

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 knowl-

edge, not been analyzed before and we investigate the existence and computation of (approximate) pure Nash equilibria 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 no unilateral deviation improves the cost of the deviating leader by more than a factor α , for some $\alpha \geq 1$ (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 for 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.

We provide a formal definition of the model in Section 2. In Sections 3 and 4, respectively, we describe the two algorithms and sketch the proofs of the corresponding results.

The complete proofs can be found in the full version of this article (Harks et al. 2021).

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). Correa et al. (2018) and Harks, Schröder, and Vermeulen (2019) 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 or equal for all leaders, respectively), 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 in a network routing game are, e.g., 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 (Babaiouff, 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 the full version of this article 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  $W_\alpha(x) := \{i \in N : x_i \neq 0 \text{ and } c_{x_i}(x) > \alpha \cdot c_r(x_{-i}, r) \text{ for a resource } r \in R\} \neq \emptyset$  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.

An illustrating example for the application of Algorithm 1 can be found in the full version.

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

In this subsection, we outline the analysis that reveals 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 $\min\{O(n^2m), O(nm^2)\}$. The complete analysis can be found in the full version.

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. The derived properties are primarily depend on the 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 . Combining these properties with further structural insights, 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.

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. The conditions on the resource loads mentioned above imply 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 application of the aforementioned conditions, we can show termination of the while-loop also for this case.

At the end of this proof sketch, we want to briefly indicate




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 .

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 $a_1 \geq 0$, 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 $\alpha = K$, independent of the budget and the cost coefficients. Ruling out several other special cases of cyclic behavior, we derive the following theorem.

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). 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 one can show that no player moves more than twice. 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, implying the desired 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

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.

at least five players and at least three resources. Thus we get the following corollary.

Corollary 5. *Exact PNE always 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 $a_1 < a_2 < a_3$; it can be shown that if there exists an α -approximate PNE, there is also one with corresponding load profile among the five listed load profiles). 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 it can be shown that if an α -PNE exists, there also exists one with decreasing loads. Also 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

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

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

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

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




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

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; see the full version for a complete proof.

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 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.
- Babaiouff, 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.
- Bild, 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.; Lianes, 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.; Henle, M.; Klimm, M.; Matuschke, J.; and Schedel, A. 2021. Multi-Leader Congestion Games with an Adversary. Technical Report arXiv:2112.07435, arxiv.org.
- 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.