Competitive Cost Sharing with Economies of Scale*

Martin Hoefer†

Abstract

We consider a general class of non-cooperative *buy-at-bulk cost sharing games*, in which k players make investments to purchase a set of resources. Each resource has a certain cost and must bought to be available to the players. Each player has a certain constraint on the number and types of resources that she needs to have available, and she can specify payments to make a resource available to her. She strives to fulfill her constraint with the smallest investment possible. Our model includes a natural economy of scale: for a subset of players capacity must be installed at the resources, and the cost increase for a resource r is composed of a fixed price $c(r)$ and a global concave capacity function g. This cost can be shared arbitrarily between players.

We consider the existence and total cost of pure-strategy exact and approximate Nash equilibria. In general, prices of anarchy and stability depend heavily on the economy of scale and are $\Theta(k/g(k))$. For non-linear functions g pure Nash equilibria might not exist, and deciding their existence is NP-hard. For subclasses of games corresponding to covering problems, primal-dual methods can be applied to derive cheap and stable approximate Nash equilibria in polynomial time. In addition, for singleton games optimal Nash equilibria exist. In this case expensive exact as well as cheap approximate Nash equilibria can be computed in polynomial time. Most of these results can be extended to games based on facility location problems.

1 Introduction

Game-theoretic aspects of large networks are a reseach area that has received much interest recently. Networks like the Internet, which are subject to the strategic behavior of various economic agents, play a crucial role in the development of modern societies. It is therefore important to understand the underlying dynamics that govern their development. In this paper we consider a general class of non-cooperative cost sharing games. They can for instance serve as model for crucial investment problems in networks like service installation, facility location or various network design problems. Our games represent a fundamental model to study the results of economic competition in a variety of investment tasks, which (telecommunication) companies and other parties concerned with the development of the Internet face today, such as topology creation, installation of amplification technology, server placement etc.

In particular, we consider games for k players that strive to obtain a number of resources with minimum investment. There is a set of resources, and each resource has a cost. Each player picks as a strategy a function that specifies her offer to each resource. If the sum of offers made by a set of players exceeds the resource cost, it is considered available for these

*An extended abstract of this work has appeared in LATIN 2008 [26].

†Department of Computer Science, RWTH Aachen, Germany. mhoefer@cs.rwth-aachen.de. Supported by DFG through UMIC Research Centre at RWTH Aachen University.

players. For each player there is a constraint on the number and types of resources that must be available for her. She strives to fulfill this constraint with minimum total investment in her strategy. This class of games, which is sometimes referred to as *arbitrary* cost sharing games [6,19], has been studied intensively recently for a variety of scenarios involving network design [3,4,6,19,27], facility location [24], and covering problems [9,24]. We will refer to such games as *regular cost sharing games*. In this paper, we consider a more realistic formulation called *buy-at-bulk cost sharing games*, which incorporates the assumption that for increased usage of resources players have to install a larger capacity. This aspect is modeled using resource costs with economies of scale. Formally, a resource becomes more expensive when it shall be available to a larger set of players. If resource r is available to a set of i players, the cost is $c(r,i) = c(r) \cdot g(i)$, in which $c(r)$ is a fixed cost and g is a non-decreasing and concave function, which is used for every resource r. We will restrict our attention to pure states in which players do not randomize over strategies. We consider existence and cost of pure Nash equilibria and leave a study of mixed Nash equilibria for future work. In this paper the term *Nash equilibrium* refers to pure ones throughout.

Regular cost sharing games have recently attracted significant research interest to model self-interested agents that have to agree upon a cost sharing of a joint investment in unregulated settings. This approach is complementary to a number of recent works on *designing* cost sharing games to obtain favorable equilibrium properties [15]. In contrast to the design perspective, we do not assume the existence of a central authority that designs and maintains the solution and dictates cost shares for each player. Instead we consider a scenario with payment functions that allows players to freely specify their payment for each resource. On the one hand, such an assumption is necessary when there is very little control over players and their bargaining options, e.g., when considering cost sharing of global investments in the Internet. Our model thus has the advantage of more freedom of choice and less required control. On the other hand, these advantages come with a cost, which is the fact that Nash equilibria might not exist (unlike, e.g., for fixed cost sharing mechanisms such as Shapley value cost sharing [5,15]). Furthermore, prices of anarchy [30] and stability [5], which are the ratios of the cost of the worst and best Nash equilibrium over the cost of a socially optimum solution, respectively, can in general be as large as the number of players k. Nevertheless, in some interesting special cases specified below, we can derive existence of optimal Nash equilibria, which does not hold for Shapley cost sharing [5], and we show how to obtain an arbitrary Nash equilibrium in polynomial time.

Our games include a variety of cases, in which computing a best response is NP-hard, which is the case in many realistic optimization contexts. Then a player must use a heuristic to find a good strategy, which will only be an approximation to the best response. Hence, we consider *approximate Nash equilibria* – pure states that are as cheap and stable as possible – and assess them with respect to their approximation to the social cost and the incentives that they give players to deviate. The assumption of using approximation algorithms to find best responses motivates *relative* performance ratios, which are standard in approximation algorithms, rather than the more common notion of *additive* approximate Nash equilibria in game theory (see also [14] for a discussion). In this paper an (α, β)-approximate Nash equilibrium is a state, in which each player can decrease his cost by unilateral deviation by at most a factor α, and which represents a β-approximation to the socially optimum cost. We refer to α as the *stability ratio* and β as the *approximation ratio*. For classes of set multi-cover and facility location games we will show the existence of approximate Nash equilibria with small stability and approximation ratios, which mitigates non-existence of pure exact Nash equilibria. Our main interest, however, is to investigate the influence of the function g on the efficiency and computational complexity of exact and approximate Nash equilibria.

In our games we explicitly incorporate different costs for different coalitions of players that strive to obtain a resource. This gives rise to slightly subtle and potentially problematic issues. Given a set of contributions of the players to a resource r, it can happen that the contributions suffice to pay the cost $c(r,i)$ for several distinct coalitions of players, but not for their union. Consider for instance one player p that contributes $c(r,2)$ to a resource r, whereas all other players do not contribute to r at all. In this case r can be available to one pair of players (p,q), for another player $q \neq p$. Every player might now believe that she is q and the resource is available to her, but clearly the cost for the complete set of players is not paid for. In general, this issue could be resolved by centralized mechanisms or separate bargaining between players to determine which coalition gets preferred access. The presence of such coordination features, however, is unlikely in unregulated settings, and making assumptions about them would crucially limit our model. Instead, note that in our example the ambiguities arise from the fact that player p is overcontributing to r, she only needs to pay $c(r,1)$ to make the resource available for herself. It is easy to observe that when such irrational overcontribution is absent, the economies of scale guarantee that there is always a unique maximal set of players to which r is available. Hence, separate conflict resolution methods are unnecessary, because, as a byproduct, individual rationality resolves all such ambiguities and conflicts. While we have mentioned above that computing a best response strategy for obtaining a subset of resources can be NP-hard, our argument here requires only that a player can determine the minimum contribution necessary to make a single resource available to her. This can be done efficiently, and a simple procedure is implicitly given below in Section 2 when we formally discuss availability of resources.

Related Work. In the area of *competitive location* there has been a high research activity on game-theoretic models for facility location during the last decades [18,33]. In these models facility owners are players that decide where to open a facility. Clients are mostly behavioral, e.g., assumed to connect to the closest facility. A recent example of this kind of location game is also found in [37]. According to our knowledge, however, none of these models considers a situation where clients are in charge of independently creating connections and facilities.

Cooperative games have been studied quite intensively in the past (see [16, 21] and the references therein). In [16] the authors prove that the core of cooperative games based on covering and packing integer programs is non-empty if and only if the integrality gap is 1. They also show results on polynomial time computability of core solutions in a number of special cases. In [21] similar results are shown for class of cooperative facility location games. Some of these games have also been analyzed with respect to mechanism design. In addition, cost sharing mechanisms have been considered for games based on set cover and facility location. Every player corresponds to a single item and has a private utility 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 [17] presented strategyproof mechanisms for set cover and facility location games. For set cover games this work was extended [32, 35] to different fairness aspects and formulations with items or sets being agents, for facility location games computing cross-monotonic cost sharing schemes was considered in [34], and in [28] lower bounds on their budget-balance were provided. In contrast, our approach is an extension of non-cooperative games, which were first studied in [6] and recently in [3,4,19,27] in a network design context. Our recent work [9, 24] provided results for exact and approximate Nash equilibria in covering and facility location games. Prices of anarchy and stability in these games are generally as large as $\Theta(k)$. None of these previous models, however, considers the influence of different economies of scale.

Buy-at-bulk problems are a vivid recent research area in the analysis of network design and facility location problems. In particular, starting with [7] network design problems with economies of scale were considered. Typically, there are a number of source-sink pairs with demands that must be routed by an unsplittable flow. Edge and/or vertex costs increase with the demand routed over them. Recently, in [12, 13] polylogarithmic approximation algorithms were given. As a lower bound, logarithmic hardness results for general resource costs were derived [2]. For special cases, e.g., single-souce or rent-or-buy problems [22] there exist constant-factor approximation algorithms. This is also the case for unit-demand metric facility location [23], for which an adjustment of recently proposed greedy algorithms [29] yields the same approximation guarantees for the buy-at-bulk as for the regular problem.

Our Contribution. Buy-at-bulk cost sharing games studied in this paper are a new general model to consider cost sharing in optimization problems with economies of scale. In addition, as an extension they address a frequent criticism to regular cost sharing games. In regular games there is only a fixed cost for each resource. As soon as this cost is paid for, the resource is available to every player, *no matter* whether she contributes or not. Hence, the game inherently allows *free riders* who can obtain a resource for free. This problem has been addressed frequently [1, 5, 8, 10, 11, 14, 20, 31] by fixing a Shapley sharing of resource cost. In contrast, our model allows smaller groups of players to obtain the resource at cheaper costs. This ensures that every player is eventually forced to contribute for availability. The severeness of this force depends on the number of players that request a resource and is dynamically affected by g. Some undesirable properties of the game like a high price of anarchy are directly influenced, the price of anarchy is exactly $\frac{k}{g(k)}$. Other properties are independent of this adjustment, e.g., for any non-linear g there are games without Nash equilibria. The price of stability is as large as $\Theta\left(\frac{k}{g(k)}\right)$, and it is NP-hard to decide the existence of Nash equilibria. Interestingly, some existence and optimality conditions for regular games can be extended to hold for buy-at-bulk games. If each player wants to cover exactly one element, optimal Nash equilibria exist, and $(1+\epsilon, \beta)$-approximate Nash equilibria can be obtained in polynomial time by a local search from any β-approximate starting state. In addition, we provide a procedure to find an (arbitrary) exact Nash equilibrium in polynomial time, which was not known before even for regular singleton games. These results are shown for set multi-cover games, which were not studied before. In general, there are (f, f)-approximate Nash equilibria even for non-singleton set cover games, where f is the maximum frequency of any element in the sets. With the exception of the last result, all our proofs translate more or less canonically to buy-at-bulk cost sharing games for facility location.

2 Model and Basic Properties

In a buy-at-bulk cost sharing game there is a set $[k]$ of k non-cooperative players and a set R of resources. Each resource $r \in R$ has a *fixed cost* $c(r) \geq 0$. In addition, there is a function

$g : \mathbb{N} \to \mathbb{R}_+^0$, which is non-negative, non-decreasing, concave, and has $g(0) = 0$ and $g(1) > 0$. We normalize the function to obey $g(1) = 1$. For convenience, we use $\mu(i) = g(i) - g(i-1)$, which is non-increasing and non-negative for all $i \geq 1$. The *bundle cost* of resource r for a set of i players is $c(r, i) = c(r) \cdot g(i)$. A strategy s_p of a player p is a function $s_p : R \to \mathbb{R}_+^0$ to specify her non-negative payment to each resource. A *state* is a vector $s = (s_1, \dots, s_k)$ with a strategy for each player. We denote by s_{-p} the same vector without s_p.

A resource r is *available* to a player p if there exists a subset $Q \subset [k]$ of players with $p \in Q$ that pays the corresponding bundle cost, i.e. $\sum_{p \in Q} s_p(r) \geq c(r, |Q|)$. In particular, a player can easily determine whether a given contribution $s_p(r)$ suffices to make r available given s_{-p} of other players. Assume players $q \in [k] - p$ are numbered in non-increasing order of $s_q(r)$. Then r is available to p if and only if

$$s_p(r) \geq \min_{i=0,\dots k-1} \{c(r, i+1) - \sum_{q=1}^{i} s_q(r)\} \ .$$

Note that when $i = 0$, player p is assumed to purchase the resource by herself, and we count only the fixed cost $c(r)$. For p we use $\rho_p(s)$ to denote the set of her available resources, and we drop the argument whenever context allows.

Each player p has a player-specific *constraint* on ρ_p. In this paper we consider constraints that inherit a free-disposal property. A constraint can never be violated by having *additional* resources available to the ones required. If ρ_p does not fulfill the constraint, we assume that the player is penalized with a prohibitively large cost, i.e., for her *individual cost* $c_p = +\infty$. Otherwise, if her constraint is satisfied, the individual cost is her total investment $c_p(s_p, s_{-p}) = \sum_{r \in R} s_p(r)$. A player wants to minimize her individual cost, so she strives to fulfill her constraint with ρ_p at the least possible investment. A *Nash equilibrium* (denoted NE) is a state, in which no player can reduce her individual cost by changing her strategy. We restrict our attention to pure states in this paper and leave a deeper study of mixed NE for future work. As *social cost* of a state s of the game we use the sum of individual costs $c(s) = \sum_{p \in [k]} c_p(s)$. A social optimum state minimizing social cost will be denoted s^* throughout. A (α, β)-approximate Nash equilibrium (denoted (α, β)-NE) is a state, in which no player can reduce her individual cost by a factor of more than α, and for which the social cost is a β-approximation to the minimum social cost over all states of the game. More formally, for a (α, β)-NE we have $c_p(s) \leq \alpha c_p(s'_p, s_{-p})$ for every possible strategy s'_p and $c(s) \leq \beta c(s^*)$.

Observe that in any NE and any social optimum state s^* the available resources for each player satisfy her constraints. Due to concavity of g, in any NE there is a unique maximal set of players (denoted Q_r), for which the resource is available. This set includes as subsets all other sets of players, for which the resource is available. No subset of i players will contribute more than $c(r, i)$ to any resource r. The strategies exactly purchase the bundle cost $c(r, |Q_r|)$ of every resource. Thus, a NE s represents a cost sharing of the set of resources. This property can be assumed for s^* as well, because in this case the cost distribution is irrelevant. Finding s^* is equivalent to finding a solution to the underlying buy-at-bulk minimization problem given by satisfying all player constraints at minimum total cost. In this problem, a feasible solution is a vector that indicates for each player, which resources are available to her such that all constraints are satisfied.

Finally, the function $g(i) \in [1, i]$ for all $i \geq 1$. Previously considered regular cost sharing games for covering [9, 24], facility location [24], and network design [3, 4, 6, 27] were buy-at-bulk cost sharing games with $g(i) = 1$ for all $i \geq 1$. When speaking of games tied to optimization problems in this paper – for instance vertex cover games – we generally refer to the buy-at-bulk version. It is explicitly pointed out when regular cost sharing games are under consideration.

2.1 Covering and Facility Location

The definition of cost sharing games allows a number of general classes of games to be defined in this framework. A simple class, which we will repeatedly consider, is a (buy-at-bulk) vertex cover game on an undirected graph $G = (V, E)$. The resources $R = V$, and each player corresponds to a subset of edges $E_p \subset E$. Her constraint is satisfied if for each edge there is at least one incident vertex available to her. In this way we generalize to set multi-cover games. There is a set of elements E, and the resources are given by $R = \mathcal{M} \subseteq 2^E$ of subsets $M \in \mathcal{M}$, such that $M \subseteq E$. Each player corresponds to a subset $E_p \subseteq E$ of elements. In addition, there is a number $b(e) > 0$ for each element $e \in E$. A player is satisfied with her choice if for each element $e \in E_p$ there are at least $b(e)$ sets available to her that include e.

Note that for regular set multi-cover games the underlying optimization problem is sometimes termed *constrained* set multi-cover [36, chapter 13.2] as we do not allow to purchase multiple copies of a set M. For general functions g we only assume that a set can be *used* by multiple players paying the corresponding bundle cost. However, each set can only contribute 1 towards the covering requirement of each of the contained elements.

Facility location games can be obtained as follows. We are given two sets T of terminals and F of facilities. The resources are facilities and connections, i.e., $R = F \cup (T \times F)$. A player p corresponds to a subset of terminals $T_p \subseteq T$. She strives to connect her terminals to facilities. As both the connections and the facilities are resources, they both generate a cost. We will refer to them as connection and opening costs, respectively. The constraint of a player p is satisfied if for each of her terminals $t \in T_p$ at least one connection (t, f) and the corresponding facility $f \in F$ are available to her. If connection costs satisfy the triangle inequality, we refer to these games as *metric* facility location games. Furthermore, a large variety of buy-at-bulk variants of cost minimization problems can be formulated as a game within this framework.

3 Cost and Complexity of Nash Equilibria

In this section we consider the behavior of prices of anarchy and stability in the game and the existence of NE. Our first result concerns the price of anarchy.

Theorem 1. *The price of anarchy in the buy-at-bulk cost sharing game is exactly $k/g(k)$.*

Proof. First, we prove the lower bound. Consider a vertex cover game on a star network, in which every player owns a single edge and each vertex v has fixed cost $c(v) = 1$. If every player contributes exactly the cost of the leaf node incident to her edge, a NE of cost k evolves. The optimum solution, however, consists of the center vertex v and has bundle cost $c(v, k) = g(k)$. This proves that the price of anarchy is at least $k/g(k)$.

For the upper bound consider any NE s of any buy-at-bulk cost sharing game with strategies s_p. In addition, let ρ_p^- be a set of resources for player p, which has minimum total fixed cost. Now consider a social optimum state s^*. Denote by ρ_p^* a subset of minimum fixed cost of the available resources of player p in s^*, which suffices to satisfy her constraint. As in terms of fixed cost the set ρ_p^- is optimal for p, it follows that

$$\sum_{r \in \rho_p^-} c(r) \leq \sum_{r \in \rho_p^*} c(r) \ . \tag{1}$$

The concavity of g ensures that with increasing demands for resources in ρ_p^-, the cost to be paid for player p can only decrease. Hence, it becomes ever more attractive for p to deviate to a strategy, which contributes only to ρ_p^-. However, as s is a NE, the fixed cost of ρ_p^- is an upper bound on current total contribution of p in s, because player p can always deviate and purchase all resources in ρ_p^- by herself:

$$\sum_{r \in R} s_p(r) \leq \sum_{r \in \rho_p^-} c(r) \ .$$

Since s is a NE, the cost of the purchased resources must be fully paid for. Using the bound from (1) we get

$$\sum_{p \in [k]} \sum_{r \in R} s_p(r) \leq \sum_{p \in [k]} \sum_{r \in \rho_p^-} c(r) \leq \sum_{p \in [k]} \sum_{r \in \rho_p^*} c(r). \tag{2}$$

Consider the following procedure of constructing a lower bound on the cost of the social optimum solution. Iteratively add players and the cost of their available resources ρ_p^* to the solution. The presence of the i-th player on ρ_i^* adds at least a cost $\mu(i) \sum_{r \in \rho_i^*} c(r)$ to the cost of s^*. As μ is monotonic decreasing, we can lower bound $c(s^*)$ by

$$\sum_{i=1}^{k} \mu(i) \sum_{r \in \rho_i^*} c(r) \leq c(s^*). \tag{3}$$

Note that the cost of the resources is determined by the final set Q_r, and this is independent of the ordering in which players are considered. Hence, the value of this lower bound is the same for any ordering of the players chosen. By making $k-1$ cyclic rotations of an initial ordering of players, we ensure that each player appears at each position i exactly once. Adding all resulting inequalities (3) we get

$$\sum_{p \in [k]} \sum_{i=1}^{k} \mu(i) \sum_{r \in \rho_p^*} c(r) = g(k) \sum_{p \in [k]} \sum_{r \in \rho_p^*} c(r) \leq k c(s^*).$$

Together with (2) this yields

$$c(s) = \sum_{p \in [k]} \sum_{r \in R} s_p(r) \leq \sum_{p \in [k]} \sum_{r \in \rho_p^*} c(r) \leq \frac{k}{g(k)} \cdot c(s^*),$$

which proves the theorem. □

In fact, the lower bound follows similarly for all classes of games that have been considered in the literature and have the free disposal property. If for a game $g(k) = k$, the game exhibits a decomposition property that allows for optimal NE. The previous theorem states that every NE is a social optimum. In fact, a similar argumentation yields the reverse statement, i.e., for linear g an optimum NE is guaranteed to exist.

Corollary 1. *If $g(k) = k$, then each Nash equilibrium is a social optimum state, and for each optimum solution to the underlying buy-at-bulk problem there is a Nash equilibrium purchasing it.*

This observation highlights the main effect of g. The price of anarchy quantifies the severity of coordination failure in a game. Steeper functions g, however, yield smaller savings for a player from payments made by others. Intuitively, this decouples player incentives and leads to the extreme case with $g(k) = k$, in which no player can expect savings through contributions of others. In this case everyone optimizes "in her own world" and has to pay fully for every resource obtained. Players do not need to coordinate anymore, and the result is the absence of coordination failures. On the other hand, if $g(k) = 1$ for $k \geq 1$, the coordination failure can be severe and the price of anarchy as large as k. The theorem yields an exact characterization for all intermediate ranges of g.

Once g is sublinear, then players start profiting in larger coalitions. In this case, we show that for a vertex cover game with sufficiently large number of players there is no NE.

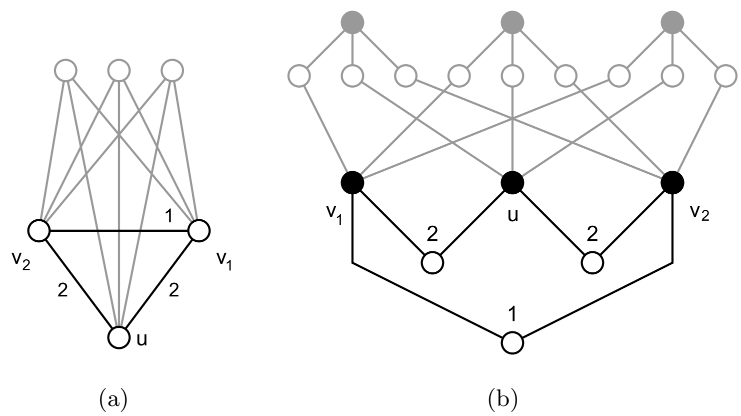

Figure 1: (a) Vertex cover game without a NE. Edge labels indicate player ownership. Gray parts are introduced when considering auxiliary players to deal with arbitrary values of k_0. (b) Transformation into a facility location game. Filled vertices are facilities, empty vertices are terminals. Labels of terminals indicate player ownership. Gray parts are introduced when considering auxiliary players.

Lemma 1. *If $g(i) = i$ for $i \leq k_0$ and $g(i) < i$ for $i > k_0$, then for any $k > k_0$ there is a vertex cover game with k players without a Nash equilibrium.*

Proof. Consider the game in Figure 1(a). We first prove the lemma for $k_0 = 1$ and then describe how to adjust the construction to any k_0. Player 1 has a single edge between vertices v_1 and v_2, player 2 has two edges connecting u to v_1 and v_2. The fixed cost $c(v_1) = c(v_2) = 1$, and, with foresight, $c(u) = 1 + x$ for any $1 > x > \mu(2)$. Now consider the possible covers that players can use to satisfy their covering requirement. Let player 1 use v_1 and player 2 use u. In this case player 2 can use v_1 and v_2 with cheaper payments, because $x > \mu(2)$. Now consider the case player 2 uses v_1 and v_2. Then she can contribute payments of at most $1 + x$. Player 1 could switch to use v_2, so she can at most contribute $\mu(2)$. Together this yields $1+x+\mu(2) < c(v_1)+c(v_2, 2) = 2+\mu(2)$, because $x < 1$. The case for player 1 choosing v_2 is symmetric, hence there is no NE in this game.

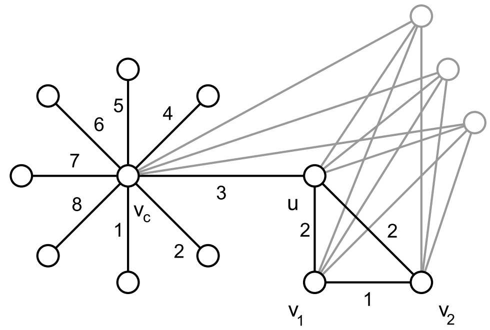

Figure 2: Vertex cover game with a price of stability of $\Theta\left(\frac{k}{g(k)}\right)$. Edge labels indicate player ownership. Gray parts are introduced when considering auxiliary players to deal with arbitrary values of k_0

For a game with $k_0 > 1$ we simply add $k_0 - 1$ auxiliary players in the following way. Each auxiliary player p owns edges of a star with a center w_p of cost $c(w_p) = 5$. The leaves are the vertices v_1, v_2, and u. Now we set $c(v_1) = c(v_2) = 1$ and $c(u) = 1 + x$ with $1 > x > \mu(k_0 + 1)$. Hence, in a NE every auxiliary player will use all three leaf vertices, in this way "boosting" the instance into the desired range where g becomes sublinear. Note, however, that the $k_0 - 1$ players together will never contribute more than $g(k_0 - 1)c(u) = (k_0 - 1)c(u)$ to vertex u (similarly for v_1 and v_2). This allows to rework the analysis above. Let player 1 use v_1 and player 2 use u. Even when joining the $k_0 - 1$ auxiliary players on u, player 2 still has to contribute $c(u)$ to get availability of u, because $g(k_0)c(u) = k_0 c(u)$ and the other players pay at most $(k_0 - 1)c(u)$. In this case, it is better to join the $k_0 - 1$ auxiliary players and player 1 on v_1 and v_2, because $x > \mu(k_0 + 1)$. Now consider the case player 2 uses v_1 and v_2. Then she can contribute payments of at most $1 + x$. Player 1 could switch to use v_2, where the cheapest way to obtain availability is to join all k_0 other players, yielding a required contribution of at most $\mu(k_0 + 1)$. Hence, the maximum payment of all players is $(k_0 - 1)(c(v_1) + c(v_2)) + (1 + x) + \mu(k_0 + 1) = 2k_0 - 1 + x + \mu(k_0 + 1)$, whereas the cheapest way to guarantee all availabilities yields a cost of $c(v_1, k_0) + c(v_2, k_0 + 1) = 2k_0 + \mu(k_0 + 1)$. This is a contradiction, because $x < 1$. The case for player 1 choosing v_2 is symmetric, hence there is no NE in this game. Finally, note that for a game with $k > k_0 + 1$ players we can simply add players that posses edges, which do not interfere with the described game. $\square$

This game can be embedded into a game that results in a price of stability in $\Theta\left(\frac{k}{g(k)}\right)$.

Theorem 2. *For vertex cover games the price of stability is in* $\Theta\left(\frac{k}{g(k)}\right)$.

Proof. Consider the game in Figure 2. It is a combination of the triangle game of Figure 1(a) and the star yielding maximum price of anarchy. Suppose every leaf vertex of the star and the star center v_c have constant fixed cost of $1 + \mu(k_0 + 1)$. The fixed costs of v_1 and v_2 are 1, for u it is $2 > c(u) > 1 + \mu(k_0 + 1)$. There are $k_0 - 1$ auxiliary players $k - k_0 + 2, \ldots, k$ that strive to cover a star centered at an additional vertex w_p. The cost $c(w_p)$ is prohibitively high, say $c(w_p) = 10$. Similar to the previous proof, in any NE the auxiliary players will choose to make v_c, v_1, v_2, and u available, while their total contribution to each vertex is at most $k_0 - 1$ times the fixed cost, e.g., for u at most $(k_0 - 1)c(u)$. So the contribution of these players boosts the game to a range where g becomes sublinear. Now suppose there is a NE, in which player 3 does not contribute to u. This leaves players 1, 2 and the auxiliary players with exactly the same game, for which we showed in the previous lemma that there is no NE.

In particular, the only possibility to stabilize players 1 and 2 is to let player 3 contribute sufficiently to vertex u such that player 2 can lower his contribution to u, and this becomes a best response. Hence, in any NE, player 3 contributes to u, which means u must be available to her, but not v_c. Now consider a NE s, in which v_c is available to at least one of the players $1, 2, 4, \ldots, k - k_0 + 1$. Then the total contribution of all players towards v_c is at least $\sum_{p \in [k]} s_p(v_c) \geq c(v_c, k_0)$. As v_c is not available to player 3, we have $s_3(v_c) = 0$. In turn, this implies player 3 pays $s_3(u) \leq c(v_c)\mu(k_0 + 1) = (1 + \mu(k_0 + 1))\mu(k_0 + 1) < c(u)\mu(k_0 + 1)$, because otherwise it would be cheaper to join a coalition purchasing v_c. Now consider player 2. Note that in any NE there are at most $k_0 + 1$ players that have an incentive to pay for u - the auxiliary players and players 2 and 3. Given the previous payment restrictions, we have $\sum_{p \in [k], p \neq 2} s_p(u) \leq c(u, k_0 + 1) - c(u)$. Thus, player 2 must invest at least $c(u) > 1 + \mu(k_0 + 1)$ to make vertex u available. This implies that she still has an incentive to join the auxiliary players and player 1 on the vertices v_1 and v_2 for a total cost of $1 + \mu(k_0 + 1)$. This is a contradiction to s being a NE.

Hence, in any NE none of the players $1, 2, 4, \ldots, k - k_0 + 1$ has the star center v_c available, which in turn means they all purchase their corresponding leaf vertices completely. We show that in this case a NE s can be obtained. The cheapest way for player 3 to obtain v_c would be to join the auxiliary players (or equivalently purchase v_c by herself), which would mean to pay $c(v_c)$. Hence, we can assume $s_3(u) = c(v_c)$, and the total contribution of the auxiliary players and player 3 to u is $\sum_{p \in [k], p \neq 2} s_p(u) = c(u, k_0 - 1) + c(v_c) > c(u, k_0 + 1) - 1 - \mu(k_0 + 1)$. Thus, player 2 can contribute $s_2(u) < 1 + \mu(k_0 + 1)$, and the bundle cost $c(u, k_0 + 1)$ will be paid for by auxiliary players and players 2 and 3. Player 1 can join the auxiliary players and purchase either v_1 or v_2, e.g., $s_1(v_1) = 1$. This gives no player an incentive to switch his strategy.

The fact that a NE evolves only when the players $1, 2, 4, \ldots, k - k_0 + 1$ purchase their corresponding leaf vertices implies that every NE has cost of at least $(1 + \mu(k_0 + 1))k + 1 + (\mu(k_0 + 1))^2 + \mu(k_0 + 1) + 3(k_0 - 1)$. Note that there is at least one such NE. In the social optimum, however, every incident player is covered by v_c yielding a cost of at most $(1 + \mu(k_0 + 1))g(k) + 2 + \mu(k_0 + 1) + 4(k_0 - 1)$. For fixed g, parameter k_0 is a constant, and the ratio grows with $k/g(k)$. □

The non-existence raises the question about the sets of games with and without NE. Can the players efficiently determine whether they could possibly agree upon a state or they are doomed to cycle? We answer this question in the negative. In particular, we derive a construction to show that given any fixed, non-linear function g, there is a class of games with sufficiently many players, in which determining existence of a NE is NP-hard.

Theorem 3. *Given any non-linear function g, for which $g(i) = i$ for $i \leq k_0$ and $g(i) < i$ for $k > k_0$, then for each $k > k_0$ there is a class of vertex cover games with g and k players, for which it is* NP*-hard to determine the existence of a Nash equilibrium.*

Proof. We again restrict to the case $k_0 = 1$ and then extend it to arbitrary k_0. Our construction reworks a proof for regular vertex cover games [9] and extends it to any sublinear function g. Consider the extended triangle game in Figure 3. Assign fixed costs of $c(u_1) = c(u_2) = 1$,

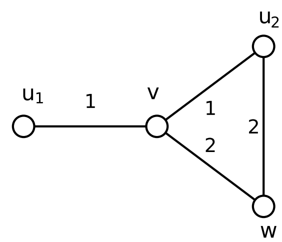

Figure 3: Extended triangle gadget, edge labels indicate player ownership. Vertex costs are defined in a way such that the game allows no NE. If a variable player purchases the fixed cost of u_1, or a clause player contributes a sufficiently large amount to the cost of w, a NE is possible.

$c(v) = 1.5 + (\mu(2)/2)$, and $c(w) = 1.25(1 + \mu(2)) + (\mu(2)^2/2)$. In similar way as for the game of Figure 1(a) we can argue that the game has no NE. This game is used in a reduction from 3SAT as follows. For each instance of 3SAT we introduce a *variable player* for each variable and a *clause player* for each clause. Each variable player strives to cover two star networks, a *true star* and a *false star*. The true (false) star has a number of leaves that equals the number of non-negated (negated) appearances of the variable in the clauses. The star centers are connected by an additional edge. They have a fixed cost that corresponds to the number of leaves. All leaves have fixed cost 1.

At each leaf of a variable gadget we attach a different extended triangle in the way that vertex u_1 of the triangle becomes the leaf vertex of the star. There are two *gadget players* who own the edges of all the extended triangles. Finally, for each clause we introduce a *clause player*. She owns a star with center vertex w_c. With $\nu = c(w)(1 + \mu(2)) - 1 - c(v)\mu(2)$ the cost of $c(w_c) = 2\nu$. This center is connected by exactly three edges to three different extended triangles. In particular, we attach the edges at vertex w of the extended triangles. If a variable occurs non-negated (negated) in the clause, we connect to an extended triangle attached to the true (false) star of the variable gadget. Note that there are sufficiently many gadgets installed such that in the final instance each extended triangle is connected to exactly one variable gadget and one clause gadget.

First suppose the 3SAT instance has a satisfying assignment. If a variable is set to true (false) in this assignment, we assign the variable player to pay the fixed cost for the center of the false (true) and all the leaves of the true (false) star. In this way the extended triangles for the true (false) star allow a stable cost sharing, in which player 1 pays the additional cost of $\mu(2)c(u_1)$ and $c(u_2)$ to make u_1 and u_2 available. Player 2 pays $c(w)$ to w. These are both best responses as $1 + \mu(2) < c(v)$ and $c(w) < \mu(2) + c(v)$. Thus, all those triangle gadgets are stabilized. As the assignment is satisfying, this implies that for each clause there is at least one stabilized triangle gadget. The remaining gadgets are stabilized as follows. The cost of the center vertex w_c of a clause is large enough to allow each clause player to pay a sufficiently large share of the cost of the vertices w for two remaining gadgets. In particular, the clause player can pay $\nu = c(w)(1+\mu(2)) - 1 - c(v)\mu(2)$ for w of each of the remaining gadgets. This means that player 2 must contribute only a cost of $1 + c(v)\mu(2)$ to w, and then w is available to both the clause player and player 2. Player 1 can pay $c(v)$ to v. This obviously allows player 2 to play a best response by contributing the remaining cost to vertex w, which in turn allows player 1 to stick to v in these gadgets. A NE evolves.

Now suppose there is a NE. The only way that extended triangles are stabilized is through contribution of variable players at u or clause players at w. The cost of the center vertex of a clause gadget allows the clause player to contribute at most 2ν to all vertices w of incident gadgets. This budget is sufficient to stabilize *at most two* extended triangles. Stabilization occurs only when the contribution is at least ν, because in this case the payment of player 2 to w is at most the cost of the optimal deviation to v and u_2. Hence, there can be at most two triangles which are stabilized by payments of the clause player. The remaining triangles must be stabilized with contributions of the variable players. Observe that a variable player can either purchase both star centers, or she can contribute to the set of leaves from *at most one*, the true or the false star. In a NE for each clause at least one variable player must stabilize an extended triangle by contributing to the corresponding leaf vertex. This induces a decision for each variable player and translates directly into a satisfying assignment for the variables.

To reduce the number of players we can merge the variable players with player 1 from the extended triangles into one global player. In similar fashion, we can merge the clause players with player 2 into another global player. This does not alter the incentives, because the edges owned by these players are in disconnected parts of the graph, respectively. Each global player will still have the same preferences and payoffs as the set of players she accumulates. This proves the theorem for the case of $k_0 = 1$.

For larger values of k_0 we introduce auxiliary players as before. For each auxiliary player we attach a vertex w_p with a prohibitively large cost such as $c(w_p) = 100nm$, where n and m denote the number of variables and clauses in the 3SAT instance, respectively. Each of the $k_0 - 1$ auxiliary players then owns a star with w_p as center and all vertices from the orignal game as leaves. In this way, in every NE all auxiliary players always strive to make the complete instance available instead of w_p. Each such player is willing to contribute the fixed cost to every vertex from the original gadget. On the basis of this contribution the arguments for the remaining players can be reworked as above. Note that k_0 is considered to be a fixed constant that does not grow with the 3SAT instance. The cost of w_p does grow, however, only logarithmically in representation. Hence, the reduction is polynomial and the theorem follows. □

3.1 Extension to Facility Location Games

Note that for every vertex cover game there is a metric facility location game that is equivalent in terms of the structure of NE. We replace each edge $e = (u, v)$ by a terminal t_e and two connections (t_e, u) and (t_e, v) of connection cost $c_{max} = \max_{v \in V} c(v)$. This creates the set of terminals. The former set of vertices becomes the set of facilities. For the remaining connections between facilities and terminals we assume a cost given by the shortest path metric, i.e. these costs are at least $3c_{max}$ (see Figure 1(b) for an example). In a NE of this game every player will purchase completely the corresponding connections for her terminals. In addition, no player will purchase a connection of cost larger than $2c_{max}$, in particular, no edges with costs determined by the shortest path metric will be bought in a NE. A NE for the facility location game provides a NE for the corresponding vertex cover game and vice versa.

Corollary 2. *If $g(i) = i$ for $i \leq k_0$ and $g(i) < i$ for $i > k_0$, then for any $k > k_0$ there is a metric facility location game for k players without a Nash equilibrium.*

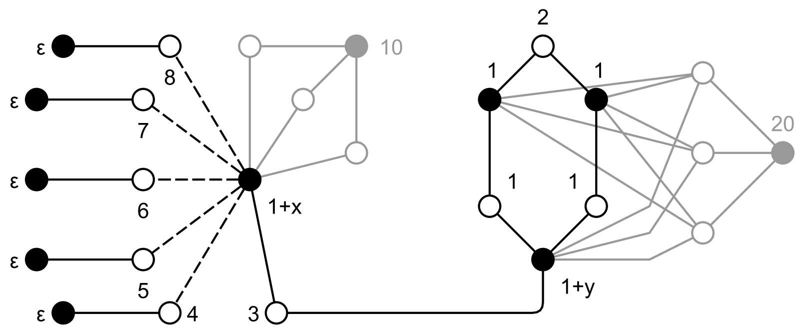

Figure 4: Metric facility location game with a price of stability of $\Theta\left(\frac{k}{g(k)}\right)$. Terminals are depicted as empty vertices, facilities are filled vertices. Terminal labels indicate player ownership, facility labels specify opening costs. In the depicted example $k_0 = 4$, $x = \mu(5)$, and $x < y < 1$. Gray parts are introduced by $2(k_0 - 1) = 6$ auxiliary players, and each such player holds exactly one terminal. Solid black edges have cost $1 + x$, dashed black edges cost $\epsilon > 0$, and gray edges cost 3. All other connection costs are given by the shortest path metric.

Corollary 3. *If $g(i) = i$ for $i \leq k_0$ and $g(i) < i$ for $i > k_0$, then for each $k > k_0$ there is a class of metric facility location games with g and k players, for which it is* NP*-hard to determine the existence of a Nash equilibrium.*

While these corollaries can be derived directly using the transformation, the next corollary requires slighlty more careful adjustment.

Corollary 4. *For metric facility location games the price of stability is in $\Theta\left(\frac{k}{g(k)}\right)$.*

The upper bound follows from Theorem 1. A construction for the lower bound is depicted in Figure 4. The bound can be established with similar arguments as for Theorem 2. We assume that $1/k > \epsilon > 0$ is a small constant and use two sets consisting of $k_0 - 1$ auxiliary players that boost the facility costs into the region where g becomes sublinear. If the center facility of cost $1 + x$ is unavailable to all regular players, player 3 can contribute sufficiently to the side gadget such that players 1 and 2 stabilize. In this case a NE of cost in $\Omega(k)$ evolves. If the center facility is available to at least one player, then with $\mu(k_0 + 1) < y < 1$ the contribution of player 3 is not sufficient to stabilize players 1 and 2 in the side gadget. In this case there can be no NE. However, if the center facility is available to all incident regular players, a solution of cost $O(g(k))$ can be obtained. This establishes the lower bound.

4 Single Element Players

The previous section clarified that vertex cover games might have no NE if there are players that own two edges. In this section we consider singleton set multi-cover games in which each player has only a single element. For these games a NE always exists and can be found in polynomial time.

Theorem 4. *Algorithm 1 returns an exact Nash equilibrium for singleton set multi-cover games in polynomial time.*

Algorithm 1: Exact NE for singleton set multi-cover games

$d_M \leftarrow 1$ for all sets M
Construct $G_s = (\mathcal{M}, A)$ with $(M_1, M_2) \in A$ iff $M_1 \cap M_2 \neq \emptyset$ and $c(M_1) \cdot \mu(d_{M_1}) < c(M_2) \cdot \mu(d_{M_2})$
while *there are remaining players* **do**
  for *every remaining player p* **do**
    if *element e of p is included in exactly $b(e)$ sets $\mathcal{M}_e$* **then**
      Assign p to contribute $s_p(M) \leftarrow c(M) \cdot \mu(d_M)$ to all these sets $M \in \mathcal{M}_e$
      Increase $d_M \leftarrow d_M + 1$ and drop p from consideration
      Adjust the arc set of G_s for the new values of d_M
  Find a sink in G_s and drop the corresponding set from consideration

Proof. Clearly, Algorithm 1 can be implemented to run in polynomial time. A set M_1 *dominates* a set M_2 iff there is a player who prefers M_1 over M_2 with the bundle costs given with d_{M_1} and d_{M_2}. The algorithm constructs and maintains a directed acyclic graph G_s, which contains a directed edge between sets M_1 and M_2 iff M_1 dominates M_2. A set M that is dropped from consideration represents a sink in G_s. Then for each remaining player with $e \in M$ it is dominated by all remaining sets that contain her element. None of these players will contibute to M, as they have a cheaper alternative to cover their element. As no contribution will be assigned to M after it has been dropped, no player wants to contribute to sets that were dropped before she was dropped. When player p gets dropped, she is left with the set $\mathcal{M}_e$ of exactly $b(e)$ sets to cover e. This reveals that she cannot profit from contributing to any other sets that contain her element. This is also true for the sets in $\mathcal{M}_e$. Consider another player q, who is assigned to contribute to $M \in \mathcal{M}_e$ after p has been dropped. q will only pay a cost representing the concave increase in bundle cost with p already counted towards d_M. Hence, there is no subset of players whose payments allow p to lower her contribution to $\mathcal{M}_e$. Thus, each player plays a best response. This proves the theorem. ☐

Unfortunately, the proposed algorithm can compute worst-case NE. Reconsider the singleton vertex cover game with a star network used to obtain a lower bound for the price of anarchy. Suppose the fixed cost for the leaf vertices is 1, and for the center vertex it is $1 + \epsilon$. Algorithm 1 will assign each player to purchase the leaf vertex incident to her edge. This obviously yields a NE, whose cost is a factor arbitrarily close to $\frac{k}{g(k)}$ worse than $c(s^*)$. In contrast, we show that there are optimal NE in every singleton set multi-cover game.

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

Proof. Consider an arbitrary feasible solution $\mathcal{R}$ for the underlying optimization problem. $\mathcal{R}$ must yield for every player at least $b(e)$ available sets that include her element e. Consider one of these sets M. For this set consider the player set Q_M, which includes every player p whose element $e \in M$ and who has exactly $b(e)$ available sets. If $Q_M = \emptyset$, the state can be improved by dropping all contributions to M without hurting feasibility of the player constraints. Otherwise, consider each $p \in Q_M$ individually, and suppose that M is unavailable. In particular, suppose that all contributions of all players to M are 0. Now player p has only $b(e) - 1$ sets available and is willing to make another set available to her. We consider the cheapest contribution that she needs to do this. The new set could be a different set $M' \neq M$, which she does not yet have available. It could also be the set M, for which she would need to pay the fixed cost $c(M)$, as we assume that nobody else contributes to it. We name the minimum contribution the *budget* of player p, denote it by $c_M(p)$, and define formally

$$c_M(p) = \min(c(M), \min_{M' \neq M} \{c(M') \cdot \mu(|Q_{M'}| + 1)\}) \ . \tag{4}$$

Note that we might not yet have determined a cost sharing for sets M'. However, we assume temporarily that the players in $Q_{M'}$ will pay for it. Then player p must at most contribute the cost that is incurred due to his additional availability. In particular, this means $c_M(p) \leq c(M')$. If player $p \notin Q_M$, we set $c_M(p) = 0$. Note that this holds in particular for players that are overcovered in $\mathcal{R}$.

Intuitively, the budget $c_M(p)$ is the maximum a player p is willing to pay for M without an incentive to deviate. We now strive to determine the coalition of players that is willing to purchase the largest bundle of M without having an incentive to deviate. For this task we order the players in non-increasing order of budgets, i.e., we let p_i be the player with the i-th largest budget, for $i = 1, \ldots, |Q_M|$. We assign a contribution for each p_i to M in this order and ensure that it never exceeds $c_M(p_i)$. In addition, we strive to obtain a largest set of players $Q^{max} \subseteq Q_M$ that can pay for a bundle cost and has no incentive to remove payments. More formally, we define a payment maximal set of players Q^{max} as follows.

Definition 1. *The set $Q^{max} \subseteq Q_M$ is called* payment maximal *if it is a largest set of players such that there exist payments s^{max} with s_p^{max} for each $p \in Q^{max}$ satisfying*

- $s_p^{max} \leq c_M(p)$ *for each* $p \in Q^{max}$,
- $\sum_{p \in Q} s_p^{max} \leq c(M, |Q|)$ *for each* $Q \subseteq Q^{max}$, *and*
- $\sum_{p \in Q^{max}} s_p^{max} = c(M, |Q^{max}|)$.

We denote $k^{max} = |Q^{max}|$.

To obtain such a set and the corresponding payments, we use an iterative procedure that considers players in the specified order and assigns

$$s_{p_i}(M) = \min\left(c_M(p_i), c(M, i) - \sum_{j=1}^{i-1} s_{p_j}(M)\right) \ . \tag{5}$$

Every player contributes at most the cost $c_M(p_i)$. In addition, the i-th player contributes at most the remaining cost share that is needed to purchase the bundle cost $c(M, i)$ using contributions of all $i - 1$ previously considered players. The procedure determines exactly one largest set Q' of k' players that achieves to pay for $c(M, k')$. The players with the largest possible contributions are considered earliest in the procedure and thus are also assigned to pay the largest shares. Note that the cost shares assigned are monotonically decreasing with increasing i. We next show that Q' represents a payment maximal set.

Lemma 2. *Q' is payment maximal and the payments according to Equation* (5) *represent a set of payments s^{max}.*

Proof. Consider a payment maximal set Q^{max} and contributions s^{max}. Without loss of generality we can assume that the k^{max} players in Q^{max} are the ones with largest budgets $c_M(p)$. Let us order the players p_i for $i = 1, \ldots, k^{max}$ with non-increasing budgets. Again, without loss of generality we can redistribute the assigned cost shares for the players such that they become non-increasing for non-increaseing budgets, i.e., that $s^{max}_{p_i} \leq s^{max}_{p_{i+1}}$ for any $i \in [1, k^{max} - 1]$. In particular, this yields $Q' \subseteq Q^{max}$.

We know that every subset of i players of Q^{max} contributes at most $c(M, i)$. If this holds for the players with largest payments $p_1, \ldots, p_i$, it continues to hold for all players. In particular, this shows that the property is also satisfied for our payments $s_p(M)$. Hence, the previous observations show that Q' and $s_p(M)$ satisfy all the conditions in the definition of payment maximal except for the maximality of k'.

Finally, we show that indeed $k' = k^{max}$ by induction on the the ordering in the assignment. In particular, our invariant is that for any $i \leq k^{max}$

$$\sum_{j=1}^{i} s_{p_j}(M) \geq \sum_{j=1}^{i} s^{max}_{p_j} \quad ,$$

i.e., our routine extracts the maximum payment possible under the conditions of payment maximal, which proves the result. Suppose that $k^{max} = 1$, then by previous observations $s_{p_1}(M) = s^{max}_{p_1}$. Consider iteration i, in which the payment $s_{p_i}(M)$ is determined and assume the invariant holds up to $i - 1$. If both $s^{max}_{p_i} = s_{p_i}(M) = c_M(p_i)$, the invariant is preserved. Otherwise, if $s_{p_i}(M) < c_M(p_i)$, then we know that

$$\sum_{j=1}^{i} s_{p_j}(M) = c(M, i) \geq \sum_{j=1}^{i} s^{max}_{p_j}$$

by definition. This proves the lemma. □

Hence, our procedure allows us to find Q^{max}. If $Q_M = Q^{max}$, then we call a set *stabilized*. Otherwise, we know

$$\sum_{p \in Q_M - Q^{max}} c_M(p) < c(M, |Q_M|) - c(M, k^{max}) \quad , \tag{6}$$

so the budgets of players of $Q_M - Q^{max}$ do not allow them to pay the additional cost that is generated by availability of M for them. Suppose each of these players switches, i.e., it is dropped from Q_M and instead is included in $Q_{M'}$ for the set M' which generated its budget $c_M(p)$. Then a new solution $\mathcal{R}'$ to the underlying covering problem evolves, in which again all constraints are satisfied. In addition, for the change in cost we derive with Equation (6) that

$$c(\mathcal{R}') - c(\mathcal{R}) \leq \sum_{p \in Q_M - Q^{max}} c_M(p) - (c(M, |Q_M|) - c(M, k^{max})) < 0 \quad , \tag{7}$$

the new solution $\mathcal{R}'$ is cheaper. If there is an unstabilized set, then there is another solution with smaller social cost. Hence, an optimum solution $\mathcal{R}^*$ consists only of stabilized sets. Finally, for stabilized sets we can use the assignment in Equation (5) to derive a cost sharing. If a player p would in a deviation strive for availability of a different set M, she must indeed contribute at least the difference in bundle cost that her availability generates. The budgets $c_M(p)$ for the sets are defined under the assumption, and they are used to bound the contributions in s_p. In addition, by the adaptive assignment of payments for s based on the ordering, she must pay at least $s_p(M)$ for any bundle that makes a set M available to her. Thus, s is a NE and a cost sharing of $\mathcal{R}^*$. This proves the theorem. $\square$

Computing optimal Nash equilibria is clearly NP-hard. However, the improvement step by removal of players from unstabilized sets allows to construct a local search procedure to obtain (α, β)-approximate NE that are near-stable and near-optimal. Recall that the stability ratio α specifies the relative incentive to deviate and β the approximation factor of the resulting solution. Using an efficient approximation algorithm [36, chapter 13.2] to compute a starting solution and the fact that each player holds exactly one element, the next theorem shows that we can obtain $(1+\epsilon, O(\log k))$-NE in polynomial time.

Theorem 6. *For any constant $\epsilon > 0$, a $(1+\epsilon, \beta)$-approximate Nash equilibrium in singleton set multi-cover games can be obtained in polynomial time from any state representing a β-approximation to the optimum social cost.*

Proof. The proof relies on the local improvement step outlined before. We use a technique from [6] for a minimum cost improvement to bound the number of improvement steps by a polynomial in n, k and ϵ^{-1}. In particular, we again start with a solution $\mathcal{R}$ to the underlying covering problem, which represents a β-approximation of the optimum social cost. We denote the cost of this solution by $c(\mathcal{R})$. For every set that is in the cover, we reduce the cost of the respective bundle that is to be bought by $\kappa = \frac{\epsilon c(\mathcal{R})}{(1+\epsilon)n\beta}$. We appropriately adjust the cost of smaller bundles for this set in order to keep the ordering of bundle cost increasing with the size. In particular, we derive an adjusted bundle cost, which is

$$c'(M, i) = \min(c(M, i),\ c(M, |Q_M|) - \kappa) \qquad \text{for } i = 1, \ldots, |Q_M| \ .$$

Hence, when considering paying for the current bundle Q_M of M the players must pay for all but a cost of κ. For computing the budgets $c_M(p)$, however, we stick to the original definition in Equation (4) with cost function c. In this way a set is stabilized even if a cost of κ of the set remains unpaid. In addition, if a set is not stabilized, the budgets guarantee a minimum improvement, i.e., Equation (7) now reads

$$c(\mathcal{R}') - c(\mathcal{R}) \leq \sum_{p \in Q_M - Q^{max}} c_M(p) - (c'(M, |Q_M|) - c(M, k^{max})) \leq -\kappa < 0 \ .$$

The solution cost decreases by at least κ. This yields a maximum of at most $\frac{(1+\epsilon)n\beta}{\epsilon}$ improvement steps and proves polynomial running time.

Now suppose the algorithm has run to completion. We denote by n' the number of sets in the final solution $\mathcal{R}^*$. There remains an unpaid cost of κ for each set. These costs of $n'\kappa$ are collectively paid for. Each player p contributes a share that corresponds to the relative amount of investment that she was assigned in s_p. Thus, we simply scale up the assigned total contributions of players until the extra cost of at most $n'\kappa$ is paid for. Using $|s_p| = \sum_{M \in \mathcal{M}} s_p(M)$ the additional payment of player p is $\delta_p = (n'\kappa) \cdot (|s_p| / \sum_q |s_q|)$. This extra contribution of all the players is allocated arbitrarily to pay for κ at each of the n' sets. A player p might contribute δ_p to sets that are not needed to cover her elements. However, even if p removes this payment, the ratio of improvement for her is still small enough. We denote by $c(\mathcal{R}')$ the cost of the final solution and derive

$$\delta_p = |s_p| \frac{\kappa n'}{c(\mathcal{R}') - \kappa} \leq \frac{\epsilon c(\mathcal{R}) |s_p|}{\beta(1+\epsilon)(1-\epsilon)c(\mathcal{R}')} \leq \epsilon |s_p| \ .$$

This resembles an argument in [6] and establishes that the amount added by scaling of payments allows a player to reduce her contribution by at most a factor of $(1+\epsilon)$. This proves the theorem. □

The arguments in the proofs of Theorems 5 and 6 require only straightforward modifications for the case of metric and non-metric facility location games. For instance, we can efficiently obtain $(1+\epsilon, 1.861)$-NE in metric facility location games [23]. We leave the details as an exercise for the interested reader.

Corollary 5. *For singleton facility location games the price of stability is 1. For any constant $\epsilon > 0$, a $(1+\epsilon, \beta)$-approximate Nash equilibrium can be obtained in polynomial time from any state representing a β-approximation to the social cost.*

5 Approximate Nash Equilibria

In this section we consider set cover games, which are set multi-cover games with $b(e) = 1$ for all elements $e \in E$. While the lower bounds shown for vertex cover games extend to this case, it is possible to obtain (f, f)-NE in polynomial time, in which $f = \max_{e \in E} |\{M \in \mathcal{M}, e \in M\}|$ denotes the maximum *frequency* of any element in the sets.

Algorithm 2: (f, f)-NE for set cover games

1 $s_p(M) \leftarrow 0$ for all players p and sets M
2 $\gamma_p(e) \leftarrow 0$ for all players p and elements e
3 **for** *every player* $p = 1, \ldots, k$ **do**
4   Set $c^p(M) = \min_Q \{c(M, |Q|+1) - \sum_{q \in Q} s_q(M)\}$ for $Q \subseteq [p-1]$ and all M
5   **while** *there is an uncovered element* $e \in E_p$ **do**
6   Let $\gamma_p(e) \leftarrow \min_{e \in M} c^p(M)$
7   Increase payments: $s_p(M) \leftarrow s_p(M) + \gamma_p(e)$ for all M with $e \in M$
8   Add all purchased sets to the cover
9   Reduce set costs: $c^p(M) \leftarrow c^p(M) - \gamma_p(e)$ for all M with $e \in M$

Theorem 7. *Algorithm 2 returns a (f, f)-approximate Nash equilibrium for set cover games in polynomial time.*

The algorithm is an adjustment of the primal-dual algorithm for minimum set cover (see for instance [36, chapter 15]). It can be implemented to run in polynomial time. In line 4 it determines the minimum cost player p has to contribute in order to make set M available for her. In particular, we take all previous contributions into account and determine a set of players $Q \cup p$, for which the missing contribution to the bundle cost is minimal. The set Q can naturally be restricted to subsets of $[p-1] = 1, \ldots, p-1$ of previouly considered players, because for all other players all contributions are still 0. We can start with $Q = \emptyset$ and add players $q < p$ in non-increasing order of the contributions $s_q(M)$. In this process the desired set Q with minimal missing contributions for the bundle cost will be found.

In Lemma 3 we show that the approximation ratio obtained by the algorithm is bounded by f. This has most likely been observed before, for completeness a proof can be found in the Appendix.

Lemma 3. *Algorithm 2 returns a state with approximation ratio f.*

Finally, we prove the bound on the stability ratio, and Theorem 7 follows.

Lemma 4. *Algorithm 2 returns a state with stability ratio f.*

Proof. We consider the p-th player after the execution of the algorithm and her best move taking into account the payments of all other players $q \neq p$. For that purpose, we consider for each set M the cost $c'(M) = \min_{Q \subset [k], p \notin Q} c(M, |Q| + 1) - \sum_{q \in Q} s_q(M)$. We have to show that the sum of the payments of player p is not greater than f times the cost of the cheapest set cover of E_p with respect to the costs c'. From the algorithm and the fact that bundle costs are concave we know that $s_p(M) \leq c'(M)$. Also from the algorithm, we know that for any set M that includes one or more elements of E_p, we have $s_p(M) = \sum_{e \in M \cap E_p} \gamma_p(e)$, so for any such M we have $\sum_{e \in M \cap E_p} \gamma_p(e) \leq c'(M)$. Now let us consider a minimum cost set cover $\mathcal{R}_p^*$ of E_p with respect to c'. We have:

$$\sum_{M \in \mathcal{R}_p^*} \sum_{e \in M \cap E_p} \gamma_p(e) \leq \sum_{M \in \mathcal{R}_p^*} c'(M) = c'(\mathcal{R}_p^*).$$

Since $\mathcal{R}_p^*$ is a set cover of E_p, the charge $\gamma_p(e)$ of each element e in E_p is counted at least once in the left-hand side above. Hence

$$\sum_{e \in E_p} \gamma_p(e) \leq \sum_{M \in \mathcal{R}_p^*} \sum_{e \in M \cap E_p} \gamma_p(e) \leq c'(\mathcal{R}_p^*).$$

Now we can conclude

$$\sum_{M \in \mathcal{M}} s_p(M) \leq f \sum_{e \in E_p} \gamma_p(e) \leq f c'(\mathcal{R}_p^*),$$

which proves the lemma and Theorem 7. □

In the special case of vertex cover ratios of $f = 2$ are tight even on regular vertex cover games [9]. In contrast to the analysis for prices of anarchy and stability the analysis of the algorithm cannot be strengthened to a ratio depending on g. The ratio can be as large as f for any function g in games without sets that contain elements of more than one player.

For linear g the well-known greedy algorithm [36, chapter 2] achieves logarithmic stability and approximation ratio simply by optimizing the cover independently for each player. In contrast, for regular set cover games with $g(i) = 1$ for all $i \geq 1$, the greedy algorithm yields an unbounded stability ratio [25, Lemma 4.5]. It is an interesting open problem to obtain a procedure with improved bounds for intermediate functions g.

References

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

[2] Matthew Andrews. Hardness of buy-at-bulk network design. In *Proc. 45th Symp. Foundations of Computer Science (FOCS)*, pages 115–124, 2004.

[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)*, 2009. To appear.

[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 and Yossi Azar. Buy-at-bulk network design. In *Proc. 38th Symp. Foundations of Computer Science (FOCS)*, pages 542–547, 1997.

[8] Vittorio Biló, Angelo Fanelli, Michele Flammini, and Luca Moscardelli. When ignorance helps: Graphical multicast cost sharing games. In *Proc. 33rd Intl. Symp. Mathematical Foundations of Computer Science (MFCS)*, pages 108–119, 2008.

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

[10] 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.

[11] 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.

[12] Chandra Chekuri, MohammadTaghi Hajiaghayi, Guy Kortarz, and Mohammad Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design. In *Proc. 47th Symp. Foundations of Computer Science (FOCS)*, pages 677–686, 2006.

[13] Chandra Chekuri, MohammadTaghi Hajiaghayi, Guy Kortsarz, and Mohammad Salavatipour. Approximation algorithms for node-weighted buy-at-bulk networks. In *Proc. 18th Symp. Discrete Algorithms (SODA)*, 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] Xiaotie Deng, Toshihide Ibaraki, Hiroshi Nagamochi, and Wenan Zhang. Totally balanced combinatorial optimization games. *Math. Prog.*, 87(3):441–452, 2000.

[17] 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.

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

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

[20] Angelo Fanelli, Michele Flammini, Giovanna Melideo, and Luca Moscardelli. Multicast transmissions in non-cooperative networks with a limited number of selfish moves. In *Proc. 31st Intl. Symp. Mathematical Foundations of Computer Science (MFCS)*, pages 363–374, 2006.

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

[22] Anupam Gupta, Amit Kumar, Martin Pál, and Tim Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *J. ACM*, 54(3):11, 2007.

[23] Mohammad Taghi Hajiaghayi, Mohammand Mahdian, and Vahab Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.

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

[25] Martin Hoefer. *Cost Sharing and Clustering under Distributed Competition*. PhD thesis, Lehrstuhl Algorithmik, Universität Konstanz, 2007.

[26] Martin Hoefer. Competitive cost sharing with economies of scale. In *Proc. 8th Latin American Theoretical Informatics Conference (LATIN)*, pages 339–349, 2008.

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

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

[29] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*, 50(6):795–824, 2003.

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

[31] Stefano Leonardi and Piotr Sankowsi. Network formation games with local coalitions. In *Proc. 26th Symp. Principles of Distributed Computing (PODC)*, pages 299–305, 2007.

[32] 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.

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

[34] 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.

[35] Zheng Sun, Xiang-Yang Li, Weizhao Wang, and Xiaowen Chu. Mechanism design for set cover games when elements are agents. In *Proc. 1st Intl. Conf. Algorithmic Applications in Management (AAIM)*, pages 360–369, 2005.

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

[37] 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.

Appendix

Proof of Lemma 3

Proof. We consider for each set M the number k_M of players that have some of their elements in M. If an element e is owned by more than one player, we introduce a separate copy of e for each player. Now we turn the orignal instance of the buy-at-bulk set cover problem into a *bundle instance*, in which we replace each set M by the set of all possible bundles, i.e., all possible 2^{k_M} subsets that are obtained by deleting a subset of the players and their elements. For each such bundle we set the corresponding bundle cost. The bundle instance is an instance of the regular set cover problem. Algorithm 2 is an adjustment of the primal-dual f-approximation algorithm for minimum set cover. If this algorithm is run directly on the bundle instance, we get a factor as large as $f \cdot 2^{k-1}$ as the dual payment for each element is offered towards all (now potentially exponentially many) bundles. However, if two bundles introduced by a single set M are bought, we can always feasibly lower the solution cost by purchasing the bundle corresponding to the union of the two. Hence, the optimum solution purchases at most one bundle for each set M. Thus, the bundle instance has the same optimum solution as the original instance. The payments $\gamma_p(e)$ computed by Algorithm 2 compose a feasible dual solution for the bundle instance, hence $\sum_{p\in[k]} \sum_{e\in E_p} \gamma_p(e)$ is a lower bound on the optimum solution for both the bundle and the orignal instance.

The property that for each set there is only one purchased bundle is also used by Algorithm 2. It restricts the attention only to one bundle using $c^p(M)$. In this way, it guarantees that the dual feasible payment $\gamma_p(e)$ is invested only at most f times for each player p and each element e. This allows to rework the analysis of Algorithm 2 on the regular set cover problem and to derive f as the approximation ratio. □