Competitive Cost Sharing with Economies of Scale

Martin Hoefer*

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

Abstract. We consider a general class of non-cooperative *buy-at-bulk cost sharing games*, in which k players must contribute to purchase a number of resources. The resources have costs and must be paid for to be available to players. Each player can specify payments and has a certain constraint on the number and types of resources that she needs to have available. She strives to fulfill this 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. The cost increase for larger sets of players is composed of a fixed price $c(r)$ for each resource r and a global concave capacity function g. This cost can be shared arbitrarily between players. We consider the quality and existence 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. Some of our results can be extended to games based on facility location problems.

1 Introduction

In this paper we consider a general class of non-cooperative *buy-at-bulk cost sharing games*, which can for instance be used to model crucial competitive cost sharing aspects of networks like the Internet, e.g. service installation, facility location or various network design problems. The formulation captures a realistic aspect of networks by including costs with economies of scale. In particular, we consider a game 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 certain cost. Each player picks as a strategy a function that specifies their 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 players. For each player there is a constraint on the number and types of resources that must be available for her. She

* Supported by DFG through German excellence cluster UMIC at RWTH Aachen

strives to fulfill this constraint with minimum total investment in her strategy. A resource becomes more expensive when it shall be available to a larger set of players. In particular, if resource r is available to a set of i players, the cost is $c(r,i) = c(r)g(i)$, in which $c(r)$ is a fixed cost and g is a non-decreasing and concave function, which is used for every resource r. A variety of problems, e.g. buy-at-bulk variants of set cover, facility location, and network design, can be turned into a game with the help of this model.

We first study our games with respect to the existence of pure-strategy exact Nash equilibria. We characterize prices of anarchy [17] and stability [2], which measure the cost of the worst and best Nash equilibria in terms of the cost of a socially optimum solution, respectively. We also consider a situation, in which a central institution with some means to influence agent behavior tries to induce a state that is as cheap and stable as possible. This poses a two-parameter optimization problem captured by the notion of (relative) (α, β)-approximate Nash equilibria. These are states, in which the equilibrium condition is relaxed by 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*. In accordance with previous work we consider properties of games, in which player constraints are equivalent to well-known covering and facility location problems. Our interest is to investigate the influence of the function g on the efficiency and computational complexity of exact and approximate Nash equilibria.

Related Work. There are a number of related game-theoretic models. Cooperative games have been studied quite intensively in the past (see [9, 11] and the references therein). In [9] 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 [11] 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. The authors in [10] presented strategyproof mechanisms for set cover and facility location games. For set cover games this work was extended [18, 20] 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 [19], and in [16] lower bounds on their budget-balance were provided. In contrast, our approach is an extension of non-cooperative games, which were first studied in [3] in a Steiner forest network design context. Recent work [5,14,15] presented extended 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 $\Theta(k)$. For singleton games, in which each player is interested only in a single element, however, optimal Nash equilibria exist. In [5, 14] we proved the applicability of primal-dual methods to derive cheap and stable approximate Nash equilibria. None of these previous models, however, considers the influence of different economies of scale.

Starting with [4] network design problems with economies of scale became a vivid area of research. 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, polylogarithmic approximation algorithms [6, 7] and logarithmic hardness results for general resource costs [1] were derived. For special cases, e.g. single-souce or rent-or-buy problems [12] there exist constant-factor approximation algorithms. This is also the case for unit-demand metric facility location [13].

Our Contribution. Buy-at-bulk investment 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 previous cost sharing games [3, 5, 14, 15], which in the following we will call *regular cost sharing games*. In regular games, only a fixed cost for each resource must be paid for to make it 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 e.g. in [2, 8] by fixing a Shapley-value cost sharing. In contrast, our model allows smaller groups of players to obtain the resource at cheaper costs. This creates a force on every player to contribute for availability. The severeness of this force depends on the number of players that request a resource and is dynamically adjusted by g. Some undesirable properties of the game like a high price of anarchy are directly affected by this, 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 upper bounds on approximate Nash equilibria for regular games can be extended to hold for buy-at-bulk games. There are (f, f)-approximate Nash equilibria for set cover games, where f is the maximum frequency of any element in the sets. 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 exact Nash equilibrium in polynomial time, which was not known before even for regular singleton games. A number of these results directly translate to a class of buy-at-bulk investment games for facility location. Due to space limitations some of the proofs are shortened or omitted.

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 $g : \mathbb{N} \rightarrow \mathbb{R}^0_+$, 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 is $c(r,i) = c(r)\,g(i)$. A strategy s_p of a player p is a function $s_p : R \rightarrow \mathbb{R}^0_+$ to specify her non-negative payment to each resource. A *state* is a vector $s = (s_1, \ldots, s_k)$ 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 is a subset $p \in Q \subset [k]$ of players such that they purchase the corresponding bundle cost, i.e. $\sum_{q \in Q} s_q(r) \geq c(r, |Q|)$. For a player 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, which has a covering aspect in the sense that it can never be violated by having *additional* resources available to the ones required. If ρ_p in the current state s does not fulfill the constraint, we assume that the player is penalized with a prohibitively large cost, i.e. for her *individual cost* $c_p(s) = +\infty$. 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 (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. A social optimum state minimizing social cost will be denoted s^* throughout.

In a NE and in a social optimum state s^* the available resources for each player satisfy her constraints. Also, in NE and s^* due to concavity of g, 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. In NE no subset of i players will contribute more than $c(r,i)$ to any resource r. The strategies exactly purchase the bundle cost $c(r, |Q|)$ 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 here 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 him, 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 were buy-at-bulk games with $g(i) = 1$ for all $i \geq 1$ [3, 5, 14, 15]. When referring to games in this paper - e.g. vertex cover games - we generally mean the buy-at-bulk version. It is explicitly mentioned when regular games are under consideration.

2.1 Covering and Facility Location

The definition allows a variety of games to be defined in this framework. A simple class 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, and there is a number $b(e) > 0$ for each $e \in E$. Player p is satisfied if for each $e \in E_p$ there are at least $b(e)$ sets available to her that include e.

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. In *metric* games the connection costs satisfy the triangle inequality.

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 hardness of finding 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 cost of the available resources of player p in s^*, which suffices to satisfy her constraint. It is obvious that for the fixed cost

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

The concavity of g ensures that with increasing demands for resources in ρ_p^-, the cost to be paid for by 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, i.e. $\sum_{r \in R} s_p(r) \leq \sum_{r \in \rho_p^-} c(r)$. 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 $\sum_{p \in [k]} \sum_{i=1}^{k} \mu(i) \sum_{r \in \rho_p^*} c(r) = g(k) \sum_{p \in [k]} \sum_{r \in \rho_p^*} c(r) \leq kc(s^*)$, and together with (2) this proves the theorem:

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

In fact, our proof bounds the price of anarchy for both, pure and mixed NE. 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. The reverse is also true, i.e. in this case there is always an optimum NE. However, once g is sublinear, then for a vertex cover game with sufficiently large number of players, there is no NE.

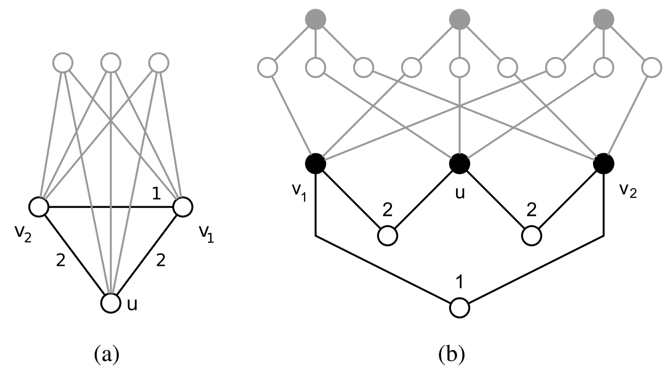

Fig. 1. (a) Vertex cover game without a NE. Edge labels indicate player ownership. Grey parts are introduced when considering auxiliary players to deal with arbitrary values of k_0. (b) Transformation into a facility location game. Filled vertices are facilities, empty vertices are terminals. Labels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

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.*

Consider the game in Figure 1(a) and $k_0 = 1$. Intuitively, whenever player 1 contributes the fixed cost to some vertex v_1 or v_2, player 2 is motivated to contribute to bundles of v_1 and v_2. In particular, she will purchase the fixed cost of the other vertex. This gives player 1 an incentive to remove payments, which gives player 2 an incentive to purchase vertex u. While this is not a formal argument, it can be verified that for each possible feasible solution no stable cost-sharing can be obtained. The transition to arbitrary k_0 uses additional $k_0 - 1$ *auxiliary players*. These players own a star with an expensive center and u, v_1, v_2 as leaves. They never contribute to the center and simply serve to "boost" the dynamics on the original game into the region, where the drop in function g occurs. Hence, as soon as players can profit from the investment of other players, they might not be able to agree upon a set of resources to purchase. Based on this observation we can 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. In addition, the price of stability can be as high as $\Theta\left(\frac{k}{g(k)}\right)$.

Theorem 2. *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.*

Theorem 3. *For vertex cover games the price of stability is in $\Theta\left(\frac{k}{g(k)}\right)$.*

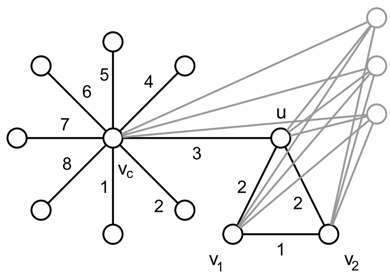

Fig. 2. Vertex cover game, in which the price of stability is $\Theta\left(\frac{k}{g(k)}\right)$. Edge labels indicate player ownership. Grey parts are introduced when considering auxiliary players to deal with arbitrary values of k_0

Proof. Consider the game in Figure 2. 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 cost 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 $k - k_0 + 2, \ldots, k$, and each has a star centered at an additional vertex w_p. The cost $c(w_p)$ is prohibitively high, so these players will boost the game to a range where g becomes sublinear. Now suppose there is at least one of the players $1, 2, 4, \ldots, k - k_0 + 1$, who strives to make v_c available to her. Then there are at least k_0 players, who pay a cost of $c(v_c, k_0)$ for v_c. Player 3 will contribute at most $c(v_c)\mu(k_0+1) = (1+\mu(k_0+1))\mu(k_0+1) < c(u)\mu(k_0+1)$ to u. Thus, player 2 has to invest at least $c(u)$ to make vertex u available. As previously noted there can be no NE in this case. Thus, none of the players $1, 2, 4, \ldots, k_0 + 1$ shall make star center v_c available to her. Then player 3 can contribute $c(v_c)$ to a bundle cost of vertex u. Player 2 can add less than $1 + \mu(k_0 + 1)$ to u, and together with the auxiliary players this purchases the bundle cost of $c(u, k_0+1)$. Note that player 1 sticks to purchasing one of v_1 and v_2, and the remaining edges of the star can be covered by purchasing the leaf vertices. A NE of cost at least $(1+\mu(k_0+1))k+1+(\mu(k_0+1))^2+\mu(k_0+1)+3(k_0-1)$ evolves. In the social optimum, however, all players $1,\ldots,k_0+1$ contribute to v_c yielding a state of cost 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 $k/g(k)$. □

Note that any vertex cover game can be translated easily to a metric facility location game, which 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. they are at least $3c_{max}$ (see Figure 1(b) for an example). Observe that a NE for the facility location game provides a NE for the corresponding vertex cover game and vice versa.

Corollary 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 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.*

4 Approximate Nash Equilibria

In this section we consider set 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 1: (f,f)-NE for set cover games

```
s_p(M) ← 0 for all players p and sets M
γ_p(e) ← 0 for all players p and elements e
for every player p = 1,...,k do
    Set c^p(M) = min_Q{c(M,|Q|+1) − Σ_{q∈Q} s_q(M)} for Q ⊆ [p−1] and all M
    while there is an uncovered element e ∈ E_p do
        Let γ_p(e) ← min_{e∈M} c^p(M)
        Increase payments: s_p(M) ← s_p(M) + γ_p(e) for all M with e ∈ M
        Add all purchased sets to the cover
        Reduce set costs: c^p(M) ← c^p(M) − γ_p(e) for all M with e ∈ M
```

Theorem 4. *Algorithm 1 returns a (f,f)-approximate Nash equilibrium for set cover games in polynomial time.*

Proof. The algorithm can be implemented to run in polynomial time. In line 4 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 is a subset of $[p-1] = 1, \ldots, p-1$, because for all other players all contributions are still 0. We start with $Q = \emptyset$ and add players $q < p$ in non-increasing order of the contributions $s_q(M)$. This yields the desired set Q.

Our algorithm represents an adjustment of the primal-dual algorithm for minimum set cover (see for instance [21, chapter 15]). An approximation guarantee of f for the buy-at-bulk set cover problem has most likely been observed before, so a proof is omitted. For the stability ratio, 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 $c'(M) = \min_{Q \subset [k], p \neq Q} c(M, |Q|+1) - \sum_{q \in Q} s_q(M)$. 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 theorem. □

In the special case of vertex cover ratios of $f = 2$ is tight even for regular vertex cover games [5]. The analysis cannot be strengthened to a ratio depending on g, because stability and approximation ratio coincide for single player games. For linear g the greedy algorithm achieves logarithmic stability and approximation ratio, but for regular set cover games this algorithm has an unbounded stability ratio [14]. It is an interesting open problem to obtain a procedure with improved bounds for intermediate functions g.

5 Single Element Players

In the previous section we showed that vertex cover games, in which each player owns at most two edges, might have no NE. Now 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 5. *Algorithm 2 returns an exact Nash equilibrium for singleton set multi-cover games in polynomial time.*

Algorithm 2: 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 $c(M_1) \cdot \mu(d_{M_1}) < c(M_2) \cdot \mu(M_2, d_{M_2})$
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 $s_p(M) = c(M) \cdot \mu(d_M)$ 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 2 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. The previous arguments show 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 repsonse. This proves the theorem. $\square$

Unfortunately, the proposed algorithm can compute worst-case NE, whose cost is a factor arbitrarily close to $\frac{k}{g(k)}$ worse than $c(s^*)$. In contrast, there are social optimal NE in every singleton set multi-cover game. Computing them is NP-hard, but with a local search procedure we can obtain near-stable and near-optimal approximate NE. The arguments can transfered to buy-at-bulk versions of *connection-restricted facility location* (CRFL) games [14]. Proofs are omitted due to space limitations.

Theorem 6. *For singleton set multi-cover and singleton CRFL 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 optimum social cost.*

References

1. M. Andrews. Hardness of buy-at-bulk network design. In *Proc. 45th FOCS*, pages 115–124, 2004.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, T. Roughgarden, É. Tardos, and T. Wexler. The price of stability for network design with fair cost allocation. In *Proc. 45th FOCS*, pages 295–304, 2004.
3. E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. 35th STOC*, pages 511–520, 2003.
4. B. Awerbuch and Y. Azar. Buy-at-bulk network design. In *Proc. 38th FOCS*, pages 542–547, 1997.
5. J. Cardinal and M. Hoefer. Selfish service installation in networks. In *Proc. 2nd WINE*, pages 174–185, 2006.
6. C. Chekuri, M.T. Hajiaghayi, G. Kortarz, and M.-Salavatipour. Approximation algorithms for non-uniform buy-at-bulk network design. In *Proc. 47th FOCS*, pages 677–686, 2006.
7. C. Chekuri, M.T. Hajiaghayi, G. Kortsarz, and M. Salavatipour;. Approximation algorithms for node-weighted buy-at-bulk networks. In *Proc. 18th SODA*, 2007.
8. H.-L. Chen and T. Roughgarden. Network design with weighted players. In *Proc. 18th SPAA*, pages 29–38, 2006.
9. X. Deng, T. Ibaraki, and H. Nagamochi. Combinatorial optimization games. In *Proc 8th SODA*, pages 720–729, 1997.
10. N. Devanur, M. Mihail, and V. Vazirani. Strategyproof cost-sharing mechanisms for set cover and facility location problems. *Decision Support Systems*, 39(1):11–22, 2005.
11. M. Goemans and M. Skutella. Cooperative facility location games. *J Algorithms*, 50(2):194–214, 2004.
12. A. Gupta, A. Kumar, M. Pál, and T. Roughgarden. Approximation via cost sharing: Simpler and better approximation algorithms for network design. *J ACM*, 54(3):11, 2007.
13. M.T. Hajiaghayi, M. Mahdian, and V. Mirrokni. The facility location problem with general cost functions. *Networks*, 42(1):42–47, 2003.
14. M. Hoefer. Non-cooperative facility location and covering games. In *Proc. 17th ISAAC*, pages 369–378, 2006.
15. M. Hoefer. Non-cooperative tree creation. In *Proc. 31st MFCS*, pages 517–527, 2006.
16. N. Immorlica, M. Mahdian, and V. Mirrokni. Limitations of cross-monotonic cost sharing schemes. In *Proc. 16th SODA*, pages 602–611, 2005.
17. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. 16th STACS*, pages 404–413, 1999.
18. X. Li, Z. Sun, and W. Wang. Cost sharing and strategyproof mechanisms for set cover games. In *Proc. 22nd STACS*, pages 218–230, 2005.
19. M. Pál and É. Tardos. Group strategyproof mechanisms via primal-dual algorithms. In *Proc. 44th FOCS*, pages 584–593, 2003.
20. Z. Sun, X. Li, W. Wang, and X. Chu. Mechanism design for set cover games when elements are agents. In *Proc. 1st AAIM*, pages 360–369, 2005.
21. V. Vazirani. *Approximation Algorithms*. Springer Verlag, Berlin, 2000.