

Covering Games: Approximation through Non-Cooperation ^{*}

Martin Gairing

Department of Computer Science, University of Liverpool, U.K.
m.gairing@liverpool.ac.uk

Abstract. We propose approximation algorithms under game-theoretic considerations. We introduce and study the *general covering problem* which is a natural generalization of the well-studied *max-n-cover problem*. In the general covering problem, we are given a universal set of weighted elements E and n collections of subsets of the elements. The task is to choose one subset from each collection such that the total weight of their union is as large as possible. In our game-theoretic setting, the choice in each collection is made by an independent player. For covering an element, the players receive a payoff defined by a non-increasing *utility sharing function*. This function defines the fraction that each covering player receives from the weight of the elements.

We show how to construct a utility sharing function such that every Nash Equilibrium approximates the optimal solution by a factor of $1 - \frac{1}{e}$. We also prove that any sequence of unilateral improving steps is polynomially bounded. This gives rise to a polynomial-time local search approximation algorithm whose approximation ratio is best possible.

1 Introduction

Motivation and Framework. Large scale distributed systems, like the Internet, usually lack a centralized control authority. Instead, they are operated and controlled in a distributed fashion by competing entities – modeled as *players* – which make their decisions in order to optimize their own private *utility*. Such systems are assumed to end up in a *Nash equilibrium* [20] – a state in which no player wishes to unilaterally leave her own strategy in order to improve the value of her private *utility*. However, Nash equilibria are often suboptimal solutions with respect to the *social objective function*. The *price of anarchy* [18] is a measure for the performance degradation. It is defined as the worst-case ratio between the values of a social objective function in a Nash equilibrium and in an optimum solution.

As the designer of a distributed system we are faced with the main challenge of how to design the distributed system in order to optimize this social objective function even in the presence of myopic players. However, even if all players

^{*} This work was supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

adhere to some centralized authority, finding the optimum solution with respect to the social objective might be an \mathcal{NP} -hard optimization problem. Approximation algorithms [23] are a powerful tool for coping with intractable optimization problems. In general, they compute (in polynomial time) suboptimal solutions but with a provable performance guarantee. However, approximation algorithms usually presume a centralized authority. In this paper we propose to consider approximation algorithms that take the selfish user behavior into account. More precisely, we propose to design the distributed system in a way that the price of anarchy is optimized and the system is guaranteed to converge to a Nash equilibrium in polynomial time.

For our study of such approximation algorithms we consider a very general covering problem. In covering problems a finite set of elements has to be covered with subsets of the elements. Such problems arise in many contexts: Covering problems can be used to model service installation problems in distributed systems. Moreover, many packing problems and fixed parameter optimization problems can be modeled as a covering problem (see [13]). A well-studied representative is the *max- n -cover* problem (see e.g. [9, 14]): Given a finite set of weighted elements E , choose n subsets of the elements (from a given collection of subsets) such that the total weight of their union is as large as possible. We consider a generalization of the *max- n -cover* problem that we call *general covering problem*. Here, we are given not 1 but n collections of subsets and we have to choose one subset from each collection. Although this generalization seems natural, we are not aware of any previous work on it.

We study the general covering problem as a *covering game*, where the choice for the subset in each of the n collections is made by an independent player. Covering games are a subclass of the *congestion games* introduced by Rosenthal [21]. For covering an element, the players receive a payoff defined by a *utility sharing function*. This function defines the fraction that each covering player receives from the weight of the element and depends only on local parameters. Those parameters are the *number of players* covering an element and the *cardinality* of an element (i.e. the number of players that have this element in at least one of their strategies). We only make two natural assumptions on the utility sharing function: First, we assume that it is non-increasing in the number of players covering the element. And second, we want that the payoff to the players for covering an element does not exceed the weight of the element.

The focus of this paper is to design utility sharing functions, such that:

1. In any Nash equilibrium the total weight of the covered elements is as large as possible. Or more precisely, the price of anarchy is maximized.
2. A Nash equilibrium is reached in polynomial time.

Obviously, each utility sharing function that fulfills both of these properties gives rise to a local search approximation algorithm. In fact, we will show that this approach yields essentially the best possible approximation ratio.

Contribution. In this paper, we introduce and study *covering games*. Such games are congestion games that have the *general covering problem* as underlying structure.

In the first part of the paper, we focus on the design of *utility sharing functions* that maximize the price of anarchy. In particular, we construct a utility sharing function which achieves a price of anarchy of $1 - \frac{1}{e}$ (Theorem 4) and show that no utility sharing function performs better (Theorem 1). To show this, we first prove a corresponding result on the price of anarchy, that depends on the maximum cardinality k of the elements (Theorem 3). Surprisingly, we get matching bounds for each fixed k . All our results on the price of anarchy hold for pure and mixed Nash equilibria.

In the second part of the paper, we show how to use our results on the price of anarchy to construct a *local search approximation algorithm* for the general covering problem (Theorem 6), which runs in polynomial time, if the weights of the elements are polynomially bounded. Our hardness result in Theorem 7 shows that this restriction on the weights is necessary. For the general case, we also present a (centralized) approximation algorithm, which is based on LP-rounding and generalizes an algorithm for MAXSAT [12].

Related Work. Congestion games and variants thereof have long been used to model non-cooperative resource sharing among selfish players. Rosenthal [21] showed that congestion games always possess pure Nash equilibria. However, computing such a pure Nash equilibrium is \mathcal{PLS} -complete [8]. The price of anarchy in congestion games has been studied extensively (see e.g. [1, 6, 10, 22]).

The *general covering problem* is a natural generalization of the well-studied *max-n-cover* problem. For the *max-n-cover* problem, the greedy approach yields a $(1 - \frac{1}{e})$ -approximation [14] and no polynomial time algorithm can do better, unless $NP \subseteq TIME(n^{O(\log \log n)})$ [9]. Applying the greedy approach to our more general problem guarantees only a $\frac{1}{2}$ -approximation. For an overview on approximation algorithms for covering problems, we refer to [13, Chapter 3].

The MAXSAT problem is a special case of the *generalized covering* problem, where each of the n collections consists of at most 2 subsets (corresponding to true/false). The power of local search for approximating MAXSAT has been studied in [2, 17]. MAXSAT has also been considered in a game-theoretic setting as a SAT-game [4, 11], which is itself a special case of our covering games. Bilò [4] mainly focuses on the expressiveness of SAT-games. Mavronicolas et al. [19] concentrate on structural properties and complexity questions for a generalization of SAT-games, called *weighted boolean formula games*. Here, each player controls a set of variables and aims to maximize the total weight of his satisfied formulas. Giannakos et al. [11] study the price of anarchy (for pure Nash equilibria) of SAT-games under different utility sharing functions and point out the relation to approximation algorithms. Our work generalizes their results in two perspectives. First, we consider a far more general class of games, and second, we allow for mixed Nash equilibria. Moreover, in Example 1 we show that their main result [11, Thm. 5] is incorrect.

Under certain conditions on the utility sharing functions, our games fall in the class of *valid utility games* [24]. For such games, Vetta [24] shows that each Nash equilibrium is a $\frac{1}{2}$ -approximation. Our result improves this ratio to $1 - \frac{1}{e}$.

Coordination mechanisms have been introduced in [7] as a notion to improve the price of anarchy. The idea is to define local policies such that the corresponding price of anarchy is as small as possible. A few other papers follow this approach, e.g. [3, 5, 15]. Our task of designing utility sharing functions can be seen as such a coordination mechanism. However, we take the idea of coordination mechanisms one step further. In the design of our utility sharing function, the goal is not only to optimize the price of anarchy, but also to ensure that the system converges to a Nash equilibrium in polynomial time. Azar et al. [3] pursue a similar approach for the unrelated scheduling problem. However, the price of anarchy of their best coordination mechanism increases significantly by requiring a polynomial convergence.

Roadmap. The rest of the paper is organized as follows. In Section 2, we introduce *covering games*. Section 3 comprises our results on the price of anarchy, while Section 4 presents our approximation algorithms. We conclude in Section 5. Due to lack of space, some proofs are omitted.

2 Model

For any two integers $l \leq m$, denote $[m] = \{1, \dots, m\}$ and $[l, m] = \{l, \dots, m\}$. For a vector $\mathbf{v} = (v_1, \dots, v_n)$, let $\mathbf{v}_{-i} = (v_1, \dots, v_{i-1}, v_{i+1}, \dots, v_n)$ and $(\mathbf{v}_{-i}, v'_i) = (v_1, \dots, v_{i-1}, v'_i, v_{i+1}, \dots, v_n)$.

The general covering problem. In the *general covering problem* we are given a finite set of elements E and a weight function $w : E \mapsto \mathbb{N}$ that assigns a positive integer weight w_e to each element $e \in E$. Moreover, we are given n collections S_1, \dots, S_n of subsets of E where for each $i \in [n]$, the collection $S_i \subset 2^E$ is a subset of the power-set of the elements. Given such an instance our task is to choose one subset s_i from each collection S_i such that their union $\cup_{i \in [n]} s_i$ has maximum total weight, i.e. $\sum_{e \in \cup_{i \in [n]} s_i} w_e$ is maximized.

Covering Games. Each *covering game* has a general covering problem as an underlying structure. Here, each of the n collections of subsets is controlled by a rational *player*, that is player $i \in [n]$ has S_i as her strategy set. Denote $S = S_1 \times \dots \times S_n$. As for the general covering problem, each element $e \in E$ has a weight $w_e \in \mathbb{N}$. For any subset of the elements $E' \subseteq E$ denote $W(E') = \sum_{e \in E'} w_e$. Let $W = W(E)$. For each element $e \in E$ denote by $k_e = |\{i \in [n] : e \in s_i \text{ for some } s_i \in S_i\}|$ the *cardinality* of e which is the number of players that can possibly cover e . Let $k = \max_{e \in E} k_e$ and $k_{\min} = \min_{e \in E} k_e$. A covering game is a *SAT-game* if $|S_i| \leq 2$ for all player $i \in [n]$. In this case, elements correspond to clauses and players correspond to variables which can be set to **true** or **false**.

Strategies and Strategy Profiles. A *pure strategy* for player i is some specific strategy $s_i \in S_i$, while a *mixed strategy* $P_i = (p(i, s_i))_{s_i \in S_i}$ is a probability distribution over S_i , where $p(i, s_i)$ denotes the probability that player i chooses the pure strategy s_i .

A *pure strategy profile* is an n -tuple $\mathbf{s} = (s_1, \dots, s_n)$ whereas a *mixed strategy profile* $\mathbf{P} = (P_1, \dots, P_n)$ is represented by an n -tuple of mixed strategies. For a

mixed strategy profile \mathbf{P} , denote by $p(\mathbf{s}) = \prod_{i \in [n]} p(i, s_i)$ the probability that the players choose the pure strategy profile \mathbf{s} .

Utility Sharing Functions. For each element $e \in E$ there is a *payoff function* f_e that describes how much a player receives for covering e . In this paper we consider payoff functions that come from a common *utility sharing function* f . We want this function to depend only on local parameters. We consider two different kinds of utility sharing functions:

- A *cardinality dependent* utility sharing function depends on the number of players covering an element and the cardinality of the element, i.e. $f : [k] \times [k] \mapsto \mathbb{N}$ and for all elements $e \in E$ and $j \in [k]$, $f_e(j) = f(j, k_e) \cdot w_e$.
- A *symmetric* utility sharing function depends only on the number of players covering an element, i.e. $f : [k] \mapsto \mathbb{N}$ and for all elements $e \in E$ and $j \in [k]$, $f_e(j) = f(j) \cdot w_e$.

For both cases we assume that f is *non-increasing* in the number of players. Moreover, we assume that f does not overpay the players, i.e. $j \cdot f(j) \leq 1$ for all $j \in [k]$ in the symmetric case (and $j \cdot f(j, l) \leq 1$ for all $j \in [l], l \in [k]$ in the cardinality dependent case).

Load and Player Utilities. For a pure strategy profile \mathbf{s} , let $\delta_e(\mathbf{s}) = |\{i \in [n] : e \in s_i\}|$ denote the *load* on element $e \in E$, i.e. the number of players covering e .

Fix a pure strategy profile \mathbf{s} . The *utility* $u_i(\mathbf{s})$ of player $i \in [n]$ is defined by the *payoff* from the elements she covers. Thus, $u_i(\mathbf{s}) = \sum_{e \in s_i} f_e(\delta_e(\mathbf{s}))$. For a mixed strategy profile \mathbf{P} , the *utility* of player $i \in [n]$ is $u_i(\mathbf{P}) = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}) \cdot u_i(\mathbf{s})$.

Social Utility. Fix a pure strategy profile \mathbf{s} . Denote by $E_{\mathbf{s}}$ the subset of elements that are covered by at least one player in \mathbf{s} , i.e. $E_{\mathbf{s}} = \{e \in E : \delta_e(\mathbf{s}) > 0\}$. The *social utility* in \mathbf{s} is the total weight $W(E_{\mathbf{s}})$ of the covered elements. We abuse notation and denote this value also as $W(\mathbf{s})$. For a mixed strategy profile \mathbf{P} the social utility $W(\mathbf{P}) = \sum_{\mathbf{s} \in \mathcal{S}} p(\mathbf{s}) \cdot W(E_{\mathbf{s}})$ is the expected total weight of the covered elements. Throughout denote by \mathbf{s}^* a pure strategy profile that maximizes the total weight of the covered elements, thus, $\mathbf{s}^* = \arg \max_{\mathbf{s} \in \mathcal{S}} W(E_{\mathbf{s}})$.

Nash Equilibria and Potential Function. A mixed strategy profile \mathbf{P} is a *Nash equilibrium* if and only if no player can increase her utility by unilaterally changing her strategy, that is, $u_i(\mathbf{P}) \geq u_i(\mathbf{P}_{-i}, s_i)$ for all $i \in [n]$ and $s_i \in S_i$. Depending on the type of strategy profile we distinguish between *pure* and *mixed* Nash equilibria. Given a pure strategy profile \mathbf{s} , a *selfish step* of player $i \in [n]$ is a deviation to a strategy profile (\mathbf{s}_{-i}, s'_i) where $u_i(\mathbf{s}_{-i}, s'_i) > u_i(\mathbf{s})$, that is player i increases her utility.

For covering games, Rosenthal's [21] *exact potential function* Φ implies the existence of a pure Nash equilibrium. For every pure strategy profile \mathbf{s} , the potential $\Phi(\mathbf{s})$ is defined by $\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{\delta_e(\mathbf{s})} f_e(i)$. If a player performs a selfish step and increases her utility by Δ , then $\Phi(\mathbf{s})$ also increases by Δ .

Price of Anarchy. Let $\mathcal{G}(k)$ be the class of covering games where $k_e \leq k$ for all $e \in E$. Fix a utility sharing function f . The *Price of Anarchy* for f , denoted

by PoA_f , is the infimum, over all instances $\Gamma \in \mathcal{G}(k)$ and Nash equilibria \mathbf{P} , of the ratio $\frac{W(\mathbf{P})}{W(\mathbf{s}^*)}$. Thus, $\text{PoA}_f(k) = \inf_{\Gamma \in \mathcal{G}(k), \mathbf{P}} \frac{W(\mathbf{P})}{W(\mathbf{s}^*)}$. Similarly, define PoA_f by dropping the restriction on k .

3 Price of Anarchy Results

In this section, we study the price of anarchy for different utility sharing functions. We start with an upper bound that holds for all utility sharing functions.

Theorem 1. *Consider the class of covering games $\mathcal{G}(k)$ with unweighted elements. Then, $\text{PoA}_f(k) \leq 1 - \frac{1}{(k-1)(k-1)! + \sum_{j \in [0, k-1]} \frac{1}{j!}}$.*

This holds (a) for every cardinality dependent utility sharing function f , (b) even for SAT-games, if we restrict ourselves to symmetric utility sharing functions f .

For SAT-games Giannakos et al. [11, Thm. 5] claim that using the cardinality dependent utility sharing function defined by $f(1, l) = 1$ and $f(j, l) = \frac{1}{2(l-1)}$ for $j \geq 2$ achieves a price of anarchy of $\frac{2}{3}$. The following example shows that this does not hold:

Example 1. Given the maximum cardinality k , define a SAT-game with k players and $k+1$ elements. We have $w_e = 1$ for all $e \in [k-1]$ and $w_k = w_{k+1} = 2(k-1)$. Each player $i \in [k-1]$ can either cover element i or element k , while player k can choose between the elements k and $k+1$.

Let \mathbf{s} be the strategy profile where each player $i \in [k]$ covers element i . It's not hard to see that \mathbf{s} is a Nash equilibrium with $W(E_{\mathbf{s}}) = 3(k-1)$. On the other hand there is a strategy profile \mathbf{s}^* (where only element 1 is not covered) with $W(\mathbf{s}^*) = 5(k-1) - 1$. For $k \geq 4$, this stands in conflict to the claim in [11, Thm. 5]; and for $k \rightarrow \infty$, we might only cover $\frac{3}{5}$ of the optimum total weight. In fact, we believe that our upper bound in Theorem 1 holds also for SAT-games with cardinality dependent utility sharing functions.

We proceed by introducing a parameter χ_f of the utility sharing function f . This parameter is a measure on how fast the utility sharing function decreases. We will use χ_f in Theorem 2 to prove a general lower bound on the price of anarchy that depends on χ_f .

Definition 1. *Given a cardinality dependent utility sharing function f , define χ_f as the minimum value such that for all cardinalities $l \in [k]$ we have*

$$j \cdot f(j, l) - f(\min\{j+1, l\}, l) \leq \chi_f \cdot f(1, l) \quad \text{for all } j \in [l].$$

Theorem 2. *Consider the class of covering games $\mathcal{G}(k)$. Let f be a cardinality dependent utility sharing function, where $\alpha_{\min} = \min_{l \in [k]} f(1, l)$ and $\alpha_{\max} = \max_{l \in [k]} f(1, l)$. Then, $\text{PoA}_f(k) \geq \frac{1}{\chi_f + 1} \cdot \frac{\alpha_{\min}}{\alpha_{\max}}$.*

Proof. Let \mathbf{P} be an arbitrary mixed Nash equilibrium and \mathbf{s}^* be an optimum pure strategy profile. Since \mathbf{P} is a Nash equilibrium, it follows that $u_i(\mathbf{P}) - u_i(\mathbf{P}_{-i}, \mathbf{s}_i^*) \geq 0$ for all player $i \in [n]$. So,

$$0 \leq \sum_{i \in [n]} u_i(\mathbf{P}) - \sum_{i \in [n]} u_i(\mathbf{P}_{-i}, \mathbf{s}_i^*) = \sum_{\mathbf{s} \in S} p(\mathbf{s}) \left(\sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) \right) \quad (1)$$

By definition of player utility, for any pure strategy profile \mathbf{s} ,

$$\sum_{i \in [n]} u_i(\mathbf{s}) = \sum_{i \in [n]} \sum_{e \in s_i} f_e(\delta_e(\mathbf{s})) = \sum_{e \in E} \delta_e(\mathbf{s}) \cdot f_e(\delta_e(\mathbf{s})) = \sum_{j \in [k]} \sum_{\substack{e \in E_s, \\ \delta_e(\mathbf{s})=j}} j \cdot f_e(j). \quad (2)$$

Moreover,

$$\begin{aligned} \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) &= \sum_{i \in [n]} \sum_{e \in s_i^*} f_e(\delta_e(\mathbf{s}_{-i}, \mathbf{s}_i^*)) \geq \sum_{i \in [n]} \sum_{e \in s_i^*} f_e(\min\{k_e, \delta_e(\mathbf{s}) + 1\}) \\ &\geq \sum_{e \in E_{\mathbf{s}^*}} f_e(\min\{k_e, \delta_e(\mathbf{s}) + 1\}) = \sum_{j=0}^k \sum_{\substack{e \in E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} f_e(\min\{k_e, j + 1\}) \quad (3) \end{aligned}$$

where the first inequality follows since f_e is a non-increasing function and the second inequality follows since $\delta_e(\mathbf{s}^*) \geq 1$ for all $e \in E_{\mathbf{s}^*}$.

So, for any pure strategy profile \mathbf{s} , (2) and (3) imply:

$$\begin{aligned} \sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) &\leq \sum_{j \in [k]} \sum_{\substack{e \in E_s, \\ \delta_e(\mathbf{s})=j}} j \cdot f_e(j) - \sum_{j=0}^k \sum_{\substack{e \in E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} f_e(\min\{k_e, j + 1\}) \\ &= \sum_{j \in [k]} \sum_{\substack{e \in E_s \setminus E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} j f_e(j) + \sum_{j \in [k]} \sum_{\substack{e \in E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} [j f_e(j) - f_e(\min\{k_e, j + 1\})] - \sum_{e \in E_{\mathbf{s}^*} \setminus E_s} f_e(1) \end{aligned}$$

By Definition 1 and the fact that f_e is non-increasing, we have $j \cdot f_e(j) \leq (\chi_f + 1)f_e(1)$ for all $e \in E$ and $j \in [k_e]$. Using this and Definition 1, we get:

$$\begin{aligned} &\sum_{i \in [n]} u_i(\mathbf{s}) - \sum_{i \in [n]} u_i(\mathbf{s}_{-i}, \mathbf{s}_i^*) \\ &\leq \sum_{j \in [k]} \sum_{\substack{e \in E_s \setminus E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} (\chi_f + 1) \cdot f_e(1) + \sum_{j \in [k]} \sum_{\substack{e \in E_{\mathbf{s}^*}, \\ \delta_e(\mathbf{s})=j}} \chi_f \cdot f_e(1) - \sum_{e \in E_{\mathbf{s}^*} \setminus E_s} f_e(1) \\ &= \sum_{e \in E_s \setminus E_{\mathbf{s}^*}} (\chi_f + 1) \cdot f_e(1) + \sum_{e \in E_{\mathbf{s}^*} \cap E_s} \chi_f \cdot f_e(1) - \sum_{e \in E_{\mathbf{s}^*} \setminus E_s} f_e(1) \\ &= \sum_{e \in E_s \setminus E_{\mathbf{s}^*}} (\chi_f + 1) \cdot f_e(1) + \sum_{e \in E_{\mathbf{s}^*} \cap E_s} (\chi_f + 1) \cdot f_e(1) - \sum_{e \in E_{\mathbf{s}^*}} f_e(1) \\ &\leq (\chi_f + 1) \cdot \alpha_{\max} \cdot W(E_s) - \alpha_{\min} \cdot W(E_{\mathbf{s}^*}). \end{aligned}$$

With (1) we get $0 \leq \sum_{s \in S} p(s) ((\chi_f + 1) \cdot \alpha_{\max} \cdot W(E_s) - \alpha_{\min} \cdot W(E_{s^*})) = (\chi_f + 1) \cdot \alpha_{\max} \cdot \sum_{s \in S} p(s) \cdot W(E_s) - \alpha_{\min} \cdot W(E_{s^*})$. Rearranging terms yields $\frac{W(\mathbf{P})}{W(s^*)} = \frac{\sum_{s \in S} p(s) \cdot W(E_s)}{W(E_{s^*})} \geq \frac{\alpha_{\min}}{(\chi_f + 1) \cdot \alpha_{\max}}$. The theorem follows since \mathbf{P} is an arbitrary Nash equilibrium. \square

In the following, we construct a utility sharing function such that the corresponding lower bound in Theorem 2 is maximized. Observe, that χ_f is independent of the values for α_{\min} and α_{\max} . So, without loss of generality, we can restrict our attention to a symmetric utility sharing function f , where $\alpha_{\min} = \alpha_{\max}$. Our task is to construct a symmetric utility sharing function that solves the following optimization problem:

$$\begin{aligned} & \text{minimize } \chi & (4) \\ & \text{subject to } j \cdot f(j) - f(j+1) \leq \chi \cdot f(1) \quad \forall j \in [k-1] \\ & & (k-1)f(k) \leq \chi \cdot f(1) \end{aligned}$$

Replacing " \leq " with " $=$ " yields a homogeneous system of linear equations. The values for χ and the utility sharing function f in the following theorem correspond to the solution of this system where $f(1) = 1$.

Theorem 3. *Given k we can construct a symmetric utility sharing function f , such that $\text{PoA}_f(k) \geq 1 - \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=0}^{k-1} \frac{1}{i!}}$.*

Proof. Given k , let f be the symmetric utility sharing function defined by

$$f(j) = (j-1)! \frac{\frac{1}{(k-1)(k-1)!} + \sum_{i=j}^{k-1} \frac{1}{i!}}{\frac{1}{(k-1)(k-1)!} + \sum_{i=1}^{k-1} \frac{1}{i!}} \quad \text{for all } j \in [k]. \quad (5)$$

It is not hard to check that f is a valid utility sharing function, i.e. f is non-increasing and $j \cdot f(j) \leq 1$ for all j . Moreover, f satisfies the constraints in (4) for $\chi = \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=1}^{k-1} \frac{1}{i!}}$. Recall that $\alpha_{\min} = \alpha_{\max}$ for symmetric utility sharing functions. The claim follows by applying Theorem 2. \square

In order to construct the utility sharing function f in Theorem 3 we need to know the maximum cardinality k over all elements. However, since the value for χ from the proof of Theorem 3 is increasing with k , we can get the same lower bound if each element $e \in E$ is only aware of her own cardinality k_e . For this case, the cardinality dependent utility sharing function is defined by replacing k with k_e in (5). This implies:

Corollary 1. *There exists a cardinality dependent utility sharing function f , such that $\text{PoA}_f(k) \geq 1 - \frac{1}{\frac{1}{(k-1)(k-1)!} + \sum_{i=0}^{k-1} \frac{1}{i!}}$.*

Observe that the lower bounds on the price of anarchy in Theorem 3 and Corollary 1 match exactly the upper bound in Theorem 1.

There might also be cases where we want to use a utility sharing function that works for all k . For example, neither k is known a priori nor the elements can determine their own cardinality. For such cases, we get:

Theorem 4. *There exists a symmetric utility sharing function f with $\text{PoA}_f(k) \geq 1 - \frac{1}{e}$, which works for arbitrary k .*

Proof. This follows by applying Theorem 3 for $k \rightarrow \infty$. In this case, f reduces to $f(j) = \frac{(j-1)!}{e-1} \left[e - \sum_{i=0}^{j-1} \frac{1}{i!} \right]$ for all $j \in \mathbb{N}$, and $\chi = \frac{1}{e-1}$. \square

We close this section with an alternative lower bound on the price of anarchy that depends on the maximum dimension $d_{\max} = \max_{i \in [n]} |S_i|$ over all players and the minimum cardinality of an element k_{\min} . For certain cases (e.g. if $k_{\min} > d_{\max}$) this is better than the bound in Theorem 4.

Theorem 5. *Consider the class of covering games \mathcal{G} where $k_e \geq k_{\min}$ for all elements $e \in E$ and $d_{\max} = \max_{i \in [n]} |S_i|$. Let f be a cardinality dependent utility sharing function with $f(1, l) = 1$ for all cardinalities $l \in [k]$. Then, $\text{PoA}_f(\mathcal{G}) \geq \frac{k_{\min}}{d_{\max} - 1 + k_{\min}}$.*

4 Approximation Algorithm

In the previous section, we have shown results on the price of anarchy for covering games. In this section, we want to use those results to construct a distributed, local-search approximation algorithm for the covering problem.

The idea of the algorithm is simple:

- Choose an appropriate utility sharing function,
- start with an arbitrary strategy profile, and
- let the players unilaterally perform selfish steps until a pure Nash equilibrium is reached.

The approximation ratio is the price of anarchy for the chosen utility sharing function. Rosenthal's potential function [21] can be used to bound the number of selfish step until a pure Nash equilibrium is reached. Unfortunately, the utility sharing functions in Theorem 3 and Theorem 4 do not provide a sub-exponential bound on the number of selfish steps, since the increase in the potential due to a single selfish step can be arbitrary small.

To overcome this, we will design a new symmetric utility sharing function f , where for each element $e \in E$ the players receive strictly positive payoff only if at most a constant number k' of players cover this element, i.e. $f(j) = 0$ for all $j > k'$.

We will show that the right choice of f yields a $(1 - \frac{1}{e} - \varepsilon)$ -approximation algorithm, where $\varepsilon = \varepsilon(k') = o(1)$.

Theorem 6. *For every constant $\varepsilon > 0$ there exists a local-search approximation algorithm with approximation ratio $(1 - \frac{1}{e} - \varepsilon)$ that uses at most $\mathcal{O}(\frac{1}{\varepsilon} \cdot \log \log \frac{1}{\varepsilon}) \cdot W$ selfish steps.*

Proof. Let $k' \in \mathbb{N}$ be some positive integer (to be determined later) and construct f as a solution to the following optimization problem:

$$\begin{aligned} & \text{minimize } \chi & (6) \\ & \text{subject to } j \cdot f(j) - f(j+1) \leq \chi \cdot f(1) \quad \forall j \in [k' - 1] \\ & \quad \quad \quad k' \cdot f(k') \leq \chi \cdot f(1) \end{aligned}$$

For the solution of the corresponding homogeneous system of linear equations with $f(1) = 1$ we get: $f(j) = (j-1)! \frac{\sum_{i=j}^{k'} \frac{1}{i!}}{\sum_{i=1}^{k'} \frac{1}{i!}}$ for all $j \in [k']$ and $f(j) = 0$ for $j > k'$. It is not hard to check that f is a valid utility sharing function, i.e. f is non-increasing and $j \cdot f(j) \leq 1$ for all j .

Observe, that $f(1) = 1, k' \cdot f(k') = \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$ and $f(j+1) = j! \frac{\sum_{i=j+1}^{k'} \frac{1}{i!} - \frac{1}{j!}}{\sum_{i=1}^{k'} \frac{1}{i!}} = j \cdot f(j) - \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$. Thus, f satisfies the constraints in (6) for $\chi = \frac{1}{\sum_{i=1}^{k'} \frac{1}{i!}}$. Applying Theorem 2 yields

$$\text{PoA}_f \geq 1 - \frac{1}{\sum_{i=0}^{k'} \frac{1}{i!}} = 1 - \frac{1}{e} - \varepsilon(k'),$$

where $\varepsilon(k') = \frac{1}{\sum_{i=0}^{k'} \frac{1}{i!}} - \frac{1}{e} = \frac{\sum_{i=k'+1}^{\infty} \frac{1}{i!}}{e \cdot \sum_{i=0}^{k'} \frac{1}{i!}} = \frac{\sum_{i=k'+1}^{\infty} \frac{k'!}{i!}}{e \cdot \sum_{i=0}^{k'} \frac{k'!}{i!}} = \Theta\left(\frac{1}{(k'+1)!}\right)$. So for every constant ε we choose $k' = \Theta\left(\frac{\log(\frac{1}{\varepsilon})}{\log \log(\frac{1}{\varepsilon})}\right)$. We will now bound the maximum number of selfish steps. To do so, we will first show an upper bound on Rosenthal's potential function $\Phi(\mathbf{s})$. Afterwards, we show a lower bound on the increase in Φ due to a selfish step. For any pure strategy profile \mathbf{s} we have

$$\Phi(\mathbf{s}) = \sum_{e \in E} \sum_{i=1}^{\delta_e(\mathbf{s})} f_e(i) \leq \sum_{e \in E} \sum_{i=1}^{k'} f_e(i) \leq \sum_{e \in E} \sum_{i=1}^{k'} \frac{1}{i} \cdot w_e = H(k') \cdot W,$$

where $H(k')$ is the harmonic number of order k' .

Recall, that $w_e \in \mathbb{N}$ for all $e \in E$. Moreover, for all $j \in [k']$, $f(j)$ is an integer multiple of $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}}$. To see this multiply the numerator and denominator by $k'!$ and observe that both become integer. So if a player improves, then she improves by at least $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}}$. Using the property of Rosenthal's *exact* potential function (cf. Sec. 2), it follows that each selfish step increases Φ by at least $\frac{1}{\sum_{i=1}^{k'} \frac{k'!}{i!}} = \frac{1}{\lfloor (e-1)k'! \rfloor}$. So the number of selfish steps is upper bounded by $\lfloor (e-1)k'! \rfloor \cdot H(k') \cdot W = \mathcal{O}(k'! \cdot \log k') \cdot W = \mathcal{O}\left(\frac{1}{\varepsilon} \cdot \log \log \frac{1}{\varepsilon}\right) \cdot W$. \square

Theorem 6 implies a polynomial time $(1 - \frac{1}{e} - \varepsilon)$ -approximation algorithm for the case that the total weight W is polynomially bounded. This includes the important case of unweighted elements, where $w_e = 1$ for all $e \in E$. In the following theorem we show that this restriction on W is necessary, since for arbitrary weights the problem of computing a pure Nash equilibrium is \mathcal{PLS} -complete (see [16] for an introduction to the complexity class \mathcal{PLS}).

Theorem 7. *Consider the class of covering games with arbitrary weights. Then, for every symmetric utility sharing function with $f(1) > f(2)$, it is \mathcal{PLS} -complete to compute a pure Nash equilibrium. This holds even for SAT-games.*

For the general case, we present a (centralized) approximation algorithm, which is based on LP-rounding and generalizes an algorithm for MaxSat [12].

Theorem 8. *There exists a (centralized) polynomial-time $(1 - \frac{1}{e})$ -approximation algorithm for the general covering problem.*

5 Conclusion

In this paper we use game theoretic concepts for the design of new local search approximation algorithms for a very general covering problem. Our approach is to design player payoff functions that minimize the price of anarchy and guarantee that any sequence of unilateral improvements by the players is of polynomial length. For the covering problem this yields essentially the best possible approximation ratio.

For future work, we propose to study how far such ideas can be utilized to get new local search approximation algorithms also for other interesting optimization problems. Certainly, our approach will not always yield the best possible approximation ratio. This gives rise to the new interesting concept of *selfish approximation ratio*, i.e. the best possible approximation ratio that can be achieved by selfish players.

Acknowledgement

We thank Adrian Vetta and Thomas Sauerwald for many helpful discussions and comments on the topic.

References

1. Aland, S., Dumrauf, D., Gairing, M., Monien, B., Schoppmann, F.: Exact Price of Anarchy for Polynomial Congestion Games. In Durand, B., Thomas, W. (eds.) STACS 2006. LNCS, vol. 3884, pp. 218–229, Springer, Heidelberg (2006)
2. Alimonti, P.: New Local Search Approximation Techniques for Maximum Generalized Satisfiability Problems. In Bonuccelli, M.A., Crescenzi, P., Petreschi, R. (eds.) CIAC 1994. LNCS, vol. 778, pp. 40–53, Springer, Heidelberg (1994)
3. Azar, Y., Jain, K., Mirrokni, V.S.: (Almost) Optimal Coordination Mechanisms for Unrelated Machine Scheduling. In: Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08). pp. 323–332 (2008)
4. Bilò, V.: On Satisfiability Games and the Power of Congestion Games. In Kao, M.-Y., Li, X.-Y. (eds.) AAIM 2007. LNCS, vol. 4508, pp. 295–304, Springer, Heidelberg (2007)
5. Chen, H.-L., Roughgarden, T., Valiant, G.: Designing Networks with Good Equilibria. In: Proc. of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'08). pp. 854–863 (2008)

6. Christodoulou, G., Koutsoupias, E.: The Price of Anarchy of Finite Congestion Games. In: Proc. of the 37th Annual ACM Symposium on Theory of Computing (STOC'05), pp. 67–73 (2005)
7. Christodoulou, G., Koutsoupias, E., Nanavati, A.: Coordination Mechanisms. In Diaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 345–357, Springer, Heidelberg (2004)
8. Fabrikant, A., Papadimitriou, C.H., Talwar, K.: The Complexity of Pure Nash Equilibria. In: Proc. of the 36th Annual ACM Symposium on Theory of Computing (STOC'04), pp. 604–612 (2004)
9. Feige, U.: A Threshold of $\ln n$ for Approximating Set Cover. *Journal of the ACM* **45**, pp. 314–318 (1998)
10. Gairing, M., Schoppmann, F.: Total Latency in Singleton Congestion Games. In Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 381–387, Springer, Heidelberg (2007)
11. Giannakos, A., Gourvès, L., Monnot, J., Paschos, V.T.: On the Performance of Congestion Games for Optimum Satisfiability Problems. In Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 220–231, Springer, Heidelberg (2007)
12. Goemans, M.X., Williamson, D.P.: New $\frac{3}{4}$ -Approximation Algorithms for the Maximum Satisfiability Problem. *SIAM Journal on Discrete Mathematics* **7**, pp. 656–666 (1994)
13. Hochbaum, D.S.: *Approximation Algorithms for NP-Hard Problems*. PWS Publishing Company (1997)
14. Hochbaum, D.S., Pathria, A.: Analysis of the Greedy Approach in Problems of Maximum k -coverage. *Naval Research Logistics* **45**(6), pp. 615–627 (1998)
15. Immorlica, N., Li, L., Mirrokni, V.S., Schulz, A.S.: Coordination Mechanisms for Selfish Scheduling. Deng, X., Ye, Y. (eds.), WINE 2005. LNCS, vol 3828, pp. 55–69, Springer, Heidelberg (2005)
16. Johnson, D.S., Papadimitriou, C.H., Yannakakis, M.: How Easy is Local Search? *Journal of Computer and System Sciences* **37**(1), pp. 79–100 (1988)
17. Khanna, S., Motwani, R.: On Syntactic versus Computational Views of Approximability. *SIAM Journal on Computing* **28**(1), pp. 164–191 (1998)
18. Koutsoupias, E., Papadimitriou, C.H.: Worst-Case Equilibria. In Meinel, C., Tison, S. (eds.) STACS 1999. LNCS, vol. 1563, pp. 404–413, Springer, Heidelberg (1999)
19. Mavronicolas, M., Monien, B., Wagner, K.W.: Weighted Boolean Formula Games. In Deng, X., Graham, F.C. (eds.) WINE 2007. LNCS, vol. 4858, pp. 469 – 481, Springer, Heidelberg (2007)
20. Nash, J.F.: Non-Cooperative Games. *Annals of Mathematics* **54**(2), pp. 286–295 (1951)
21. Rosenthal, R.W.: A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory* **2**, pp. 65–67 (1973)
22. Roughgarden, T.: Intrinsic Robustness of the Price of Anarchy. In: Proc. of the 41st Annual ACM Symposium on Theory of Computing (STOC'09), pp. 513–522 (2009)
23. Vazirani, V.V.: *Approximation Algorithms*. Springer, Heidelberg (2003)
24. Vetta, A.: Nash Equilibria in Competitive Societies, with Applications to Facility Location, Traffic Routing and Auctions. In: Proc. of the 43rd Annual Symposium on Foundations of Computer Science (FOCS'02), pp. 416–425 (2002)