

The Impact of Collusion on the Price of Anarchy in Nonatomic and Discrete Network Games

Tobias Harks

Institute of Mathematics, Technical University Berlin, Germany
harks@math.tu-berlin.de

Abstract. Hayrapetyan, Tardos and Wexler recently introduced a framework to study the impact of collusion in congestion games on the quality of Nash equilibria. We adopt their framework to network games and focus on the well established price of anarchy as a measure of this impact. We first investigate nonatomic network games with coalitions. For this setting, we prove upper bounds on the price of anarchy for polynomial latencies, which improve upon the current best ones except for affine latencies.

Second, we consider discrete network games with coalitions. In discrete network games, a finite set of players is assumed each of which controlling a discrete amount of flow. We present tight bounds on the price of anarchy for polynomial latencies, which improve upon the previous best ones except for the affine and linear case. In particular, we show that these upper bounds coincide with the known upper bounds for weighted congestion games.

As we do not use the network structure for any of our results but only rely on variational inequalities characterizing a Nash equilibrium, the derived bounds are also valid for the more general case of nonatomic and weighted congestion games with coalitions. Additionally, all our results imply bounds on the price of collusion proposed by Hayrapetyan et al.

Keywords: collusion in network games, congestion games, price of anarchy.

1 Introduction

Over the past years, the impact of the behavior of selfish, uncoordinated users in congested networks has been investigated intensively in the theoretical computer science literature. In this context, *network routing games* have proved to be a reasonable means of modeling selfish behavior in networks. The basic idea is to model the interaction of selfish network users as a *noncooperative game*. We are given a directed graph with latency functions on the arcs and a set of origin-destination pairs, called *commodities*. Every commodity is associated with a *demand*, which specifies the rate of flow that needs to be sent from the respective origin to the destination.

In the nonatomic variant, every demand represents a continuum of agents, each controlling an infinitesimal amount of flow. The latency that an agent experiences to traverse an arc is given by a (non-decreasing) function of the total

flow on that arc. Agents are assumed to act selfishly and route their flow along a minimum-latency path from their origin to the destination; this corresponds to a common solution concept for noncooperative games, that of a *Nash equilibrium*. In the discrete variant, we assume a finite set of players each of which controlling a discrete amount of flow. In both variants a Nash flow is given when no agent or player can improve his own cost by unilaterally switching to another path.

Our focus is to study the impact of coalitions of subsets of agents or players, respectively, on the price of anarchy. To study such coalitions we adopt the framework of Hayrapetyan et al. [10]: Consider a nonatomic network game and a subset of agents forming a coalition. We assume that this coalition aims at minimizing the average delay experienced by this coalition. In this setting, we study Nash equilibria: a stable point, where no coalition can unilaterally improve its cost by rerouting its flow. Note that the assumption that a coalition aims at minimizing its average delay does not imply that *all* members of the coalition experience lower individual cost compared to a Nash equilibrium in which they act independently. Still, one might hope that the total cost of the coalition improves. As discovered by Cominetti et al. [5], even this is not true in general. They presented an instance of a nonatomic network game, where a coalition of nonatomic agents belonging to the same source-destination pair experiences higher cost at Nash equilibrium compared to the case, where every agent acts independently. In the following, we review the known positive and negative results in this setting and comment on issues, which are still open.

1.1 Related Work

Koutsoupias and Papadimitriou [11] initiated the investigation of the efficiency loss caused by selfish behavior. They introduced a measure to quantify the inefficiency of Nash equilibria which they termed the *price of anarchy*. The price of anarchy is defined as the worst-case ratio of the cost of a Nash equilibrium over the cost of a system optimum. In a seminal work, Roughgarden and Tardos [16] showed that the price of anarchy for network routing games with nonatomic players and linear latency functions is $4/3$; in particular, this bound holds independently of the underlying network topology. The case of more general families of latency functions has been studied by Roughgarden [13] and Correa et al. [6]. (For an overview of these results, we refer to the book by Roughgarden [14].)

For discrete network games, where players control a discrete amount of flow, Roughgarden and Tardos examined the price of anarchy for the unsplittable variant [16]. Awerbuch et al. [2], Christodoulou and Koutsoupias [4], and Aland et al. [1] studied the price of anarchy for weighted and unweighted congestion games with polynomial latency functions.

Closest to our work are the papers by Hayrapetyan et al. [10] and Cominetti et al. [5]. The former presented a general framework for studying congestion games with colluding players. Their goal is to investigate the price of collusion: the factor by which the quality of Nash equilibria can deteriorate when coalitions form. Their results imply that for symmetric nonatomic load balancing games with coalitions the price of anarchy does not exceed that of the game without

coalitions. For weighted congestion games with coalitions and polynomial latencies they proved upper bounds of $\mathcal{O}(2^d d^{d+1})$. They also presented examples showing that in discrete network games, the price of collusion may be strictly larger than 1, i.e., coalitions may strictly increase the social cost. Cominetti et al. [5] studied the atomic splittable selfish routing model, which is a special case of the nonatomic congestion game with coalitions. They observed that the price of anarchy of this game may exceed that of the standard nonatomic selfish routing game. Based on the work of Catoni and Pallotino [3], they presented an instance with affine latency functions, where the price of anarchy is 1.34. Using a variational inequality approach, they presented bounds on the price of anarchy for concave and polynomial latency functions of degree two and three of 1.5, 2.56, and 7.83, respectively. For polynomials of larger degree, their approach does not yield bounds.

Fotakis, Kontogiannis, and Spirakis [9] studied algorithmic issues in the setting of atomic congestion games with coalitions.

1.2 Our Results

Albeit the above mentioned advances in the context of bounding the price of anarchy in nonatomic and discrete network games with coalitions, several questions remain open. Can the upper bound for the nonatomic network game with coalitions especially for higher degree polynomials be improved? Consider discrete network games with coalitions. Are there similar counterintuitive phenomena possible as in the nonatomic variant? In particular, does the price of anarchy in discrete network games with coalitions exceed that of discrete network games without coalitions?

In this paper, we contribute to partially answering these questions. For nonatomic network games with coalitions, we present new bounds on the price of anarchy for polynomial latency functions with nonnegative coefficients that improve upon all previous bounds, except for the affine latency case.

For discrete network games with coalitions and latency functions in the sets $\{\ell(z) : \ell(cz) \geq c^d \ell(z), c \in [0, 1]\}$, $d \in \mathbb{N}$, we prove upper bounds on the price of anarchy of Φ_d^{d+1} , where Φ_d is the real solution of the equation $(x+1)^d = x^{d+1}$. Among others, the above sets of latencies contain polynomials with nonnegative coefficients and degree d and concave latencies for $d = 1$. Since classical weighted network games are a special case of the respective games with coalitions, the matching lower bounds for polynomials due to Aland et al. [1] show that our upper bounds are tight. For an overview of the results see Table 1.

We obtain our results by extending the variational inequality approach of Correa et al. [7], Cominetti et al. [5] and Roughgarden [13]. We introduce an additional nonnegative parameter $\lambda \in \mathbb{R}^+$, which allows to bound the cost of a Nash flow in terms of λ times the cost of an optimal flow plus an $\omega(\lambda)$ fraction of the Nash flow itself. Then, the idea is to vary λ so as to find the best possible representation, i.e. one that minimizes the price of anarchy given by $\frac{\lambda}{1-\omega(\lambda)}$. Aland et al. [1] have used a similar technique for bounding the price of anarchy

Table 1. Comparison of our results to [5], [10] and [1]. We show upper bounds on the price of anarchy for nonatomic network games with coalitions (NNGC) and discrete network games with coalitions (DNGC). The matching lower bounds of [1] for weighted network games are also valid lower bounds for discrete network games with coalitions. Considered are polynomial latency functions with nonnegative coefficients and degree d . The results in the first row are with respect to linear latencies $\{a_1 x : a_1 \geq 0\}$. The previous best upper bounds for DNGC are $\mathcal{O}(2^d d^{d+1})$ proved by [10].

d	NNGC		DNGC	
	LB [5] and [13]	Our UB	UB [5]	Our UB-LB [1]
1^*	1.17	1.46	1.5	2.16
1	1.34	1.5	1.5	2.16
2	1.67	2.55	2.56	9.90
3	1.981	5.06	7.83	47.82
4	2.287	11.09	–	277
5	2.40	26.32	–	1858
6	2.63	66.89	–	14,099
7	2.86	180.27	–	118,926
8	3.08	512	–	1,101,126
9	3.30	1,524	–	11,079,429
10	3.51	4,634	–	120,180,803
d	$\Omega(d)$	$(\frac{1}{2}\sqrt{d} + \frac{1}{2})^d \frac{(d+\sqrt{d}+1)}{\sqrt{d}+1}$	–	Φ_d^{d+1}

in weighted congestion games. Their approach, however, requires polynomial latencies while our technique does not need this assumption.

As we do not use the network structure for any of our results but only rely on variational inequalities characterizing a Nash equilibrium, the derived bounds are also valid for the more general case of nonatomic and weighted congestion games with coalitions. Additionally, all our results imply bounds on the price of collusion proposed by Hayrapetyan et al. [10].

2 Nonatomic Network Games with Coalitions

In a network routing game we are given a directed network $G = (V, A)$ and k origin-destination pairs $(s_1, t_1), \dots, (s_k, t_k)$ called *commodities*. For every commodity $i \in [k]$, a demand $r_i > 0$ is given that specifies the amount of flow with origin s_i and destination t_i . Let \mathcal{P}_i be the set of all paths from s_i to t_i in G and let $\mathcal{P} = \cup_i \mathcal{P}_i$. A *flow* is a function $f : \mathcal{P} \rightarrow \mathbb{R}_+$. The flow f is *feasible* (with respect to r) if for all i , $\sum_{P \in \mathcal{P}_i} f_P = r_i$. For a given flow f , we define the flow on an arc $a \in A$ as $f_a = \sum_{P \ni a} f_P$. Moreover, each arc $a \in A$ has an associated variable *latency* denoted by $\ell_a(\cdot)$. For each $a \in A$ the latency function ℓ_a is assumed to be nonnegative, nondecreasing and differentiable. If not indicated otherwise, we also assume that ℓ_a is defined on $[0, \infty)$ and that $x\ell_a(x)$ is a convex function of x . Such functions are called *standard* [13]. The latency of a path P with respect

to a flow f is defined as the sum of the latencies of the arcs in the path, denoted by $\ell_P(f) = \sum_{a \in A} \ell_a(f_a)$. The *cost* of a flow f is $C(f) = \sum_{a \in A} f_a \ell_a(f_a)$. The feasible flow of minimum cost is called *optimal*.

In a nonatomic network game, infinitely many agents are carrying the flow and each agent controls only an infinitesimal fraction of the flow. The continuum of agents of type j (traveling from s_j to t_j) is represented by the interval $[0, d_j]$. It is well known that for this setting flows at Nash equilibrium exist and their total latency is unique, see [14]. Furthermore, the price of anarchy, which measures the worst case ratio of the cost of any Nash flow and that of an optimal flow is well understood, see Roughgarden and Tardos [16], Correa et al. [6, 7], Perakis [12], and Roughgarden [14].

In this paper, we study the impact of coalition formation of subsets of agents on the price of anarchy. Consider a discrete quantity of flow that forms a coalition. We assume that this coalition aims at minimizing the average delay experienced by this coalition. Let $[C] = \{c_1, \dots, c_m\}$ denote the set of coalitions. Each coalition $c_i \in [C]$ is characterized by a tuple (c_{i1}, \dots, c_{ik}) , where every c_{ij} is a subset of the continuum $[0, r_j]$ of agents of type j . We assume that every agent of type j belongs to exactly one coalition c_i , i.e, the disjoint union of c_{ij} , $i \in [m]$ represents the entire continuum of $[0, r_j]$. The tuple (G, r, ℓ, C) is called an *instance* of the nonatomic network game with coalitions. Note that this model includes the special case, where we have exactly k coalitions each of which controlling the flow for commodity k . This case corresponds to the atomic splittable selfish routing model studied by Roughgarden and Tardos [16], Correa et al. [7], and Cominetti et al. [5].

We denote by f^{c_i} any feasible flow of coalition c_i . The cost for coalition c_i is defined as

$$C(f^{c_i}; f^{-c_i}) := \sum_{a \in A} \sum_{j \in [k]} \sum_{h \in [c_{ij}]} \ell_a(f_a) f_a^h,$$

where f^{-c_i} denotes the flow of all other coalitions. It is straightforward to check that at Nash equilibrium, every coalition routes its flow so as to minimize $C(f^{c_i}; f^{-c_i})$ with the understanding that coalition c_i optimizes over f^{c_i} while the flow f^{-c_i} of all other coalitions is fixed.

If latencies are restricted to be standard, minimizing $C(f^{c_i}; f^{-c_i})$ is a convex optimization problem, see for example Roughgarden [15]. The following conditions are necessary and sufficient to characterize a flow at Nash equilibrium for a nonatomic network game with coalitions. We define $f_a^{c_i} := \sum_{j \in [k]} \sum_{h \in [c_{ij}]} f_a^h$ to be the aggregated flow of coalition c_i on arc a .

Lemma 1. *A feasible flow f is at Nash equilibrium for a nonatomic network game with m coalitions if and only if for every $i \in [m]$ the following inequality is satisfied:*

$$\sum_{a \in A} (\ell_a(f_a) + \ell'_a(f_a) f_a^{c_i}) (f_a^{c_i} - x_a^{c_i}) \leq 0 \text{ for all feasible flows } x^{c_i}. \quad (1)$$

The proof is based on the first order optimality conditions and the convexity of $C(f^{c_i}; f^{-c_i})$, see Dafermos and Sparrow [8].

3 Bounding the Price of Anarchy

For every arc a , latency function ℓ_a , and nonnegative parameter λ we define the following nonnegative value:

$$\omega(\ell_a; m, \lambda) := \sup_{f_a, x_a \geq 0} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \ell'_a(f_a) \left(\sum_{i \in [m]} (f_a^{c_i} x_a^{c_i} - (f_a^{c_i})^2) \right)}{\ell_a(f_a) f_a}. \quad (2)$$

We assume $0/0 = 0$ by convention. For a given class of functions \mathcal{L} , we further define $\omega_m(\mathcal{L}; \lambda) := \sup_{\ell_a \in \mathcal{L}} \omega(\ell_a; m, \lambda)$. Moreover, we define the following set:

Definition 1. *Given a class of latency functions \mathcal{L} . The set of feasible $\lambda \geq 0$ is defined as*

$$\Lambda(\mathcal{L}) := \{\lambda \in \mathbb{R}^+ \mid (1 - \omega_m(\mathcal{L}; \lambda)) > 0\}.$$

Theorem 1. *Let \mathcal{L} be a class of continuous, nondecreasing, and standard latency functions. Then, the price of anarchy for the nonatomic network game with coalitions is at most*

$$\inf_{\lambda \in \Lambda(\mathcal{L})} \left[\lambda (1 - \omega_m(\mathcal{L}; \lambda))^{-1} \right].$$

Proof. Let f be a flow at Nash equilibrium, and x be any feasible flow. Then,

$$\begin{aligned} C(f) &\leq \sum_{a \in A} (\ell_a(f_a) f_a + \sum_{i \in [m]} (\ell_a(f_a) + \ell'_a(f_a) f_a^{c_i}) (x_a^{c_i} - f_a^{c_i})) \quad (3) \\ &= \sum_{a \in A} (\ell_a(f_a) x_a + \sum_{i \in [m]} \ell'_a(f_a) f_a^{c_i} (x_a^{c_i} - f_a^{c_i})) \\ &= \sum_{a \in A} (\lambda \ell_a(x_a) x_a + (\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \sum_{i \in [m]} \ell'_a(f_a) f_a^{c_i} (x_a^{c_i} - f_a^{c_i})) \\ &\leq \lambda C(x) + \omega_m(\mathcal{L}; \lambda) C(f). \quad (4) \end{aligned}$$

Here, (3) follows from the variational inequality stated in Lemma 1. The last inequality (4) follows from the definition of $\omega_m(\mathcal{L}; \lambda)$. Taking x as the optimal flow the claim is proven. \square

Note that whenever $\Lambda(\mathcal{L}) = \emptyset$ or $\Lambda(\mathcal{L}) = \{\infty\}$, the approach does not yield a finite price of anarchy. Our definition of $\omega_m(\mathcal{L}; \lambda)$ is closely related to the parameter $\beta^m(\mathcal{L})$ in Cominetti, Correa, and Stier-Moses [5] and $\alpha^m(\mathcal{L})$ in Roughgarden [15] for the atomic splittable selfish routing model. For $\lambda = 1$ we have $\beta^m(\mathcal{L}) = \omega_m(\mathcal{L}; 1)$ and $\alpha^m(\mathcal{L}) = (1 - \omega_m(\mathcal{L}; 1))^{-1}$.

As we show in the next section, the generalized value $\omega_m(\mathcal{L}; \lambda)$ implies improved bounds for a large class of latency functions, e.g., polynomial latency functions. The previous approaches with $\beta^m(\mathcal{L})$ (or $\alpha^m(\mathcal{L})$) failed for instance

to generate upper bounds for polynomials of degree $d \geq 4$ since this value exceeds 1 (or is infinite). The advantage of Theorem 1 is that we can tune the parameter λ and, hence, $\omega_m(\mathcal{L}; \lambda)$ so as to minimize the price of anarchy.

We make use of a result due to Cominetti, Correa, and Stier-Moses [5].

Theorem 2 (Cominetti et al. [5]). *The value $\beta^m(\ell_a) = \omega(\ell_a; m, 1)$ is at most*

$$\sup_{x_a, f_a \geq 0} \frac{(\ell_a(f_a) - \ell_a(x_a)) x_a + \ell'_a(f_a) [(x_a)^2/4 - (f_a - x_a/2)^2/m]}{\ell_a(f_a) f_a}.$$

Since the necessary calculations to prove the above claim only affect the last term in (2), which is the same for $\omega(\ell_a; m, \lambda)$ and $\beta^m(\ell_a)$, this bound carries over for arbitrary nonnegative values of λ .

Corollary 1. *If $\lambda \geq 0$, the value $\omega(\ell_a; m, \lambda)$ is at most*

$$\sup_{x_a, f_a \geq 0} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \ell'_a(f_a) [(x_a)^2/4 - (f_a - x_a/2)^2/m]}{\ell_a(f_a) f_a}.$$

3.1 Linear and Affine Linear Latency Functions

Cominetti et al. [5] proved an upper bound of 1.5 for affine latencies. In the following, we present a stronger result for linear latencies. We also show that for affine latencies the best bound can be achieved by setting $\lambda = 1$. In this case, we have $\beta^m(\mathcal{L}) = \omega_m(\mathcal{L}; 1)$.

Theorem 3. *Consider linear latency functions in $\mathcal{L}_1^* = \{a_1 z : a_1 \geq 0\}$ and $m \geq 2$ coalitions. Then, the price of anarchy is at most*

$$P(m) = \frac{(2m + \sqrt{2} \sqrt{m(m+1)}) (m + 1 + \sqrt{2} \sqrt{m(m+1)}) \sqrt{2}}{8 \sqrt{m(m+1)} (m+1)}.$$

Furthermore, $\lim_{m \rightarrow \infty} P(m) = \frac{3}{4} + \frac{1}{2} \sqrt{2} \approx 1.46$.

Proof. For proving the first claim, we start with the bound on $\omega(\ell_a; m, \lambda)$ given in Corollary 1.

$$\omega(\ell_a; m, \lambda) \leq \frac{a_1 (f_a - \lambda x_a) x_a + a_1 ((x_a)^2/4 - (f_a - x_a/2)^2/m)}{a_1 (f_a)^2}.$$

We define $\mu := \frac{x_a}{f_a}$ for $f_a > 0$ and 0, otherwise, and replace $x_a = \mu f_a$. This yields $\omega(\ell_a; m, \lambda) \leq \max_{\mu \geq 0} (\mu^2 (\frac{m-1-\lambda 4m}{4m}) + \mu (\frac{m+1}{m} - \frac{1}{m}))$. For $\lambda > \frac{m-1}{4m}$ this is a strictly convex program with a unique solution given by $\mu^* = \frac{-2(m+1)}{m-1-\lambda 4m}$. Inserting the solution, yields $\omega(\ell_a; m, \lambda) \leq \frac{m+3-4\lambda}{4\lambda m+1-m}$. To determine an appropriate λ , we also have to analyze the feasible set $\Lambda(\mathcal{L}_1^*)$. The condition $\lambda \in \Lambda(\mathcal{L}_1^*)$ is equivalent to $\lambda > \max \left\{ \frac{m-1}{4m}, \frac{m-2}{2m-2} \right\}$. We define $\lambda = \frac{1}{2} + \frac{1}{4} \sqrt{2(m+1)/m} \in \Lambda(\mathcal{L}_1^*)$. Applying Theorem 1 with this value proves the claim. \square

The proof for affine latencies is similar and leads to $C(f) \leq \min_{\lambda \geq 1} \frac{4\lambda^2 - \lambda}{4\lambda - 2} C(x)$ showing that best bound can be achieved, when $\lambda = 1$.

3.2 Polynomial Latency Functions

We start this section with bounding the value $\omega(\ell_a; m, \lambda)$ for standard latency functions. Some of the following results (Lemma 2, and Proposition 1) are based on results obtained by Cominetti et al. [5]. In addition to their approach that is based on analyzing the parameter $\beta^m(\mathcal{L})$, we need to keep track of restrictions on the parameter λ . For this reason, we also present complete proofs. We focus in the following on the general case $m \in \mathbb{N} \cup \{\infty\}$. Therefore, we define

$$\omega(\ell_a; \infty, \lambda) := \sup_{x_a, f_a \geq 0} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \ell'_a(f) (x_a)^2/4}{\ell_a(f_a) f_a}. \quad (5)$$

Then, it follows from Theorem 2 that $\omega(\ell_a; m, \lambda) \leq \omega(\ell_a; \infty, \lambda)$, since the square is nonnegative and $\lim_{m \rightarrow \infty} (f_a - x_a/2)^2/m = 0$.

Lemma 2. *If $\lambda \geq 1$ and $\ell_a(f_a) f_a$ is a convex function, then the value $\omega(\ell_a; \infty, \lambda)$ is at most:*

$$\omega(\ell_a; \infty, \lambda) \leq \sup_{0 \leq x_a \leq f_a} \frac{(\ell_a(f_a) - \lambda \ell_a(x_a)) x_a + \ell'_a(f) (x_a)^2/4}{\ell_a(f_a) f_a}. \quad (6)$$

Proof. Consider the function $h(x_a)$ defined as the numerator of the supremum in (5). To prove that the solution satisfies $x_a \leq f_a$, we show that $h'(x_a) \leq 0$ if $x_a \geq f_a$. Using that $h'(x_a) = \ell_a(f_a) - \lambda \ell_a(x_a) - \lambda x_a \ell'_a(x_a) + \frac{x_a}{2} \ell'_a(f_a)$, the derivative is negative if and only if $\ell_a(f_a) + \frac{x_a}{2} \ell'_a(f_a) \leq \lambda (\ell_a(x_a) + x_a \ell'_a(x_a))$. By assumption $\ell_a(f_a) f_a$ is convex, hence, $\ell_a(f_a) + \ell'_a(f_a) f_a \leq \ell_a(x_a) + \ell'_a(x_a) x_a$ for $x_a \geq f_a$. Since furthermore $\lambda \geq 1$, the proof is complete. \square

The following characterization of $\omega(\ell_a; m, \lambda)$ via a differentiable function $s(z)$ can be proved with techniques from Cominetti et al. [5].

Proposition 1. *Let \mathcal{L} be a class of continuous, nondecreasing and convex latency functions. Furthermore, assume that $\lambda \geq 1$ and $\ell_a(\kappa f_a) \geq s(\kappa) \ell_a(f_a)$ for all $\kappa \in [0, 1]$, where $s : [0, 1] \rightarrow [0, 1]$ is a differentiable function with $s(1) = 1$. Then,*

$$\omega(\ell_a; \infty, \lambda) \leq \max_{0 \leq u \leq 1} u (1 - \lambda s(u) + s'(1) u/4).$$

In the following, we consider polynomial latency functions of the form

$$\mathcal{L}_d := \{a_d x^d + \dots + a_1 x + a_0 : a_s \geq 0, s = 0, \dots, d\}.$$

Corollary 2. *For latency functions in \mathcal{L}_d , $d \geq 1$, the price of anarchy is at most*

$$\inf_{\lambda \in \Lambda(\mathcal{L}_d) \cap \mathbb{R}^{\geq 1}} \left[\lambda \left(1 - \max_{0 \leq u \leq 1} u (1 - \lambda u^d + d u/4) \right)^{-1} \right]. \quad (7)$$

Proof. All assumptions of Proposition 1 are satisfied with $s(f) = f^d$. Therefore, $s'(1) = d$ and

$$\omega(\ell_a; \infty, \lambda) \leq \max_{0 \leq u \leq 1} u (1 - \lambda u^d + d u/4). \quad (8)$$

Applying Theorem 1 yields the claim. \square

In Table 1, we present upper bounds for latency functions in $\mathcal{L}_1, \dots, \mathcal{L}_{10}$. The results have been obtained by optimizing the expression in (7) over the parameter $\lambda \in \Lambda(\mathcal{L}_d) \cap \mathbb{R}^{\geq 1}$. An asymptotic approximation for general d is provided in the next theorem.

Theorem 4. *Given are latency functions in \mathcal{L}_d , $d \geq 2$. Then, the price of anarchy is at most $(\frac{1}{2}\sqrt{d} + \frac{1}{2})^d \frac{d+\sqrt{d+1}}{\sqrt{d+1}}$.*

Proof. We define $\lambda(d)$ as follows $\lambda(d) := (\frac{1}{2}\sqrt{d} + \frac{1}{2})^d \frac{d+\sqrt{d+1}}{(\sqrt{d+1})(d+1)}$. The proof proceeds by proving a claim, which yields a bound on $\omega(\ell_a; \infty, \lambda(d))$ for $\ell_a \in \mathcal{L}_d$.

Claim. $\max_{0 \leq u \leq 1} [T(u) := u(1 - \lambda(d)u^d + d\frac{u}{4})] = \frac{d}{d+1}$, for all $d \geq 2$.

Proof. To prove the claim it is convenient to write $\lambda(d)$ as $\lambda(d) = (du_1(d) + 4u_1(d) - 4d/(d+1))/(4u_1(d)^{d+1})$ with $u_1(d) := 2/(\sqrt{d} + 1)$. Then, the claim is proven by verifying the following facts:

1. $T'(u_1(d)) = 0$, $T''(u_1(d)) < 0$ and $T''(u)$ has at most one zero in $(0, 1)$
2. $T(0) = 0$, $T(1) \leq 0$ and $T(u_1(d)) = \frac{d}{d+1}$.

Before we prove these facts, we show how it implies the claim. The first fact implies that $u_1(d)$ is the *only* local maximum of $T(u)$ in the open interval $(0, 1)$. Then, by comparing $T(u_1(d))$ to the boundary values $T(0)$ and $T(1)$ it follows that $T(u_1(d)) = d/(d+1)$ is the global maximum.

We start by proving the first fact. The first derivative $T'(u)$ evaluates to $T'(u_1(d)) = -(d(u_1(d))^2 d + 4u_1(d) - u_1(d)^2 - 4)/(4u_1(d))$. Then, it is easy to check that $u_1(d)$ solves the equality $x^2 d + 4x - x^2 - 4 = 0$ proving $T'(u_1(d)) = 0$. Furthermore, $T''(u_1(d)) = -d(d+1)(\frac{d}{4} + \frac{1}{u_1(d)} - \frac{d}{(d+1)u_1(d)^2}) + \frac{d}{2}$. For $T''(u_1(d)) < 0$, it is sufficient to show that $(d+1)(\frac{d}{4} + \frac{1}{u_1(d)} - \frac{d}{(d+1)u_1(d)^2}) > \frac{1}{2}$. Inserting the definition of $u_1(d)$ and rewriting yields $\frac{2d+2\sqrt{d+2}}{4} > \frac{1}{2}$, which holds for all $d \geq 1$. Furthermore, it is straightforward to check that $T''(u)$ has at most one zero in $(0, 1)$, so we omit the details. \square

Using the above claim we can now invoke Corollary 2, which implies that $\omega(\ell_a; \infty, \lambda(d))$ is bounded by $\frac{d}{d+1}$. Hence, $\lambda(d) \in \Lambda(\mathcal{L}_d)$ so we can use Theorem 1 to obtain the claimed bound of $(d+1)\lambda(d)$. \square

In the following we analyze the growth of the derived upper bound for large d , ($d \geq 4$). Th proof consists of standard calculus and is omitted.

Corollary 3. $(\frac{1}{2}\sqrt{d} + \frac{1}{2})^d \frac{d+\sqrt{d+1}}{\sqrt{d+1}} \leq \sqrt{d}^d$ for $d \geq 4$.

4 Discrete Network Games with Coalitions

In this section, we study discrete network games with coalitions. We adopt the same notation as in the previous section and only point out the differences. The main distinction between nonatomic network games and discrete network games is that in the latter one we are given a finite set of players denoted by $[P] = \{1, \dots, P\}$. Each player $p \in [P]$ controls a discrete amount of flow, which has to be routed along a single path. We call a player p of type $j \in [k]$ if player p controls a fraction of commodity j .

A coalition $c_i \in [C]$ is characterized by a tuple (c_{i1}, \dots, c_{ik}) , where every c_{ij} is a subset of players of type j . As before, we assume that every player p of type j belongs to exactly one coalition c_i . The tuple (G, r, ℓ, P, C) is called an *instance* of the discrete network game with coalitions.

We denote by f^p and f^{c_i} the flow of player p and coalition i , respectively. The cost for coalition i is given by

$$C(f^{c_i}; f^{-c_i}) = \sum_{a \in A} \sum_{p \in [c_i]} \ell_a(f_a) f_a^p.$$

We define $f_a^{c_i} = \sum_{p \in [c_i]} f_a^p$ to be the aggregated flow of coalition i on arc a . It is straightforward to check that at Nash equilibrium every coalition routes its flow so as to minimize $C(f^{c_i}; f^{-c_i})$. Note that while the coalition determines a flow minimizing $C^{c_i}(f^{c_i}; f^{-c_i})$, still all individual flows f^p have to be routed along a single path.

The following lemma states that at Nash equilibrium every coalition minimizes $C^{c_i}(f^{c_i}; f^{-c_i})$, i.e. $C^{c_i}(f^{c_i}; f^{-c_i}) \leq C^{c_i}(x^{c_i}; f^{-c_i})$ for any other feasible flow x^{c_i} for the demands of coalition i .

Lemma 3. *Let f be a feasible flow at Nash equilibrium for a discrete network game with coalitions C . Then, for every $[c_i] \in [C]$, $i \in [m]$ the following inequality is satisfied:*

$$\sum_{a \in A} \ell_a(f_a) f_a^{c_i} \leq \sum_{a \in A} \ell_a \left(\sum_{\substack{p \in [P] \\ p \notin [c_i]}} f_a^p + x_a^{c_i} \right) x_a^{c_i} \text{ for all feasible flows } x^{c_i}. \quad (9)$$

Using a similar technique as in the previous section, we define for a given latency function $\ell(z)$, $\lambda \geq 1$, and nonnegative values f, x , the following values

$$\bar{\omega}(\ell; \lambda) := \sup_{x, f \geq 0} \begin{cases} \frac{(\ell(f+x) - \lambda \ell(x)) x}{\ell(f) f} & \text{if } \ell(f) f > 0 \\ 0 & \text{if } \ell(f) f = 0. \end{cases} \quad (10)$$

For a class \mathcal{L} of latency functions, we further define $\bar{\omega}(\mathcal{L}; \lambda) := \sup_{\ell \in \mathcal{L}} \bar{\omega}(\ell; \lambda)$. Figure 1 shows an illustration of $\bar{\omega}(\ell; \lambda)$.

Definition 2. *For latency functions in \mathcal{L} , the feasible scaling set for λ is defined as $\bar{\Lambda}(\mathcal{L}) := \Lambda_1(\mathcal{L}) \cap \Lambda_2(\mathcal{L})$, where $\Lambda_1(\mathcal{L}) := \{\lambda \geq 1 : 1 - \bar{\omega}(\mathcal{L}; \lambda) > 0\}$ and $\Lambda_2(\mathcal{L}) := \{\lambda \geq 1 : (\ell(f+x) - \lambda \ell(x)) x \leq 0 \text{ for all } f, x \in \mathbb{R}^+, \ell \in \mathcal{L}, \ell(f) = 0\}$.*




Fig. 1. Illustration of the value $\bar{\omega}(\ell; \lambda)$ in (10) with $1 < \lambda < \frac{\ell(f+x)}{\ell(x)}$. The light-gray shaded area corresponds to the numerator of $\bar{\omega}(\ell; \lambda)$ and the dark-gray shaded area to the denominator.

Theorem 5. *Let \mathcal{L} be a class of continuous and nondecreasing latency functions. Then, the price of anarchy for a discrete network game with coalitions is at most*

$$\inf_{\lambda \in \bar{A}(\mathcal{L})} \left[\lambda (1 - \bar{\omega}(\mathcal{L}; \lambda)^{-1}) \right].$$

Proof. Let f be a flow at Nash equilibrium, and x be any feasible flow. Then,

$$\begin{aligned} C(f) &= \sum_{a \in A} \sum_{i \in [m]} \ell_a(f_a) f_a^{c_i} \\ &\leq \sum_{a \in A} \sum_{i \in [m]} \ell_a \left(\sum_{\substack{p \in [P] \\ p \neq [c_i]}} f_a^p + x_a^{c_i} \right) x_a^{c_i} \end{aligned} \quad (11)$$

$$\leq \sum_{a \in A} \ell_a(f_a + x_a) x_a \quad (12)$$

$$\begin{aligned} &= \sum_{a \in A} \lambda \ell_a(x_a) x_a + (\ell_a(f_a + x_a) - \lambda \ell_a(x_a)) x_a \\ &\leq \lambda C(x) + \bar{\omega}(\mathcal{L}; \lambda) C(f). \end{aligned} \quad (13)$$

Here, (11) follows from Lemma 3. The second inequality (12) is valid since latency functions are nondecreasing. The bound involving $\bar{\omega}(\ell_a; \lambda)$ holds because of the following: If $\ell_a(f) = 0$ then it follows $(\ell_a(f+x) - \lambda \ell_a(x)) x \leq 0$ since $\lambda \in \Lambda_2(\mathcal{L})$. Thus, the bound is true. If $f = 0$, then $(\ell_a(f) - \lambda \ell_a(x)) \leq 0$ since $\lambda \geq 1$ and $\ell_a(\cdot)$ is nondecreasing. The case $\ell_a(f) f > 0$ follows from the definition. Taking x as the optimal offline solution and since $\lambda \in \bar{A}(\mathcal{L})$, the claim is proven. \square

Note that our only assumptions on feasible latencies are continuity and monotonicity. As an implication of Theorem 5, we show that the price of anarchy for continuous and nondecreasing latencies satisfying $\ell(cz) \geq \ell(z) c^d$ for all $c \in [0, 1]$ is at most Φ_d^{d+1} , where Φ_d is the solution to the equation $(x+1)^d = x^d$.

Theorem 6. *Let \mathcal{L} be a class of continuous and nondecreasing latency functions satisfying $\ell(cz) \geq \ell(z)c^d$ for all $c \in [0, 1]$. Then, the price of anarchy of a discrete network game with coalitions is bounded by Φ_d^{d+1} , where Φ_d is the solution to the equation $(x+1)^d = x^{d+1}$.*

Proof. We present a complete prove for $d = 1$ and sketch the proof for larger d . Key for proving the claim is the following lemma.

Lemma 4. *Let \mathcal{L} be a class of continuous and nondecreasing latency functions satisfying $\ell(cz) \geq \ell(z)c$ for all $c \in [0, 1]$. Then, $\bar{\omega}(\ell; \lambda) \leq \frac{1}{4(\lambda-1)}$.*

Proof. We assume $f > 0$ and $x > 0$ since otherwise the claim is trivially true. We consider two cases: (1) $f \geq x$. In this case we define $\mu = \frac{x}{f}$ with $\mu \in [0, 1]$. Then, we have

$$\begin{aligned} \bar{\omega}(\ell; \lambda) &= \sup_{f \geq 0, \mu \in [0, 1]} \frac{(\ell((1+\mu)f) - \lambda \ell(\mu f)) \mu f}{\ell(f) f} \\ &\leq \sup_{f \geq 0, \mu \in [0, 1]} \frac{(1+\mu)\ell(f) - \lambda \mu \ell(f)}{\ell(f)} \mu \\ &= \sup_{\mu \in [0, 1]} (1 + \mu - \lambda \mu) \mu. \end{aligned}$$

This problem is a concave program with optimal value $\frac{1}{4(\lambda-1)}$. The second case is similar: (2) $f \leq x$. In this case we define $\mu = \frac{f}{x}$ with $\mu \in [0, 1]$. Then, using similar arguments we arrive at

$$\bar{\omega}(\ell; \lambda) = \sup_{x \geq 0, \mu \in [0, 1]} \frac{(\ell((1+\mu)x) - \lambda \ell(\mu x)) x}{\ell(\mu x) \mu x} \leq \sup_{\mu \geq 0} \frac{1 + \mu - \lambda}{\mu^2}.$$

This problem is again a concave program with optimal value $\frac{1}{4(\lambda-1)}$. \square

By Lemma 4, $\bar{\omega}(\mathcal{L}; \lambda)$ is bounded by $\frac{1}{4(\lambda-1)}$. Thus, Theorem 5 implies that the price of anarchy is bounded by $\frac{\lambda(4\lambda-1)}{4\lambda-5}$. Choosing $\lambda = \frac{5+\sqrt{5}}{4}$ as the minimizer of the price of anarchy, the claim for $d = 1$ is proven. For general $d \in \mathbb{N}$, the idea is to bound $\bar{\omega}(\ell; \lambda)$ by $\sup_{\mu \in [0, 1]} ((1+\mu)^d - \lambda \mu^d) \mu$. Then, we can write $(1+\mu)^d$ as $\kappa^{1-d} + \mu^d (1-\kappa)^{1-d}$ for some $\kappa \in (0, 1)$ and solve the supremum. Finally, we define λ and κ as a function of Φ_d leading to the desired bound. \square

Corollary 4. *For continuous, nondecreasing and concave latency functions, the price of anarchy of a discrete network game with coalitions is bounded by $\frac{3+\sqrt{5}}{2}$.*

Remark 1. As weighted network games are a special case of discrete network games with coalitions, the matching lower bounds for polynomials due to Aland et al. [1] show that our upper bounds are tight.

Acknowledgment

We thank Nicholas E. Stier-Moses and Marc E. Pfetsch for fruitful discussions and suggestions.

References

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *STACS*, volume 3884 of *Lecture Notes in Computer Science*, pages 218–229. Springer, 2006.
2. B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 57–66, New York, NY, USA, 2005. ACM Press.
3. S. Catoni and S. Pallotino. Traffic equilibrium paradoxes. *Transportation Science*, 25:240–244, 1991.
4. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC '05: Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 67–73, New York, NY, USA, 2005. ACM Press.
5. R. Cominetti, J. R. Correa, and N. E. Stier-Moses. Network games with atomic players. In *Proceedings of the 33rd International Colloquium of Automata, Languages and Programming (ICALP'06)*, 2006.
6. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Math. Oper. Res.*, 29:961–976, 2004.
7. J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. On the inefficiency of equilibria in congestion games. In M. Jünger and V. Kaibel, editors, *Proc. of 11th International Conference on Integer Programming and Combinatorial Optimization, (IPCO), Berlin*, volume 3509 of *LNCS*, pages 167–181. Springer, Berlin, 2005.
8. S. Dafermos and F. Sparrow. The traffic assignment problem for a general network. *J. Res. Natl. Bur. Stand., Sect. B*, 73:91–118, 1969.
9. D. Fotakis, S. C. Kontogiannis, and P. G. Spirakis. Atomic congestion games among coalitions. In *Proceedings of the 33rd International Colloquium of Automata, Languages and Programming (ICALP'06)*, pages 572–583, 2006.
10. A. Hayrapetyan, Tardos, and T. Wexler. The effect of collusion in congestion games. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 89–98, New York, NY, USA, 2006. ACM Press.
11. E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In C. Meines and S. Tison, editors, *Proc. of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *LNCS*, pages 404–413. Springer, 1999.
12. G. Perakis. The price of anarchy when costs are non-separable and asymmetric. In D. Bienstock and G. Nemhauser, editors, *Proc. of 10th International Conference on Integer Programming and Combinatorial Optimization, (IPCO), New York*, volume 3064 of *Lecture Notes in Computer Science*, pages 46–58. Springer, Berlin Heidelberg, 2004.
13. T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Science*, 67:341–364, 2002.
14. T. Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.
15. T. Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 973–974, 2005.
16. T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.