

Multi-Leader Congestion Games with an Adversary

Tobias Harks¹, Mona Henle², Max Klimm³, Jannik Matuschke⁴, Anja Schedel¹

¹ University of Augsburg, 86159 Augsburg, Germany

² University of Applied Sciences Augsburg, 86161 Augsburg, Germany

³ TU Berlin, 10623 Berlin, Germany

⁴ KU Leuven, 3000 Leuven, Belgium

tobias.harks@math.uni-augsburg.de, mona.henle@hs-augsburg.de, klimm@math.tu-berlin.de,
jannik.matuschke@kuleuven.be, anja.schedel@math.uni-augsburg.de

Abstract

We study a multi-leader single-follower congestion game where multiple users (leaders) choose one resource out of a set of resources and, after observing the realized loads, an adversary (single-follower) attacks the resources with maximum loads, causing additional costs for the leaders. For the resulting strategic game among the leaders, we show that pure Nash equilibria may fail to exist and therefore, we consider approximate equilibria instead. As our first main result, we show that the existence of a K -approximate equilibrium can always be guaranteed, where $K \approx 1.1974$ is the unique solution of a cubic polynomial equation. To this end, we give a polynomial time combinatorial algorithm which computes a K -approximate equilibrium. The factor K is tight, meaning that there is an instance that does not admit an α -approximate equilibrium for any $\alpha < K$. Thus $\alpha = K$ is the smallest possible value of α such that the existence of an α -approximate equilibrium can be guaranteed for *any* instance of the considered game. Secondly, we focus on approximate equilibria of a given fixed instance. We show how to compute efficiently a best approximate equilibrium, that is, with smallest possible α among all α -approximate equilibria of the given instance.

1 Introduction

Hierarchical leader-follower games have received considerable attention in the artificial intelligence community, especially, because several real-world applications related to the protection of vulnerable systems can be modeled within this framework. Applications include the security domain (Kiekintveld et al. 2009; Marchesi, Castiglioni, and Gatti 2019; Sinha et al. 2018; Gan, Elkind, and Wooldridge 2018), where a leader aims at protecting a set of valuable targets and moves first by applying a defender strategy such as controls or fortification of resources. The adversary acts as a follower and, after observing the leader’s defensive strategy, chooses a strategy incurring maximum damage. The leader anticipates the followers’ strategy. Thus, the computation of the defender strategy takes the follower reaction into account.

While most works in this literature consider the case of a single leader, the case of *multiple leaders* playing a simultaneous-move strategic game subject to one or more followers has received much less attention and only very few results with respect to the existence and computational complexity of equilibria are known. Note that the multi-leader

case applies to several scenarios, for example, in the analysis of deregulated electricity markets in which some of the large energy producers are the leaders and the smaller energy producers and independent system operator are the followers; see Leyffer and Munson (2010) and references therein. Also in the security domain related to transport and communication networks, there are usually multiple leaders that compete over the network resources subject to a followers’ response; see Kulkarni and Shanbhag (2015). One well-known obstacle in the analysis of multi-leader games—even in the realm of continuous formulations with convex action spaces for the leaders and single-valuedness of the followers’ response—is the inherent non-convexity of the best-response correspondence that results in non-existence of pure Nash equilibria; see Kulkarni and Shanbhag (2015).

In this paper, we consider a class of multi-leader single-follower games on discrete strategy spaces that are motivated by security applications with congestion effects. Consider a standard singleton congestion game where multiple users (leaders) choose one resource out of a set of resources. After observing the realized loads, an adversary (single-follower) attacks the resources with maximum loads, causing additional disutilities for the users on the attacked resources. The adversary may be thought of as either being a malicious player attacking the resources in order to maximize the caused damage or as controls by a central authority to counter tax or fare evasion; see Correa et al. (2017) for a related mathematical model of fare evasion without any congestion or load balancing effects. In both applications it is sensible to assume that the adversary has limited resources, modeled by a fixed budget for his interventions that can be distributed freely on the resources, and that he acts rationally, investing the budget only on resources with maximum load. The users anticipate this strategy. From their perspective, every maximum-load resource is equally likely under attack, that is, they assume the budget to be spent evenly among the resources with maximum load (this can be interpreted as a randomized strategy of the adversary choosing the uniform distribution over maximum-load resources). For the users, the additional cost term corresponds to the expected additional damage cost due to an attack.

This fundamental model has, to the best of our knowledge, not been analyzed before and we investigate the existence and computation of (approximate) pure Nash equilibria.

ria of this multi-leader single-follower game.

1.1 Our Results and Proof Techniques

We first observe that pure Nash equilibria do not always exist in the introduced game, not even for linear congestion costs. This motivates the analysis of approximate pure Nash equilibria, where any unilateral deviation cannot improve the cost of the deviating leader by more than a factor α , for some $\alpha \geq 1$. (Note that the adversary is still assumed to act optimally.) We analyze existence and efficient computation of approximate equilibria for the introduced game with linear congestion costs.

As our first main result, we show that $\alpha = K \approx 1.1974$ is the smallest possible value of α such that the existence of an α -approximate pure Nash equilibrium can always be guaranteed (K is the unique solution of some cubic polynomial equation). For the proof, we give an efficient algorithm which computes a K -approximate equilibrium. The basic approach is to start with an empty game, and add the players one after another, always placing them on a best-response resource. If the addition of a new player makes some of the already added players “unhappy”, meaning that there is a unilateral deviation decreasing their cost by more than a factor α , we let the unhappy players deviate one after another until all players are happy again, that is, an α -approximate pure Nash equilibrium is reached for the subset of players already added. Only then the next player is added. By choosing the possible deviations carefully, we can show that this procedure terminates after a polynomial number of steps for $\alpha = K$, showing existence of K -approximate pure Nash equilibria and giving an efficient way of computing them. A similar approach has been used before to compute *exact* pure Nash equilibria in the context of weighted congestion games (Milchtaich 1996; Ackermann, Röglin, and Vöcking 2009), but we are not aware of any results regarding approximate equilibria applying this technique. We furthermore provide an instance which does not admit an α -approximate pure Nash equilibrium for any $\alpha < K$. This shows that $\alpha = K$ is tight in the sense that it is the smallest possible value such that the existence of an α -approximate equilibrium can be guaranteed for *all* instances of the introduced game.

However, for a single given instance, better approximate equilibria might exist, that is, there may be α -approximate equilibria with $\alpha < K$. We show how to compute efficiently, for a given instance, a *best* approximate equilibrium, that is, an α -approximate equilibrium for the smallest value of α for which such an equilibrium exists. Note that this in particular implies that we can decide efficiently whether a given instance admits an exact pure Nash equilibrium, and in case of existence, we can also compute such an equilibrium. Our algorithm is based on a careful analysis of the structure of optimal approximate equilibria, which allows us to enumerate a polynomially-sized set of possible resource-load configurations, from which an optimal approximate equilibrium can then be found using a simple linear program.

1.2 Related Work

The game that we analyze in this paper constitutes a Stackelberg game with multiple leaders and a single follower. The leaders’ game is a singleton congestion game and we assume symmetric strategies, meaning that all leaders have the same strategy space. Stackelberg games with an underlying congestion game for (a subset of) the players have received considerable attention in the literature. Castiglioni et al. (2019) and Marchesi, Castiglioni, and Gatti (2019) consider a game with a *single* leader and multiple followers where all players participate in a congestion game (but the leader’s congestion cost functions may be different from the followers’). Depending on the structure of strategy spaces and congestion cost functions, they analyze the computational complexity of computing exact equilibria. In particular, they find that efficient algorithms are only possible for singleton strategy spaces (unless $P = NP$), and derive such algorithms for singleton strategy spaces where either all followers have the same strategies (Castiglioni et al. 2019), or the followers can be divided in “classes” having the same strategies (Marchesi, Castiglioni, and Gatti 2019).

There are several works analyzing hierarchical situations with a subsequent nonatomic network routing game, where a set of infinitesimally small players chooses paths in a network, and each player aims to minimize the (load-dependent) length of her chosen path leading to a Wardrop equilibrium (Wardrop 1952). For works analyzing situations where a single leader determines capacities or prices in order to reduce the total congestion (plus investments for the case of capacities) of the Wardrop equilibria in a subsequent network routing game, we refer to Marcotte (1986); Gairing, Harks, and Klimm (2017) for setting capacities, and Beckmann, McGuire, and Winsten (1956) and Yang and Huang (2004) for setting prices. Labbé, Marcotte, and Savard (1998) study a model where a single leader sets prices in order to maximize her profit in a subsequent network routing game (but without congestion effects). Harks, Schröder, and Vermeulen (2019) and Correa et al. (2018) consider a game where multiple leaders set prices in order to maximize their own profits achieved in a network routing game. The prices that the leaders are allowed to choose are upper-bounded by price caps (leader-specific in Correa et al. (2018), equal for all leaders in Harks, Schröder, and Vermeulen (2019)), and the two papers consider the (three-level) problem of a system designer who chooses the cap(s) in order to minimize total congestion. Finally, models where multiple leaders choose prices and capacities to maximize their individual profits achieved in a network routing game are for example analyzed by Johari, Weintraub, and Van Roy (2010), Liu, Chen, and Huang (2011), and Harks and Schedel (2019).

Regarding the computation of approximate equilibria in atomic congestion games, we refer to Caragiannis et al. (2011, 2015). Finally, we also mention here congestion games with an adversarial structure such as agent or resource failures, see Bilò, Moscardelli, and Vinci (2018); Meir et al. (2012); Li et al. (2017) or games with malicious players (Babaiouf, Kleinberg, and Papadimitriou 2009).

load profile	deviation (of some player using r to r' , notation $r \rightarrow r'$)	resulting cost improvement
(5, 0, 0)	$r_1 \rightarrow r_2$	$5a_{r_1} + B = 6 > 2 = a_{r_2}$
(4, 1, 0)	$r_1 \rightarrow r_2$	$4a_{r_1} + B = 6 > 4 = 2a_{r_2}$
(3, 2, 0)	$r_1 \rightarrow r_3$	$3a_{r_1} + B = 6 > 5 = a_{r_3}$
(3, 1, 1)	$r_3 \rightarrow r_2$	$a_{r_3} = 5 > 4 = 2a_{r_2}$
(2, 2, 1)	$r_2 \rightarrow r_1$	$2a_{r_2} + B/2 = 7 > 6 = 3a_{r_1} + B$

Table 1: Improving deviations for the candidate profiles for the game in Example 1.

2 The Model

For an integer $k \in \mathbb{Z}_{\geq 0}$, let $[k] := \{1, \dots, k\}$. Let $N = [n]$ be a finite set of players (leaders) and $R = \{r_1, \dots, r_m\}$ be a finite set of m resources. For each player i , the set of strategies available to player i is $X_i = R$. We call $x = (x_1, \dots, x_n)$ with $x_i \in X_i$ for all $i \in N$ a *strategy profile*, and $X = X_1 \times \dots \times X_n$ the *strategy space*.

We use standard game theory notation; for a strategy profile $x \in X$, we write $x = (x_i, x_{-i})$ meaning that x_i is the strategy that player i plays in x and x_{-i} is the partial strategy of all players except i . Every strategy profile $x = (x_1, \dots, x_n) \in X$ induces a *load* or *congestion* on the resources given by $\ell_r(x) := |\{i \in N \mid x_i = r\}|$, $r \in R$. We are further given linear cost functions $a_r \ell_r(x)$, $r \in R$ with nonnegative coefficients $a_r \geq 0$. In classical congestion games, the private cost of player i under strategy profile $x \in X$ is defined as $\pi_i(x) = a_{x_i} \ell_{x_i}(x)$. Now we model the actions of an *adversary* (follower) after the leaders have chosen their joint strategy profile $x \in X$. Formally, given $x \in X$ the adversary solves

$$\max \sum_{r \in R} \ell_r(x) \kappa_r \quad \text{s.t.:} \quad \sum_{r \in R} \kappa_r \leq B, \quad \kappa \geq 0. \quad (LP)$$

The linear program (LP) has the interpretation that the adversary has a budget of $B > 0$ that can be freely distributed among the resources. For each unit of budget spent on a resource, the adversary receives a utility equal to the number of players on that resource since any interaction with a leader on a resource is equally beneficial for the follower. Thus, the adversary strategically selects those resources that are used by the maximum number of players in order to maximize the caused damage. It is not hard to see that these are precisely the optimal solutions to (LP). While (LP) may have multiple optimal solutions, a reasonable selection among the optimal solutions is the following, where we use the notation $M(x) := \max_{r \in R} \{\ell_r(x)\}$ together with $M^{-1}(x) := \arg \max_{r \in R} \{\ell_r(x)\}$:

$$\kappa_r^*(x) = \begin{cases} \frac{B}{|M^{-1}(x)|}, & \text{if } r \in M^{-1}(x), \\ 0, & \text{else.} \end{cases} \quad (1)$$

Clearly, $\kappa_r^*(x)$ is an optimal solution to (LP) and has the intuitive interpretation that, assuming that every maximum-load resource is equally likely to be under attack by the adversary, from the perspective of the players it represents the expected additional resource cost due to an attack.

The multi-leader congestion game with an adversary can be defined as the game $G = (N, X, B, \pi)$ in strategic form,

where

$$\pi_i(x) := a_{x_i} \ell_{x_i}(x) + \kappa_{x_i}^*(x). \quad (2)$$

We furthermore define

$$c_r(x) := a_r \ell_r(x) + \kappa_r^*(x) \quad (3)$$

for all $r \in R$ (this is useful in case we do not want to consider a specific player using r).

A strategy profile $x \in X$ is called a *pure Nash equilibrium* (PNE) of G if for all $i \in N$:

$$\pi_i(x) \leq \pi_i(y_i, x_{-i}) \quad \text{for all } y_i \in X_i.$$

Let us give an example of a multi-leader congestion game with an adversary showing that pure Nash equilibria need not exist in general.

Example 1. Consider the game with $m = 3$ resources, $n = 5$ players, budget $B = 6$ and resource cost coefficients

$$a_{r_1} = 0, \quad a_{r_2} = 2, \quad \text{and } a_{r_3} = 5.$$

We proceed to show that this game has no pure Nash equilibrium. To this end, we show for each strategy profile x that there exists a player who can decrease her cost by a unilateral deviation from x . Since the players are symmetric, it suffices to analyze all “load” profiles $(\ell_{r_1}(x), \ell_{r_2}(x), \ell_{r_3}(x))$, that is, all vectors $(\ell_{r_1}, \ell_{r_2}, \ell_{r_3}) \in \mathbb{N}^3$ with $\ell_{r_1} + \ell_{r_2} + \ell_{r_3} = n = 5$. Since $a_{r_1} < a_{r_2} < a_{r_3}$, the only candidates for a PNE are those strategy profiles x with $\ell_{r_1}(x) \geq \ell_{r_2}(x) \geq \ell_{r_3}(x)$. Thus, it suffices to show that there is no PNE among the five load profiles (5, 0, 0), (4, 1, 0), (3, 2, 0), (3, 1, 1) and (2, 2, 1). Table 1 provides an improving deviation for each of these candidate profiles, showing that the game has no PNE.

This example motivates the analysis of approximate equilibria defined as follows.

Definition 2. A strategy profile $x \in X$ is an α -approximate pure Nash equilibrium (α -PNE) of G for some $\alpha \geq 1$, if for all $i \in N$:

$$\pi_i(x) \leq \alpha \cdot \pi_i(y_i, x_{-i}) \quad \text{for all } y_i \in X_i.$$

A unilateral deviation which decreases the cost of the deviating player more than a factor α is called an α -improving deviation, or an α -improving move.

For $\alpha = 1$, we obtain the standard PNE. For general $\alpha \geq 1$, the interpretation is that no player can improve her cost by a unilateral deviation gaining more than a factor α . We remark that, while one can similarly define *additively* approximate equilibria, no existence guarantees can be given for such equilibria for any additive constant due to the scale-invariance of the games studied in this article (see Appendix F for details).

Algorithm 1: Computation of an α -approximate PNE.

Input: Player set $N = [n]$, resource set $R = \{r_1, \dots, r_m\}$, resource cost coefficients $0 \leq a_1 \leq \dots \leq a_m$, and $\alpha \geq 1$.

Output: α -approximate pure Nash equilibrium x .

```
1  $x \leftarrow (0, 0, \dots, 0)$ ;  
2 for  $k = 1, 2, \dots, n$  do  
3    $k' \leftarrow \min\{i \in [m] : c_{r_i}(x_{-k}, r_i) \leq c_r(x_{-k}, r) \forall r \in R\}$ ;  
4    $x_k \leftarrow r_{k'}$ ;  
5   while  $W_\alpha(x) \neq \emptyset$  do  
6      $R_\alpha(x) \leftarrow \{r \in R : x_i = r \text{ for some } i \in W_\alpha(x)\}$ ;  
7      $i^- \leftarrow \max\{j \in [m] : r_j \in R_\alpha(x) \text{ and } c_{r_j}(x) \geq c_r(x) \forall r \in R_\alpha(x)\}$ ;  
8      $i \leftarrow$  some player with  $x_i = r_{i^-}$ ;  
9      $i^+ \leftarrow \min\{j \in [m] : c_{r_j}(x_{-i}, r_j) \leq c_r(x_{-i}, r) \forall r \in R\}$ ;  
10     $x_i \leftarrow r_{i^+}$ ;
```

3 Computing K -approximate PNE

In this section, we analyze approximate PNE of the introduced multi-leader congestion games with an adversary. As Example 1 shows, existence of (exact) PNE can not be guaranteed for these games. We show that $\alpha = K \approx 1.1974$ is the smallest possible α such that the existence of an α -PNE can be guaranteed for any instance, where

$$K := \frac{1}{6} \left(1 + \sqrt[3]{109 - 6\sqrt{330}} + \sqrt[3]{109 + 6\sqrt{330}} \right) \quad (4)$$

is the unique solution of the equation $-x^3 + x^2/2 + 1 = 0$. To this end, we provide Algorithm 1 which efficiently computes a K -approximate PNE (see Theorem 3). This result is complemented by an instance where no α -approximate PNE with $\alpha < K$ exists (see Theorem 6). As an easy consequence of the proof, we also get that exact PNE are guaranteed to exist if $n \leq 4$ or $m \leq 2$ (see Corollary 5).

3.1 An Algorithm for Computing α -Approximate Equilibria

For computing an α -approximate PNE, we use the following basic approach. Starting with an empty game, we add the players one after another to the game, where a newly added player is always placed on a best response. If the addition of a new player makes some of the earlier added players “unhappy”, meaning that they now have an α -improving move, we let unhappy players deviate one after another to a best response until all players are happy again. Only then we add the next player, etc.

Note that this approach has been used before to compute *exact* PNE, for example for player-specific costs or weighted congestion games on matroids (Milchtaich 1996; Ackermann, Röglin, and Vöcking 2009). However, to the best of our knowledge, it has not been utilized in the context of approximate PNE. We believe that this technique will be useful for showing existence and computing approximate equilibria beyond the class of games that we analyze here.

For the formal description of our algorithm see Algorithm 1. We assume that the resource set $R = \{r_1, \dots, r_m\}$ is ordered such that $a_1 \leq \dots \leq a_m$ (where $a_j := a_{r_j}$). Since we need to consider strategy profiles for subsets of

the players, let us extend the notion of a strategy profile to the set of vectors $x \in (R \cup \{0\})^n$, where $x_i = 0$ means that player i has not yet been added to the game. Given such a vector $x \neq 0$, the cost $c_{x_i}(x)$ incurred to player i with $x_i \neq 0$ is defined as the cost which is experienced by her in the game where only the players j with $x_j \neq 0$ are present. Now define, for a strategy profile $x \in (R \cup \{0\})^n$, the following set of players $W_\alpha(x)$ who are “unhappy” with their strategy, meaning that they have an α -improving move:

$$W_\alpha(x) := \{i \in N : x_i \neq 0 \text{ and } c_{x_i}(x) > \alpha \cdot c_r(x_{-i}, r) \text{ for a resource } r \in R\}.$$

In each iteration of the for-loop in line 2 of Algorithm 1, a new player k is added to the game. We place k on a best response (with respect to the strategies of the players $\{1, \dots, k-1\}$ already added to the game). If there is more than one best response, we choose the one with smallest index, see lines 3 and 4. After having added player k , it may be the case that some players are not happy with their strategy, that is, $W_\alpha(x) \neq \emptyset$, where x denotes the current strategy profile. In the while-loop starting in line 5, we iteratively choose a player who is “maximally unhappy”, meaning that she experiences maximum cost among all unhappy players, and let her deviate to a best response, until all players in $\{1, \dots, k\}$ are happy. If there is more than one resource with maximum cost among the resources used by unhappy players, we choose a player on the maximum cost resource with largest index, and if there is more than one best response, we choose the one with smallest index (see lines 7, 8 and 9). After this, we return to line 2 where the next player $k+1$ is added to the game.

It is clear that if Algorithm 1 terminates, the computed strategy profile is an α -PNE. In the next subsection, we show that Algorithm 1 terminates for $\alpha = K$. Note that the special choices made by the algorithm in case of non-unique best responses or non-unique most expensive resources, together with the fact that resources are ordered such that $a_1 \leq a_2 \leq \dots \leq a_m$, ensure that the loads are always decreasing along the resources, that is, $\ell_{r_1}(x) \geq \ell_{r_2}(x) \geq \dots \geq \ell_{r_m}(x)$ always holds during the algorithm.

Appendix A contains an illustrating example for the application of Algorithm 1.

3.2 Termination of Algorithm 1 for $\alpha = K$

In this subsection, we show that Algorithm 1 terminates for $\alpha = K$, and thus computes a K -approximate PNE. We also show that the running time of Algorithm 1 can be bounded by $\min\{O(n^2m), O(nm^2)\}$.

We need to prove that the while-loop terminates in each iteration $k \in \{1, \dots, n\}$ of the for-loop. For $k = 1$, the while-loop obviously terminates since we only have one player, and she is placed on the best response r_1 . Therefore, $W_K(x) = W_K(r_1, 0, \dots, 0) = \emptyset$ and the while-loop terminates without changing anything. For $k = 2$, the second player is placed on r_1 if $2a_1 + B \leq a_2 + B/2$, and is placed on r_2 otherwise. It is easy to see that in both cases, both players are happy and the while-loop immediately terminates.

For iteration $k \geq 3$, it may be the case that some players change their strategy during the while-loop. We show inductively that the while-loop also terminates in iteration $k \geq 3$. Note that since the while-loop terminated in iteration $k - 1$, the players in $\{1, \dots, k - 1\}$ were happy before the next player k was added to the game. That is, with respect to the profile x directly before player k is added, no player $i \in \{1, \dots, k - 1\}$ has a K -improving deviation, i.e., an alternative resource r with $c_{x_i}(x) > K \cdot c_r(x_{-i}, r)$. But it may be the case that some players are unhappy after the addition of player k . Additionally, the deviation of one of these players may cause further players to be unhappy. For the termination of the while-loop, we need to show that after finitely many deviations, all players in $\{1, \dots, k\}$ are happy again. We thus need to keep track of the set of unhappy players during the course of the while-loop. To this end, we derive necessary properties for players who either become unhappy due to the addition of player k , or due to the subsequent deviation of some other player during the while-loop (see Lemma B.5 and Lemma B.6 in the appendix). The derived properties are mostly in terms of loads, e.g., conditions on the load of an unhappy player i 's current strategy x_i , and on the load of a corresponding K -improving deviation r . Using these two lemmas, as well as some further structural insights (see Appendix B), we then proceed by a careful case distinction regarding the sequence of deviating players in iteration k , and show that this sequence terminates in all cases. Let us now briefly sketch the mentioned case distinction (the complete proof of Theorem 3 can be found in Appendix C).

Let x and x' denote the profiles directly before and after the new player k is added. If all players are happy with their strategy in x' , the statement follows; thus assume that player i changes from $x'_i = x_i$ to r in the while-loop. Using Lemma B.5, we know that there are three possible cases regarding the load of player k 's current strategy x'_k . Namely, the load $\ell_{x'_k}(x)$ of x'_k (before player k is added) needs to be in $\{M - 2, M - 1, M\}$, where $M := M(x)$ denotes the maximum load with respect to x . We then analyze all three cases. As it turns out, the cases $\ell_{x'_k}(x) = M - 2$ and $\ell_{x'_k}(x) = M - 1$ are very simple, whereas $\ell_{x'_k}(x) = M$ is more complicated and requires further subcases. However, by repeated use of the lemmas contained in Appendix B, in particular Lemma B.6, we can show termination of the while-loop also for this case.




Figure 1: A cycling sequence of deviations which might occur during the while-loop of Algorithm 1: First, some player moves from r_3 to r_2 , then another player changes her strategy from r_1 to r_3 , and finally, a player using r_2 moves to r_1 .

At the end of this proof sketch, we want to briefly indicate the role of the constant K . To this end, consider Figure 1 which shows a simplified version of a subcase occurring in the proof. In particular, the displayed sequence of deviations might occur during the while-loop of the algorithm, and this implies that the following three inequalities need to hold:

$$\begin{aligned} a_3 &> \alpha \cdot 2a_2 \\ 3a_1 + B &> \alpha \cdot a_3 \\ 2a_2 \cdot + B/2 &> \alpha \cdot (3a_1 + B) \end{aligned}$$

From this, one can derive $(1 + \alpha^2/2 - \alpha^3)B > (\alpha^3 - 1)3a_1$, which yields a contradiction for $\alpha = K$ and any $a_1 > 0$, as the left-hand side is equal to 0, whereas the right-hand side is non-negative. Thus, cycles of this form cannot occur during the algorithm.

For a complete proof of Theorem 3, see Appendix C.

Theorem 3. For $\alpha = K$, where $K = 1/6 \cdot (1 + \sqrt[3]{109 - 6\sqrt{330}} + \sqrt[3]{109 + 6\sqrt{330}}) \approx 1.1974$ is the unique solution of the equation $-x^3 + x^2/2 + 1 = 0$, Algorithm 1 computes a K -approximate PNE.

The proof of Theorem 3 also yields the following upper bound on the running time of Algorithm 1.

Corollary 4. For $\alpha = K$, the running time of Algorithm 1 can be bounded by $\min\{O(n^2m), O(nm^2)\}$.

Proof. First note that each iteration of the while-loop can be implemented in $O(m)$ (note that although the cost of deviating to a resource r is in general player-specific, since it depends on the load of the deviating player's current resource, it can in fact only be different for two players if one of these players is using a resource with maximum load, and the other not, cf. Lemma B.1). Furthermore, in the k th iteration of the for-loop, we can bound the number of iterations of the while-loop either by $O(m)$, or alternatively by $2k$ since no player moves more than twice, see the proof of Theorem 3 in Appendix C. Since there are n iterations of the for-loop, we get $\min\{O(n^2), O(nm)\}$ as an upper bound for the total number of iterations of the while-loop and this yields the given bound on the total running time. \square

The proof of Theorem 3 also reveals that $n \geq 5$ and $m \geq 3$ need to hold for any instance which has no exact PNE. This follows from the fact that the case displayed in Figure 1 essentially is the only situation where the while-loop might not terminate for $\alpha = 1$, and this case requires at

load profile	α -improving deviation (of some player using r to r' , notation $r \rightarrow r'$)	conditions on α
(5, 0, 0)	$r_1 \rightarrow r_2$	$5a_1 + B > \alpha a_2 \Leftrightarrow \alpha < 1/a_2 = \frac{2}{K-1/2} \approx 2.86$
(4, 1, 0)	$r_1 \rightarrow r_2$	$4a_1 + B > \alpha 2a_2 \Leftrightarrow \alpha < 1/(2a_2) = \frac{1}{K-1/2} \approx 1.43$
(3, 2, 0)	$r_1 \rightarrow r_3$	$3a_1 + B > \alpha a_3 \Leftrightarrow \alpha < 1/a_3 = K$
(3, 1, 1)	$r_3 \rightarrow r_2$	$a_3 > \alpha 2a_2 \Leftrightarrow \alpha < \frac{1}{K(K-1/2)} = K$
(2, 2, 1)	$r_2 \rightarrow r_1$	$2a_2 + B/2 > \alpha(3a_1 + B) \Leftrightarrow \alpha < 2a_2 + 1/2 = K$

Table 2: α -improving deviations for the candidate profiles.

least five players and at least three resources (for a complete proof, see Appendix C). Thus we get the following corollary.

Corollary 5. *Exact PNE are guaranteed to exist if $n \leq 4$ or $m \leq 2$.*

3.3 Tightness of $\alpha = K$

In this subsection, we provide an instance where no α -approximate PNE with $\alpha < K$ exists, showing that $\alpha = K$ is the smallest possible value such that the existence of an α -approximate equilibrium can be guaranteed.

Theorem 6. *There exists an instance with three resources and five players such that there is no α -approximate PNE for any $\alpha < K$, where $K \approx 1.1974$ is as specified in (4).*

Proof. Consider the instance with $m = 3$ resources, $n = 5$ players, budget $B = 1$, and resource cost coefficients

$$a_1 := a_{r_1} = 0, \quad a_2 := a_{r_2} = K/2 - 1/4 \approx 0.3487, \quad \text{and} \\ a_3 := a_{r_3} = 1/K \approx 0.8351.$$

We proceed to show that there is no α -approximate PNE for $\alpha < K$. To this end, note that it suffices to show that there is no α -approximate PNE among the five load profiles (5, 0, 0), (4, 1, 0), (3, 2, 0), (3, 1, 1), (2, 2, 1) (since $a_1 < a_2 < a_3$; if there exists an α -approximate PNE, there is also one with corresponding load profile among the five listed load profiles, see Lemma E.1). Let $\alpha < K$. We show that for any of the five load profiles, there exists an α -improving deviation, showing the claim. To this end, consider Table 2, where we provide a deviation for each candidate profile which is α -improving if the given conditions on α are satisfied. It is easy to check that these conditions are indeed fulfilled for $\alpha < K$. \square

4 Computing Optimal Approximate Equilibria

In the last section, we showed that $\alpha = K$ is the smallest possible value for α such that the existence of an α -approximate PNE can be *guaranteed* for any instance of a multi-leader congestion game with an adversary. However, there are clearly instances where α -PNE with $\alpha < K$ exist (in particular, all instances exhibiting an exact PNE). We show in this section how to compute efficiently a best approximate PNE for a given instance, that is, with smallest possible α such that an α -PNE exists for *this* instance.

To this end, consider a multi-leader congestion game with an adversary with resource set $R = [m]$ and $a_1 \leq \dots \leq a_m$. We can restrict our attention to strategy profiles x with decreasing loads, that is, with $\ell_1(x) \geq \dots \geq \ell_m(x)$, since if an α -PNE exists, there also exists one with decreasing loads (see Lemma E.1 and note that we can assume $\alpha \in [1, K]$ since we want to find the smallest possible α). Thus let x be a strategy profile with $\ell_1(x) \geq \dots \geq \ell_m(x)$. Note that, clearly, $M(x) = \max\{\ell_r(x) : r \in R\} \geq \lceil \frac{n}{m} \rceil$ holds. Furthermore, if $M(x)m = n$ (which is equivalent to $\ell_r(x) = M(x)$ for all $r \in R$), x is an α -PNE if and only if $a_m \cdot M(x) + B/m \leq \alpha(a_1 \cdot (M(x) + 1) + B)$ holds. Thus we can assume in the following that there are resources with load $< M(x)$. We denote by $k = k(x) < m$ the largest resource having maximum load $M = M(x)$. Similarly, $k' = k'(x)$ denotes the smallest resource with load strictly smaller than $M - 1$, and $k'' = k''(x)$ denotes the smallest resource with load strictly smaller than $M - 2$. In other words, $\ell_r(x) = M$ for all $r \in \{1, \dots, k\}$, $\ell_r(x) = M - 1$ for all $r \in \{k+1, \dots, k'-1\}$, $\ell_r(x) = M - 2$ for all $r \in \{k', \dots, k''-1\}$, and $\ell_r(x) \leq M - 3$ for all $r \in \{k'', \dots, m\}$, see Figure 2 for illustration. Note that $k' = k + 1$ or $k' = k''$ are possible, in which case there are no resources with load $M - 1$ or $M - 2$, respectively.

We now define the following values $\bar{c}_M = \bar{c}_M(x)$ and $\bar{c}_{<M} = \bar{c}_{<M}(x)$, which essentially describe the cost of a best alternative resource for a player using a resource with load M , and load smaller than M , respectively:

$$\bar{c}_M = \begin{cases} \min\{a_{k+1} \cdot M + B/k \text{ if } k' \geq k + 2, \\ a_{k'} \cdot (M - 1) + B/k' \text{ if } k' < k'', \\ a_r \cdot (\ell_r(x) + 1) \forall r \geq k''\}, & \text{if } k = 1, \\ \min\{a_1 \cdot (M + 1) + B, \\ a_{k+1} \cdot M + B/k \text{ if } k' \geq k + 2, \\ a_{k'} \cdot (M - 1) \text{ if } k' < k'', \\ a_r \cdot (\ell_r(x) + 1) \forall r \geq k''\}, & \text{if } k \geq 2, \end{cases}$$

and

$$\bar{c}_{<M} = \min\{a_1 \cdot (M + 1) + B, \\ a_{k+1} \cdot M + B/(k + 1) \text{ if } k' \geq k + 2, \\ a_{k'} \cdot (M - 1) \text{ if } k' < k'', \\ a_r \cdot (\ell_r(x) + 1) \forall r \geq k''\}.$$




Figure 2: Illustration for the definition of k , k' and k'' .

Using these definitions, x is an α -approximate PNE iff

$$a_r \cdot \ell_r(x) \leq \alpha \cdot \bar{c}_{<M} \text{ for all } r > k, \text{ and} \quad (5)$$

$$a_k \cdot M + B/k \leq \alpha \cdot \bar{c}_M. \quad (6)$$

Regarding this, note that there are some cases in which \bar{c}_M or $\bar{c}_{<M}$ do not denote the cost of a best alternative for some players, but only provide a lower bound. This happens in two cases. Firstly, if $k \geq 2$ and the minimum in the definition of \bar{c}_M is uniquely attained at $a_1 \cdot (M+1) + B$, then \bar{c}_M is in general not the cost of a best alternative for players currently using resource $r = 1$. However, it is then also clear that these players are satisfied with their strategies anyways. Similarly, if the minimum in the definition of $\bar{c}_{<M}$ is uniquely attained for some resource $r > k$, then $\bar{c}_{<M}$ does not denote the cost of a best alternative for the players using r . But it is again clear that these players do not want to deviate.

Also note that the minimum properties of \bar{c}_M and $\bar{c}_{<M}$, as well as the Nash condition (5), yield upper and lower bounds for the loads $\ell_r(x)$ of all resources $r \geq k''$ (for $r < k''$, the load $\ell_r(x)$ is uniquely determined by the definitions of M, k, k' and k''). By using this, as well as the fact that for given values of M, k, k' and k'' , there are only polynomially many possible values for \bar{c}_M and $\bar{c}_{<M}$, we can show the following result (for a complete proof, see Appendix D):

Lemma 7. *Given $M \in \{\lceil \frac{n}{m} \rceil, \dots, n\}$, $k \in \{1, \dots, m-1\}$, $k' \in \{k+1, \dots, m\}$ and $k'' \in \{k', \dots, m\}$, as well as $\alpha \in [1, K]$, we can decide efficiently if there exists an α -approximate PNE x such that $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$ holds.*

In case of existence, we can furthermore compute a corresponding load vector efficiently.

Using Lemma 7, we get the main result of this section:

Theorem 8. *We can efficiently compute the smallest possible α such that an α -approximate PNE exists, as well as a corresponding load vector.*

Proof. First note that if x is an α -PNE corresponding to the best possible α , there need to be resources $r, r' \in R$ (as well as a player i with $x_i = r$) such that $a_r \ell_r(x) + \kappa_r^*(x) = \alpha \cdot (a_{r'} \ell_{r'}(r', x_{-i}) + \kappa_{r'}^*(r', x_{-i}))$ (otherwise α cannot be smallest possible), or, equivalently,

$$\alpha = \frac{a_r \ell_r(x) + \kappa_r^*(x)}{a_{r'} \ell_{r'}(r', x_{-i}) + \kappa_{r'}^*(r', x_{-i})}.$$

Since there are only $O(nm^2)$ many possible values for $a_r \ell_r(x) + \kappa_r^*(x)$, as well as for $a_{r'} \ell_{r'}(r', x_{-i}) + \kappa_{r'}^*(r', x_{-i})$, there are only $O(n^2 m^4)$ many possible values for α .

The result then follows by applying the underlying procedure of Lemma 7 for all $M \in \{\lceil \frac{n}{m} \rceil, \dots, n\}$, $k \in \{1, \dots, m-1\}$, $k' \in \{k+1, \dots, m\}$ and $k'' \in \{k', \dots, m\}$, as well as the $O(n^2 m^4)$ many possible values for α (thus in total $O(n^3 m^7)$ many times). \square

5 Conclusion

We introduced a multi-leader congestion game with an adversary which is motivated by security applications with congestion effects. Since PNE do not exist in general, we studied approximate equilibria. Our first main result shows that a K -approximate PNE always exists, where $K \approx 1.1974$ is the unique solution of a cubic polynomial equation. To this end, we presented an efficient algorithm which computes a K -approximate PNE. Furthermore, we showed that the factor K is tight by providing an instance where no α -approximate PNE with $\alpha < K$ exists. However, for a *specific* instance there might be a better α -approximate PNE, i.e., with $\alpha < K$. We presented an efficient procedure that computes a *best* approximate PNE of a given instance.

Our work also suggests several interesting directions for further research regarding multi-leader congestion games with an adversary. For example, one could analyze whether the results from Section 3 continue to hold if one allows more general strategy spaces in the leaders' congestion game. A first natural generalization in this regard would be to consider asymmetric strategies, and/or bases of matroids. It would furthermore be interesting to analyze the quality of approximate PNE. For example, one could measure the social cost of a strategy profile by the total cost of all players, and then compare a (best or worse) approximate PNE to a social optimum.

Acknowledgements We thank the anonymous referees for their comments that helped to improve the presentation of the paper. This work was supported by Deutsche Forschungsgemeinschaft (DFG – German Research Foundation) under grants HA 8041/4-1, MA 8439/1-1, and under Germany’s Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

References

- Ackermann, H.; Röglin, H.; and Vöcking, B. 2009. Pure Nash equilibria in player-specific and weighted congestion games. *Theoret. Comput. Sci.*, 410(17): 1552–1563.
- Babaioff, M.; Kleinberg, R.; and Papadimitriou, C. H. 2009. Congestion games with malicious players. *Games Econ. Behav.*, 67(1): 22–35.
- Beckmann, M.; McGuire, C.; and Winsten, C. 1956. *Studies in the Economics and Transportation*. New Haven, CT, USA: Yale University Press.
- Bilò, V.; Moscardelli, L.; and Vinci, C. 2018. Uniform Mixed Equilibria in Network Congestion Games with Link Failures. In Chatzigiannakis, I.; Kaklamanis, C.; Marx, D.; and Sannella, D., eds., *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic*, volume 107 of *LIPICs*, 146:1–146:14. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Caragiannis, I.; Fanelli, A.; Gravin, N.; and Skopalik, A. 2011. Efficient Computation of Approximate Pure Nash Equilibria in Congestion Games. In Ostrovsky, R., ed., *IEEE 52nd Annual Symposium on Foundations of Computer Science, FOCS 2011, Palm Springs, CA, USA, October 22-25, 2011*, 532–541. IEEE Computer Society.
- Caragiannis, I.; Fanelli, A.; Gravin, N.; and Skopalik, A. 2015. Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure. *ACM Trans. Economics and Comput.*, 3(1): 2:1–2:32.
- Castiglioni, M.; Marchesi, A.; Gatti, N.; and Coniglio, S. 2019. Leadership in singleton congestion games: What is hard and what is easy. *Artif. Intell.*, 277: 103177.
- Correa, J. R.; Guzmán, C.; Lianeas, T.; Nikolova, E.; and Schröder, M. 2018. Network Pricing: How to Induce Optimal Flows Under Strategic Link Operators. In *Proc. 19th ACM Conf. Electronic Commerce (EC)*, 375–392.
- Correa, J. R.; Harks, T.; Kreuzen, V. J. C.; and Matuschke, J. 2017. Fare Evasion in Transit Networks. *Oper. Res.*, 65(1): 165–183.
- Gairing, M.; Harks, T.; and Klimm, M. 2017. Complexity and Approximation of the Continuous Network Design Problem. *SIAM J. Optim.*, 27(3): 1554–1582.
- Gan, J.; Elkind, E.; and Wooldridge, M. J. 2018. Stackelberg Security Games with Multiple Uncoordinated Defenders. In André, E.; Koenig, S.; Dastani, M.; and Sukthankar, G., eds., *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2018, Stockholm, Sweden*, 703–711. International Foundation for Autonomous Agents and Multiagent Systems Richland, SC, USA / ACM.
- Harks, T.; and Schedel, A. 2019. Capacity and Price Competition in Markets with Congestion Effects. In *Proc. 15th Internat. Conference on Web and Internet Econom.*, 341.
- Harks, T.; Schröder, M.; and Vermeulen, D. 2019. Toll caps in privatized road networks. *Eur. J. Oper. Res.*, 276(3): 947–956.
- Johari, R.; Weintraub, G. Y.; and Van Roy, B. 2010. Investment and Market Structure in Industries with Congestion. *Oper. Res.*, 58(5): 1303–1317.
- Kiekintveld, C.; Jain, M.; Tsai, J.; Pita, J.; Ordóñez, F.; and Tambe, M. 2009. Computing Optimal Randomized Resource Allocations for Massive Security Games. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, AAMAS ’09*, 689–696. Richland, SC: International Foundation for Autonomous Agents and Multiagent Systems.
- Kulkarni, A. A.; and Shanbhag, U. V. 2015. An Existence Result for Hierarchical Stackelberg v/s Stackelberg Games. *IEEE Transactions on Automatic Control*, 60(12): 3379–3384.
- Labbé, M.; Marcotte, P.; and Savard, G. 1998. A Bilevel Model of Taxation and Its Application to Optimal Highway Pricing. *Management Science*, 44(12): 1608–1622.
- Leyffer, S.; and Munson, T. 2010. Solving multi-leader-common-follower games. *Optim. Methods Softw.*, 25(4): 601–623.
- Li, Y.; Jia, Y.; Tan, H.; Wang, R.; Han, Z.; and Lau, F. C. M. 2017. Congestion Game With Agent and Resource Failures. *IEEE Journal on Selected Areas in Communications*, 35(3): 764–778.
- Liu, T.-L.; Chen, J.; and Huang, H.-J. 2011. Existence and efficiency of oligopoly equilibrium under toll and capacity competition. *Transportation Research Part E: Logistics and Transportation Review*, 47(6): 908–919.
- Marchesi, A.; Castiglioni, M.; and Gatti, N. 2019. Leadership in Congestion Games: Multiple User Classes and Non-Singleton Actions. In Kraus, S., ed., *Proc. 28th Internat. Joint Conf. Artif. Intell. (IJCAI)*, 485–491.
- Marcotte, P. 1986. Network Design Problem with Congestion Effects: A Case of Bilevel Programming. *Math. Program., Ser. A*, 34: 142–162.
- Meir, R.; Tennenholtz, M.; Bachrach, Y.; and Key, P. B. 2012. Congestion Games with Agent Failures. In Hoffmann, J.; and Selman, B., eds., *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence, July 22-26, 2012, Toronto, Ontario, Canada*. AAAI Press.
- Milchtaich, I. 1996. Congestion Games with Player-Specific Payoff Functions. *Games Econom. Behav.*, 13(1): 111–124.
- Sinha, A.; Fang, F.; An, B.; Kiekintveld, C.; and Tambe, M. 2018. Stackelberg Security Games: Looking Beyond a Decade of Success. In Lang, J., ed., *Proc. 27th Internat. Joint Conf. Artif. Intell. (IJCAI)*, 5494–5501.


Wardrop, J. 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers*, 1(Part II): 325–378.

Yang, H.; and Huang, H.-J. 2004. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Res.*, 38(B): 1–15.


A An Illustrating Example for Algorithm 1

Example A.1. Consider a game with seven players and five resources with $a_1 = 1$, $a_2 = 4$, $a_3 = 4$, $a_4 = 10$, and $a_5 = 10$. Let $B = 9$ and $\alpha = K \approx 1.1974$ as specified in (4). We now describe the steps performed by Algorithm 1, see also Figure 3. Note that in the first six iterations of the for-loop, no players change during the while-loop. Only after player 7 is added, the players 5 and 6 become unhappy. Since the players 5 and 6 currently experience the same cost and player 6 uses the resource with a larger index, player 6 changes from r_5 to r_2 due to the tie-breaking rule of the algorithm. Note that the resources r_2 and r_3 both constitute a best response but r_2 has the smaller index. The deviation of player 6 causes the players on r_1 to be unhappy since they can improve their cost from 12 to 10 (more than a factor K) by deviating to r_3 . Furthermore player 5 can improve her cost from 10 to 8. Since $12 > 10$ a player on r_1 changes due to the tie-breaking rule of the algorithm. The resulting strategy profile if one of these players, say player 7, deviates to r_3 , is a K -approximate PNE since no player has a K -improving deviation anymore. (In fact, the profile even is a $12.5/12 = 25/24$ -PNE.)


Note that if some player now deviates from r_2 to r_1 (which is a best response), the resulting strategy profile is not a K -approximate PNE anymore—the situation is essentially as in Figure 3 (a). This also shows that, starting from a K -PNE, a sequence of best responses does not necessarily yield a K -PNE.



(a) Situation after player 7 is added. In the while-loop, player 6 wants to deviate to r_2 since $10 > K \cdot 8 \approx 9.58$.



(b) Situation after player 6 deviated to r_2 . Player 7 now wants to deviate to r_5 since $12 > K \cdot 10 \approx 11.97$.



(c) Situation after player 7 deviated to r_5 . The while-loop terminates since no player has a K -improving deviation.

Figure 3: Example for Algorithm 1 with $n = 7$ players, and five resources r_1, \dots, r_5 . Player $i \in [n]$ is represented by a square containing number i .

B Omitted Lemmas from Section 3

This section contains various lemmas which are used to prove Theorem 3. Before we state the Lemmas B.5 and B.6 which are the basis of the case distinction used in the proof of Theorem 3, we introduce the following notation, and provide some simple lemmas.

Notation. Given a strategy profile x , a player i with $x_i \neq 0$ and a resource $r \neq x_i$, we write $\text{dev}(i, r, x) := c_r(x_{-i}, r)$ for the cost experienced by player i if she unilaterally deviates from her current strategy x_i to a different resource r . We also call $\text{dev}(i, r, x)$ the deviation cost of player i for r (with respect to x).

Lemma B.1. Let x be a strategy profile, i and j two players with strategies $x_i, x_j \in R$, and r a resource different from x_i and x_j . Furthermore assume that $\ell_{x_i}(x) \geq \ell_{x_j}(x)$, and define $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$. Then:

$$\text{dev}(i, r, x) \neq \text{dev}(j, r, x) \Leftrightarrow$$

One of the following two cases holds:

$$\ell_r(x) = M - 1, \ell_{x_i}(x) = M, \text{ and } \ell_{x_j}(x) < M, \text{ or}$$

$$\ell_r(x) = M - 2 \text{ and } x_i \text{ the only resource with load } M \text{ in } x.$$

Proof. If a player i with $\ell_{x_i}(x) = M$ changes from her current strategy x_i to $r \neq x_i$ with $\ell_r(x) = M - 1$, she has to pay $\text{dev}(i, r, x) = a_r(\ell_r(x) + 1) + \frac{B}{p}$, where $p := |M^{-1}(x)|$ denotes the number of resources with load M in x . If a player j with $\ell_{x_j}(x) < M$ changes from her current strategy x_j to $r \neq x_j$ with $\ell_r(x) = M - 1$, she has to pay $\text{dev}(j, r, x) = a_r(\ell_r(x) + 1) + \frac{B}{p+1}$. Thus, $\text{dev}(i, r, x) \neq \text{dev}(j, r, x)$. If x_i is the only resource with load M in x and player i changes from her current strategy x_i to $r \neq x_i$ with $\ell_r(x) = M - 2$, she has to pay $\text{dev}(i, r, x) = a_r(\ell_r(x) + 1) + \frac{B}{\tilde{p}+1}$, where \tilde{p} is the number of resources with load $M - 1$ in x . If a player j changes from her current strategy $x_j \neq x_i$ to $r \neq x_j$ with $\ell_r(x) = M - 2$, she has to pay $\text{dev}(j, r, x) = a_r(\ell_r(x) + 1)$. Thus, $\text{dev}(i, r, x) \neq \text{dev}(j, r, x)$.

Now we will show that if $\text{dev}(i, r, x) \neq \text{dev}(j, r, x)$, then either $\ell_r(x) = M - 1$, $\ell_{x_i}(x) = M$ and $\ell_{x_j}(x) < M$ or $\ell_r(x) = M - 2$ and x_i is the only resource with load M in x . We consider a deviation to $r \neq x_i, x_j$ for player i and j in x , respectively. It is clear that the congestion part of the deviation cost of player i or j for r , which is $a_r(\ell_r(x) + 1)$, is the same. Since we assume $\text{dev}(i, r, x) \neq \text{dev}(j, r, x)$, the adversary part, which is $\kappa_r^*(x_{-i}, r)$ and $\kappa_r^*(x_{-j}, r)$, must be different. In particular, the adversary parts cannot be 0 for both. We write $M_{i,r} := \max\{\ell_{r'}(x_{-i}, r) : r' \in R\}$ for the maximum load after player i deviated from her current strategy x_i to r . We now distinguish between the possible values for $M_{i,r} \in \{M - 1, M, M + 1\}$.

Case $M_{i,r} = M + 1$: This case implies $\ell_r(x) = M$ and $\ell_r(x_{-i}, r) = \ell_r(x_{-j}, r) = M + 1$. Thus, both players have to pay the whole adversary budget B after changing to r in x . Remark that the other resources have load at most M after player i or j deviates. This contradicts our assumption that the adversary parts must be different for the two players.

Case $M_{i,r} = M - 1$: This case implies that x_i is the only resource with load M in x and $\ell_r(x) \leq M - 2$. If player j changes from her current strategy $x_j \neq x_i$ to r the adversary part is 0 (the resource x_i has still load M). Since the adversary part should be different after the deviations to the resource r , we conclude that the adversary part of the deviation cost of player i must be positive. Therefore, $\ell_r(x) + 1 = M - 1$ which is equivalent to $\ell_r(x) = M - 2$.

Case $M_{i,r} = M$: This case implies $\ell_r(x) \leq M - 1$ (if $\ell_r(x) = M$ we would get $M_{i,r} = M + 1$). If $\ell_r(x) \leq M - 2$ and player i changes to r the adversary part of player i 's cost is 0. Note that x_i can not be the only resource with load M which would imply $M_{i,r} = M - 1$. Since the adversary part should be different after the deviations to the resource r , we conclude that the adversary part of the deviation cost of player j must be positive. Therefore, x_j must be the only resource with load M in x and $\ell_r(x) = M - 2$. Thus, $\ell_{x_j}(x) > \ell_{x_i}(x)$ which contradicts the assumption that $\ell_{x_j}(x) \leq \ell_{x_i}(x)$. Hence the case $\ell_r(x) \leq M - 2$ can not occur and we know $\ell_r(x) = M - 1$. This implies $\ell_r(x_{-i}, r) = \ell_r(x_{-j}, r) = M$ and both adversary parts of the deviation costs of player i and j for the resource r must be positive. Since the adversary part must be different the number of resources with maximal load must be different after the respective deviations to r . Let $p := |M^{-1}(x)|$ be the number of resources with load M in x , and p_i and p_j the number of resources with load M in (x_{-i}, r) and (x_{-j}, r) , respectively. As mentioned above the adversary part must be different, that is, $p_i \neq p_j$. If $\ell_{x_i}(x) = M$ we get $p_i = p$ (note that $\ell_r(x) = M - 1$). If, additionally, $\ell_{x_j}(x) = M$, we get $p_j = p = p_i$ which is a contradiction to our assumption. Since $\ell_{x_j}(x) \leq \ell_{x_i}(x)$ we can conclude $\ell_{x_j}(x) \leq M - 1$. Finally, we exclude the case $\ell_{x_i}(x) \leq M - 1$. Since we assume $\ell_{x_j}(x) \leq \ell_{x_i}(x) \leq M - 1$, we get $p_i = p + 1 = p_j$ (note that $\ell_r(x) = M - 1$), which is a contradiction to our assumption. \square

Lemma B.2. *Let x be a strategy profile occurring in Algorithm 1, $i \in N$ a player with strategy $x_i \in R$, and $r \in R$ a resource with load $\ell_r(x) = \ell_{x_i}(x) - 1$. Then player i does not want to deviate to r .*

Proof. We show that a deviation to r can only increase player i 's cost. Clearly, the adversary part of the cost remains unchanged. Furthermore, the congestion part can only be larger, since $a_{x_i} \leq a_r$ holds: Recall that loads are always decreasing along the resources, that is, $\ell_{r_1}(x) \geq \ell_{r_2}(x) \geq \dots \geq \ell_{r_m}(x)$ holds, and furthermore $a_{r_1} \leq a_{r_2} \leq \dots \leq a_{r_m}$. Altogether, the deviation is not beneficial for player i . \square

Lemma B.3. *Let x and x' be two strategy profiles occurring in Algorithm 1, where x' results from x by a deviation in the while-loop of player i from x_i to x'_i . Then, if a player $j \in N$ wants to deviate to a resource $r \in R$ with $\ell_r(x') = \ell_{x_i}(x')$, she prefers x_i .*

Proof. Note that since player i deviates from x_i , we know that $\ell_{x_i}(x) > \ell_{x_i}(x') = \ell_r(x') \geq \ell_r(x)$ for $r \neq x_i$. Thus, since the loads are decreasing along the resources, $a_{x_i} \leq$

a_r holds. Consequently, $\text{dev}(j, x_i, x') \leq \text{dev}(j, r, x')$ for $r \in R$ with $\ell_{x_i}(x') = \ell_r(x')$. Furthermore, the resource x_i has a smaller index than the resource r since the loads are decreasing. Thus, by construction, player j prefers the resource x_i . \square

Lemma B.4. *Let x and x' be two strategy profiles occurring in Algorithm 1, where x' results from x either due to the addition of some new player, or by a deviation of a player in the while-loop. Denote the new or deviating player by i . Then any player j with $\ell_{x'_j}(x') = \ell_{x'_i}(x')$ is happy with respect to x' .*

Proof. Note that since the loads are decreasing along the resources, $a_{x'_j} \leq a_{x'_i}$ holds. Therefore, if $x'_j \neq x'_i$, player j does not want to deviate to x'_i (her cost can only increase). Now let $r \neq x'_j$ be some possible alternative for player j . We show $c_{x'_j}(x') \leq \text{dev}(j, r, x')$, that is, player j does not want to deviate to r . Due to the above, we may assume that $r \neq x'_i$. First note that $c_{x'_j}(x') \leq c_{x'_i}(x')$ holds since $a_{x'_j} \leq a_{x'_i}$. Secondly, $c_{x'_i}(x') \leq \text{dev}(i, r, x')$ since x'_i was a best response for player i . Finally, $\text{dev}(i, r, x') = \text{dev}(j, r, x')$ holds since the deviation cost can only be different for two players if they have different load (see Lemma B.1). Altogether, we conclude $c_{x'_j}(x') \leq \text{dev}(j, r, x')$ and thus player j does not want to deviate to r . \square

We can now state and prove the Lemmas B.5 and B.6 which are the basis of the case distinction used in the proof of Theorem 3.

Lemma B.5. *Assume that $i \in \{1, \dots, k-1\}$ is a player who becomes unhappy due to the addition of the new player k , that is, player i has a K -improving deviation r with respect to the strategy profile $x' := (x'_k, x_{-k})$ which results from x due to the addition of the new player k . Then, regarding the loads of the resources $x'_k, x'_i = x_i$, and r , one of the cases shown in Table 3 needs to hold, where $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ denotes the maximum load with respect to x , that is, before player k is added. Also note that $x_i \neq x'_k$ holds since all players using x'_k are happy with respect to x' .*

Proof. Let $i \in \{1, \dots, k\}$ be a player who wants to change from her current strategy x_i to a resource $r \neq x_i$.

We first show that $c_{x_i}(x') \leq c_{x_i}(x)$, that is, the addition of player k to the game did not increase the cost experienced by player i . Clearly, $\ell_{x_i}(x') = \ell_{x_i}(x)$, since the load on resources different from x'_k has not changed. Therefore, the congestion-part of the cost experienced by player i has not changed. We now turn to the adversary part and show that it can only be smaller than before player k was added: Assume, by contradiction, that the adversary part has increased. In particular, the adversary part with respect to x' has to be positive. This implies that $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\} \geq \ell_{x_i}(x) = \ell_{x_i}(x') = \max\{\ell_{r'}(x') : r' \in R\} \geq M$ holds. Thus x_i is a maximum load resource both before and after player k was added, and clearly the number of resources with load M can not decrease by the addition of

$\ell_{x'_k}(x)$	$\ell_{x_i}(x') = \ell_{x_i}(x)$	$\ell_r(x')$	further conditions
$M - 2$	M	$M - 2$	x_i is the only resource with load M in x (also in x')
$M - 1$	$\leq M - 1$	$M - 1$	
M	$\leq M$	M	
	$\leq M - 1$	$M - 1$	

Table 3: Cases for the loads of the resources x'_k, x_i and r , where r is a K -improving deviation for player i who becomes unhappy due to the addition of new player k on x'_k (each row corresponds to one possible case).

player k . This contradicts our assumption that the adversary part experienced by player i has increased, and altogether shows that the total cost of player i can only be smaller than before player k was added.

Before player k was added, player i was happy with her strategy x_i , thus she in particular did not want to change to resource r then. Furthermore, as we already showed, player i 's own cost has not increased due to the addition of player k . Since we assume that player i now wants to change to r , we conclude that the deviation cost of player i for r has decreased, that is, $\text{dev}(i, r, x') < \text{dev}(i, r, x)$ holds. Clearly, the congestion part of the deviation cost did not decrease, since $\ell_r(x') \geq \ell_r(x)$ holds. Therefore, the adversary part needs to be smaller than before, and this is only possible if it was positive with respect to x . We thus get $\ell_r(x) \geq M - 2$, where $\ell_r(x) = M - 2$ is only possible if x_i is the only resource with load M in x . Furthermore, the adversary part of the deviation cost with respect to x' needs to be smaller than B . Moreover $r = x'_k$ is not possible since the adversary part of the deviation cost needs to decrease. Therefore, we can assume in the following that $r \neq x'_k$ and consequently $\ell_r(x) = \ell_r(x')$. We now distinguish between the three possible values for $\ell_r(x) \in \{M - 2, M - 1, M\}$. If $\ell_r(x) = M$, then the adversary part of the deviation cost with respect to x' can only be smaller than B if $\ell_{x'_k}(x) = M$. If $\ell_r(x) = M - 1$, we get that $\ell_{x'_k}(x) \in \{M - 1, M\}$ needs to hold. Finally if $\ell_r(x) = M - 2$ and x_i is the only resource with load M in x , we conclude $\ell_{x'_k}(x) = M - 2$ (note that $x_i \neq x'_k$, and players having load $\ell_{x'_k}(x) + 1$ are happy according to Lemma B.4). Using that $\ell_r(x) \neq \ell_{x_i}(x) - 1$ (see Lemma B.2), this yields the conditions stated in Table 3, completing the proof. \square

Lemma B.6. *Assume that player $i \in \{1, \dots, k\}$ deviates in some iteration of the while-loop of Algorithm 1, and denote the strategy profiles before and after this deviation by x and $x' := (x_{-i}, x'_i)$, respectively. Assume further that player $j \in \{1, \dots, k\} \setminus \{i\}$ becomes unhappy due to player i 's deviation, that is, j was happy with respect to x , but she has a K -improving deviation $r \neq x'_j = x_j$ with respect to x' . Then, regarding the loads of the resources x_i, x'_i, x_j and r , one of the cases given in Table 4 has to hold, where $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ denotes the maximum load with respect to x , that is, before player i deviated. Also note that $x_j \notin \{x_i, x'_i\}$, since all players using x_i are already unhappy with respect to x , and all players*

using x'_i are happy with respect to x' .

Proof. Assume that player i is unhappy with her strategy x_i in x and changes to x'_i . Denote the resulting strategy profile by $x' := (x_{-i}, x'_i)$. Assume that player j is happy with her strategy x_j in x and wants deviate now, with respect to x' , to $r \neq x_j$. Note that $x_j \neq x_i$ since player j is happy with her strategy in x . Furthermore $x_j = x'_j$ since only player $i \neq j$ deviates in x . Additionally $x'_j \neq x'_i$ because player i is happy in x' . If player j is happy with respect to the profile x and wants to deviate now in x' from $x_j = x'_j$ to a resource $r \neq x_i$, then either her cost must be increase, that is $c_{x_j}(x') > c_{x_j}(x)$ or her deviation cost for r must decrease, that is $\text{dev}(j, r, x') < \text{dev}(j, r, x)$. These two cases are discussed in Claim 1 and Claim 2. The situations where a player j wants to deviate to x_i in x' are analyzed in Claim 3.

Claim 1. $c_{x_j}(x') > c_{x_j}(x) \Leftrightarrow$ *One of the two cases displayed in Table 5 needs to hold, where $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ denotes the maximum load with respect to x .*

Proof of Claim 1. Assume that player i changes from the only resource with load M in x to x'_i with $\ell_{x'_i}(x) \leq M - 2$ and player j is on a resource $x'_j = x_j$ with $\ell_{x_j}(x) = M - 1$. Since $x_j \neq x_i, x_j \neq x'_i$ (see above) we know $\ell_{x_j}(x') = M - 1$. Then, player j 's cost increases from $c_{x_j}(x) = a_{x_j} \ell_{x_j}(x)$ to $c_{x_j}(x') = a_{x_j} \ell_{x_j}(x) + \frac{B}{\tilde{p}}$, where $\tilde{p} \geq 2$ is the number of resources with load $M - 1$ in x' . Assume now, player i changes from a resource $x_i \neq x_j$ with load M to x'_i with $\ell_{x'_i}(x) \leq M - 2$ and player j is on a resource x_j with $\ell_{x_j}(x) = M$. Then, player j 's cost increases from $c_{x_j}(x) = a_{x_j} \ell_{x_j}(x) + \frac{B}{p}$ to $c_{x_j}(x') = a_{x_j} \ell_{x_j}(x) + \frac{B}{p-1}$, where $p := |M^{-1}(x)| \geq 2$ is the number of resources with load M in x .

Now, we will show that if $c_{x_j}(x') > c_{x_j}(x)$, one of the two cases displayed in Table 5 needs to hold. Since $x_j = x'_j, x_j \neq x_i$ and $x_j \neq x'_i$ the congestion-part of the cost experienced by player j does not change, comparing the profile x' with x . Thus, if $c_{x_j}(x') > c_{x_j}(x)$ holds, the adversary-part must be increase. Therefore, the adversary-part with respect to x' can not be 0, which implies x_j has maximum load in x' .

If x_j has not maximum load in x , then x_i was the only resource with maximum load in x , which is M and deviates to x'_i with $\ell_{x'_i}(x) \leq M - 2, \ell_{x_j}(x) = M - 1$.

$\ell_{x_i}(x)$	$\ell_{x'_i}(x)$	$\ell_{x_j}(x') = \ell_{x_j}(x)$	$\ell_r(x')$	further conditions
M	M	$\leq M - 1$	$M - 1 (x_i)$	
		$\leq M$	M	
	$\leq M - 2$	M	M	x_i is not the only resource with load M in x
		M	$\leq M - 2$	x_i is not the only resource with load M in x
	$\leq M - 3$	$M - 1$	$M - 1 (x_i)$	x_i is the only resource with load M in x
		$M - 1$	$\leq M - 3$	x_i is the only resource with load M in x
$\leq M - 1$	M	$\leq M$	M	
		$\leq M - 1$	$M - 1$	
	$M - 1$	$\leq M - 1$	$M - 1$	
$\leq M - 2$	$M - 2$	M	$\leq M - 3 (x_i)$	x_j is the only resource with load M in x
		M	$M - 2$	x_j is the only resource with load M in x

Table 4: Cases for the loads of the resources x_i, x'_i, x_j and r , where r is a K -improving deviation for player j who becomes unhappy due to player i 's deviation from x_i to x'_i . Each row corresponds to one possible case. In some cases, we derive that if player j is chosen as the next deviating player in the while-loop, she deviates to x_i ; and whenever this is the case, x_i is added in brackets.

$\ell_{x_i}(x)$	$\ell_{x'_i}(x)$	$\ell_{x_j}(x)$	further conditions
M	$\leq M - 2$	M	
		$M - 1$	x_i is the only resource with load M in x

Table 5: Conditions for a player j who is happy with his strategy in x and wants to change from x_j to r considering the profile x' , which results from x after player i moved from x_i to x'_i (each row corresponds to one possible case).

If x_j has maximum load in x , which is M , the number of resources with maximum load must be decrease since the adversary-part must be increase. This implies that player i deviates from x_i with $\ell_{x_i}(x) = M$ to x'_i with $\ell_{x'_i}(x) \leq M - 2$. ■

Claim 2. $\text{dev}(j, r, x') < \text{dev}(j, r, x)$ for $r \neq x_i \Leftrightarrow r \neq x'_i$ and one of the four cases displayed in Table 6 needs to hold, where $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ denotes the maximum load with respect to x .

Proof of Claim 2. Assume that player i changes from x_i with $\ell_{x_i}(x) \leq M$ to x'_i with $\ell_{x'_i}(x) = M$. Then, with respect to the profile x' , a deviation to a resource $r \neq x_i, x'_i$ with $\ell_r(x) = M$ would now be better for player j . That is $\text{dev}(j, r, x') = a_r(\ell_r(x) + 1) + \frac{B}{2} < a_r(\ell_r(x) + 1) + B = \text{dev}(j, r, x)$. Let us consider now the situation where player i changes from a resource x_i with load $\leq M$ to x'_i with $\ell_{x'_i}(x) = M$, or from a resource x_i with load $\leq M - 1$ to x'_i with $\ell_{x'_i}(x) = M - 1$ and $r \neq x_i, x'_i$ is a resource with load $M - 1$. Then, player j 's deviation cost for r decreases from $\text{dev}(j, r, x) = a_r(\ell_r(x) + 1) + \frac{B}{p+1}$ to $\text{dev}(j, r, x') = a_r(\ell_r(x) + 1)$, or from $\text{dev}(j, r, x) = a_r(\ell_r(x) + 1) + \frac{B}{p+1}$ to $\text{dev}(j, r, x') = a_r(\ell_r(x) + 1) + \frac{B}{p+2}$ respectively, where $p := |M^{-1}(x)|$ is the number of resources with load M in

x . If player i changes from x_i with load $\leq M - 1$ to x'_i with load $M - 1$, or from x_i with load $\leq M - 2$ to x'_i with load $M - 2$ and player j is on the only resource with load M in x , then a deviation to $r \neq x_i, x'_i$ with $\ell_r(x) = M - 2$ is now better than before. In the first mentioned case we have $\text{dev}(j, r, x') = a_r(\ell_r(x) + 1) < a_r(\ell_r(x) + 1) + \frac{B}{\tilde{p}+2} = \text{dev}(j, r, x)$; in the second case we have $\text{dev}(j, r, x') = a_r(\ell_r(x) + 1) + \frac{B}{\tilde{p}+3} < a_r(\ell_r(x) + 1) + \frac{B}{\tilde{p}+2} = \text{dev}(j, r, x)$, where \tilde{p} is the number of resources with load $M - 1$ in x .

Now we will show that if $\text{dev}(j, r, x') < \text{dev}(j, r, x)$, one of the two cases displayed in Table 6 needs to hold and $r \neq x'_i$. First, we discuss the case $r = x'_i$. Since player i changes from x_i to x'_i in x resulting the profile x' , we thus get $\ell_r(x) + 1 = \ell_r(x')$. This implies that the congestion part of the deviation cost to $r = x'_i$ does not decrease. The same holds for the adversary part since the load from x'_i increases while the load of the other resources do not increase. This implies $\text{dev}(j, x'_i, x') \geq \text{dev}(j, x'_i, x)$ which is a contraction to our assumption.

Assume now that $r \neq x'_i$. Since $r \neq x_i$ we get $\ell_r(x) = \ell_r(x')$ which implies that the congestion part of the deviation cost to r is the same in x and x' . Thus, the adversary part must decrease comparing x to x' to achieve $\text{dev}(j, r, x') < \text{dev}(j, r, x)$. Therefore, the adversary part of the deviation cost needs to be positive with respect to x . Consequently,

$\ell_{x_i}(x)$	$\ell_{x'_i}(x)$	$\ell_r(x)$	further conditions
$\leq M$	M	M	
		$M - 1$	
$\leq M - 1$	$M - 1$	$M - 1$	
		$M - 2$	x_j is the only resource with load M in x
$\leq M - 2$	$M - 2$	$M - 2$	x_j is the only resource with load M in x

Table 6: Conditions for a player j who is happy with his strategy in x and wants to change from x_j to r considering the profile x' , which results from x after player i moved from x_i to x'_i (each row corresponds to one possible case).

$\ell_r(x) + 1 = M(x_{-j}, r) =: M_{j,r} \in \{M - 1, M, M + 1\}$ needs to hold. We now distinguish between the three possible values for $M_{j,r}$.

Case $M_{j,r} = M + 1$: This implies $\ell_r(x) = M$. If player j changes to r in x she has to pay the whole adversary budget B (since r is the only resource with load $M + 1$ after this deviation). Now, with respect to x' , the adversary part of the deviation cost must decrease. Therefore, there must be another resource with load $M + 1$ in x' . This implies that player i 's move must be from a resource with load $\leq M$ to a resource with load M .

Case $M_{j,r} = M$: This implies $\ell_r(x) = M - 1$ since we assume $\ell_r(x) + 1 = M_{j,r} = M$ (see above). The adversary part of the deviation cost to r with respect to x is either $\frac{B}{p}$ or $\frac{B}{p+1}$, where $p := |M^{-1}(x)|$ is the number of resources with load M in x . The first mentioned case occurs if $\ell_{x_j}(x) = M$, the second if $\ell_{x_j}(x) < M$. Now, with respect to x' , the adversary part needs to be smaller. Thus, there are two possible situations. Either the adversary part gets 0 or $\frac{B}{\bar{p}}$, where \bar{p} is the number of resources with load $M_{j,r} = M$ in (x'_{-j}, r) . If the the adversary part gets 0 we know with $\ell_r(x') + 1 = \ell_r(x) + 1 = M$ that the maximum load in x' is $> M$. Thus, player i 's move must be from a resource with load $\leq M$ to a resource with load M . If the adversary part gets $\frac{B}{\bar{p}}$ we know with the assumption that the adversary part needs to decrease, that player i 's move must be from a resource with load $\leq M - 1$ to a resource with load $M - 1$.

Case $M_{j,r} = M - 1$: This implies that x_j must be the only resource with load M in x and $\ell_r(x) \leq M - 2$. Since we assumed that $\ell_r(x) + 1 = M_{j,r} = M - 1$ (see above), we know $\ell_r(x) = M - 2$. The adversary part of the deviation cost for r with respect to x is $\frac{B}{\tilde{p}+2}$, where \tilde{p} is the number of resources with load $M - 1$ in x . Note that x_j and r have load $M - 1$ after the change. Now, with respect to x' , the adversary part needs to be smaller. Thus there are two possible situations. Either the adversary part gets 0 or $\frac{B}{\hat{p}}$, where \hat{p} is the number of resources with load $M_{j,r} = M - 1$ in (x'_{-j}, r) . If the the adversary part gets 0 we know with $\ell_r(x') + 1 = M - 1$ that the maximum load in x' is $> M - 1$. Thus, player i 's move must be from a resource with load $\leq M - 1$ to a resource with load $M - 1$. Note that x_i and x'_i have maximum load $M - 1$ because $x_j \neq x_i, x'_i$ (see above) is the only resource with load M in x . If the adversary part gets $\frac{B}{\hat{p}}$ we know with

the assumption that the adversary part needs to decrease, that player i 's move must be from a resource with load $\leq M - 2$ to a resource with load $M - 2$. ■

Claim 3. Assume player i changed in x from x_i to x'_i resulting the profile x' . Then, with respect to x' , a player j might want to deviate to x_i in x' if one of the six cases displayed in Table 7 holds, where $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ denotes the maximum load with respect to x .

Proof of Claim 3. Obviously, player j wants to change from her current strategy x_j to x_i in x' if and only if

$$c_{x_j}(x') > K \text{dev}(j, x_i, x'). \quad (7)$$

Moreover, player i wants to change from x_i to x'_i with respect to x if and only if

$$c_{x_i}(x) > K \text{dev}(i, x'_i, x). \quad (8)$$

We first show by contradiction that it is not possible that all of the three following inequalities are fulfilled: $\text{dev}(j, x_i, x') \geq c_{x_i}(x)$, $\text{dev}(i, x'_i, x) \geq \text{dev}(j, x'_i, x)$, and $c_{x_j}(x) \geq c_{x_j}(x')$. Assume that these three inequalities are fulfilled. Together with inequality (7) and inequality (8), we thus get:

$$\begin{aligned} c_{x_j}(x) &\geq c_{x_j}(x') > K \text{dev}(j, x_i, x') \geq K c_{x_i}(x) \\ &> K^2 \text{dev}(i, x'_i, x) \geq K^2 \text{dev}(j, x'_i, x). \end{aligned}$$

With $K^2 \text{dev}(j, x'_i, x) > K \text{dev}(j, x'_i, x)$ we can conclude that player j wants to deviate to x'_i with respect to x and $c_{x_j}(x) > K c_{x_i}(x) > c_{x_i}(x)$. This contradicts our assumption that in the algorithm player i deviates in x and not player j . Thus, at least one of these three mentioned inequalities is not fulfilled. We now distinguish between which of these inequalities does not hold.

Case $\text{dev}(j, x_i, x') < c_{x_i}(x)$: We will show that in this situation one of the two following cases holds. Either $\ell_{x_j}(x) \leq M$ and player i deviates from x_i with load M to x'_i with load M in x , or $\ell_{x_j}(x) \leq M - 1$ and player i deviates from x_i with load M to x'_i with load $M - 1$ in x . Obviously, the congestion part of the cost experienced by player i on x_i in x and the congestion part of the deviation cost experienced by player j by changing to x_i in x' is the same since $\ell_{x_i}(x) = \ell_{x_i}(x'_{-j}, x_i)$. Thus, to achieve $\text{dev}(j, x_i, x') < c_{x_i}(x)$, the adversary part needs to be

$\ell_{x_i}(x)$	$\ell_{x'_i}(x)$	$\ell_{x_j}(x)$	further conditions
M	M	$\leq M$	
	$M - 1$	$\leq M - 1$	
	$\leq M - 2$	M	
		$M - 1$	x_i is the only resource with load M in x
$\leq M - 1$	$M - 1$	M	
	$M - 2$	M	x_j is the only resource with load M in x

Table 7: Conditions for a player j who wants to change to x_i considering the profile x' , which results from x after player i moved from x_i to x'_i (each row corresponds to one possible case).

smaller. This implies that the adversary part of the cost experienced by player i on x_i in x needs to be positive. Therefore, $\ell_{x_i}(x) = M$ and the adversary part of the cost is $\frac{B}{p}$, where $p := |M^{-1}(x)|$ is the number of resources with load M in x . Since the adversary part of the deviation cost experienced by player j by changing to x_i in x' must be smaller, it can be 0 or $\frac{B}{\tilde{p}}$, where \tilde{p} is the number of resources with load M in (x'_{-j}, x_i) , and $\tilde{p} > p$. If the adversary part is 0, we know with $\ell_{x_i}(x) = \ell_{x_i}(x'_{-j}, x_i) = M$ that the maximum load is $M + 1$ in (x'_{-j}, x_i) . This implies that player i changes from x_i with load M to x'_i with load M in x . Now we discuss the case that the adversary part is $\frac{B}{\tilde{p}}$. In this situation player i changes from x_i with load M to x'_i with load $\leq M - 1$ in x . Let p' be the number of resources with load M in x' . If $\ell_{x_j}(x) = M$ we get $\tilde{p} = p' \leq p$ which contradicts $\tilde{p} > p$. Thus, $\ell_{x_j}(x) \leq M - 1$. If $\ell_{x'_i}(x) \leq M - 2$ we get $p = p' + 1$, and with $\ell_{x_j}(x) \leq M - 1$ we get $\tilde{p} = p' + 1$. Together we get $p = p' + 1 = \tilde{p}$ which contradicts $\tilde{p} > p$. Thus, $\ell_{x'_i}(x) = M - 1$.

Case $\text{dev}(i, x'_i, x) < \text{dev}(j, x'_i, x)$: We show that in this situation one of the two following cases holds. Either $\ell_{x_j}(x) = M$ and player i deviates from x_i with load $\leq M - 1$ to x'_i with load $M - 1$ in x , or x_j is the only resource with load M in x and player i deviates from x_i with load $\leq M - 1$ to x'_i with load $M - 2$ in x . According to Lemma B.1 we know that either $\ell_{x'_i}(x) = M - 1$ and one of the two resources x_j, x_i has load M while the other has load $< M$ in x , or $\ell_{x'_i}(x) = M - 2$ and one of the two resources x_j, x_i is the only resource with load M in x . To make that more precise, we have to know the relation between the load of x_j and x_i in x . We prove by contradiction that $\ell_{x_j}(x) \geq \ell_{x_i}(x)$. Assume that $\ell_{x_j}(x) < \ell_{x_i}(x)$. This implies $\ell_{x_i}(x) = M$ and $\ell_{x_j}(x) \leq M - 1$ with Lemma B.1. Furthermore, we know $\ell_{x'_i}(x) \in \{M - 1, M - 2\}$ according to Lemma B.1. If $\ell_{x'_i}(x) = M - 1$ we get $\text{dev}(i, x'_i, x) = a_{x'_i}M + \frac{B}{p} > a_{x'_i}M + \frac{B}{p+1} = \text{dev}(j, x'_i, x)$, where $p := |M^{-1}(x)|$ is the number of resources with load M in x . This contradicts our assumption that $\text{dev}(i, x'_i, x) < \text{dev}(j, x'_i, x)$. If $\ell_{x'_i}(x) = M - 2$ and x_i is the only resource with load M in x , we get $\text{dev}(i, x'_i, x) = a_{x'_i}(M - 1) + \frac{B}{\tilde{p}+2} > a_{x'_i}(M - 1) = \text{dev}(j, x'_i, x)$, where \tilde{p} is the number of

resources with load $M - 1$ in x . This contradicts our assumption, too. Thus, $\ell_{x_j}(x) \geq \ell_{x_i}(x)$ needs to hold and with Lemma B.1 we know that either $\ell_{x'_i}(x) = M - 1$, $\ell_{x_j}(x) = M$ and $\ell_{x_i}(x) \leq M - 1$, or $\ell_{x'_i}(x) = M - 2$ and x_j is the only resource with load M in x .

Case $c_{x_j}(x) < c_{x_j}(x')$: We show that in this situation one of the following two cases holds. Either $\ell_{x_j}(x) = M - 1$ and player i deviates from x_i which is the only resource with load M in x to x'_i with load $\leq M - 2$ in x , or $\ell_{x_j}(x) = M$ and player i deviates from x_i with load M to x'_i with load $\leq M - 2$ in x . Clearly, the congestion part of the cost experienced by player j in x and x' does not change since $\ell_{x_j}(x) = \ell_{x_j}(x')$. Note that $x_j \neq x_i$ and $x_j \neq x'_i$ (see above). Thus, the adversary part of player j 's cost must increase comparing x' to x . This implies that the adversary part of player j 's cost is positive with respect to x' . Therefore, the maximum load in x' must be $\leq M$, since $\ell_{x_j}(x') = \ell_{x_j}(x) \leq M$. In particular, $\ell_{x'_i}(x') \leq M$ and thus $\ell_{x'_i}(x) \leq M - 1$ needs to hold. Moreover $\ell_{x_j}(x') \in \{M - 1, M\}$ is the maximum load in x' . We now distinguish between the possible values for the maximum load in x' . Assume first that $\ell_{x_j}(x') = M - 1$. This implies that player i changes from x_i which is the only resource with load M in x to x'_i with load $\leq M - 2$ in x . Thus, the adversary part of player j 's cost is 0 with respect to x and $\frac{B}{\tilde{p}}$ with respect to x' , where $\tilde{p} \geq 2$ is the number of resources with load $M - 1$ in x' . This shows an increase of the adversary part. Assume now that $\ell_{x_j}(x') = M$. Thus, the adversary part of player j 's cost is $\frac{B}{p}$ with respect to x and $\frac{B}{p'}$ with respect to x' , where p and p' denote the numbers of resources with load M in x and x' , respectively. Since the adversary part needs to increase we know $p > p'$. Therefore, player i deviates from x_i with load M to x'_i with load $\leq M - 2$ in x . ■

Altogether, we can conclude the conditions for a player j who is happy with her strategy in x and wants to change from $x'_j = x_j$ to r considering the profile x' , which results from x after player i moved from x_i to x'_i . These conditions are displayed in Table 8.

Using Lemma B.4, Lemma B.2 and Lemma B.3 we can reduce Table 8 to Table 4, completing the proof of Lemma B.6. □

$\ell_{x_i}(x)$	$\ell_{x'_i}(x)$	$\ell_{x_j}(x') = \ell_{x_j}(x)$	$\ell_r(x')$	further conditions	
M	M	$\leq M$	M		
		$\leq M$	$M - 1$		
		$\leq M$	$M - 1(x_i)$		
	$M - 1$	$\leq M - 1$	$M - 1(x_i)$		
	$M - 2$	M	$\leq M$		x_i is not the only resource with load M in x
		M	$\leq M - 1(x_i)$		x_i is not the only resource with load M in x
		$M - 1$	$M - 1(x_i)$		x_i is the only resource with load M in x
		$M - 1$	$\leq M - 1$		x_i is the only resource with load M in x
	$\leq M - 3$	M	$\leq M$		x_i is not the only resource with load M in x
		M	$M - 1(x_i)$		x_i is not the only resource with load M in x
$M - 1$		$M - 1(x_i)$		x_i is the only resource with load M in x	
$M - 1$		$\leq M - 1$		x_i is the only resource with load M in x	
$\leq M - 1$	M	$\leq M$	M		
		$\leq M$	$M - 1$		
	$M - 1$	M	$M - 2$		x_j is the only resource with load M in x
		M	$\leq M - 2(x_i)$		
	$M - 2$	M	$M - 1$		
		M	$\leq M - 2(x_i)$		x_j is the only resource with load M in x
$\leq M - 2$	$M - 2$	M	$M - 2$	x_j is the only resource with load M in x	
		M	$\leq M - 3(x_i)$	x_j is the only resource with load M in x	

Table 8: Conditions for a player j who is happy with her strategy in x and wants to change from $x'_j = x_j$ to r considering the profile x' , which results from x after player i moved from x_i to x'_i (each row corresponds to one possible case). Whenever there is an explicit resource r where player i wants to deviate to, we add this resource in brackets.

C Omitted Proofs for the Main Results from Section 3

In this section, we provide complete proofs for Theorem 3 and Corollary 5.

Proof of Theorem 3. It suffices to show that in iteration $k \geq 3$ of the for-loop, the while-loop terminates. Let x and x' denote the profiles directly before and after the new player k is added. If all players are happy with their strategy in x' , the statement follows; thus assume that player i changes from $x'_i = x_i$ to r in the while-loop. Due to Lemma B.5, we know that one of the three cases displayed in Table 3 needs to hold. We now analyze each of these cases, and make repeated use of Lemma B.6. To this end, let $M := M(x) = \max\{\ell_{r'}(x) : r' \in R\}$ be the maximum load in x , that is, before player k is added, and recall that x'_k denotes the resource to which player k is added, thus $\ell_{x'_k}(x)$ denotes the corresponding load before player k is added. Figure 4 illustrates the complete case distinction that we carry out in the proof, and we start with the analysis of the three different cases regarding $\ell_{x'_k}(x)$.

Case $M - 2$:

If $\ell_{x'_k}(x) = M - 2$, then $M(x') = M$ and $x_i = x'_i$ is the only resource with load M in x' and $\ell_r(x') = M - 2$ holds. Using Lemma B.6, we conclude that there are no players who become unhappy due to player i 's change, and the players on $x'_i = x_i$ (the only unhappy players w.r.t. x') are now happy due to Lemma B.4. Therefore, after player i deviated, all players are happy with their strategy and the while-loop terminates after one iteration.

Case $M - 1$:

If $\ell_{x'_k}(x) = M - 1$, then $M(x') = M$, $\ell_{x_i}(x') \leq M - 1$ and $\ell_r(x') = M - 1$ hold. Lemma B.6 yields that after this deviation, the only players who might be unhappy are players using resources with load $\leq M - 1$, and they want to change to resources with load $M - 1$. But since there are at most $m - 2$ resources with load $M - 1$ in x' , we conclude that the while-loop terminates after $O(m)$ iterations.

Case M :

Now turn to the case that $\ell_{x'_k}(x) = M$. Note that $x'_k = r_1$ and $M(x') = M + 1$ hold, and that r_1 is the only resource with load $M + 1$ in x' . We have to consider the different possibilities regarding the loads of x_i and r as given in

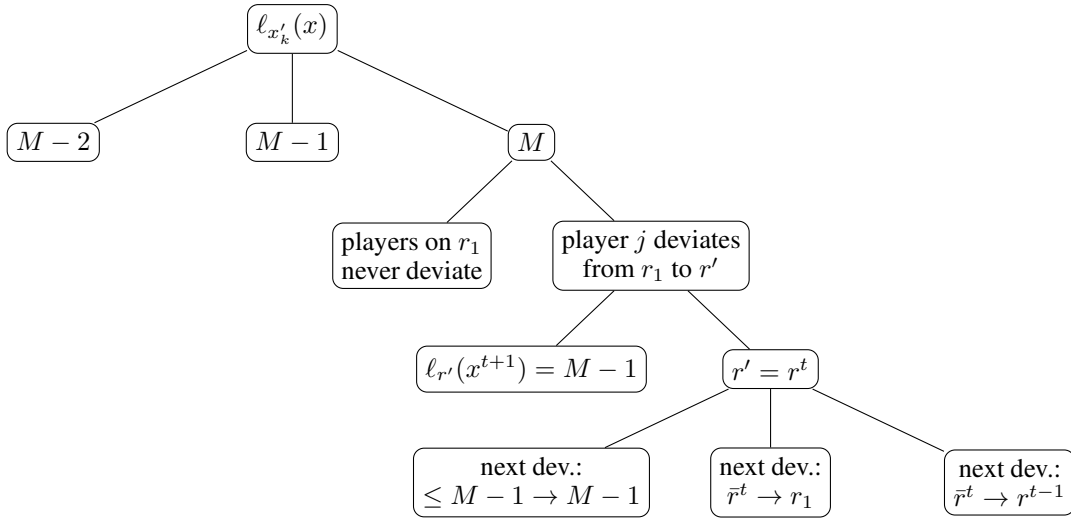


Figure 4: Illustration for the case distinction in the proof of Theorem 3.

Table 3, namely that player i changes from load $\leq M$ to load M or from $\leq M-1$ to $M-1$. First note that if $\ell_r(x') = M$, the only players who might want to change in $x'' := (x'_{-i}, r)$ are using resources with load $\leq M$ and want to change to load M , or want to change from load $\leq M-1$ to load $M-1$. This follows from Lemma B.6, where one should note that the maximum load in x' is $M+1$. For the case that $\ell_r(x') = M-1$ (and consequently, $\ell_{x_i}(x') \leq M-1$), the only new unhappy players could be players using $x'_k = r_1$, and they might want to change to a resource with load $M-1$, or to x'_i (see Lemma B.6).

Using the above, we can argue that if the players on r_1 never deviate during the while-loop, then the only changes are from resources having load $\leq M$ to M , or from $\leq M-1$ to $M-1$. But this terminates after $O(m)$ deviations (there can be at most $m-1$ deviations to a resource with load M , and at most $2m-2$ deviations to a resource with load $M-1$).

Thus we can now assume that in some iteration of the while-loop, a player j deviates from r_1 to a resource r' . Consider the first such change, that is, all changes before have been from load $\leq M$ to M or from $\leq M-1$ to $M-1$. More exactly, note that the changes before cannot include a change to load M , because after such a change, all players having load $M+1$ are happy (see Lemma B.4) and never become unhappy afterwards (see Lemma B.6 and note that r_1 is not the only resource with load $M+1$ anymore). Therefore, all changes before were moves from load $\leq M-1$ to $M-1$ (and there was at least one such change). Let i_1, \dots, i_t be the corresponding sequence of deviating players, where $i_1 = i$ is the first, and i_t the last player deviating before player j . Let r^s and \bar{r}^s denote player i_s 's resources before and after her change, for $s = 1, \dots, t$, and let x^{s+1} be the strategy profile resulting from i_s 's change. Finally, let $x^{t+2} := (x_{-j}^{t+1}, r')$ be the strategy profile resulting from x^{t+1} due to player j 's deviation from r_1 to r' . Figure 5 shows an illustration for the

situation before player j deviates.

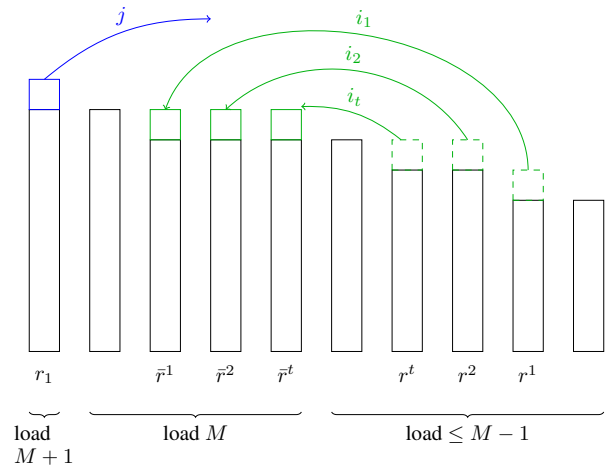


Figure 5: Situation before player j deviates (with $t = 3$).

Note that r_1 is the only resource with load $M+1$ in x^{t+1} . Furthermore, all players having load M are happy in x^{t+1} (see Lemma B.4), thus the only unhappy players in x^{t+1} (except from the players using r_1) might want to change from load $\leq M-1$ to load $M-1$ or to load M . Now consider the situation after player j deviated. Recall that player j either deviated to a resource with load $M-1$, or to r^t (smallest resulting cost and smallest index among all resources where the players i_1, \dots, i_t deviated from).

Subcase $\ell_{r'}(x^{t+1}) = M-1$:

If player j deviates to a resource with load $M-1$, then all players having load M are happy in x^{t+2} . Furthermore, there are no new unhappy players in x^{t+2} (see Lemma B.6). Therefore, after player j 's change, the only possible further changes are from load $\leq M-1$ to load

$M-1$ or load M . Even more, we can also exclude changes from $\leq M-1$ to M , since the cost for deviating to r_1 (the best resource with load M) would be the same as it was before the new player k was added to the game, and in x no one wanted to change to r_1 . Furthermore, the cost on a resource with load $\leq M-1$ can now only be smaller than it was in x . Altogether, the only further changes can be from load $\leq M-1$ to $M-1$, and this terminates after $O(m)$ iterations.

Subcase $r' = r^t$:

It remains to analyze the case that player j deviates from r_1 to r^t (the resource where player i_t deviated from). New unhappy players might want to change from load M to r_1 or to a resource with load $\leq M-2$ (see Lemma B.6). Among these players, a player on \bar{r}^t moves first (biggest index among the resources with load M). Former unhappy players who are not included above want to change from $\leq M-1$ to $M-1$ or to M . More exactly, we can again exclude changes from $\leq M-1$ to M , since the cost for deviating to r_1 (the best resource with load M) would be the same as player j had to pay before her change, and if this is strictly smaller than the current cost of some unhappy player with load $\leq M-1$, we get a contradiction to the fact that player j moved first (the other player could improve 'more' by deviating to r^t). Therefore, the next deviation may be from \bar{r}^t to r_1 , from \bar{r}^t to $\leq M-2$, or from $\leq M-1$ to $M-1$.

In the last case, all players with load M are happy after the change, and the only further moves can be from load $\leq M-1$ to load $M-1$. This terminates after $O(m)$ iterations.

The case $\bar{r}^t \rightarrow r_1$ (see Figure 6 for an illustration) can be excluded since it leads to the following contradiction. Let ℓ_{r^t} denote the load of resource r^t directly before player i_t deviated from it. Furthermore note that there are at least two resources with load M in x^{t+2} (namely r_1 and \bar{r}_1). Then, the following three inequalities hold.

$$a_{r^t} \cdot \ell_{r^t} > K \cdot a_{\bar{r}^t} \cdot M \quad (9)$$

$$a_1 \cdot (M+1) + B > K \cdot a_{r^t} \cdot \ell_{r^t} \quad (10)$$

$$a_{\bar{r}^t} \cdot M + B/2 > K \cdot (a_1 \cdot (M+1) + B) \quad (11)$$

Using that (11) is equivalent to $Ka_{\bar{r}^t} \cdot M > K^2a_1 \cdot (M+1) + (K^2 - K/2)B$ and combining this with (9) and (10) yields

$$\begin{aligned} K^2a_1 \cdot (M+1) + (K^2 - K/2)B &< a_{r^t} \cdot \ell_{r^t} \\ &< \frac{a_1 \cdot (M+1) + B}{K}. \end{aligned}$$

From this we conclude the contradiction

$$0 \leq (K^3 - 1)a_1 \cdot (M+1) < (-K^3 + K^2/2 + 1)B = 0,$$

where we additionally used that $K^3 > 1$ and $-K^3 + K^2/2 + 1 = 0$ by definition of K .

Therefore, the only remaining case is that a player on \bar{r}^t changes to a resource r'' with load $\leq M-2$. We will first show that $r'' = r^{t-1}$ holds. Note that the cost for deviating to r'' is strictly smaller than the cost for deviating to r_1 (otherwise, one would change to r_1). Therefore, r'' needs to have strictly smaller load now than




Figure 6: Illustration for the case $\bar{r}^t \rightarrow r_1$ (with $t = 3$).

it had before player k was added (otherwise, player k would be placed on r'' rather than on r_1). This shows that $r'' \in \{r^1, \dots, r^{t-1}\}$ needs to hold, and since deviating to r^{t-1} is cheapest and r^{t-1} has the smallest index among the cheapest deviations, the assertion follows. Consider the situation after the change from \bar{r}^t to r^{t-1} (of some player j'). It may be that a player on \bar{r}^{t-1} (biggest index among the resources with load M) wants to change to r_1 (the best resource with load M) or to a resource with load $\leq M-2$. Players with load $\leq M-1$ do not want to change, since a deviation to \bar{r}^t (the best resource with load $M-1$) cannot be beneficial (player j' moved first). Furthermore, the case $\bar{r}^{t-1} \rightarrow r_1$ leads to a contradiction (note that in this case, (9)-(11) hold since $a_{\bar{r}^t} \geq a_{\bar{r}^{t-1}}$). Thus we are in the situation that a player on \bar{r}^{t-1} changes to a resource with load $\leq M-2$. By repeating the argumentation of the last paragraph, we conclude that there can be no more than $t-3$ further deviations. \square

Proof of Corollary 5. Note that the only case where the proof of Theorem 3 fails if we consider $\alpha = 1$ instead of $\alpha = K$ is the case corresponding to Figure 6. Since r_1 , \bar{r}^t and r^t are three different resources, and with respect to x' there are at least 3 players using r_1 , at least 1 player using \bar{r}^t and at least 1 player using r^t , the statement follows. \square

D Omitted Proofs from Section 4

Proof of Lemma 7. Given $\bar{c}_M \geq 0$ and $\bar{c}_{<M} \geq 0$, the following procedure determines a vector $\ell = \ell(x)$ corresponding to an α -PNE x with $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$, as well as $\bar{c}_M(x) \geq \bar{c}_M$ and $\bar{c}_{<M}(x) \geq \bar{c}_{<M}$, if such a strategy profile x with the additional property that $\bar{c}_M(x) = \bar{c}_M$ and $\bar{c}_{<M}(x) = \bar{c}_{<M}$ exists.

1. Test first whether all of the following inequalities are satisfied (if at least one inequality is not satisfied, there is no

α -PNE with the desired properties and we can stop):

$$\begin{aligned}
& a_k \cdot M + B/k \leq \alpha \cdot \bar{c}_M \\
& \text{if } k' \geq k + 2: a_{k'-1} \cdot (M - 1) \leq \alpha \cdot \bar{c}_{<M} \\
& \quad \text{if } k' < k'': a_{k''-1} \cdot (M - 2) \leq \alpha \cdot \bar{c}_{<M} \\
& \quad \text{if } k \geq 2: a_1 \cdot (M + 1) + B \geq \bar{c}_M \\
& \text{if } k' \geq k + 2: a_{k+1} \cdot M + \frac{B}{k} \geq \bar{c}_M \\
& \text{if } k = 1 \text{ and } k' < k'': a_{k'} \cdot (M - 1) + \frac{B}{k'} \geq \bar{c}_M \\
& \quad \text{if } k \geq 2 \text{ and } k' < k'': a_{k'} \cdot (M - 1) \geq \bar{c}_M \\
& \quad \quad a_1 \cdot (M + 1) + B \geq \bar{c}_{<M} \\
& \quad \text{if } k' \geq k + 2: a_{k+1} \cdot M + \frac{B}{k+1} \geq \bar{c}_{<M} \\
& \quad \quad \text{if } k' < k'': a_{k'} \cdot (M - 1) \geq \bar{c}_{<M}
\end{aligned}$$

Note that the first three inequalities come from the Nash conditions (players using resources with load $M, M - 1$ or $M - 2$ do not want to deviate). The remaining inequalities are due to the minimum properties of \bar{c}_M and $\bar{c}_{<M}$.

2. Set $\ell_r := M \forall r \in \{1, \dots, k\}$, $\ell_r := M - 1 \forall r \in \{k+1, \dots, k'-1\}$, $\ell_r := M - 2 \forall r \in \{k', \dots, k''-1\}$. If $\ell_r < 0$ for some of the above assigned values, we can stop since there is no α -PNE with the desired properties.
3. Let $n' := n - k \cdot M - (k' - k - 1) \cdot (M - 1) - (k'' - k') \cdot (M - 2)$ be the number of players which are not assigned yet.
4. Test for all $r \geq k''$ whether the following inequalities lead to a contradiction (if this is the case, there is no α -PNE with the desired properties and we can stop):

$$\begin{aligned}
\ell_r & \geq 0 \\
\ell_r & \leq M - 3 \\
\ell_r & \leq \lfloor \frac{\alpha \cdot \bar{c}_{<M}}{a_r} \rfloor \\
\ell_r & \geq \lfloor \frac{\bar{c}_{<M}}{a_r} \rfloor - 1 \\
\ell_r & \geq \lfloor \frac{\bar{c}_M}{a_r} \rfloor - 1
\end{aligned}$$

Note that the third inequality ensures that no player on resource r wants to deviate. The fourth and fifth inequality are due to the minimum properties of \bar{c}_M and $\bar{c}_{<M}$.

5. For all $r \geq k''$, let b_r and b'_r be the lower and upper bounds on ℓ_r which are induced by the inequalities in 4., that is, $b_r := \max\{0, \lfloor \frac{\bar{c}_{<M}}{a_r} \rfloor - 1, \lfloor \frac{\bar{c}_M}{a_r} \rfloor - 1\}$ and $b'_r := \min\{M - 3, \lfloor \frac{\alpha \cdot \bar{c}_{<M}}{a_r} \rfloor\}$.
6. If $n' \notin [\sum_{r \geq k''} b_r, \sum_{r \geq k''} b'_r]$, there is no α -PNE with the desired properties and we can stop.
7. Else set $\ell_r := b_r$ for all $r \geq k''$ and update $n' \leftarrow n' - \sum_{r \geq k''} b_r$.
8. If $n' > 0$, then consider the resources $r \geq k''$ from k'' to m and update $\ell_r \leftarrow \min\{b'_r, b_r + n'\}$ and $n' \leftarrow n' - (\ell_r - b_r)$, until $n' = 0$.

It is clear that if an α -PNE x with $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$, as well as $\bar{c}_M(x) = \bar{c}_M$ and $\bar{c}_{<M}(x) = \bar{c}_{<M}$ exists, the above procedure terminates with a vector ℓ . We now show that ℓ has the desired properties. Clearly, $\ell_r \geq 0$ for all $r \in R$ and $\sum_{r \in R} \ell_r = n$. Thus there is a strategy profile x with load vector $\ell(x) = \ell$. We now argue that x is an α -approximate PNE with the desired properties. Clearly, $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$. Furthermore, $\bar{c}_M \leq \bar{c}_M(x)$ and $\bar{c}_{<M} \leq \bar{c}_{<M}(x)$. Finally, $a_k \cdot M + B/k \leq \alpha \bar{c}_M$ and $a_r(\ell_r(x)) \leq \alpha \bar{c}_{<M}$ for all $r > k$. Thus x is in fact an α -PNE.

Note that for a strategy profile x with $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$, there are at most $3 + (m - k'' + 1) \cdot (M - 2)_+ = O(mn)$ many possible values for $\bar{c}_{<M}(x)$, and also at most $3 + (m - k'' + 1) \cdot (M - 2)_+ = O(mn)$ many possible values for $\bar{c}_M(x)$. Thus if we apply the above procedure for all possible choices of $\bar{c}_M(x)$ and $\bar{c}_{<M}(x)$, we get an α -PNE x with $M(x) = M$, $k(x) = k$, $k'(x) = k'$ and $k''(x) = k''$, if such an α -PNE exists. Since the number of times that we need to apply the above procedure (for given M, k, k', k'') is bounded by $O(n^2 m^2)$, and the procedure itself is efficient, the overall algorithm is efficient, too. \square

E Existence of an approximate PNE with decreasing load profile

Lemma E.1. *For a multi-leader congestion game with an adversary where the resource set $R = [m]$ is ordered such that $a_1 \leq \dots \leq a_m$, the following holds for any $\alpha \in [1, 2]$: If there exists an α -approximate PNE, then there also exists an α -approximate PNE x with decreasing loads, that is, with $\ell_1(x) \geq \dots \geq \ell_m(x)$.*

Proof. Assume there is a α -approximate Nash equilibrium y with $a_r \geq a_{r'}$ and $\ell_r(y) > \ell_{r'}(y)$ for some $r, r' \in R$. Then, we show that the profile x with $\ell_r(x) = \ell_{r'}(y)$, $\ell_{r'}(x) = \ell_r(y)$ and $\ell_{\bar{r}}(x) = \ell_{\bar{r}}(y)$ for all $\bar{r} \in R \setminus \{r, r'\}$ is a α -approximate Nash equilibrium too. The profile x results from y by exchanging all players from r and r' . Denote $z_{u,v}$ for the profile which results from a profile $z \in X$ by moving one player from a resource $u \in R$ to a resource $v \in R$. Since y is a α -approximate Nash equilibrium we have:

$$a_r \ell_r(y) + \kappa_r^*(y) \leq \alpha(a_{\bar{r}}(\ell_{\bar{r}}(y) + 1) + \kappa_{\bar{r}}^*(y_{r,\bar{r}})) \quad (12)$$

for all $\bar{r} \in R \setminus \{r\}$. First, we show, that a player on the resource r is satisfied with his strategy considering the profile x . Let $M := M(y) = \max\{\ell_{\bar{r}}(y) : \bar{r} \in R\}$ denotes the maximum load with respect to y . If $\ell_r(y) = M$, $\ell_{r'}(y) \leq M - 1$ and $\ell_{\bar{r}}(y) = M - 1$ for a resource $\bar{r} \in R \setminus \{r, r'\}$ or $\ell_r(y) = M$, $\ell_{r'}(y) \leq M - 1$ and $\ell_{\bar{r}}(y) = M - 2$ for a resource $\bar{r} \in R \setminus \{r, r'\}$ and r is the only resource with load M regarding y we know with Lemma B.1 that the cost of a deviation from r to \bar{r} must be not the same than the costs of a deviation from r' to \bar{r} . In all the other cases the cost are equal and with inequality (12)

we get:

$$\begin{aligned}
c_r(x) &= a_r \ell_r(x) + \kappa_r^*(x) \\
&\leq a_r \ell_r(y) \\
&\leq \alpha(a_{\bar{r}}(\ell_{\bar{r}}(y) + 1) + \kappa_{\bar{r}}^*(y_{r,\bar{r}})) \\
&= \alpha(a_{\bar{r}}(\ell_{\bar{r}}(x) + 1) + \kappa_{\bar{r}}^*(x_{r,\bar{r}})),
\end{aligned}$$

for all $\bar{r} \in R \setminus \{r', r\}$. Note that $\ell_r(x) = \ell_{r'}(y) < \ell_r(y) = \ell_{r'}(x)$ and $\kappa_r^*(x) = 0$ since r has not maximal load regarding x . Now, let us discuss the other two cases. First let $\ell_r(y) = M$, $\ell_{r'}(y) \leq M - 1$ and $\ell_{\bar{r}}(y) = M - 1$ for a resource $\bar{r} \in R \setminus \{r, r'\}$. We prove by contradiction, that in this situation a player on r does not want to deviate to \bar{r} considering the profile x . Let us assume that a player on r wants to deviate to \bar{r} considering the profile x , that means:

$$c_r(x) = a_r \ell_r(x) > \alpha(a_{\bar{r}}M + \frac{B}{q}), \quad (13)$$

where $q := |M^{-1}(y)| \geq 1$ denotes the number of resources with load M in y . Since y is a Nash equilibrium we know:

$$c_r(y) = a_r M + \frac{B}{q} \leq \alpha(a_{\bar{r}}M + \frac{B}{q+1}), \quad (14)$$

where, again, $q := |M^{-1}(y)| \geq 1$ denotes the number of resources with load M in y . With inequality (13) and (14) we get:

$$\frac{a_r M + \frac{B}{q}}{\alpha} - \frac{B}{q} < \frac{a_r(M-1) + \frac{B}{q}}{\alpha} - \frac{B}{q+1}.$$

Thus, we have

$$B(q+1-\alpha) < -a_r q(q+1) \leq 0.$$

But with $B > 0$ and $q+1-\alpha \geq 1+1-2 = 0$ we get a contradiction. If $\ell_r(y) = M$, $\ell_{r'}(y) \leq M-1$ and $\ell_{\bar{r}}(y) = M-2$ for a resource $\bar{r} \in R \setminus \{r, r'\}$ and r is the only resource with load M regarding y we also can show by contradiction, that in this situation a player on r does not want to deviate to \bar{r} considering the profile x . Let us assume that a player on r wants to deviate to \bar{r} considering the profile x , that means:

$$c_r(x) = a_r \ell_r(x) > \alpha(a_{\bar{r}}(M-1)). \quad (15)$$

Since y is a Nash equilibrium we know:

$$c_r(y) = a_r M + B \leq \alpha(a_{\bar{r}}(M-1) + \frac{B}{p+2}), \quad (16)$$

where p denotes the number of resources with load $M-1$ in y . With inequality (15) and (16) we get:

$$\frac{a_r M + B}{\alpha} - \frac{B}{p+2} + a_{\bar{r}} < \frac{a_r(M-1)}{\alpha} + a_{\bar{r}},$$

where it follows

$$B(1 - \frac{\alpha}{p+2}) < -a_r \leq 0.$$

But with $B > 0$ and $1 - \frac{\alpha}{p+2} \geq 1 - \frac{2}{2} = 0$ we get a contradiction. Altogether, we know that a player on the resource r

does not want to change to a resource $\bar{r} \in R \setminus \{r, r'\}$ considering the profile x . It remains to show that no player on r wants to deviate to r' . With inequality (12) we get:

$$\begin{aligned}
c_r(x) &= a_r \ell_r(x) + \kappa_r^*(x) \\
&\leq a_r \ell_r(y) \\
&\leq \alpha(a_{r'}(\ell_{r'}(y) + 1) + \kappa_{r'}^*(y_{r,r'})) \\
&\leq \alpha(a_{r'}(\ell_{r'}(x) + 1) + \kappa_{r'}^*(x_{r,r'})).
\end{aligned}$$

Thus, the players on r are satisfied with their strategy considering x . Now, let us prove the same for the players on r' . Since $a_r \geq a_{r'}$ we have:

$$\begin{aligned}
c_{r'}(x) &= a_{r'} \ell_{r'}(x) + \kappa_{r'}^*(x) \\
&\leq a_r \ell_r(y) + \kappa_r^*(y) \\
&\leq \alpha(a_{\bar{r}}(\ell_{\bar{r}}(y) + 1) + \kappa_{\bar{r}}^*(y_{r,\bar{r}})) \\
&\leq \alpha(a_{\bar{r}}(\ell_{\bar{r}}(x) + 1) + \kappa_{\bar{r}}^*(x_{r,\bar{r}})),
\end{aligned}$$

for all $\bar{r} \in R \setminus \{r', r\}$. The second inequality holds due to inequality (12). Furthermore with inequality (12) and $a_r \geq a_{r'}$ a player on r' does not want to change to r :

$$\begin{aligned}
c_{r'}(x) &= a_{r'} \ell_{r'}(x) + \kappa_{r'}^*(x) \\
&\leq a_r \ell_r(y) + \kappa_r^*(y) \\
&\leq \alpha(a_{r'}(\ell_{r'}(y) + 1) + \kappa_{r'}^*(y_{r,r'})) \\
&\leq \alpha(a_r(\ell_r(x) + 1) + \kappa_r^*(x_{r,r'})).
\end{aligned}$$

Thus, the players on r' do not want to deviate considering the profile x . Finally, let us show that the players on $\bar{r} \in R \setminus \{r', r\}$ are satisfied with their strategy regarding x . Clearly, they do not want to deviate to a resource which is not r, r' since they have the same costs regarding x and y . Furthermore we know, that y is a α -approximate Nash equilibrium and a deviation to an resource which is not r, r' would provide the same cost considering x and y . Thus, it remains to show that a player on \bar{r} does not want to change to r and r' . First, let us consider a deviation from \bar{r} to r . Since $a_r \geq a_{r'}$ and y is a α -approximate Nash equilibrium we have:

$$\begin{aligned}
c_{\bar{r}}(x) &= a_{\bar{r}} \ell_{\bar{r}}(x) + \kappa_{\bar{r}}^*(x) \\
&= a_{\bar{r}} \ell_{\bar{r}}(y) + \kappa_{\bar{r}}^*(y) \\
&\leq \alpha(a_{r'}(\ell_{r'}(y) + 1) + \kappa_{r'}^*(y_{\bar{r},r'})) \\
&\leq \alpha(a_r(\ell_r(x) + 1) + \kappa_r^*(x_{\bar{r},r})).
\end{aligned}$$

To complete the proof let us consider a deviation from \bar{r} to r' . Since $a_r \geq a_{r'}$ and y is a α -approximate Nash equilibrium we get:

$$\begin{aligned}
c_{\bar{r}}(x) &= a_{\bar{r}} \ell_{\bar{r}}(x) + \kappa_{\bar{r}}^*(x) \\
&= a_{\bar{r}} \ell_{\bar{r}}(y) + \kappa_{\bar{r}}^*(y) \\
&\leq \alpha(a_{r'}(\ell_{r'}(y) + 1) + \kappa_{r'}^*(y_{\bar{r},r'})) \\
&\leq \alpha(a_{r'}(\ell_{r'}(x) + 1) + \kappa_{r'}^*(x_{\bar{r},r'})).
\end{aligned}$$

Altogether, we can conclude that x is a α -approximate Nash equilibrium. \square

F Additively Approximate PNE

In addition to multiplicatively approximate equilibria, as studied in this paper, it also possible to define *additively*

approximate equilibria. A ε -additive approximate PNE is a strategy profile $x \in X$ such that

$$\pi_i(x) \leq \pi_i(y_i, x_{-i}) + \varepsilon \text{ for all } y_i \in X_i.$$

Contrasting the case of K -multiplicative approximate PNE, however, the existence of ε -additive approximate PNE for the class of games studied in this paper cannot be guaranteed for any constant $\varepsilon > 0$.

To see this, note that the considered multi-leader congestion games are invariant under scaling in the sense that multiplying all cost coefficients a_r and the adversary's budget B with the same factor $\lambda \geq 0$ results in a game in which each player's private cost for any given strategy profile is scaled by the same factor λ . As a result, given any constant $\varepsilon > 0$, any instance of a game that does not have an exact PNE (such as, e.g., the one described in Section 3.3) can be scaled in such a way that it does not allow for a ε -additive approximate PNE.

We remark that, while the existence of approximate PNE with a small additive constant cannot be guaranteed, the approach in Section 4 can easily be adjusted to compute ε -additive approximate PNE with the smallest possible ε for a given instance.