > 0$. For the metric UFL game there is an algorithm to compute $(3, 3)$-NE in polynomial time. We will extend the ideas to a more general setting coming from [26], which we term *connection-restricted facility location (CRFL)*. For a subclass of *closed* CRFL (CCRFL) games, the price of stability is 1 if the integrality gap of the

Game Class	NE	find best NE	[PoA, PoS]	apx. NE
Covering	see Set Cover			
Set Cover(SC)	NP-hard		$[k, k-1]$	(f, f)
Singleton SC	$\surd$	NP-hard	$[k, 1]$	$(1+\varepsilon, O(\log n))$
Integral SC	$\surd$	P	$[k, 1]$	$(1, 1)$
CRFL	NP-hard		$[\infty, k-2]$	
Metric UFL	NP-hard		$[k, k-2]$	$(3, 3)$
Singleton UFL	$\surd$		$[k, 1]$	$(1+\varepsilon, 1.5)$
Integral CCRFL	$\surd$	P	$[\infty, 1]$	$(1, 1)$

Table 1: Main results

corresponding LP-relaxation is 1. The best Nash equilibrium can be derived with LP-duality.

Table 1 summarizes the results obtained in this paper. Column "Game Class" denotes the class of games; Covering for the general case, set cover, UFL or CCRFL for more specific games; finally, the special cases integral, metric, and singleton as mentioned above. Column "NE" shows whether or not an exact Nash equilibrium is guaranteed to exist, and whether or not it is hard to decide existence. Column "find best NE" shows the complexity of finding the best Nash equilibrium. Column "[PoA, PoS]" shows upper bounds on the price of anarchy as well as lower bounds on the price of stability. Finally, "apx. NE" displays bounds on (α, β)-NE that can be computed in polynomial time.

Part of this work has appeared previously in two extended abstracts [11, 29]. In [11] we treated only the case of vertex cover games, whereas the remaining results of this paper are reported in [29]. In [11] many proofs are sketched or misssing (e.g., NP-hardness to determine existence of exact Nash equilibria, existence of $(2, 1)$-NE, computation of $(2, 2)$-NE in polynomial time), which are presented in full detail in this paper. In addition, instead of bipartiteness of the graph as stated in [11], the right condition for existence of Nash equilibria is an integrality gap of 1 for the LP. While this condition is considered for the more general class of set cover games in [29], a full proof of the result has not been given before this paper. In addition, we present complete proofs that were presented only as sketches in [29], for instance for the price of stability and the existence and computability of $(3, 3)$-NE in metric UFL games, and for (f, f)-NE in set cover games.

1.2 Discussion and Related Work

Our analysis uses concepts developed for non-cooperative games in the area of algorithmic game theory, in particular prices of anarchy and stability characterizing worst- and best-case Nash equilibria. The price of anarchy has been studied in a large and diverse number of games, e.g., in areas like routing and congestion [7, 35, 43], network creation [6, 24], wireless ad-hoc networks [20, 28], or facility location [45]. The price of stability [5] has been introduced more recently and studied for instance in network creation games [5, 31] or linear congestion games [16].

The non-cooperative model we consider stems from [6], who proposed a connection game based on the Steiner forest problem. They show that prices of anarchy and stability are $\Theta(k)$ and give a polynomial time algorithm for $(4.65+\varepsilon, 2)$-NE. In our uncapacitated facility location (UFL) game we assume that each of the clients must be connected directly to a facility. The connection requirement for a client cannot be satisfied by connecting it to a different client, which is (possibly over other clients) connected to a facility. Hence, it is possible to turn this into a single source connection game on a directed graph. The idea is to introduce a source node s, connect all facilities f to it and direct all edges from clients to facilities. The costs for the new edges (f, s) are given by the opening costs $c(f)$ of the corresponding facilities. If we allow connecting to facilities using longer paths, the same construction allows us to turn the game into an undirected single source connection game (SSC) considered in [6, 30]. For both the UFL and the SSC game results in [6] suggest that the price of anarchy is k and the price of stability is 1 if each player has a single client. In [30] we provided improved algorithms for $(3.1 + \varepsilon, 1.55)$-NE for the SSC game. In addition, we showed that deciding the existence of exact Nash equilibria is NP-hard. Extended results on the price of stability for Steiner network and group formation games have been obtained recently in [3, 4].

A similar model for non-cooperative cost sharing in covering scenarios was proposed by Anshelevich et al. [5], who study an egalitarian cost sharing protocol closely related to the Shapley value. The resulting game is a potential game [40], i.e., pure Nash equilibria are guaranteed to exist, and the price of stability is in $\Theta(\log k)$. Variants of this game recently have been subject to increased research interest [1, 12–14, 22]. For a class of similar protocols for network cost sharing games, Chen et al. [15] derived a characterization of the ones that allow pure Nash equilibria and showed that the Shapley protocol results in minimal prices of anarchy and stability. The main motivation for the study of these aspects is the *design* of routing protocols. The task is to set cost sharing rules in order to obtain favorable properties in the resulting game like pure Nash equilibria and low equilibrium inefficiency. Player strategies are limited by design to chosing routes. However, when considering the interplay of companies building the Internet, there is no agency that is able to prescribe a certain way of sharing the cost of an investment. Instead, we take a more general approach with payment functions as strategies, a standard in the area of cost sharing mechanisms. On the one hand, this results in non-existence of pure Nash equilibria and a high price of stability. On the other hand, in many settings (including ours) finding an optimal investment strategy is NP-hard. This makes it reasonable to consider cheap approximate equilibria, which exist in our model because of more general strategy spaces.

Another model, which is close to our approach, are cooperative games and mechanism design problems based on optimization. The set cover problem is a fundamental combinatorial optimization problem and has been studied for decades. Recently, a cooperative set cover game was studied by Immorlica et al. [32]. In this coalitional game, each item is an agent and each coalition of players is associated with a certain cost value - the cost of a minimum cover. In [32] cross-monotonic cost sharing schemes were investigated. For each coalition of players covered they distribute the cost to players in a way that every player is better off if the coalition expands. The authors showed that for vertex cover no more than $O(n^{-\frac{1}{3}})$, for set cover no more than $O(\frac{1}{n})$ and for uncapacitated facility location no more than $\frac{1}{3}$ of the cost can be charged to the agents with a cross-monotonic scheme, respectively.

Closely related to cooperative games is the study of cost sharing mechanisms. In this scenario a central authority distributes service to players and strives for their cooperation. Starting with [19] cost sharing mechanisms for set cover problems have been considered. Every player corresponds to a single item and has a private utility (i.e., a willingness to pay) for being in the cover. The mechanism asks each player for her utility value. Based on this information its goal is to pick a subset of items to be covered, to find a minimum cost cover for the subset and to distribute costs to covered item players such that no coalition can be covered at a smaller cost. A strategyproof mechanism allows no player to lower her cost by misreporting her utility value. The authors in [19] presented strategyproof mechanisms for set cover and facility location games. For set cover games this work was extended recently by [36, 37] to the consideration of different social desiderata like fairness aspects and model formulations with items or sets being agents.

More general cooperative games based on covering/packing problems were studied in [17]. It was shown that the core of such games is non-empty if and only if the integrality gap is 1. Additional results included polynomial time computability of core members in a number of games encompassed by this framework. In [26] similar results are shown for a general class of facility location games and an appropriate integer programming formulation.

Cooperative games and the mechanism design framework are used to capture situations with selfish service receivers who can either cooperate to an offered cost sharing or manipulate. Players may also be excluded from the game depending on their utility. A major goal has been to derive good cost sharing schemes that guarantee truthfulness or budget balance. Our game, however, is strategic and non-cooperative in nature and allows players a much richer set of actions. In our game each player is motivated to participate in the game. We investigate uncoordinated service installation scenarios rather than a coordinated environment with a mechanism choosing customers, providing service, and charging costs.

Under the name *competitive location* there has been a high research activity on game-theoretic models for spatial and graph-based facility location during the last decades [21, 39]. These models consider facility owners as players that selfishly decide where to open a facility. Clients are modeled as part of player utility, e.g., they are assumed to connect to the closest facility and thus represent part of the revenue a player gets from opening a facility. Recent examples of this kind of location games are also found in [18,45]. According to our knowledge, however, none of these models consider the clients as non-cooperative players that need to create connections and facilities without central coordination.

2 Covering Games

2.1 Vertex and Set Cover Games

We start by introducing covering games for the important special case of vertex cover. The *vertex cover game* for k players is defined as follows. In an undirected graph $G = (V, E)$ with $n = |V|$ and $m = |E|$ each player i owns a set $E_i \subseteq E$ of edges. We denote by $G[E_i]$ the graph induced by the edges in E_i, and by $V(G[E_i])$ the set of vertices of $G[E_i]$. Without loss of generality we assume that $E = \bigcup_i E_i$, as any edge not owned by a player is irrelevant in the game. Each player strives to establish a service at one or both endpoints of each of her edges. For each vertex v there is a nonnegative cost $c(v)$ for establishing service at this vertex. A *strategy* for a player i is a function $p_i : V \to \mathbb{R}_0^+$ specifying an offer to costs of each vertex. A *state* or *payment scheme* is a vector $p = (p_1, \ldots, p_k)$ specifying a strategy for each player. If the sum of offers of all players for vertex v exceeds its cost (i.e., if $\sum_i p_i(v) \geq c(v)$) it is considered *bought*. Bought vertices can be used by all players to cover their incident edges. The individual cost of a player is $c_i(p) = \sum_{v \in V} p_i(v)$ if there is a cover of bought vertices for E_i and $c_i(p) = \infty$ otherwise. Hence, a player tries to minimize her investment but insists on covering her edges.

A (pure strategy) exact *Nash equilibrium* is a payment scheme such that no player i can unilaterally improve her individual cost by changing her strategy, that is $c_i(p_i, p_{-i}) \leq c_i(p'_i, p_{-i})$ for any player i and strategy p_i'. Note that we consider only pure strategy Nash equilibria in this paper. The *social cost* of a state p is given by the sum of player costs $c(p) = \sum_i c_i(p)$. A *social optimum* is a strategy profile p^* such that $p^* = \arg\min_p c(p)$. It is easy to verify that in a Nash equilibrium p a vertex v is either bought exactly or not at all, i.e., $\sum_i p_i(v) = c(v)$ or $\sum_i p_i(v) = 0$. Furthermore, in every Nash equilibrium the set of bought vertices contains one incident vertex for every edge $e \in E_i$ of every player i, because otherwise the respective player is not satisfied. Thus, the set of bought vertices in a Nash equilibrium is a *vertex cover* $\mathcal{C} \subseteq V$ such that for each $e \in E$ there is at least one incident vertex in $\mathcal{C}$. For the social cost of a Nash equilibrium p this implies $c(p) = c(\mathcal{C}) = \sum_{v \in \mathcal{C}} c(v)$ with the corresponding cover $\mathcal{C}$. A similar observation holds for the social optimum p^*; it is easy to observe that in p^* the vertices of a minimum cost vertex cover $\mathcal{C}^*$ are bought exactly, and thus $c(p^*) = c(\mathcal{C}^*)$. In this way, Nash equilibria and social optima represent cost sharings of possibly different vertex covers.

To quantify the inefficiency of the covers bought in Nash equilibria, we consider the *price of anarchy* and the *price of stability*, i.e., the ratio of the cost of the worst and best Nash equilibrium over the cost of the social optimum, respectively. The price of anarchy is defined as the maximum of $c(p)/c(\mathcal{C}^*)$ for any Nash equilibrium p. Similarly, the price of stability is the minimum of $c(p)/c(\mathcal{C}^*)$ for any Nash equilibrium p. In addition, we consider approximate Nash equilibria. A *(α, β)-NE* is a payment scheme p such that $c(p_i, p_{-i}) \leq \alpha c(p_i', p_{-i})$ for any player i and strategy p_i', and that $c(p) \leq \beta c(\mathcal{C}^*)$. In such a state every player can improve her cost at most by a factor of α by switching to another strategy, and the total payments approximate the minimum cost cover by a factor β. We will refer to the factor β as the *approximation ratio*, and we term α as the *stability ratio*. Finally, we call a vertex cover game *unweighted* if all vertices have equal costs, and *weighted* otherwise. We refer to games with a planar graph G as *planar games*.

This construction is easily generalized to a *set cover game*, which is based on an instance of the set cover problem. In the set cover problem there is a set $\mathcal{S}$ of n sets S over a set E of m elements. For each $S \in \mathcal{S}$ we have $S \subseteq E$ and a cost $c(S) \geq 0$. The problem is to find a subset of $\mathcal{S}$ covering all elements with minimum total cost. Note that in the special case of vertex cover elements are edges, and vertices can be seen as sets of incident edges. Using this analogy, it is easy to generalize all the constructions above to this case. In particular, in a set cover game, each player i strives to cover a set of elements E_i. Player i chooses p_i with $p_i(S)$ being the payments offered to set S. Set S is bought if the total contributions exceed the cost $c(S)$. A Nash equilibrium p represents a cost sharing of a *set cover* $\mathcal{C} \subseteq \mathcal{S}$ of E, i.e., a set of sets such that $\bigcup_{S \in \mathcal{C}} S = E$. A social optimum p^* is a cost sharing of a set cover $\mathcal{C}^*$ of minimum total cost. By $\mathsf{f} = \max_{e \in E} |\{S \in \mathcal{S} \mid e \in S\}|$ we denote the maximum frequency of any element in the sets. Note that a vertex cover game is a set cover game with $\mathsf{f} = 2$, as each edge is incident to exactly two vertices.

2.2 Integer Covering Games

Our games have an interesting connection to the linear programming formulation of the underlying optimization problem. Consider an instance of the vertex cover problem, in which we simply strive to find $\mathcal{C}^*$, as an integer linear program. For each vertex v there is a binary variable x_v indicating whether it is in the cover or not. Furthermore, for each edge $e = (u, v)$ there is a constraint $x_u + x_v \geq 1$ ensuring that e is covered. The cost of a vertex is given by the cost coefficient $c(v)$, which appears in the objective function.

$$\begin{array}{lll} \text{Min} & \sum_{v \in V} c(v) x_v & \\ \text{subject to} & x_v + x_u \geq 1 & \forall (u,v) \in E \\ & x_v \in \{0,1\} & \forall v \in V. \end{array} \tag{1}$$

Suppose for a vertex cover game the underlying optimization problem of finding $\mathcal{C}^*$ is described by the above integer program. Then a player actually strives to satisfy the constraints corresponding to her edges. This is done by offering money to the purchase of variable units (i.e., vertices). Once the sum of offers exceeds the cost given by the coefficient in the objective function, it is considered bought and can be used by all players to satisfy their constraints (i.e., cover their edges). The remaining specification of the game and the definitions follow accordingly.

This formulation of the game allows a straightforward translation to games based on arbitrary covering integer programs. A *covering game* is based on a covering integer problem (CIP) [44, chapter 13.2] given as

$$\begin{array}{lll} \text{Min} & \sum_{f=1}^{n} c(f) x_f & \\ \text{subject to} & \sum_{f=1}^{n} a_{tf} x_f \geq b_t & \forall t = 1, \ldots, m \\ & x_f \in \mathbb{N} & \forall f = 1, \ldots, n. \end{array} \tag{2}$$

All constants are assumed to have non-negative (rational) entries a_{tf}, b_t, $c(f) \geq 0$ for all $t = 1, \ldots, m$ and $f = 1, \ldots, n$. Each player i owns a subset of the constraints $C_i \subset \{1, \ldots, m\}$, which she strives to satisfy. Integral units of a resource f have cost $c(f)$. They must be bought to be available for constraint satisfaction. Each player offers money to purchase some integral units of the resources f. A unit is considered bought if the sum of all offered payments of all players purchase its cost. If a unit is paid for, it can be used by all players no matter whether they contribute to the cost or not.

More formally, we assume that each player i chooses a strategy p_i to specify the non-negative contribution $p_i(f)$ to each resource f. In a state $p = (p_1, \ldots, p_k)$ a total of $x_f = \lfloor \sum_i p_i(f)/c_f \rfloor$ units of resource f are bought, and the contribution of player i to each unit is $p_i(f)/x_f$. The individual cost is $c_i(p) = \sum_f p_i(f)$ if $\sum_f a_{tf} x_f \geq b_t$ for all $t \in C_i$ and $c_i(p) = \infty$ otherwise. Thus, as before, a player minimizes investment but insists on satisfying her constraints. The notions of Nash equilibrium, social optimum, and (α, β)-NE follow accordingly. Note that in a Nash equilibrium and a social optimum no player contributes to an unbought unit, and the bought units are paid for exactly. Thus, we again have that a Nash equilibrium p and a social optimum p^* correspond to feasible and optimal solutions x and x^* of the CIP instance, respectively. The payments represent cost sharings of these solutions. This again implies that for a Nash equilibrium p and a social optimum p^* we have $\sum_i p_i(f) = c(f) x_f$, and $\sum_i p_i^*(f) = c(f) x_f^*$. Note that this general formulation includes vertex and set cover games as special cases.

2.3 Initial Observations

The following observations can be used to simplify a game. Suppose a constraint is not included in any of the players constraint sets. This constraint has no influence on the game. Hence, in the following w.l.o.g. we will assume that the sets C_i cover all constraints. In particular, for the vertex cover game this means $E = \bigcup_{i=1}^{k} E_i$.

For the vertex cover game we observe a decomposition property. For a player i assume the graph $G[E_i]$ induced by the players edge set E_i is not connected. The player has to cover edges in each component and her optimum strategy decomposes to cover both components independently at minimum cost. Hence, we can form another game in which the edges for each of the k_i components are owned by different subplayer $i_1, \ldots, i_{k_i}$. Suppose we adjust the game like this and derive an (approximate) Nash equilibrium under the condition that the edges of each player form only a single connected component. When we translate this payment scheme back to the original game, the stability ratio can only improve. This property can be translated to covering games, however, it does not seem to have a similar intuitive meaning.

Finally, suppose a constraint t is owned by a player i and a set of players J. Now consider a Nash equilibrium for an adjusted game in which the constraint is owned only by player i. In this equilibrium a player $j \in J$ has no better strategy to satisfy the constraints in $C_j - t$. However, t is satisfied as well, potentially by a different player. If t is added to C_j again, j has no incentive to deviate from her strategy. This is due to the fact that the individual cost of a player can only increase by adding constraints. The Nash equilibrium for the adjusted game yields a Nash equilibrium in the original game. Clearly, the other direction is not true: Consider a game with a single resource f and a single constraint $x_f \geq 1$ owned by two players. A Nash equilibrium is obtained, e.g., if the players share the cost of resource f in equal shares. This can be no Nash equilibrium for the case, in which only one player owns the constraint. Nevertheless, for computation of approximate equilibria we will for convenience use an adjusted game and assume that all constraint sets C_i are mutually disjoint. The above idea can be used to show that the derived bounds on the ratios continue to hold for the original game, in which the sets C_i overlap.

3 Quality and Existence of Nash Equilibria

In this section we consider lower bounds on the quality of Nash equilibria in vertex cover games and on the hardness of determining their existence. In general it is not possible to guarantee their existence, they can be hard to find or expensive. At first, observe that the price of anarchy in the vertex cover game is exactly k, and that this result generalizes to general covering games.

Theorem 1 *The price of anarchy in any covering game is at most k, and there exists a vertex cover game with price of anarchy at least k.*

Proof. Consider a star in which each vertex has cost 1 and each player owns a single edge. The minimum cost cover $\mathcal{C}^*$ is the center vertex of cost 1. If each player purchases the vertex of degree 1 incident to her edge, we get a Nash equilibrium of cost k. Hence, the price of anarchy is at least k. On the other hand, k is a simple upper bound. If there is a Nash equilibrium $\mathcal{C}$ with $c(\mathcal{C}) > kc(\mathcal{C}^*)$, there is at least one player i that pays more than $c(\mathcal{C}^*)$. She could unilaterally improve by purchasing $\mathcal{C}^*$ all by herself. As the argumentation for the upper bound does not use specific properties of the vertex cover game, it continues to hold for all covering games. □

Note that the price of anarchy is k even for very simple games, in which every player owns only one edge and G is a tree. Hence, we will in the following consider existence and quality of the best Nash equilibrium in a game.

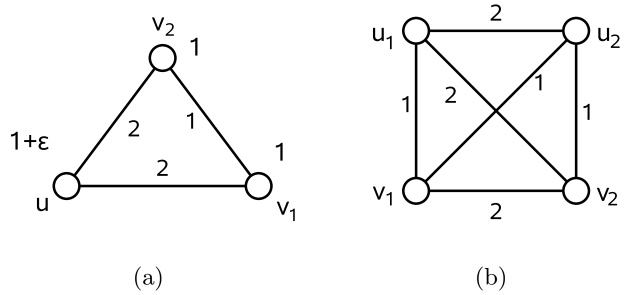

Figure 1: Vertex cover games for two players without Nash equilibria. (a) weighted game; (b) unweighted game. Edge type indicates player ownership. For the weighted game numbers at vertices indicate vertex costs.

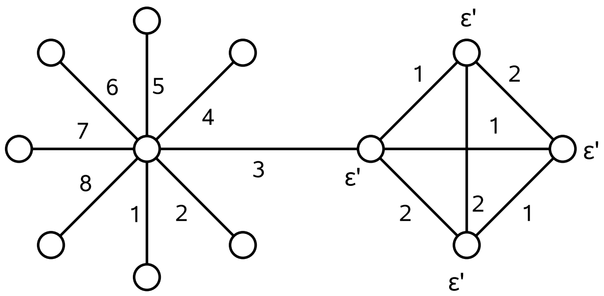

Figure 2: A game with k=8, for which the cost of any Nash equilibrium is close to $(k-1)c(\mathcal{C}^*)$. Numbering of edges indicates player ownership. Indicated vertices have cost $\varepsilon' \ll 1$, vertices without labels have cost 1.

Lemma 1 *There are planar vertex cover games for two players without Nash equilibria.*

Proof. We consider the game for two players in Figure 1(a) for an $\varepsilon > 0$. For this game we consider the four possible covers. A cover including all three vertices cannot be an equilibrium, because vertex u is not needed by any player to fulfill the covering requirement. Hence, any player contributing to the cost of u could feasibly improve by removing these payments. Suppose the cover representing an equilibrium includes v_1 and v_2. If player 1 contributes to v_1, she can remove these payments, because she only needs v_2 to cover her edge. With the symmetric statement for v_2 we can see that in equilibrium player 1 could not pay anything. Player 2, however, cannot purchase both v_1 and v_2, because buying u offers a cheaper alternative to cover her edges. Finally, suppose u and v_1 are in the cover. In equilibrium player 1 will not pay anything for u. Player 2, however, cannot purchase u completely, because v_2 offers a cheaper alternative to cover the edge (u, v_2). With the symmetric observation for the cover of u and v_2, we see that there is no feasible cover that can be purchased in a Nash equilibrium. With similar arguments we can prove that the game on K_4 depicted in Figure 1(b) has no Nash equilibria. This proves the lemma. □

Note that every game with less edges, vertices, or players than the game in Figure 1(a) is guaranteed to have a Nash equilibrium.

Theorem 2 *For any $\varepsilon > 0$ there is a weighted vertex cover game in which the price of stability is at least $(k-1)-\varepsilon$. There is an unweighted vertex cover game in which the price of stability is $\frac{k+2}{4}$.*

Proof. Consider a game as depicted in Figure 2. The social optimum cover includes the center vertex of the star and three vertices of the K_4-gadget yielding a total cost of $1+3\varepsilon'$. If the center vertex of the star is in the cover and we assume to have a Nash equilibrium, no player other than players 1 and 2 can contribute anything to vertices of the K_4-gadget incident to edges of player 1 and 2. For this network structure, however, it is easy to note that players 1 and 2 cannot agree on a set of vertices covering their edges. Hence, to allow for a Nash equilibrium, the star center must not be picked which in turn requires all other adjacent star vertices to be in the cover. Under these conditions the best feasible cover includes the vertex that connects K_4 to the star yielding a cost of $k-1+3\varepsilon'$. Note that we can derive a Nash equilibrium purchasing this cover by assigning each player to purchase a star vertex - including the vertex that also belongs to K_4. Players 1 and 2 are assigned to purchase one of the additional K_4 vertices, respectively. With $\varepsilon = \frac{3\varepsilon'(k-2)}{1+3\varepsilon'}$ the first part of the theorem follows. For the unweighted case we simply consider the game graph with all vertex costs equal to 1.

A similar analysis delivers the stated bound and proves the second part of the theorem, where we insist on the assumption that $G[E_i]$ is connected for each player. Here we drop the edges incident to the star center from E_1 and E_2, which results in connected subgraphs for these players. Then we introduce two new players, and each one owns one of the dropped edges. This results in every $G[E_i]$ being connected. In addition, the Nash equilibria remain equivalent and the numerical value of the price of stability is the same before and after the transformation (c.f. Section 2.3). However, the number of players k has increased by 2, and thus the expression now reads $k-3-\varepsilon$ and $k/4$. □

Theorem 3 *It is* NP*-hard to determine if*

1. *an unweighted vertex cover game*
2. *a weighted vertex cover game for 2 players*

has a pure strategy Nash equilibrium, even if the graphs $G[E_i]$ are forests.

Proof. We describe a reduction from 3SAT [25] for weighted games and then show how to adjust it for unweighted games. Given an instance of 3SAT for every variable we introduce a gadget with a *decision player*. This player owns the edges of two stars, denoted as the *true* star and the *false* star. The number of leaves of the (false) true star is equal to the number of (negated) nonnegated occurrences of the variable in the clauses of the instance. The cost of each center vertex is equal to the number of leaves, the cost of the leaves is equal to 1. In addition, we include a direct connection between the centers of the stars. Furthermore, we use extended triangle gadgets depicted in Figure 3. This gadget represents a game without a Nash equilibrium, which can be shown along the lines of the proof of Lemma 1. At each leaf vertex of the stars we install an extended triangle gadget. As no pair of gadgets is directly connected, this introduces only two new *triangle players*. The leaf vertices of the stars become the u_1-vertices of the gadgets. An example variable gadget is depicted in Figure 4(a). For each clause we introduce a new *clause player*. She owns a star of three edges connecting a new center vertex of cost 1 to three extended triangle gadgets. We let the edges

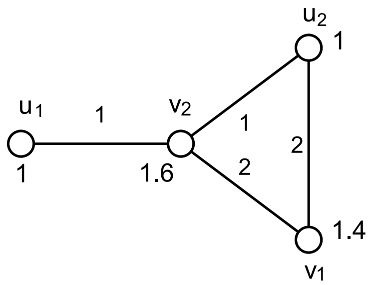

Figure 3: Extended triangle used in the proof of Theorem 3. This game does not have any Nash equilibrium. To stabilize the extended triangle in the gadgets described below, a third player can either buy u_1 completely, then the first two players can pay for u_2 and v_1; or reduce the cost of v_1 of 0.5, allowing the first two players to pay 0.9 for v_1 and 1.6 for v_2.

connect to triangles of the false or true star of a variable gadget depending on whether the variable appears negated or non-negated in the clause, respectively. In particular, the edges connect to the v_1-vertices of the extended triangles. As we have installed a sufficient number of these gadgets, we construct the network such that no two edges of different clause players are incident at the same vertex. An example of a clause gadget is depicted in Figure 4(b).

Suppose there is a satisfying assignment for the instance of 3SAT. Then we construct an equilibrium as follows. If a variable is true in the assignment, we pick the center vertex of her false star and all leaf vertices of the true star of its gadget to be in the cover and let the decision player pay for it. All extended triangles incident to the false star then allow a stable cost distribution, in which u_2 and v_1 are bought by the triangle players (see Figure 3). In the case a variable is false we pick the leaf vertices of the false star and adjust the assigned payments accordingly. As we have a satisfying assignment for the 3SAT instance, this stabilizes at least one triangle gadget per clause. So each clause player has the chance to reduce the cost of the vertices of the remaining two incident triangles by 0.5 each. The triangle players can then purchase the v_1 and v_2 vertices in the remaining unstabilized gadgets (see Figure 3). This assignment leaves no player an incentive to defect and forms a Nash equilibrium.

Now suppose there exists a Nash equilibrium. In equilibrium a decision player can either purchase one or both of the star centers. Once she purchases the center of a star, she is not willing to contribute anything to the leaf vertices of the star. Hence, if she does not contribute to the extended triangles attached to a star, the clause players must help the triangle players agree upon a cover. However, a clause player can only contribute a total cost of 1 to the triangle vertices, because otherwise she can pick her star center as a cheaper alternative. The minimum cost reduction that she can achieve at every v_1 vertex of her incident triangle gadgets is 1/3, which is not enough to allow for a stable cost assignment in all triangles. Hence, we need to have the decision players purchase the stars such that they trigger a stable cover in at least one extended triangle from each clause. Furthermore, as they can only trigger a stable cover in triangle gadgets attached to one of their stars, this naturally translates to a satisfying assignment for the 3SAT instance.

Finally, we use the transformations mentioned in Section 2 to obtain an equivalent game by merging all decision players into one player and all clause players into another player. Note that

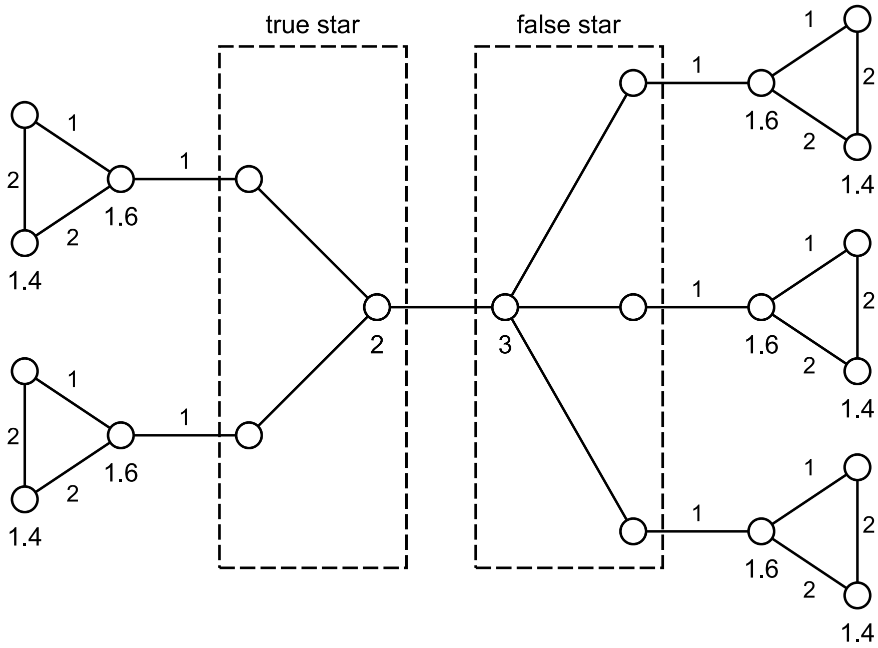

(a) A gadget for a variable occurring nonnegated in two clauses and negated in three clauses.

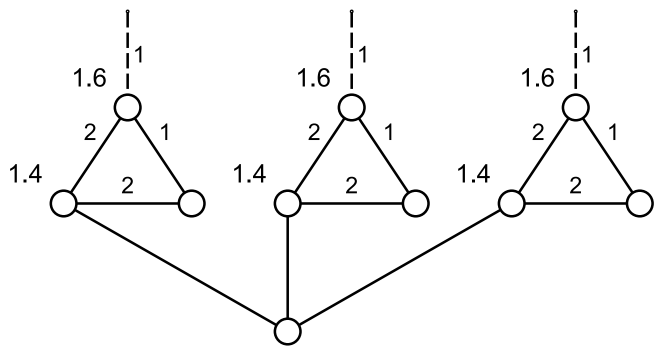

(b) A gadget for a clause.

Figure 4: Variable and clause gadgets. The edges owned by triangle players are numbered, while edges owned by decision and clause player are unlabeled. All vertices have cost 1, except for the two vertices in the extended triangles and the star centers.

we have introduced only two triangle players, whose edges form a partition of all edges from the extended triangle gadgets. The class of decision players shares endpoints only with one of the triangle players. The same is true for clause players and the other triangle player. Hence, we can merge the players again forming an equivalent game with only two players. This proves NP-hardness for weighted games and two players, even in the case where the graphs induced by the set of edges of each player are forests.

For unweighted games we replace the extended triangles by the games on K_4 depicted in Figure 1(b). Vertices labeled u_1 and v_1 indicate where to connect the decision and clause player stars, respectively. In the variable gadgets the star centers have cost 1. In addition, we introduce a number of new players such that edges of the true and false stars are each owned by a different player. For a clause player we now install two stars instead of one star. The stars have different centers, but leaf nodes from the same K_4 gadgets. Observe that in every variable gadget players in equilibrium contribute only to the leaves of at most one star. Furthermore, a clause player must invest at least a cost of 1 to stabilize a K_4 gadget. Hence, if at least one gadget per clause is stabilized by the decision players, there exists a Nash equilibrium. On the other hand, this condition is also necessary, because the centers of the clause stars allows the clause players to stabilize at most two K_4 gadgets. This proves NP-hardness.

To show that the result holds even when the graphs $G[E_i]$ are forests, we can remark that the two stars of a clause player can be shared among two distinct players and the above reasoning is still correct. It can be checked that all graphs $G[E_i]$ in this game are forests. Note that in this case it can be checked in polynomial time whether a payment scheme is a Nash equilibrium, because the minimum vertex cover problem can be solved in polynomial time on trees. Hence the problem of finding whether a Nash equilibrium exists, when restricted to these instances, is also in NP and so is NP-complete. This proves Theorem 3. □

4 Games with cheap Nash Equilibria

The results in the previous section reveal that in general pure Nash equilibria can be absent, even from very simple variants of the vertex cover game. In this section we outline two classes of set cover games that have Nash equilibria which are even social optima.

We present two classes of set cover games that always exhibit optimal Nash equilibria: integral games, in which the integrality gap of the integer programming relaxation is 1, and singleton games, in which each player owns only a single edge. For singleton games it is NP-hard to find a social optimum. Hence, we study existence and algorithmic computation of (α, β)-NE. Formally, for (α, β)-NE the stability ratio $\alpha \geq 1$ specifies the violation of the Nash equilibrium inequality, and $\beta \geq 1$ is the approximation ratio of the social cost. We show that for singleton set cover games $(1 + \varepsilon, O(\log n))$-NE can be obtained in polynomial time, for any constant $\varepsilon > 0$.

4.1 Integral Games

In this subsection we turn to integral games, in which the integrality gap is 1. The *integrality gap* of a linear program for a minimization problem is the worst-case ratio of the cost of the best integral solution over the cost of the best fractional solution. Thus, in integral games even the linear relaxation of the underlying integer program has an optimal integral solution.

Theorem 4 *For any set cover game, in which the integrality gap of the underlying CIP is 1, the price of stability is 1 and an optimal Nash equilibrium can be found in polynomial time.*

Proof. We define the payments for the players using the optimal values of the dual variables. Then the total contribution of the players suffices to pay exactly for the solution. Basically, the proof is an application of standard properties of LP-duality. In the following we provide a detailed presentation of the specific steps.

Consider the LP-relaxation of the CIP derived by setting $x_S \geq 0$ instead of $x_S \in \mathbb{N}$. For the set cover game, primal and dual of the LP-relaxation can be formulated as

$$\begin{array}{llll} \text{Min} & \sum_{S \in \mathcal{S}} c(S) x_S & & \\ \text{s.t.} & \sum_{S: e \in S} x_S \geq 1 & \forall e \in E \\ & x_S \geq 0 & \forall S \in \mathcal{S} \end{array} \qquad \begin{array}{llll} \text{Max} & \sum_{e \in E} y_e & \\ \text{s.t.} & \sum_{e \in S} y_e \leq c(S) & \forall S \in \mathcal{S} \\ & y_e \geq 0 & \forall e \in E. \end{array} \tag{3}$$

We can find the optimum primal solution x^* and the optimum dual solution y^* in polynomial time. Note that x^* yields a feasible cover $\mathcal{C}^*$ by setting $S \in \mathcal{C}^*$ if and only if $x_S^* = 1$, because all $x_S^* \in \{0, 1\}$ due to an integrality gap of 1. Both x^* for the primal and y^* for the dual have the same value by strong LP duality. Now assign each player to pay $p_i(S) = \sum_{e \in S \cap E_i} y_e^* x_S^*$. The resulting state p is a Nash equilibrium if every player plays a best response. We first show a necessary condition for this, i.e., that $\mathcal{C}^*$ is exactly paid for.

Lemma 2 *In the state p defined above the players exactly purchase the optimum cover. The contribution of player i in p is exactly $\sum_{S \in \mathcal{S}} p_i(S) = \sum_{e \in E_i} y_e^*$.*

Proof. At first, note that no player contributes to a set $S \notin \mathcal{C}^*$, because $x_S^* = 0$. For sets $S \in \mathcal{C}^*$ we use complementary slackness of optimal solutions to primal and dual linear programs, which yields $x_S^* \cdot (c(S) - \sum_{e \in S} y_e^*) = 0$ for any $S \in \mathcal{S}$. If $S \in \mathcal{C}^*$, then $x_S^* > 0$ and thus $\sum_{e \in S} y_e^* = c(S)$. Hence, for such a set the total contribution is

$$\sum_{i=1}^{k} p_i(S) = \sum_{i=1}^{k} \sum_{e \in S \cap E_i} y_e^* x_S^* = \sum_{i=1}^{k} \sum_{e \in S \cap E_i} y_e^* = \sum_{e \in S} y_e^* = c(S),$$

because w.l.o.g. we can assume that each element needs to be covered by exactly one player i. This shows that in p all sets in $\mathcal{C}^*$ get exactly paid for and no player contributes to sets $S \notin \mathcal{C}^*$. This proves the first part of the lemma.

For the second part of the lemma we again use complementary slackness, i.e., that $y_e^* \sum_{S \in \mathcal{S}} x_S^* = 0$. For every overcovered element e with $\sum_{S \in \mathcal{S}} x_S^* > 1$ this shows that $y_e^* = 0$. Thus, the total contribution of player i in p_i is

$$\sum_{S \in \mathcal{S}} p_i(S) = \sum_{S \in \mathcal{S}} x_S^* \sum_{e \in S \cap E_i} y_e^* = \sum_{e \in E_i} y_e^*.$$

This proves the second part of the lemma. □

Now we prove that the strategy assignment p_i is a best response against p_{-i} for every player i.

Lemma 3 *In the state p defined above every player plays a best response.*

Proof. Lemma 2 implies that no player has a prohibitively large individual cost, because for each player a collection of sets covering her elements is paid for. Thus, the individual cost of each player i is $c_i(p) = \sum_S p_i(S)$. To show the lemma we consider for each player i the optimization problem of finding a strategy that minimizes the function $c_i(\cdot, p_{-i})$. An optimum solution p_i^* to this problem is obviously a best response for player i. We will show that p_i is such an optimum solution. This proves that p is a Nash equilibrium.

For player i the optimization of $c_i(\cdot, p_{-i})$ has to consider the contributions of other players as fixed. In particular, player i only has to contribute exactly $c'(S) = c(S) - \sum_{j \neq i} p_j(S)$ to make set S become bought. Thus, in order to find a best response, player i faces a *reduced set cover problem*, in which she must cover elements E_i with sets $\mathcal{S}$ of cost $c'(S)$. Using Lemma 2 we note that $c'(S) \geq 0$ for every $S \in \mathcal{S}$. An optimum solution $\mathcal{C}'$ to the reduced problem yields a best response $p_i^*(S) = c'(S)$ if $S \in \mathcal{C}'$ and $p_i^*(S) = 0$ otherwise. Consider the linear relaxation of the reduced problem and the dual:

$$\begin{array}{llll} \text{Min} & \sum_{S \in \mathcal{S}} c'(S) x'_S & & \\ \text{s.t.} & \sum_{S : e \in S} x'_S \geq 1 & \forall e \in E_i \\ & x'_S \geq 0 & \forall S \in \mathcal{S} \end{array} \qquad \begin{array}{llll} \text{Max} & \sum_{e \in E_i} y'_e & \\ \text{s.t.} & \sum_{e \in S \cap E_i} y'_e \leq c'(S) & \forall S \in \mathcal{S} \\ & y'_e \geq 0 & \forall e \in E_i. \end{array} \tag{4}$$

Note that setting $x'_S = x^*_S$ is feasible, because all elements in E_i were covered using x^*_S. In addition, this solution results in the original payments $p_i(S)$ and cost $c_i(p) = \sum_{S \in \mathcal{S}} p_i(S)$. Now consider setting $y'_e = y^*_e$ in the dual. This is feasible for the dual in (4), because

$$\begin{aligned} c'(S) &= c(S) - \sum_{j \neq i} p_j(S) = c(S) - \sum_{j \neq i} \sum_{e \in S \cap E_j} x^*_S y^*_e = c(S) - \sum_{e \in S - E_i} x^*_S y^*_e \\ &\geq c(S) - \sum_{e \in S - E_i} y^*_e \geq \sum_{e \in S \cap E_i} y^*_e, \end{aligned}$$

and because y^*_e was feasible for the original dual. In addition, it results in an objective function value of $\sum_{e \in E_i} y^*_e$ for the dual. Due to LP-duality $\sum_{e \in E_i} y^*_e \leq \sum_{S \in \mathcal{S}} c'(S) x'_S$ for every feasible x'_S in the reduced primal. Lemma 2 and the observations above imply equality for $x' = x^*$. Hence, the solutions x^* and y^* have the same objective function value for the primal and dual reduced problems. By strong duality this implies that $x'_S = x^*_S$ is an integral optimal solution for the primal LP-relaxation in (4). Hence, the cover $\mathcal{C}^*$ is an optimum solution to the reduced problem. It results in the original payments $p_i(S)$, which therefore are a best response for player i. This proves the lemma. □

The previous two lemmas show that it is possible to obtain a Nash equilibrium p, which exactly pays for the cost of an optimum solution. This proves the theorem. □

For illustration of the arguments consider a vertex cover game on a bipartite graph G, which is known to have an integrality gap of 1, see Figure 5(a). To obtain an optimum dual solution one can employ a flow network using a standard construction by adding a source and a sink vertex, see Figure 5(b). Each of these two vertices is then connected by directed edges to all vertices from one partition of the graph. The additional edges are directed away from the source to the sink and

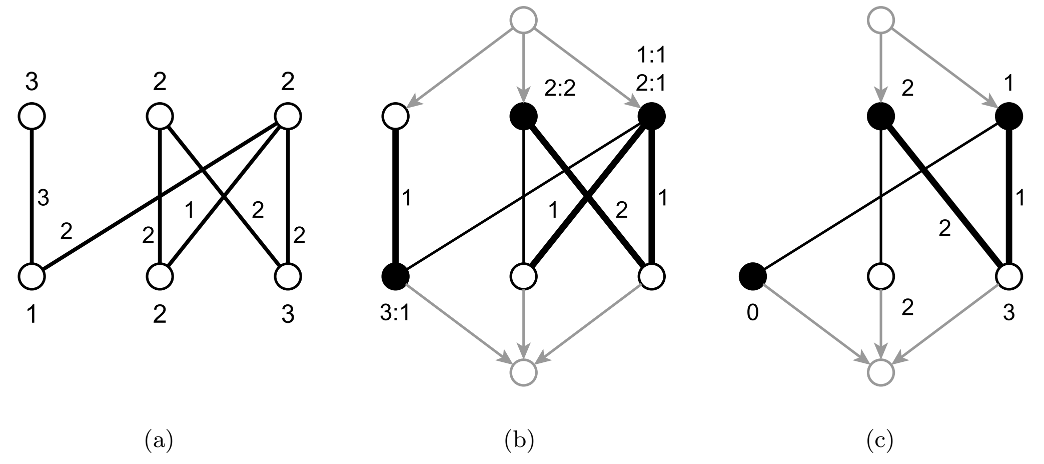

Figure 5: (a) A vertex cover game with an optimal Nash equilibrium. Edge numbering indicates player ownership, vertex numbering indicates cost. (b) Dual variables correspond to maximum flow values in an extended graph. It yields a cost sharing of an optimum cover indicated by filled vertices. Edge numbering represents top-down flow, labels of a vertex v are $i : p_i(v)$. (c) Player 2 plays a best response, because in the reduced problem the flow over her edges stays feasible. Edge numbering represents top-down flow, vertex numbering indicates cost, filled vertices constitute an optimum cover.

receive as capacity the cost of the incident vertex from G. All edges from G receive infinite capacity. A maximum flow in this network yields an optimum solution to the corresponding LP-dual of the vertex cover problem in G. The dual variables correspond to the flow values on the edges. They can be used to construct a cost sharing of an optimum vertex cover. To show that this represents a Nash equilibrium, Figure 5(c) illustrates the reduced problem of finding a best response for player 2. For the reduced problem the flow over her edges is still feasible, hence by LP-duality it lower bounds the optimum cover cost. As the flow also yields a feasible strategy, player 2 plays a best response.

4.2 Singleton Games

For singleton set cover games, in which $|E_i| = 1$ for all players i, we prove a similar result as for integral games. The proof, however, is along different lines. It does not immediately yield an efficient algorithm to compute an optimal Nash equilibrium, however, it allows us to obtain $(1 + \varepsilon, O(\log n))$-NE in time polynomial in n, k and ε^{-1}, for any constant $\varepsilon > 0$.

Note that one can prove these results in a quite simple way. We can transform a singleton set cover game into an equivalent single source connection game in a directed graph. The construction has three layers. In the first layer there is exactly one terminal node for each player. In the second layer there is a node for each set. There is an edge between a player node and a set node if and only if the element of the player is contained in this set. These edges all have cost 0. Finally, the third layer is the source node, and a node for set S is connected to the source node with an edge of cost $c(S)$. It is obvious that solutions and Nash equilibria in this connection game correspond to solutions and Nash equilibria in the singleton set cover game and vice versa. Hence, using results of [6] it follows that optimal Nash equilibria exist and that starting from an r-approximate solution we can construct a sequence of polynomially many steps improving the social cost until reaching a $(1+\varepsilon, r)$-NE. The results are stated in Theorem 5 and Theorem 6. For completeness, we also provide self-contained proofs, which apply elementary arguments along the lines of [6] directly to singleton set cover games and highlight some of the important main ideas.

Theorem 5 *For singleton set cover games the price of stability is 1.*

Proof. Consider a cover $\mathcal{C}$ and a set $S \in \mathcal{C}$, and remove S from $\mathcal{C}$. Players that are still covered can be assigned no contribution to S. Now consider the set of remaining players I_S. For each player $i \in I_S$ independently consider the case, in which S is removed, and consider her cheapest strategy to cover her element by a set different from S. Player i must purchase completely this new set. We denote the cost of this set by c_i^S. A set is called *stabilized* if $c(S) \leq \sum_{i \in I_S} c_i^S$. For a stabilized set S we assign the players to pay $p_i(S) = \frac{c(S)c_i^S}{\sum_{i \in I_S} c_i^S}$. This will not yield any incentives to remove payments.

A cover $\mathcal{C}$ is called *exchange minimal* if all sets are stabilized. Observe that an exchange minimal cover can be purchased by a Nash equilibrium specified by the above given assignment of payments. If $\mathcal{C}$ is not exchange minimal, it is possible to exchange an unstabilized set S with all the alternative sets used to define c_i^S for the players $i \in I_S$. The total cost of this alternative collection is at most $\sum_{i \in I_s} c_i^S$. Such a local exchange step results in a new, strictly cheaper cover. Thus, if $\mathcal{C} = \mathcal{C}^*$, then this is not possible. Hence, in $\mathcal{C}^*$ all sets are stabilized and there is a Nash equilibrium as cheap as $\mathcal{C}^*$. □

The exchange step mentioned in the proof suggests that with an iterated application cover and payments converge to a Nash equilibrium, which is cheaper than the starting cover. The problem with this approach is the running time, which is not guaranteed to be polynomial in the input size. Here we borrow a trick from [6]. For computing exchanges we artificially reduce the cost of currently bought sets by a value κ such that each exchange operation of the proposed algorithm guarantees a minimum improvement of the overall cost. For computing a starting solution any algorithm can be used, in particular, we can apply the greedy $O(\log n)$-approximation algorithm [44]. Using this algorithm we can compute $(1+\varepsilon, O(\log n))$-NE in time polynomial in n, k and ε^{-1}, for any constant $\varepsilon > 0$. Note that the greedy algorithm alone does not necessarily output an exchange minimal cover.[1]

Theorem 6 *For singelton set cover games a $(1+\varepsilon, O(\log n))$-NE can be computed in polynomial time, for any constant $\varepsilon > 0$.*

Proof. Consider the output $\mathcal{C}$ of any β-approximation algorithm for set cover. Starting from this cover we consider an unstabilized set S and create a new cover $\mathcal{C}'$, in which we remove S. Then we let all the players with uncovered elements simultaneously switch to their best responses and add the corresponding sets to the cover. Let $\varepsilon > 0$ and define $\kappa = \frac{\varepsilon' c(\mathcal{C})}{(1+\varepsilon')n\beta}$. We perform such an exchange step only if it improves the cost of the cover by at least κ. The algorithm terminates when no further improving exchange steps are possible. Hence, in total the algorithm makes at most $\frac{(1+\varepsilon)n'\beta}{\varepsilon}$ such exchange steps.

[1] Consider a game with 3 players and 3 sets: $S_1 = \{e_1, e_2\}$, $S_2 = \{e_2, e_3\}$, and $S_3 = \{e_1\}$, where player i owns element e_i, for $i = 1, 2, 3$. The costs are $c(S_1) = 2$ and $c(S_2) = c(S_3) = 1.5$. The greedy algorithm first picks S_1 and then S_2. However, e_2 is overcovered, and player 1 can purchase S_3. So $c_1^{S_1} = 1.5 < 2$, and S_1 is not stabilized.

If the iteration has terminated, we assign players to purchase the cost of the cover as described in the proof of Theorem 5. Note, however, that for each unit a cost of at most κ remains unpaid. Suppose in the final cover $\mathcal{C}'$ there are n' bought sets. This creates a remaining cost of at most $\kappa n'$ to be paid for. Now each player is assigned a proportional share of $\frac{p_i(\mathcal{C}')}{c(\mathcal{C}')-n'\kappa}$ of all remaining costs, which might require a player to pay at sets not containing her elements. To establish the stability ratio of $(1+\varepsilon')$ we note that the increase for player i is only

$$p_i(\mathcal{C}')\frac{\kappa n'}{c(\mathcal{C}')-\kappa} \leq \frac{\varepsilon' c(\mathcal{C})p_i(\mathcal{C}')}{\alpha(1+\varepsilon')(1-\varepsilon)c(\mathcal{C}')} \leq \varepsilon' p_i(\mathcal{C}'),$$

Hence, even if all additional payments of player i contribute to resource units she does not need for coverage, she cannot decrease her payments by more than a factor of $(1+\varepsilon)$. This yields the result. □

A possible extension of this result is the case in which each player i has a threshold M_i and would rather stay uncovered if her assigned payments exceed M_i. The presented algorithm finds a Nash equilibrium with stability ratio arbitrarily close to 1 for this case as well. Note that we can further generalize the result to games based on set multicover. It is, however, an interesting open problem to adjust this procedure to cope with general covering games with general a_{tf} and b_t.

5 Approximate Equilibria

In the previous sections we saw that exact Nash equilibria exist only in special cases, in general they can be absent or highly inefficient. In this section we focus on (α,β)-NE for general vertex and set cover games. For set cover games we show that (f,f)-NE can be computed in polynomial time. For vertex cover games we give an algorithm to divide the cost of any optimum solution into payments representing a $(2,1)$-NE. In addition, we prove a tight lower bound. We show that for any given $\alpha < 2$ and $\beta \geq 1$, there is a vertex cover game without any (α,β)-NE.

Algorithm 1: (f,f)-NE for set cover games

$p_i(S) \leftarrow 0$ for all players i and sets S
$\gamma_i(e) \leftarrow 0$ for all players i and elements e
while *there is an uncovered element* e **do**
    Let i be the player owning element e, and let $\gamma_i(e) \leftarrow \min_{S\ni e} c(S)$
    Increase payments: $p_i(S) \leftarrow p_i(S) + \gamma_i(e)$ for all S with $e \in S$
    Reduce set costs: $c(S) \leftarrow c(S) - \gamma_i(e)$ for all S with $e \in S$
    Add all purchased sets to the cover

Theorem 7 *Algorithm 1 returns an* (f,f)*-NE for set cover games in polynomial time.*

Proof. This is a variant of the well-known primal-dual f-approximation algorithm for minimum set cover for all elements in E (see for instance [44, chapter 15]). It can clearly be implemented to run in polynomial time. It accounts contributions to elements. For each element e it simultaneously raises the contribution $\gamma_i(e)$ to all covering sets until one set is bought. A player is assigned to pay all contributions of her elements. It remains to be shown that the stability ratio is equal to f as well.

Intuitively, for an element e of player i we simultaneously raise the contribution $\gamma_i(e)$ to each set S it can be covered with. When player i deviates to a best response, she has to purchase the (remaining share of the) cost of at least one set that contains e, and thus she has to pay $\gamma_i(e)$ at least once. This yields a stability ratio of at most f.

More formally, after the execution of the algorithm, we consider the ith player and her best move taking into account the payments of all other players $j \neq i$. We define the remaining costs $c'(S)$ for each set S, by letting $c'(S) = c(S) - \sum_{j\neq i} p_j(S)$. We have to show that the sum of the payments of player i is not greater than f times the cost of the cheapest set cover of E_i with respect to the costs c'.

At first, we bound the payments of player i. The algorithm does not raise payments above costs, i.e, for any set S we have $\sum_{j=1,2,\ldots,k} p_j(S) \leq c(S)$. In particular, player i does not overpay the cost $c'(S)$:

$$p_i(S) \leq c(S) - \sum_{j\neq i} p_j(S) = c'(S).$$

The algorithm simultaneously raises contributions to all sets that contain an element e. The total contribution of player i to a set S equals the $\sum_{e\in S} \gamma_i(e)$, so $p_i(S) = \sum_{e\in S\cap E_i} \gamma_i(e)$. Again, as no set is overpaid, this contribution is at most $c'(S)$

$$\sum_{e\in S\cap E_i} \gamma_i(e) \leq c'(S).$$

Now we turn to lower bounding the cost of a best response set cover $\mathcal{C}_i^*$ of E_i, which has minimum cost with respect to the cost function c'. With the previous arguments, the cost of this cover is at least

$$c'(\mathcal{C}_i^*) = \sum_{S\in\mathcal{C}_i^*} c'(S) \geq \sum_{S\in\mathcal{C}_i^*} \sum_{e\in S\cap E_i} \gamma_i(e).$$

Since $\mathcal{C}_i^*$ is a set cover of E_i, the charge $\gamma_i(e)$ of each element e in E_i is counted at least once in the right-hand side above. Hence, the left-hand side is at least $\sum_{e\in E_i} \gamma_i(e)$, which allows to conclude

$$\sum_{S\in\mathcal{S}} p_i(S) = \sum_{S\in\mathcal{S}} \sum_{e\in S\cap E_i} \gamma_i(e) \leq \mathsf{f} \cdot \sum_{e\in E_i} \gamma_i(e) \leq \mathsf{f} \cdot c'(\mathcal{C}_i^*).$$

This proves that for any best response player i has to pay for all $e \in E_i$ the share $\gamma_i(e)$, and the theorem follows. □

We remark here that our arguments are also implicitly used in [33], which considers dual payments and core solutions in cooperative games. We could as well employ their results to show that the algorithm has stability ratio f. A proof along these lines is given for the primal-dual Algorithm 3 for the UFL game below in Section 6.1.

For the case of vertex cover we can show that any socially optimum cover $\mathcal{C}^*$ can always be purchased by a $(2,1)$-NE.

Theorem 8 *For every vertex cover game there is a* $(2,1)$*-NE.*

Algorithm 2: $(2,1)$-NE for vertex cover games

$p_i(v) \leftarrow 0$ for all players i and vertices $v \in V$
$\gamma_i(e) \leftarrow 0$ for all players i and edges $e \in E$
$E^2 \leftarrow$ edges e covered twice in $\mathcal{C}^*$
Set all edges unmarked
while *there is an unmarked edge* $e \in E^2$ **do**
  Let i be the player owning edge $e = (u, v)$, and let $\gamma_i(e) \leftarrow \min\{c(u), c(v)\}$
  Increase payments: $p_i(u) \leftarrow p_i(u) + \gamma_i(e)$, $p_i(v) \leftarrow p_i(v) + \gamma_i(e)$
  Reduce vertex costs: $c(u) \leftarrow c(u) - \gamma_i(e)$, $c(v) \leftarrow c(v) - \gamma_i(e)$
  Mark all purchased vertices and their incident edges
Create $G' = (V', E')$ and $\mathcal{C}'$ by removing all marked vertices and edges from G and $\mathcal{C}^*$
Find optimum y^* for dual LP of G' that corresponds to $\mathcal{C}'$ for primal LP (3)
for *all edges* $e \in E'$ **do**
  Let i be the player owning edge $e = (u, v)$, and let $\gamma_i(e) \leftarrow y^*_e$
  Increase payments: $p_i(u) \leftarrow p_i(u) + \gamma_i(e)$, $p_i(v) \leftarrow p_i(v) + \gamma_i(e)$
$p_i(v) \leftarrow 0$ for all $v \notin \mathcal{C}^*$

Proof. Consider a vertex cover game on a graph G with vertex costs $c(v)$ and an arbitrary optimal solution $\mathcal{C}^*$. We use Algorithm 2 to obtain a cost sharing of $\mathcal{C}^*$ that is a $(2,1)$-NE. It proceeds in two phases. The first phase is a run of Algorithm 1 on a restricted set of edges. The second phase uses insights from Theorem 4 to assign the remaining costs. The key observation is that after the first phase, the remaining graph G' is bipartite, and $\mathcal{C}'$ is a minimum weight vertex cover for G'.

Lemma 4 *The graph G' is bipartite. $\mathcal{C}'$ is a minimum weight vertex cover for G'.*

Proof. In the first phase, the algorithm considers the set of edges E^2 for which both incident vertices are in $\mathcal{C}^*$. Note that every odd cycle of G has at least one edge in E^2. Thus, if E^2 is empty, then G must be bipartite, and the lemma follows. Otherwise, Algorithm 1 is called with E^2, which assigns budgets $\gamma_i(e)$ to edges $e \in E^2$ and reduces the vertex costs accordingly. In each step, the cost function c is adjusted, the costs of both incident vertices of e are reduced by $\gamma_i(e)$. As both these vertices are in $\mathcal{C}^*$, it remains an optimal cover after the adjustment. After considering all edges of E^2, the algorithm removes all purchased vertices with cost $c(v) = 0$ from $\mathcal{C}^*$ and G, along with all incident edges. This leaves $\mathcal{C}'$ optimal for the resulting graph G' and proves the second part of the lemma.

Due to the algorithm, all edges of E^2 are marked and removed. This adjustment breaks every odd cycle of G. Hence, in $\mathcal{C}'$ every vertex $v \in V'$ has frequency 1, and G' is bipartite. □

For the remaining game given by G' we show in Theorem 4 that there is an exact Nash equilibrium, in which the remaining cost of $\mathcal{C}'$ is exactly paid for. It can be obtained by assigning values of an optimal solution y^* of the LP-dual as cost shares to the edges. Instead, we create a budget for the edges $e \in E'$ with $\gamma_i(e) \leftarrow y^*_e$, which is offered to *both* incident vertices. The next lemma shows that with this adjustment players will not overpay any cost $c(v)$ and will contribute $c(v)$ for every vertex $v \in \mathcal{C}^*$.

Lemma 5 *It holds that*

$$\sum_{i=1}^{k} \sum_{e:e=(u,v)} \gamma_i(e) \leq c(v) \text{ for all } v \in V \text{ and}$$
$$\sum_{i=1}^{k} \sum_{e:e=(u,v)} \gamma_i(e) \geq c(v) \text{ for all } v \in \mathcal{C}^*.$$

Proof. Note that after the first phase, the first part of the lemma clearly holds. In the second phase, player i contributes the dual payment $p_i(e) = \gamma_i(e) = y_e^*$ for $e \in E'$ to both end vertices. Feasibility of y^* for the dual (3) means that $\sum_{e:e=(u,v)} y_e^* \leq c(v)$ for every vertex v, thus proving the first part of the lemma.

For the second part of the lemma, a subset of vertices of $\mathcal{C}^*$ is bought directly after the first phase. For the remaining vertices and the remaining costs we use Theorem 4. It shows that the primal solution x^* corresponding to $\mathcal{C}'$ under the remaining costs is completely paid for by the dual cost shares y^*. □

This property allows to use the proof idea of Theorem 7 to show that the stability ratio is bounded by 2. By dropping all payments to vertices $v \notin \mathcal{C}^*$ in the last step of Algorithm 2, we obtain a cost sharing of $\mathcal{C}^*$ with stability ratio at most 2. This dropping step is not necessary for the bound on the stability ratio, it serves only to obtain a payment scheme that exactly purchases $\mathcal{C}^*$. This proves the theorem. □

For lower bounds on the ratios we note that any algorithm to find a (α, β)-NE in the set cover game can be used as an approximation algorithm for the set cover problem with approximation ratio $\min(\alpha, \beta)$. This follows simply by considering a game with one player. This observation can be combined with recent results on the complexity status of the set cover problem. In particular, set cover cannot be approximated in polynomial time to a factor of $o(\log n)$ unless $\mathsf{P} = \mathsf{NP}$ [2]. Thus, a polynomial time algorithm for $(O(\log n), O(\log n))$-NE is all we can hope for. For the special case of the minimum weighted vertex cover problem a recent result [34] suggests that if $\mathsf{P} \neq \mathsf{NP}$ and the unique games conjecture holds, there is no polynomial time algorithm to approximate vertex cover to a factor of $2 - \varepsilon$. Thus, in this case our algorithm delivers the best factors we can hope for in polynomial time. Note that both bounds apply only to polynomial time computability. We now show that in vertex cover games the frequency $\mathsf{f} = 2$ is also a lower bound for the stability ratio, in a much stronger sense.

Theorem 9 *For any $\alpha < 2$ there is an unweighted vertex cover game without (α, β)-NE for any $\beta \geq 1$.*

Proof. The proof follows with a game on K_{4g} with $g \in \mathbb{N}$. We assume the vertices are numbered v_1 to v_{4g} and distribute the edges of the game to $2g^2 + g$ players in $g + 1$ classes as follows. In the first class there are $2g$ players. Every player i from this class owns only single edge (v_i, v_{2g+i}). Then, for each integer $j \in [1, g-1]$ there is another class of $2g$ players. A player i in one of the classes owns a cycle of four edges (v_i, v_{i+j}), (v_{i+j}, v_{2g+i}), (v_{2g+i}, v_{2g+i+j}) and (v_{2g+i+j}, v_i). Finally, there are g players in the last class. Each player i in this class also owns a cycle of four edges

(v_i, v_{g+i}), (v_{g+i}, v_{2g+i}), (v_{2g+i}, v_{3g+i}) and (v_{3g+i}, v_i). See Figure 6 for $g = 2$ and the distribution of the 10 players into 3 classes on K_8.

It is well-known that any feasible vertex cover of a complete graph is composed of either all or all but one vertices. For a cover of all $4g$ vertices we can simply drop the payments to one vertex. This reduces the payment for at least one player. In addition, it only increases the cost of some of the deviations as the players must now purchase the uncovered vertex in total. Hence, the stability ratio of the resulting payment scheme can only decrease. Hence, the minimum stability ratio is obtained by purchasing $4g - 1$ vertices.

So w.l.o.g. consider a cover of $4g-1$ vertices including all but vertex v_{4g}. Note that some player subgraphs do not include v_{4g}, and there are only two types of player subgraphs - a single edge or a cycle of length 4. First, consider a player subgraph that consists of a single edge and both end vertices are covered. If the player contributes to the cost of the incident vertices, she can drop the maximum of both contributions. Thus, if she contributes more than 0 to at least one of the vertices, her incentive to deviate is at least a factor of 2. Second, consider a player subgraph that consists of a cycle of length four. Label the four included vertices along a Euclidean tour with u_1, u_2, u_3 and u_4. Let the contributions of the player to u_j be x_j for $j = 1, 2, 3, 4$, respectively. To optimally deviate from a given payment scheme, the player picks one of the possible minimum vertex covers $\{u_1, u_3\}$ or $\{u_2, u_4\}$ and removes all payments outside this cover. A factor of r bounding her incentives to deviate must thus obey the inequalities $\sum_{j=1}^{4} x_j \leq r(x_1+x_3)$ and $\sum_{j=1}^{4} x_j \leq r(x_2+x_4)$. Note that a player might also contribute to vertices outside her cycle. These additional contributions, however, would unnecessarily tighten the bounds and require an increase in r. Hence, in order to find the minimum r that is achievable we assume the player contributes only to vertices inside her subgraph. Summing the two inequalities yields $(2 - r) \sum_{j=1}^{4} x_j \leq 0$, so either her overall contribution is 0 or $r \geq 2$. To derive a payment scheme with stability ratio of strictly less than 2, all $4g - 1$ vertices in the cover must be purchased by the $2g$ players whose subgraph includes v_{4g}.

For the rest of the proof we will concentrate on these $2g$ players. We will refer to player i, if she includes v_i in her subgraph, for $i = 1, \ldots, 2g - 1$. All these players own cycle subgraphs. The player that owns the edge (v_{2g}, v_{4g}) is labeled player $2g$. See Figure 7 for an example on K_8. We denote the contribution of player i to vertex v_j by p_{ij} for all $i = 1, \ldots, 2g$ and $j = 1, \ldots, 4g - 1$.

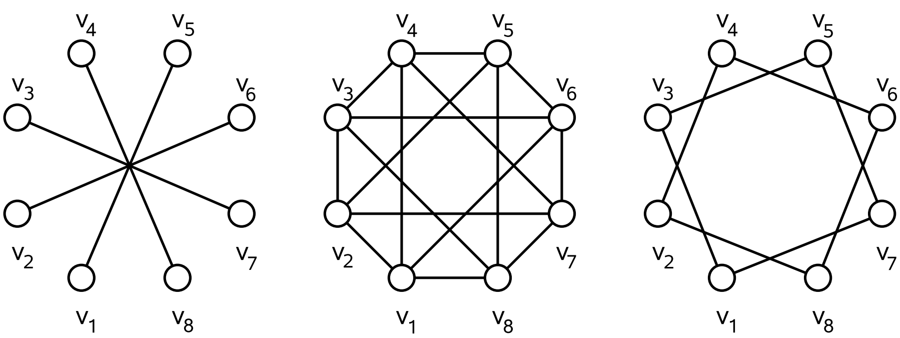

Figure 6: From left to right the edges owned by the players in the first, second, and third classes of players for K_8. The first and second class consist of four players each, the third class of two players. Players in the first class own a single edge, while players in other classes own cycles of length 4.

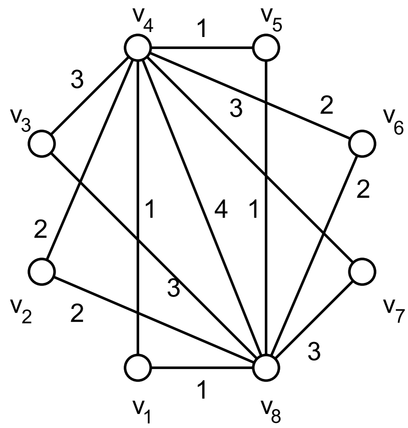

Figure 7: Players that include v_8 in their subtree. Numbering of players as described in the text. Edges labels indicate player ownership.

Observe that for each player the set $\{v_{2g}, v_{4g}\}$ forms a feasible vertex cover. To achieve a stability ratio r, we must ensure that each player can only reduce her payments by a factor of at most r when switching to this cover. In the case of player $2g$ only $\{v_{2g}\}$ is needed, so we must ensure that she can reduce her payments by at most r when dropping all payments but $p_{2g,2g}$. As v_{4g} is not part of the purchased cover its cost of 1 must be purchased completely by a player that strives to use it in a deviation. This yields the following set of $2g$ inequalities:

$$\begin{aligned} \sum_{j=1}^{4g-1} p_{ij} &\leq r(p_{i,2g}+1) \qquad \text{for } i = 1,\dots,2g-1 \\ \sum_{j=1}^{4g-1} p_{2g,j} &\leq r p_{2g,2g} \end{aligned}$$

We again strive to obtain the minimum ratio r that is possible. Note that in the minimum case no vertex gets overpaid, i.e., $\sum_{i=1}^{2g} p_{ij} = 1$ for all $j = 1,\dots,4g-1$. Using this property in the sum of all the inequalities gives

$$4g-1 = \sum_{j=1}^{4g-1}\sum_{i=1}^{2g} p_{ij} \leq r\left(2g-1+\sum_{i=1}^{2g} p_{i,2g}\right) \leq 2gr,$$

which finally yields $r \geq 2 - \frac{1}{2g}$. This proves that in the presented game no (α,β)-NE with $\alpha < 2 - \frac{1}{2g}$ exists. Thus, for every $\varepsilon > 0$ we can pick $g \geq (2\varepsilon)^{-1}$, which then yields a game without $(2-\varepsilon,\beta)$-NE for any $\beta \geq 1$. □

Although we did not consider all deviation possibilities, this lower bound is the best we can get for the presented example games. The following payment scheme is for the players including v_{4g} in their subgraph. Player i purchases $1 - \frac{1}{2g}$ of vertices v_i and v_{2g+i} for $i = 1,\dots,2g-1$. Player $2g$ purchases v_{2g} completely and the remaining $\frac{1}{2g}$ at every other vertex.

It would be interesting to see if this lower bound is connected to the integrality gap of vertex cover. Such a relation exists for approximate budget balanced core solutions in the cooperative game [33]. In a core solution each possible player coalition S contributes no more than the cost of a minimum vertex cover for S. In our game, however, players make concrete strategic investments at the vertices, which alter the cost of the minimum cover for other players. In particular, our result is mainly due to the fact that the majority of players is sufficiently overcovered leaving only a small number of contributing players. This makes a relation to the integrality gap seem more complicated to establish.

Some classes of the vertex cover problem can be approximated to a better extent. For example, there is a PTAS for the vertex cover problem on planar graphs [8]. It is therefore natural to explore whether for planar games we can find covers with approximation and stability ratio arbitrarily close 1. The bad news is that in general there are also limits to the existence of cheap approximate Nash equilibria even for planar games. In particular, Theorem 9 provides a lower bound of 1.5 on the stability ratio for unweighted planar games. For weighted planar games there is an additional trade-off between stability and approximation ratios that yields a stability ratio close to 2 for socially near-optimal covers.

Corollary 1 *There is a planar unweighted vertex cover game without (α, β)-NE for any $\alpha < 1.5$ and $\beta \geq 1$. For any $\beta < \frac{7}{6}$ there is a planar weighted vertex cover game without (α, β)-NE for $\alpha < 2/(2\beta - 1)$.*

Proof. With the planarity of K_4 and Theorem 9 the first part follows. For the second part consider a game from Figure 1(a) with $\varepsilon > 0$. Here every algorithm with $\beta < \frac{2+\varepsilon}{2}$ returns $\mathcal{C}^* = \{v_1, v_2\}$. How good can this cover be in terms of the stability ratio? If player 1 contributes, she can always drop payments to the one vertex to which she contributes the most. This shows that if her contribution is greater than 0, her deviation incentive is at least a factor of 2. If we assign player 2 to purchase the whole cover, this delivers $\alpha = \frac{2}{1+\varepsilon} < 2$ for all $\varepsilon > 0$. Hence, once an algorithm returns (α, β)-NE with $\beta < \frac{2+\varepsilon}{2}$, then for this game any such cover has $\alpha > \frac{2}{1+\varepsilon}$. Solving for ε we get the bound, which proves the second part of the corollary. □

So the better an algorithm is required to be in terms of social cost, the more it allows for selfish improvement by a factor close to 2. Note that these lower bounds apply directly to any algorithm with or without polynomial running time.

6 Facility Location Games

In this section we extend our model to facility location games. We will first restrict ourselves to one of the most simple variants, the uncapacitated facility location problem. It can be stated as follows. A complete bipartite graph $G = (T \cup F, T \times F)$ with vertex sets F of n_f *facilities* and T of n_t *clients* or *terminals* is given. Each facility $f \in F$ has nonnegative opening costs $c(f)$, and for each terminal t and each facility f there is a nonnegative connection cost $c(t, f)$. The goal is to connect each terminal to exactly one opened facility at minimum total cost. The classic integer programming formulation of the UFL problem is due to Balinski [9].

$$\begin{array}{lll} \text{Min} & \sum_{f\in F} c(f)y_f + \sum_{t\in T} c(t,f)x_{tf} & \\ \text{subject to} & \sum_{f\in F} x_{tf} \geq 1 & \forall t \in T \\ & y_f - x_{tf} \geq 0 & \forall t \in T, f \in F \\ & y_f, x_{tf} \in \{0,1\} & \forall t \in T, f \in F. \end{array} \qquad (5)$$

Note that even for this simple version the integer program is not a CIP, as there are negative coefficients in the constraints. Nevertheless, it is possible to construct a non-cooperative game based on an UFL instance as follows. Each of the k non-cooperative players holds a set $T_i \subset T$ of terminals and insists to satisfy the constraints corresponding to her terminals $t \in T_i$. She offers money to the connection and opening costs. In particular, she picks as a *strategy* a pair of two *payment functions* $p_i^c : T \times F \to \mathbb{R}_0^+$ and $p_i^o : F \to \mathbb{R}_0^+$ that specify her offers to the connection and opening costs. If the total offers of all players exceed the cost of a connection or facility, the corresponding variable is raised to 1. In this case the connection or facility is considered bought or opened, respectively. All players can use bought connections and opened facilities for free, no matter whether they contribute to the cost or not. The definitions of payment scheme, Nash equilibrium and (α, β)-NE follow directly. Furthermore, it is possible to apply some of the simplifications observed in Section 2 for covering games. Hence, for the rest of this paper, we will assume that the sets T_i form a partition of the terminal set T. In particular, this means that in a Nash equilibrium connection costs are not shared between players.

6.1 Exact and approximate Nash equilibria

In this section we present results on exact and approximate Nash equilibria for the metric UFL game. For lower bound constructions we mainly use the following transformation to turn a vertex cover game with graph $G = (V, E)$ into a metric UFL game. The set of facilities F is given by the vertex set V of the graph G. For the opening costs $c(f) = c(v)$. The terminal set T is given by the edge set E. For each terminal t corresponding to $(u, v) \in E$ we specify the connection costs for edges between t and the two facilities corresponding to u and v. These edges will be termed *basic* edges, as we explicitly specify the connection cost. All other edge costs are given by the shortest path metric over basic edges.

Even in the metric UFL game the price of anarchy is exactly k. The lower bound is derived by an instance with two facilities, f_1 with cost k and f_2 with cost 1. Each player i has one terminal t_i, and all connection costs are $\varepsilon > 0$. The argumentation follows essentially Theorem 1. Note that also the upper bound of k is easily translated to metric and non-metric UFL games. To derive a bound on the price of stability, we note that there are games without Nash equilibria.

Lemma 6 *There is a metric UFL game without Nash equilibria.*

Proof. The proof follows by translating the game of Figure 1(a) into a metric UFL game. We set the cost of vertex u to 1.5 and the cost of each basic edge to 1. Note that in equilibrium no player will consider to pay a connection cost of 3 to connect a terminal to a facility, because it is always possible to open another facility and connect the terminal with a total cost of less than 3. Hence, in equilibrium only basic edges are bought and the total connection cost is 3. Then the opened

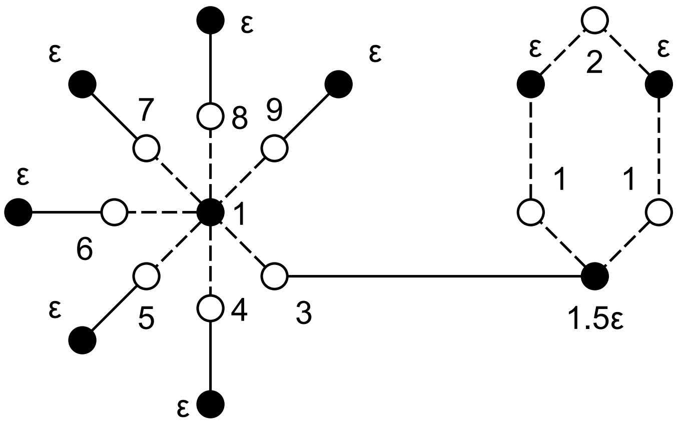

Figure 8: Construction of metric UFL games, in which the price of stability is arbitrarily close to $k-2$. Filled vertices are facilities, empty vertices are terminals. Numerical labels of terminals indicate player ownership, labels of facilities indicate opening costs. All basic solid edges have cost 1, all basic dashed edges cost ε. All other edge costs are given by the shortest path metric.

facilities must resemble a feasible vertex cover for the original instance. This proves the lemma. □

We can use this game to make the price of stability as large as $k-2$.

Theorem 10 *There is a metric UFL game with price of stability at least $k-2$.*

Proof. Consider the game in Figure 8. This game is in essence obtained by transformation from the game in Figure 2. In addition to the transformation there are two major adjustments to be made. First, instead of the unweighted game in Figure 1(b) we attach the game in Figure 1(a) to the star of players $3, \ldots, k$. Second, after the transformation we must adjust opening and connection costs to ensure the property that in equilibrium no non-basic edges are purchased. It is easy to verify that for the presented game this property holds. The argumentation then follows the proof of Theorem 2. In particular, there can be no Nash equilibrium purchasing the facility in the center of the star composed by players $3, \ldots, k$, because in this case players 1 and 2 must agree on opening some of the cheap facilities. This is not possible due to Lemma 6. If, however, the center facility of the star is not bought, each of the players $4, \ldots, k$ pays the connection and opening costs of the corresponding leaf facility resulting in a total cost of $1+\varepsilon$ for each player. Player 3 can connect to the cheap facility and contribute a cost of ε to the opening cost. Then player 1 can contribute the remaining cost of $\varepsilon/2$ and connect her 2 terminals at a cost of 2ε. Player 2 purchases one of the facilities of cost ε and the connection to it. This yields a Nash equilibrium of cost $(k-2)(1+\varepsilon)+4.5\varepsilon$. The social optimum solution has cost $1+(k+3)\varepsilon$. Thus, if ε tends to 0, the lower bound becomes arbitrarily close to $k-2$. □

The next theorem concerns the hardness of deciding Nash equilibrium existence.

Theorem 11 *It is* NP*-hard to decide if a metric UFL game has a Nash equilibrium.*

Proof. The theorem follows directly by transformation from the vertex cover game. If we transform the variable and clause gadgets of Figure 3 using a cost of 1 for each basic edge, then every non-basic edge has cost at least 3. Thus, in equilibrium no non-basic edge is purchased, as a player can always open another facility and connect a terminal with a basic edge with smaller cost. Thus, the set of opened facilities in equilibrium resembles a feasible vertex cover. This proves the theorem. □

Thus, exact Nash equilibria can be quite costly, may not exist, and existence is hard to decide. For some classes of games, however, there is a cheap Nash equilibrium. In particular, results in [6] can be used to show that UFL games with a single terminal per player allow for an iterative improvement procedure similar to the one presented in Section 4.2. Hence, the price of stability is 1, and $(1+\varepsilon, 1.5)$-NE can be found using a recent approximation algorithm [10] to compute a starting solution. In addition, there is another class of games with cheap equilibria that can be computed efficiently.

Theorem 12 *For any metric UFL game, in which the underlying UFL problem has integrality gap 1, the price of stability is 1, and an optimal Nash equilibrium can be computed in polynomial time.*

Proof. The proof works by repeating and adjusting the ideas in the proof of Theorem 4. Reconsider the IP formulation (5) and its corresponding LP-relaxation obtained by allowing $y_f, x_{tf} \geq 0$. The integrality gap is assumed to be 1, so the optimum solution (x^*, y^*) to (5) is optimal for the relaxation. Its dual is

$$\begin{array}{lll} \text{Min} & \sum_{t \in T} \gamma_t & \\ \text{subject to} & \gamma_t - \delta_{tf} \leq c(t,f) & \forall t \in T, f \in F \\ & \sum_{t \in T} \delta_{tf} \leq c(f) & \forall f \in F \\ & \gamma_t, \delta_{tf} \geq 0 & \forall t \in T, f \in F. \end{array} \tag{6}$$

We can find the optimum dual solution (γ^*, δ^*) in polynomial time. It has the same value as (x^*, y^*) for (5). Now assign each player to pay $p_i^o(f) = y_f^* \left(\sum_{t \in T_i} \delta_{tf}^*\right)$ to the opening cost of each facility f. By complementary slackness this assignment purchases every opened facility exactly. In addition, if a terminal t is connected to facility f then $x_{tf}^* = 1$ and $\gamma_t - \delta_{tf} \leq c(t,f)$ is tight. As x^* is integral, for each terminal t we have $x_{tf}^* = 1$ for exactly one opened facility f, and we can let the owning player i contribute $p_i^c(t,f) = x_{tf}^*(\gamma_t^* - \delta_{tf}^*)$ to $c(t,f)$ for every facility f. This assignment exactly pays for all opened facilities and established connections. In addition, with a similar analysis as in Lemma 2 using complementary slackness it is possible to show that $\sum_{f \in F} p_i^o(f) + \sum_{t \in T_i} p_i^c(t,f) = \sum_{t \in T_i} \gamma_t^*$.

To show that $p = (p^o, p^c)$ is a set of best responses we consider the optimization problem of finding a best response for a single player i and her set of terminals T_i. We again construct a *reduced problem* by reduction of costs through contributions in p_{-i}^o. This yields new costs $c'(f) = c(f) - \sum_{j \neq i} p_j^o(f)$. Note that in p no player contributes to the connection costs of a different player, thus $c'(t,f) = c(t,f)$. We can formulate the reduced problem as an integer program with variables (x', y'), relax this program, and consider the LP-dual of the relaxation with variables (γ', δ'). Now consider as primal solution the reduced optimal solution, i.e., $y_f' = y_f^*$ and $x_{tf}' = x_{tf}^*$ for all $t \in T_i$ and $f \in F$. This is still feasible, because (x^*, y^*) satisfies the connection requirements of all $t \in T_i$. Also, it results in the original payments of $p_i^o(f)$ and $p_i^c(t,f)$ and thus in a total cost of $\sum_{f \in F} p_i^o(f) + \sum_{t \in T_i} p_i^c(t,f)$ for the primal. Now consider the LP-dual for the reduced problem and solutions $\gamma_t' = \gamma_t^*$ and $\delta_{tf}' = \delta_{tf}^*$ for all $f \in F$ and $t \in T_i$. These solutions are feasible for the facility constraints in the reduced dual, because they were feasible for the facility constraints in the original dual:

$$c'(f) = c(f) - \sum_{j \neq i} p_j^o(f) = c(f) - y_f^* \sum_{t \in T - T_i} \delta_{tf}^* \geq \sum_{t \in T_i} y_f^* \delta_{tf}^*.$$

The same holds trivially for the connection constraints, because $c'(t,f) = c(t,f)$. This results in a total objective function value of $\sum_{t \in T_i} \gamma'_t$. By LP-duality we know that $\sum_{t \in T_i} \gamma'_t \leq \sum_{f \in F} c'(f) y'_f + \sum_{t \in T_i} c(t,f) x'_{tf}$ for every feasible solution (x', y') to the primal. For (x^*, y^*) we use observations above to show equality. Hence, (x^*, y^*) and (γ^*, δ^*) have the same value for primal and dual LPs of the reduced problem. Thus, (x^*, y^*) yields an optimum solution for the reduced problem. This proves that $p_i = (p_i^o, p_i^c)$ is a best response for player i. Hence, p is a Nash equilibrium with optimum social cost. □

For general games we consider approximate Nash equilibria.

Theorem 13 *For the metric UFL game there is an algorithm to derive* $(3,3)$*-NE in polynomial time.*

Proof. In the following algorithm we denote a terminal by t, a facility by f, and the player owning t by i_t. The algorithm raises budgets for each terminal, which are offered for purchasing the connection and opening costs. Facilities are opened if the opening costs are covered by the total budget offered, and if they are located sufficiently far away from other opened facilities.

For the approximation ratio of 3 we note that the algorithm is a primal-dual method for the UFL problem [38, 41].

For the analysis of the stability ratio consider a single player i and her payments. Note that the algorithm stops raising the budget of a terminal by the time it becomes directly or indirectly connected. We will first show that for the final budgets $\sum_{t \in T_i} B_t$ is a lower bound on the cost of any deviation for player i. For any terminal t we denote by $f(t)$ the facility t is connected to in the calculated solution.

Lemma 7 $c(t,f) \geq B_t$ *for any terminal* t *and any opened facility* $f \neq f(t)$*.*

Proof. First suppose there is such a facility for a terminal that is indirectly connected at the end of the algorithm. This is a contradiction, because then the terminal would have been tight to an opened facility during the run of the algorithm. If this happens, t gets directly connected to f.

Otherwise suppose t is directly connected to $f(t)$. Then, f and $f(t)$ are within a distance of $2B_t$, which is too close for both of them to be open. As t is directly connected to $f(t)$, either $f(t)$ or both f and $f(t)$ are opened at a time where the current budget $B \geq B_t$. If f is opened first and the algorithm tries to open $f(t)$, then with t there is a terminal $c(t,f) + c(t,f(t)) = 2B_t \leq 2B$. Thus, $f(t)$ must stay closed. Otherwise, if the algorithm tries to open f after $f(t)$, then f must be closed for the same reason. □

Hence, if a player has a deviation that improves upon B_t, she must open a new facility and connect some of her clients to it. By opening a new facility, however, the player is completely independent of the cost contributions of other players. Similar to [41] we can argue that the final budgets yield a feasible solution for the dual of the LP-relaxation. Hence, they form a 3-approximately budget balanced core solution for the cooperative game [33]. Now suppose there is a deviation for a player, which opens a new facility f and connects a subset of her terminals T_f to f thereby reducing her

Algorithm 3: Primal-dual algorithm for (3,3)-NE

In the beginning all terminals are unconnected, all budgets B_t are 0, and all facilities closed. Raise budgets of *unconnected* terminals at the same rate until one of the following events occurs. We denote the current budget of unconnected terminals by B. We call a terminal t *tight* with facility f if $B_t \geq c(t, f)$.

1. An unconnected terminal t goes tight with an opened facility f.
   In this case set t *connected* to f and assign player i_t to pay $p^c_{i_t}(t, f) = c(t, f)$.
2. For a facility f not yet definitely closed the sum of the budgets of unconnected and indirectly connected terminals t pays for opening and connection costs: $\sum_t \max(B_t - c(t, f), 0) = c(f)$. Then stop raising the budgets of the unconnected tight terminals. Also,
   (a) if there are opened facility f' and terminal t' with $c(t', f) + c(t', f') \leq 2B$, set f *definitely closed* and all unconnected terminals t tight with f *indirectly connected* to f'.
   (b) Otherwise open f and set all terminals, which are tight with f and not yet directly connected to some other facility, *directly connected* to f. For each such terminal assign player i_t to pay $p^c_{i_t}(t, f) = c(t, f)$ and $p^o_{i_t}(f) = B_t - c(t, f)$.

In the end connect all indirectly connected terminals to the closest opened facility and assign the corresponding players to pay for the connection cost.

cost below the budgets. Then the cost of $c(f) + \sum_{t \in T_f} c(t,f) < \sum_{t \in T_f} B_t$. This, however, would mean that the coalition formed by T_f in the coalitional game has a way to improve upon their budgets, which is a contradiction to B_t having the core-property. Hence, we know that $\sum_{t \in T_i} B_t$ is a lower bound on every deviation cost. Finally, note that for every directly connected terminal $t \in T_i$ player i pays B_t. A terminal t becomes indirectly connected only if it is unconnected and tight to a facility f by the time f is definitely closed. f becomes definitely closed only if there is another previously opened facility f' at distance $2B_t$ from f. Hence, there is an edge $c(t,f') \leq 3B_t$ by the metric inequality. So in the end player i pays at most $3B_t$ when connecting an indirectly connected terminal to the closest opened facility. This establishes the bound on the stability ratio. □

Our proof uses results for cooperative games to lower bound the deviation possibilities of a player i that are not influenced by the contribution of other players than i. Note that it is possible to derive a self-contained proof as for Algorithm 1 before.

Finally, we discuss some observations regarding lower bounds on the stability ratio. There is no polynomial time algorithm for the UFL problem with an approximation ratio of 1.463 unless $\mathsf{NP} \subset \mathsf{DTIME}(n^{O(\log\log n)})$ [27]. Again, this transfers to a lower bound for the stability ratio in terms of polynomial time computability. In addition there is a game giving a constant lower bound r such that the stability ratio $\alpha > r$ in every approximate Nash equilibrium. Reconsider the UFL game obtained from transforming the vertex cover game of Figure 1(a). The structure of the graph is fixed as well as all connection costs. Hence, as there is no exact Nash equilibrium, there must be a constant $r > 1$ such that any cost distribution of any feasible solution can represent only a (α,β)-NE with $\alpha > r$. By appropriate adjustment of edge costs we can obtain a small bound of $r = 1.097$, but the proof is by tedious case analysis and therefore omitted.

6.2 Connection-Restricted Facility Location Games

This section extends the UFL game to connection-restricted facility location (CRFL) problems as considered in [26]. Instead of the constraints $y_f - x_{tf} \geq 0$ there is for each facility f a set P_f of feasible subsets S of terminals that can be connected simultaneously to f. This formulation allows for instance capacity, quota, or incompatibility constraints and thus encompasses several well-known generalizations of the problem. For these games some of the previous results can be extended to hold. In particular, lower bounds on the prices of anarchy and stability follow simply by extension. The price of anarchy can be unbounded, as there might be infeasible assignments, which no player can resolve unilaterally.[2] In the following we show that for subclasses of these games cheap Nash equilibria exist in the CRFL game. We restrict to *closed* CRFL games (CCRFL games) in which the sets P_f are downward closed. In particluar, we assume that for $S \in P_f$ and $S' \subseteq S$ it holds $S' \in P_f$. Note that this class is still quite general, as it includes capacity or incompatibility constraints. An exception are quota constraints, which in general do not yield closed games.

Theorem 14 *For any CCRFL game, in which a partially conic relaxation of the underlying CRFL problem has integrality gap 1, the price of stability is 1.*

[2] Consider two players 1 and 2 owning terminal t_1 and t_2, respectively. There are two facilities f_1 and f_2, for which $P_{f_1} = \{\{t_1\}\}$ and $P_{f_2} = \{\{t_2\}\}$. All opening and connection costs are 1. Suppose t_1 pays for $c(f_2)$ and $c(t_1, f_2)$ and t_2 for $c(f_1)$ and $c(t_2, f_1)$. Then both player constraintss are violated, and they both have infinite cost. However, no player can remove payments of others, so no player can unilaterally obtain a feasible solution. Therefore, the state is a Nash equilibrium of unbounded cost.

Proof. We can formalize the CRFL problem by an integer program as follows:

$$
\begin{array}{lll}
\text{Min} & \sum_{f\in F} c(f)y_f + \sum_{t\in T} c(t,f)x_{tf} & \\
\text{subject to} & \sum_{f\in F} x_{tf} \geq 1 & \forall t\in T \\
& (y_f, x_{1f},\dots,x_{n_t f}) \in P_f & \forall f \in F \\
& y_f, x_{tf} \in \{0,1\} & \forall t\in T, f\in F.
\end{array} \tag{7}
$$

Here $P_f = \{(0,\dots,0)\}\cup\{(1,\chi_S)|S\subset T \text{ feasible for } f\}\subset\{0,1\}^{n_t+1}$, and χ_S denotes the characteristic vector of the subset S.

Following the argumentation in [26] it is possible to use the conic hull of the sets P_f to derive a linear relaxation:

$$
\begin{array}{lll}
\text{Min} & \sum_{f\in F} c(f)y_f + \sum_{t\in T} c(t,f)x_{tf} & \\
\text{subject to} & \sum_{f\in F} x_{tf} \geq 1 & \forall t\in T \\
& (y_f, x_{1f},\dots,x_{n_t f}) \in \text{cone}(P_f) & \forall f \in F.
\end{array} \tag{8}
$$

For this program a dual can be given by

$$
\begin{array}{lll}
\text{Max} & \sum_{t\in T} \gamma_t & \\
\text{subject to} & \sum_{t\in S} \gamma_t \leq c(f) + \sum_{t\in S} c(t,f) & \forall f\in F, S\in P_f.
\end{array} \tag{9}
$$

Now we will apply similar arguments given before in Theorems 4 and 12. An integral optimum solution (x^*,y^*) to the LP relaxation represents a partition of the terminal set T into a collection of feasible sets $S_f^*\in P_f$, one for each facility f. We will again construct a set of strategies $p=(p^o,p^c)$ for the players such that γ_t^* from the corresponding optimum dual solution is the total payment of a terminal t. For player i we assign an amount $p_i^c(t,f) = c(t,f)$ as connection cost to f with $t\in S_f^*$. For the opening costs $p_i^o(f) = \sum_{t\in S_f\cap T_i}\gamma_t^* - c(t,f)$. For a player i this yields $\sum_{f\in F} p_i^o(f) + \sum_{t\in T_i} p_i^d(t,f) = \sum_{t\in T_i}\gamma_t^*$ as total payment. W.l.o.g. we again assume that the sets T_i are mutually disjoint. Thus, the total payment of all players is $\sum_{t\in T}\gamma_t^*$.

We first show that for all $t\in S_f^*$ we have $\gamma_t^*\geq c(t,f)$. Thus, for every player i and every facility f we have $p_i^o(f)\geq 0$. In addition, we note that the solution is exactly paid for. This shows that state $p=(p^o,p^c)$ is feasible and all costs of the solution are exactly paid for.

Lemma 8 *For the optimum dual solution it holds $\gamma_t^*\geq c(t,f)$ for every $t\in S_f^*$. In addition, for every $f\in F$ we have $\sum_{t\in S_f^*}\gamma_t^* = c(f)+\sum_{t\in S_f^*} c(t,f)$.*

Proof. The proof of the first part follows using the closed property of the sets P_f. Note that by strong LP-duality we have $\sum_{t\in T_i}\gamma_t^* = \sum_{f\in F} c(f)y_f^* + \sum_{t\in T} x_{tf}^*$, i.e., the total contribution of all players exactly matches the cost of the optimum solution that is to be paid for. Suppose for contradiction that for some facility f there is a terminal $t'\in S_f^*$ with $\gamma_{t'}^* < c(t',f)$.

First suppose that $\sum_{t\in S_f^*}\gamma_t^* = c(f)+\sum_{t\in S_f^*} c(t,f)$. In this case, consider the set of players $S_f^* - \{t'\}$, for which we have $\sum_{t\in S_f^*, t\neq t'}\gamma_t^* > c(f) + \sum_{t\in S_f^*-\{t'\}} c(t,f)$. We have a CCRFL, thus we know that $S_f^* - \{t'\} \in P_f$. This is a contradiction, because γ^* would violate the corresponding dual constraint.

Now suppose that $\sum_{t \in S_f^*} \gamma_t^* < c(f) + \sum_{t \in S_f^*} c(t,f)$. Because primal and dual solutions have the same value, there must be at least one facility f' and a set of players $S_{f'}^* \in P_{f'}$ with $\sum_{t \in S_{f'}^*} \gamma_t^* > c(f) + \sum_{t \in S_{f'}^*} c(t,f')$. This, however, again violates the corresponding constraint in the dual. This proves the first part of the lemma. Note that the last argument also proves the second part of the lemma. □

We now show that p is a collection of best responses. Consider a player i and the problem of finding a best response strategy. This reduces to optimally solving the facility location problem for terminals $t \in T_i$ under the fixed payments p_{-i} of other players. Similarly to the previous theorems this yields a *reduced problem* to cover terminals $t \in T_i$ via allowed connections to open facilities, where opening costs are $c'(f) = c(f) - \sum_{j \neq i} p_j^o(f)$ for each $f \in F$. Due to Lemma 8 we again have $c'(t,f) = c(t,f)$ for each $f \in F$ and $t \in T_i$ as in Theorem 12. For the reduced problem we must not only fix the payments but also the connectivity conditions of all other players than i. We need to check these conditions to evaluate whether a unilateral deviation for player i is feasible or not. A terminal $t \in T - T_i$ is connected to facility f if $p_i^c(t,f) \geq c(t,f)$. Thus, the set of necessarily connected terminals of players $j \neq i$ for facility f in state p can be given by $N_f = \{t \in T - T_i \mid x_{tf}^* = 1\}$. The feasible subsets at f for player i reduce to all subsets of T_i that can be combined with N_t to yield a set in P_f. More formally, we get $P_f' = \{S \in P_f \cap T_i \mid S \in P_f, N_t \subseteq S\}$, which is the set of feasible subsets in the reduced problem. A conic relaxation of the corresponding integer program can be given by

$$\begin{aligned} \text{Min} \quad & \sum_{f \in F} c'(f) y_f' + \sum_{t \in T_i} c'(t,f) x_{tf}' \\ \text{subject to} \quad & \sum_{f \in F} x_{tf}' \geq 1 && \forall t \in T_i \\ & (y_f', x') \in \text{cone}(P_f') && \forall f \in F. \end{aligned} \tag{10}$$

For this program the dual can be given by

$$\begin{aligned} \text{Max} \quad & \sum_{t \in T_i} \gamma_t' \\ \text{subject to} \quad & \sum_{t \in S} \gamma_t' \leq c'(f) + \sum_{t \in S} c(t,f) \quad \forall f \in F, S \in P_f'. \end{aligned} \tag{11}$$

As done previously, we use (x^*, y^*) as candidate solution for the reduced problem. Note that it is trivially feasible, because it yields a feasible solution for the original problem and was used to define P_f'. Also, it results in $c_i(p) = \sum_{f \in F} p_i^o(f) + \sum_{t \in T_i} p_i^c(t,f)$ as objective function value. For the dual we again use $\gamma_t' = \gamma_t^*$. This choice of γ_t' is feasible, because for every $S \in P_f'$ we have

$S \cup N_t \in P_f$ and

$$\begin{aligned} c'(f) + \sum_{t \in S} c(t,f) &= c(f) - \sum_{j \neq i} p_j^o(f) + \sum_{t \in S} c(t,f) \\ &= c(f) - \left(\sum_{t \in N_t} \gamma_t^* - c(t,f) \right) + \sum_{t \in S} c(t,f) \\ &= c(f) + \sum_{t \in S \cup N_t} c(t,f) - \sum_{t \in N_t} \gamma_t^* \\ &\geq \sum_{t \in S} \gamma_t^*, \end{aligned}$$

because γ_t^* was feasible for the original dual. We get as value of the dual objective function $\sum_{t \in T_i} \gamma_t^* = \sum_{f \in F} p_i^o(f) + \sum_{t \in T_i} p_i^c(t,f)$, which shows that (x^*, y^*) represents an optimum solution to the reduced problem. Thus, $p_i = (p_i^o, p_i^c)$ is a best response for player i. This proves that p is a Nash equilibrium of optimum cost. □

If we can optimize over the cones in polynomial time we can also compute the best Nash equilibrium in polynomial time. In particular, this is the case if the cone can be described by a polynomial number of inequalities.

Finally, we observe that there are singleton CCRFL games (in which every player has exactly one terminal) without any pure Nash equilibrium. In particular, insights from [6] and Theorem 5 that prove existence of optimal Nash equilibria cannot be extended to this class of games. For example, consider a game with two clients $T = \{t_1, t_2\}$ as players.[3] There are two facilities f_1 and f_2. All connection costs are 0. The opening costs are $c(f_1) = 2.2$ and $c(f_2) = 1$. The set $P_{f_1} = 2^T$, whereas f_2 has a capacity constraint of $x_{t_1 f_2} + x_{t_2 f_2} \leq 1$, which yields $P_{f_2} = \{\{t_1\}, \{t_2\}\}$. In every feasible solution f_1 must be open, but f_1 can never be paid for in a Nash equilibrium. Note, however, that the integrality gap of the underlying problem is not 1, because we can assign each player fractionally $x_{t_1 f_1} = x_{t_2 f_1} = x_{t_1 f_2} = x_{t_2 f_2} = 0.5$ and open both facilities to a degree of $y_1 = y_2 = 0.5$. This satisfies all constraints and yields a total cost of 1.6, which is strictly less than for any feasible integral solution.

7 Conclusion

In this paper we have studied a general class of cost sharing games for non-cooperative agents based on combinatorial optimization problems. Exact Nash equilibria in these games are often absent, costly, and their existence is hard to decide. Nevertheless, there are non-trivial classes of games that even allow optimal exact Nash equilibria, e.g., in case agents have a single element or the integrality gap of an underlying linear programming formulation is 1. For many games we provided efficient algorithms to compute approximate Nash equilibria with small constant stability and approximation ratios. They represent an interesting trade-off between optimality of social cost and stability of agent incentives, especially in light of the fact that finding social optima and best response strategies in these games is NP-hard.

There are a number of interesting open problems that arise from our work. We have shown that certain classes of primal-dual algorithms can be used to compute approximate Nash equilibria.

[3] We thank an anonymous reviewer for pointing out this example to us.

A similar result is known to hold for cooperative games and mechanism design. However, in our non-cooperative model we need to find stable allocations of payments to resources. Because of this allocation property, the application of approximation algorithms to compute approximate Nash equilibria is much more challenging. In particular, it would be interesting to see if primal-dual algorithms for Steiner network problems can be applied to compute (α, β)-NE with small constant ratios in variants of connection games.

In addition, there are other interesting stability and fairness concepts in our model. For instance, it is an interesting open problem to characterize existence and cost of exact and approximate versions of coalitional stability concepts like strong equilibria, which have been studied recently in single source conncetion games [22]. Finally, there are many other interesting minimization problems, for which the analysis of corresponding cost sharing games can provide interesting insights into the stability of incentives.

Acknowledgements

The authors acknowledge support by the Communauté française de Belgique (ARC project), by DFG through Research Training Group 1042 "Explorative Analysis and Visualization of Large Information Spaces", UMIC Research Centre at RWTH Aachen University, and DFG grant Kr 2332/1-2 within Emmy Noether research group on "Algorithmic Game Theory".

References

[1] Susanne Albers. On the value of coordination in network design. *SIAM J. Comput.*, 38(6):2273–2302, 2009.

[2] Noga Alon, Dana Moshkovitz, and Shmuel Safra. Algorithmic construction of sets with k-restrictions. *ACM Trans. Algorithms*, 2(2):153–177, 2006.

[3] Elliot Anshelevich and Bugra Caskurlu. Exact and approximate equilibria for optimal group network formation. In *Proc. 17th European Symposium on Algorithms (ESA)*, pages 239–250, 2009.

[4] Elliot Anshelevich and Bugra Caskurlu. Price of stability in survivable network design. In *Proc. 2nd Intl. Symp. Algorithmic Game Theory (SAGT)*, pages 208–219, 2009.

[5] 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.

[6] Elliot Anshelevich, Anirban Dasgupta, Éva Tardos, and Tom Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4:77–109, 2008.

[7] Baruch Awerbuch, Yossi Azar, and Amir Epstein. The price of routing unsplittable flow. In *Proc. 37th Symp. Theory of Computing (STOC)*, pages 331–337, 2005.

[8] Brenda Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994.

[9] Michel Balinski. Integer programming: Methods, uses, computation. *Management Science*, 12(3):253–313, 1965.

[10] Jaroslaw Byrka. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. In *Proc. 10th Intl. Workshop Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 29–43, 2007.

[11] Jean Cardinal and Martin Hoefer. Selfish service installation in networks. In *Proc. 2nd Intl. Workshop Internet & Network Economics (WINE)*, pages 174–185, 2006.

[12] Moses Charikar, Howard Karloff, Claire Mathieu, Joseph Naor, and Michael Saks. Online multicast with egalitarian cost sharing. In *Proc. 20th Symp. Parallelism in Algorithms and Architectures (SPAA)*, pages 70–76, 2008.

[13] Chandra Chekuri, Julia Chuzhoy, Liane Lewin-Eytan, Joseph Naor, and Ariel Orda. Non-cooperative multicast and facility location games. *IEEE J. Selected Areas in Communications*, 25(6):1193–1206, 2007.

[14] Ho-Lin Chen and Tim Roughgarden. Network design with weighted players. *Theory Comput. Syst.*, 45(2):302–324, 2009.

[15] Ho-Lin Chen, Tim Roughgarden, and Gregory Valiant. Designing networks with good equilibria. In *Proc. 19th Symp. Discrete Algorithms (SODA)*, pages 854–863, 2008.

[16] George Christodoulou and Elias Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proc. 13th European Symposium on Algorithms (ESA)*, pages 59–70, 2005.

[17] Xiaotie Deng, Toshihide Ibaraki, and Hiroshi Nagamochi. Algorithmic aspects of the core of combinatorial optimization games. *Math. Oper. Res.*, 24(3):751–766, 1999.

[18] Nikhil Devanur, Naveen Garg, Rohit Khandekar, Vinayaka Pandit, Amin Saberi, and Vijay Vazirani. Price of anarchy, locality gap, and a network service provider game. In *Proc. 1st Intl. Workshop Internet & Network Economics (WINE)*, pages 1046–1055, 2005.

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

[20] Stefan Eidenbenz, Amit Kumar, and Sigrid Zust. Equilibria in topology control games for ad hoc networks. *Mobile Networks and Applications*, 11(2):143–159, 2006.

[21] Horst Eiselt, Gilbert Laporte, and Jacques-Francois Thisse. Competitive location models: A framework and bibliography. *Transportation Science*, 27:44–54, 1993.

[22] Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. *Games Econ. Behav.*, 67(1):51–68, 2009.

[23] Hugh Everett. Recursive games. In H. Kuhn and A. Tucker, editors, *Contributions to the theory of games III*, volume 39 of *Annals of Mathematical Studies*, pages 47–78. 1957.

[24] Alex Fabrikant, Ankur Luthera, Elitza Maneva, Christos Papadimitriou, and Scott Shenker. On a network creation game. In *Proc. 22nd Symp. Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[25] Michael Garey and David Johnson. *Computers and Intractability: A guide to the theory of NP-completeness*. W.H. Freeman and Company, New York, 1979.

[26] Michel Goemans and Martin Skutella. Cooperative facility location games. *J. Algorithms*, 50(2):194–214, 2004.

[27] Sudipto Guha and Samir Khuller. Greedy strikes back: Improved facility location algorithms. *J. Algorithms*, 31:228–248, 1999.

[28] Magnus Halldórsson, Joe Halpern, Li Li, and Vahab Mirrokni. On spectrum sharing games. In *Proc. 22nd Symp. Principles of Distributed Computing (PODC)*, pages 107–114, 2004.

[29] Martin Hoefer. Non-cooperative facility location and covering games. In *Proc. 17th Intl. Symp. Algorithms and Computation (ISAAC)*, pages 369–378, 2006.

[30] Martin Hoefer. Non-cooperative tree creation. *Algorithmica*, 53(1):104–131, 2009.

[31] Martin Hoefer and Piotr Krysta. Geometric network design with selfish agents. In *Proc. 11th Conf. Computing and Combinatorics (COCOON)*, pages 167–178, 2005.

[32] Nicole Immorlica, Mohammad Mahdian, and Vahab Mirrokni. Limitations of cross-monotonic cost sharing schemes. *ACM Trans. Algorithms*, 4(2), 2008. Special Issue SODA 2005.

[33] Kamal Jain and Vijay Vazirani. Applications of approximation algorithms to cooperative games. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 364–372, 2001.

[34] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate within 2-ϵ. *J. Comput. Syst. Sci.*, 74(3):335–349, 2008.

[35] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *Proc. 16th Symp. Theoretical Aspects of Computer Science (STACS)*, pages 404–413, 1999.

[36] Xiang-Yang Li, Zheng Sun, and Weizhao Wang. Cost sharing and strategyproof mechanisms for set cover games. In *Proc. 22nd Symp. Theoretical Aspects of Computer Science (STACS)*, pages 218–230, 2005.

[37] Xiang-Yang Li, Zheng Sun, Weizhao Wang, Xiaowen Chu, ShaoJie Tang, and Ping Xu. Mechanism design for set cover games with selfish element agents. *Theor. Comput. Sci.*, 411(1):174–187, 2010.

[38] Ramgopal Mettu and Greg Plaxton. The online median problem. *SIAM J. Comput.*, 32(3):816–832, 2003.

[39] Tan Miller, Terry Friesz, and Roger Tobin. *Equilibrium Facility Location in Networks*. Springer Verlag, 1996.

[40] Dov Monderer and Lloyd Shapley. Potential games. *Games Econ. Behav.*, 14:1124–1143, 1996.

[41] 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.

[42] Christos Papadimitriou. Algorithms, games and the internet. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 749–753, 2001.

[43] Tim Roughgarden and Éva Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[44] Vijay Vazirani. *Approximation Algorithms*. Springer Verlag, 2000.

[45] Adrian Vetta. Nash equilibria in competitive societies with application to facility location, traffic routing and auctions. In *Proc. 43th Symp. Foundations of Computer Science (FOCS)*, page 416, 2002.