

Weighted Congestion Games: Price of Anarchy, Universal Worst-Case Examples, and Tightness

Kshipra Bhawalkar^{*1}, Martin Gairing^{**2}, and Tim Roughgarden^{***1}

¹ Stanford University, Stanford, CA, USA, {kshipra,tim}@cs.stanford.edu

² University of Liverpool, Liverpool, U.K. gairing@liverpool.ac.uk

Abstract. We characterize the price of anarchy in weighted congestion games, as a function of the allowable resource cost functions. Our results provide as thorough an understanding of this quantity as is already known for nonatomic and unweighted congestion games, and take the form of universal (cost function-independent) worst-case examples. One noteworthy byproduct of our proofs is the fact that weighted congestion games are “tight”, which implies that the worst-case price of anarchy with respect to pure Nash, mixed Nash, correlated, and coarse correlated equilibria are always equal (under mild conditions on the allowable cost functions). Another is the fact that, like nonatomic but unlike atomic (unweighted) congestion games, weighted congestion games with trivial structure already realize the worst-case POA, at least for polynomial cost functions.

We also prove a new result about unweighted congestion games: the worst-case price of anarchy in symmetric games is, as the number of players goes to infinity, as large as in their more general asymmetric counterparts.

1 Introduction

The class of *congestion games* is expressive enough to capture a number of otherwise unrelated applications – including routing, network design, and the migration of species (see references in [20]) – yet structured enough to permit a useful theory. Such a game has a ground set of resources, and each player selects a subset of them (e.g., a path in a network). Each resource has a univariate cost function that depends on the load induced by the players that use it, and each player aspires to minimize the sum of the resources’ costs in its chosen strategy (given the strategies chosen by the other players).

Congestion games have played a starring role in recent research on quantifying the inefficiency of game-theoretic equilibria. They are rich enough to

* Supported by a Stanford Graduate Fellowship.

** Partly supported by a fellowship within the Postdoc-Programme of the German Academic Exchange Service (DAAD).

*** Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, an ONR PECASE Award, an AFOSR MURI grant, and an Alfred P. Sloan Fellowship.

encode the Prisoner’s Dilemma, and more generally can have Nash equilibria in which the sum of the players’ costs is arbitrarily larger than that in a minimum-cost outcome. Thus the research goal is to understand how the parameters of a congestion game govern the inefficiency of its equilibria, and in particular to establish useful sufficient conditions that guarantee near-optimal equilibria.

A simple observation is that the inefficiency of equilibria in a congestion game depends fundamentally on the “degree of nonlinearity” of the cost functions. Because of this, we identify a “thorough understanding” of the inefficiency of equilibria in congestion games with a simultaneous solution to every possible special case of cost functions. In more detail, an ideal theory would include the following ingredients.

1. For every set \mathcal{C} of allowable resource cost functions, a relatively simple recipe for computing the largest-possible *price of anarchy (POA)* – the ratio between the sum of players’ costs in an equilibrium and in a minimum-cost outcome – in congestion games with cost functions in \mathcal{C} .
2. For analytically simple classes \mathcal{C} like bounded-degree polynomials, an exact formula for the worst-case POA in congestion games with cost functions in \mathcal{C} .
3. An understanding of the “game complexity” required for the worst-case POA to be realized. Ideally, such a result should refer only to the strategy sets and be independent of the allowable cost functions \mathcal{C} .
4. An understanding of the equilibrium concepts – roughly equivalently, the rationality assumptions needed – to which the POA guarantees apply.

This paper is about the fundamental model of *weighted* congestion games [16–18], where each player i has a weight w_i and the load on a resource is defined as the sum of the weights of the players that use it. Such weights model non-uniform resource consumption among players and can be relevant for many reasons: for different durations of resource usage [22]; for modeling different amounts of traffic (e.g., by Internet Service Providers from different “tiers”); and even for collusion among several identical users, who can be thought of as a single “virtual” player with weight equal to the number of colluding players [9, 12]. Our main contribution is a thorough understanding of the worst-case POA in weighted congestion games, in the form of the four ingredients listed above.

1.1 Overview of Results

Result 1: Exact POA of general weighted congestion games with general cost functions. We provide the first characterization of the exact POA of general weighted congestion games with general cost functions. For a given set of cost functions \mathcal{C} , the properties of the functions in this class determine certain “feasible values” of two parameters λ, μ , which lead to upper bounds of the form $\lambda/(1 - \mu)$. We denote the best upper bound that can be obtained using this two-parameter approach by $\zeta(\mathcal{C})$. The hard work then lies in proving that there always exists a weighted congestion game that realizes this upper

bound. The abstract approach is to make use of the inequalities used in the upper bound proof — in the spirit of “complementary slackness” arguments in linear programming — with the cost functions and loads on the resources that make these inequalities tight employed in the worst-case example. Ultimately, we can exhibit examples with POA arbitrarily close to our upper bound of $\zeta(\mathcal{C})$.

This approach is similar to the one taken in [21] for unweighted congestion games, although non-uniform player weights create additional technical issues that we need to resolve. Our constructions here are completely different from those used for unweighted congestion games in [21].

A side effect of our upper bound proof is that $\zeta(\mathcal{C})$ is actually the “Robust POA” defined by Roughgarden [21]. As a consequence, the same upper bound also applies to the POA of the mixed Nash, correlated, and coarse correlated equilibria of a weighted congestion game (see e.g. [24, chapter 3] for definitions). Since we prove a matching lower bound for the pure POA — which obviously extends to the three more general equilibrium concepts — our bound of $\zeta(\mathcal{C})$ is the exact worst-case POA with respect to all four equilibrium concepts.

We note that exact worst-case POA formulas for weighted congestion games with polynomial cost functions and nonnegative coefficients were given previously in [1]. Characterizing the POA with respect to an arbitrary set of cost functions is more technically challenging and is also a well-motivated goal: for example, M/M/1 delay functions of the (non-polynomial) form $1/(u - x)$ for a parameter u pervade the networking literature (e.g. [4]).

Result 2: Exact POA of congestion games on parallel links with polynomial cost functions. We prove that, for polynomial cost functions with nonnegative coefficients and maximum degree d , the worst-case POA is realized on a network of parallel links (for each d). We note that even for affine cost functions, only partial results were previously known about the worst-case POA of weighted congestion games in networks of parallel links. Our work implies that, at least for polynomial cost functions, the worst-case POA of weighted congestion games is essentially independent of the allowable network topologies (in the same sense as [19]). This result stands in contrast to unweighted congestion games, where the worst-case POA in networks of parallel links is provably smaller than the worst-case POA in general unweighted congestion games [2, 8, 13].

Result 3: POA of symmetric unweighted congestion games is as large as asymmetric ones. Our final result contributes to understanding how the worst-case POA of *unweighted* congestion games depends on the game complexity. We show that the POA of symmetric unweighted congestion games with general cost functions is the same value $\gamma(\mathcal{C})$ as that for asymmetric unweighted congestion games. This fact was previously known only for affine cost functions [7]. Thus, the gap between the worst-case POA of networks of parallel links and the worst-case POA of general unweighted atomic congestion games occurs inside the class of symmetric games, and not between symmetric and asymmetric games.

1.2 Further Related Work

Koutsoupias and Papadimitriou [14] initiated the study of the POA of mixed Nash equilibria in weighted congestion games on parallel links (with a different objective function). For our objective function, the earliest results on the pure POA of unweighted congestion games on parallel links are in [15, 23]. The first results for general congestion games were obtained independently by Christodoulou and Koutsoupias [7] and Awerbuch, Azar, and Epstein [3]. Christodoulou and Koutsoupias [7, 6] established an exact worst-case POA bound of $5/2$ for unweighted congestion games with affine cost functions. They obtained the same bound for the POA of pure Nash, mixed Nash, and correlated equilibria [6]. They also provided an asymptotic POA upper bound of $d^{\Theta(d)}$ for cost functions that are polynomials with nonnegative coefficients and degree at most d .

In their concurrent paper, Awerbuch, Azar, and Epstein [3] provided the exact POA of weighted congestion games with affine cost functions ($1 + \phi \approx 2.618$, where ϕ is the golden ratio). They also proved an upper bound of $d^{\Theta(d)}$ for the POA of weighted congestion games with cost functions that are polynomials with nonnegative coefficients and degree at most d . Later, Aland et al. [1] obtained exact worst-case POA bounds for both weighted and unweighted congestion games with cost functions that are polynomials with nonnegative coefficients.

Caragiannis et al. [5] analyzed asymmetric singleton congestion games with affine cost functions. They proved lower bounds of $5/2$ and $1 + \phi$ on the worst-case POA of unweighted and weighted such games, respectively. Gairing and Schoppmann [10] provided a detailed analysis of the POA of singleton congestion games. They generalized the results in [5] to polynomial cost functions, showing that the worst-case POA in asymmetric singleton games is as large as in general games. For symmetric singleton congestion games (i.e., networks of parallel links) and polynomial cost functions they showed that the POA is bounded below by certain Bell numbers. This last result is subsumed by our second contribution.

2 Preliminaries

Congestion Games. A *general weighted congestion game* Γ is composed of a set of N players \mathcal{N} , a set of resources E and a set of non-negative and non-decreasing cost functions from \mathbb{R}^+ to \mathbb{R}^+ . For each player $i \in \mathcal{N}$ a weight w_i and a set $S_i \subseteq 2^E$ of strategies are specified. The congestion on a resource is the total weight of all players using that resource and the associated congestion cost is specified by a function c_e of the congestion on that edge.

An outcome is a choice of strategies $\mathbf{s} = (s_1, s_2, \dots, s_N)$ by players with $s_i \in S_i$. The congestion on a resource e is given by $x_e = \sum_{i: e \in s_i} w_i$. A player's cost is $C_i(\mathbf{s}) = w_i \sum_{e \in s_i} c_e(x_e)$. The *social cost* of an outcome is the sum of the players' costs: $C(\mathbf{s}) = \sum_{i=1}^N C_i(\mathbf{s})$. The total cost can also be written as $C(\mathbf{s}) = \sum_{e \in E} x_e c_e(x_e)$.

In *unweighted congestion games*, all players have unit weight. A congestion game is *symmetric* if all players have the same set of strategies: $S_i = S \subseteq 2^E$ for

all i . In a *singleton* congestion game, every strategy of every players consists of a single resource. Symmetric singleton games are also called *congestion games on parallel links*.

Pure Nash Equilibrium. A *pure Nash equilibrium* is an outcome in which no player has an incentive to deviate from its current strategy. Thus the outcome \mathbf{s} is a pure Nash equilibrium if for each player i and alternative strategy $s_i^* \in S_i$, $C_i(\mathbf{s}) \leq C_i(s_i^*, \mathbf{s}_{-i})$. Here, (s_i^*, \mathbf{s}_{-i}) denotes the outcome that results when player i changes its strategy in \mathbf{s} from s_i to s_i^* . We refer to this inequality as the *Nash condition*.

A pure Nash equilibrium need not have the minimum-possible social cost. The *price of anarchy (POA)* captures how much worse Nash equilibria are compared to the cost of the best social outcome. The POA of a congestion game is defined as the largest value of the ratio $C(\mathbf{s})/C(\mathbf{s}^*)$, where \mathbf{s} ranges over pure Nash equilibria and \mathbf{s}^* ranges over all outcomes. The POA of a class of games is the largest POA among all games in the class.

Weighted congestion games might not admit pure Nash equilibria [16]; see Harks and Klimm [11] for a detailed characterization. Fortunately, our lower bound constructions produce games in which such equilibria exist; and, our upper bounds apply to more general equilibrium concepts that *do* always exist (mixed Nash, correlated, and coarse correlated equilibria, see e.g.. [24, chapter 3]). The POA with respect to these other equilibrium concepts is defined in the obvious way, as the worst-case ratio between the expected social cost of an equilibrium and the minimum-possible social cost.

3 Upper Bounds for Weighted Congestion Games

In this section we provide an upper bound on the POA of weighted congestion games with general cost functions. To explain the upper bound, consider a set \mathcal{C} of cost functions and the set of parameter pairs

$$\mathcal{A}(\mathcal{C}) = \{(\lambda, \mu) : \mu < 1; \quad x^*c(x + x^*) \leq \lambda x^*c(x^*) + \mu xc(x)\}, \quad (1)$$

where the constraints range over all functions $c \in \mathcal{C}$ and real numbers $x \geq 0$ and $x^* > 0$. Each parameter pair (λ, μ) in $\mathcal{A}(\mathcal{C})$ yields the following upper bound on the POA of weighted congestion games with cost functions in \mathcal{C} .

Proposition 1. *For every class \mathcal{C} of cost functions and every pair $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$, the POA of every weighted congestion game with cost functions in \mathcal{C} is at most $\lambda/(1 - \mu)$.*

The upper bound in Proposition 1 is proved via a “smoothness argument” in the sense of [21]. Many earlier works also proved upper bounds on the POA in congestion games in this way (e.g. [1, 7]). As proved in [21], every such upper bound automatically applies to (among other things) the POA of mixed Nash equilibria, correlated equilibria, and coarse correlated equilibria [21].

The best upper bound implied by Proposition 1 is denoted by $\zeta(\mathcal{C})$.

Definition 1. For a class of functions \mathcal{C} , define $\zeta(\mathcal{C}) := \inf \left\{ \frac{\lambda}{1-\mu} : (\lambda, \mu) \in \mathcal{A}(\mathcal{C}) \right\}$. Define $\zeta(\mathcal{C}) := +\infty$ if $\mathcal{A}(\mathcal{C})$ is empty.

4 Lower Bounds for Weighted Congestion Games

In this section we describe two different lower bounds that match the upper bound $\zeta(\mathcal{C})$ given in Definition 1. The first lower bound applies to every class \mathcal{C} of allowable cost functions satisfying a mild technical condition, and makes use of asymmetric congestion games. The second lower bound applies only to polynomial cost functions with nonnegative coefficients, but uses only networks of parallel links.

For each lower bound, we are given a class of cost functions \mathcal{C} , and we will describe a series of games with POA approaching $\zeta(\mathcal{C})$. For each game we specify player weights, player strategies, and resource cost functions. Additionally, we describe two outcomes \mathbf{s} and \mathbf{s}^* . Justifying the example as a proper lower bound requires checking that \mathbf{s} is a (pure) Nash equilibrium and that the ratio of costs of the outcomes \mathbf{s} and \mathbf{s}^* is close to $\zeta(\mathcal{C})$.

Our lower bound constructions are guided by the aspiration to satisfy simultaneously all of the inequalities in the proof of Proposition 1 exactly. This goal translates to the following conditions.

- (a) In the outcome \mathbf{s} , each player is indifferent between its strategy s_i and the deviation s_i^* .
- (b) For each player i , the strategies s_i and s_i^* are disjoint.
- (c) Cost functions and resource congestion in the outcomes \mathbf{s} and \mathbf{s}^* form tuples (c, x, x^*) that correspond to binding constraints in the infimum in (1).
- (d) In the outcome \mathbf{s}^* , each resource is used by a single player.

We believe that satisfying all of these conditions simultaneously is impossible (and can prove it for congestion games on parallel links). Nevertheless, we are able to “mostly” satisfy these conditions, which permits an asymptotic lower bound of $\zeta(\mathcal{C})$ as the number of players and resources tend to infinity.

4.1 Weighted Congestion Games with General Cost Functions

Now we present the lower bound examples that obtain POA arbitrarily close to $\zeta(\mathcal{C})$ for most classes of cost functions \mathcal{C} . We assume that the class \mathcal{C} is *closed under scaling and dilation*, meaning that if $c(x) \in \mathcal{C}$ and $r \in \mathbb{R}^+$, then $rc(x)$ and $c(rx)$ are also in \mathcal{C} . Standard scaling and replication tricks (see [19]) imply that closure under scaling is without loss of generality. Closure under dilation is not without loss but is satisfied by most natural classes of cost functions.

To generically construct lower bound examples, we examine the set of constraints to find ones that are binding, and use them to get close to the values of (λ, μ) that yield $\zeta(\mathcal{C})$.

First, we observe that scaling and dilation does not change the set of constraints in the definition of $\mathcal{A}(\mathcal{C})$. The set of cost functions \mathcal{C} can then be seen

as composed of a number of disjoint equivalence classes (where the relation is differing by scaling and dilation). For simplicity, we assume that the number of equivalence classes of functions in \mathcal{C} is countable.

The set $\mathcal{A}(\mathcal{C})$ is restricted by a number of constraints – one for each cost function $c \in \mathcal{C}$ and non-negative real numbers x, x^* . In the next lemma we establish that the uncountable set of constraints on the set $\mathcal{A}(\mathcal{C})$ can be refined to a countable set of constraints that yield the same set. This proves useful later when we need to reason about constraints that are “most binding”.

Lemma 1. *Given a set of cost functions \mathcal{C} containing a countable number of equivalence classes, the set $\mathcal{A}(\mathcal{C})$ can be represented as constrained by a countable set of constraints.*

Order this countable set of constraints and let A_n denote the set of (λ, μ) pairs that satisfy the first n constraints. Let ζ_n be the minimum value of $\lambda/(1-\mu)$ on this set. We define $\zeta_n = +\infty$ if the set A_n is empty. Observe that ζ_n is a nondecreasing sequence with limit $\zeta(\mathcal{C})$.

For every finite set of constraints we can identify precisely the point where the least value of $\lambda/(1-\mu)$ is obtained. We next note some properties of this infimum that are crucial in building our lower bound example.

Lemma 2. *Fix n , and let A_n, ζ_n be defined as above. Suppose that there exist λ_n, μ_n such that $\zeta_n = \lambda_n/(1-\mu_n)$. Then there exist z_1, z_2 s.t. for every $w > 0$ there exist $c_1, c_2, \eta \in [0, 1]$ s.t.*

$$c_j(w \cdot (z_j + 1)) = \lambda_n c_j(w) + \mu_n c_j(w \cdot z_j) \cdot z_j \text{ for } j \in \{1, 2\} \text{ and ,} \quad (2)$$

$$\begin{aligned} \eta \cdot c_1(w \cdot (z_1 + 1)) + (1 - \eta) \cdot c_2(w \cdot (z_2 + 1)) \\ = \eta \cdot c_1(w \cdot z_1) \cdot z_1 + (1 - \eta) \cdot c_2(w \cdot z_2) \cdot z_2. \end{aligned} \quad (3)$$

For a fixed index n , if ζ_n is attained for some λ, μ , the lemma above provides z_1, z_2 and c_1, c_2 that correspond to the most binding constraints. We use these in our lower bound construction. The following example serves as a lower bound.

Lower Bound 1. For some parameter $k \in \mathbb{N}$ (chosen later) we construct a weighted congestion game with player set \mathcal{N} and resource set E as follows (see Figure 1).

Player strategies: Organize the resources in a tree of depth k , which is a complete binary tree of depth $k - 1$ with each leaf extended by a path of length 1. For each non-leaf node i in the tree there is a player i with 2 strategies: either choose node i or all children of i .

Player weights: If i is the root then $w_i = 1$ otherwise if node i is the left (right) child of some node j then $w_i = w_j \cdot z_1$ ($w_i = w_j \cdot z_2$), where z_1, z_2 are chosen for (λ_n, μ_n) as in Lemma 2. Let $\mathcal{N}_L \subset \mathcal{N}$ be the set of players connected to a leaf.

Cost functions: The cost functions of the resources are defined recursively:
– For the root we can choose any cost function $c \in \mathcal{C}$ with $c(1) = 1$. By a scaling argument such a cost function exists.

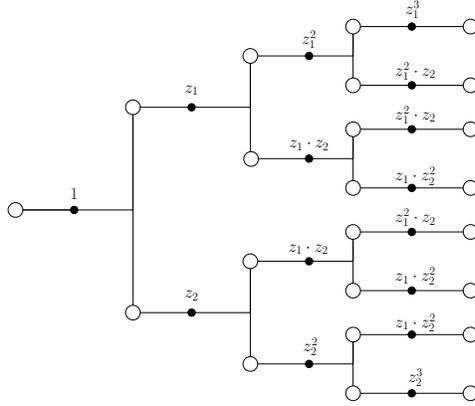


Fig. 1. Construction for the Asymmetric Lower Bound ($k = 4$). Hollow circles denote resources and solid circles denote players.

- Any leaf resource gets the same cost function as its parent.
- Consider an arbitrary resource e which is not a leaf nor its children are leaves. Let c_e be its cost function and let w_e be the weight of the corresponding player e . Let l, r be the left and right child of e respectively. By construction the corresponding players have weights $w_l = z_1 \cdot w_e$ and $w_r = z_2 \cdot w_e$. Under all pairs of cost functions c_1, c_2 that satisfy Lemma 2 for $x^* = w_e$ choose those that also satisfy

$$c_j(w_e \cdot z_j) \cdot z_j = c_e(w_e) \quad (4)$$

for $j \in \{1, 2\}$. By a scaling argument such a pair always exists. Let η_e be the corresponding value for η in Lemma 2 and define

$$c_l = \eta_e \cdot c_1 \quad \text{and} \quad c_r = (1 - \eta_e) \cdot c_2. \quad (5)$$

Nash strategy: The Nash outcome in this example is the outcome \mathbf{s} where each player chooses the resource closer to the root.

Optimal strategy: the optimal outcome \mathbf{s}^* is the outcome where each player chooses its strategy further from the root.

We have to show that, in the construction above, the outcome \mathbf{s} is indeed a Nash equilibrium and that the ratio $C(\mathbf{s})/C(\mathbf{s}^*)$ is at least $\lambda_n/(1 - \mu_n)$. We prove these claims in the following lemma:

Lemma 3. *The POA of the congestion game in Lower Bound 1 is $\lambda_n/(1 - \mu_n)$.*

The analysis above still leaves out the case when ζ_n is not obtained by any feasible values of λ_n, μ_n . A slightly modified example serves as the lower bound in this case. Combining all of the analysis above we obtain the following result,

Theorem 1. *For every class \mathcal{C} of allowable cost functions that is closed under dilation, the worst-case POA of weighted congestion games with cost functions in \mathcal{C} is precisely $\zeta(\mathcal{C})$.*

4.2 A Lower Bound for Parallel-Link Networks

In the following, we focus on weighted congestion games with cost functions that are polynomials with nonnegative coefficients and maximum degree d . For such games we show that the POA does not change if we restrict to (symmetric) weighted congestion games on parallel links. Recall that for the general case the POA was shown to be ϕ_d^{d+1} [1] where ϕ_d is such that $(\phi_d + 1)^d = \phi_d^{d+1}$. We establish the following theorem.

Theorem 2. *For weighted congestion games on parallel-link networks with cost functions that are polynomials with nonnegative coefficients and maximum degree d , the worst-case POA is precisely ϕ_d^{d+1} .*

We reiterate that the lower bound in Theorem 2 is new even for affine cost functions. We describe the example below.

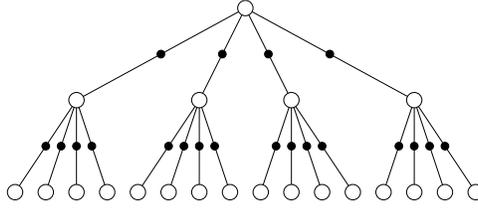


Fig. 2. Structure of the symmetric lower bound example for branching factor $\alpha = 4$ and number of levels $k = 3$. Hollow circles denote resources and solid circles denote players.

Lower Bound 2. Let k be an integer. We construct the following congestion game on parallel links. (See Figure 2 for reference.) Let ϕ_d satisfy $(\phi_d + 1)^d = \phi_d^{d+1}$. Let α be an integer satisfying $\alpha^d \geq \phi_d^{2d+2}$ which implies, $\alpha^d \geq \phi_d^{d+1}$.

Player strategies: Player strategies are single resources and all players have access to all resources.

Cost functions: Group the resources in groups $A_0, A_1 \dots A_k$. For each $i = 0, 1, \dots, k - 1$, group A_i contains α^i resources with cost function $c_i(x) = (\alpha^d / \phi_d^{d+1})^i x^d$. The last group A_k contains α^k resources with the cost function $c_k(x) = \alpha^d (\alpha^d / \phi_d^{d+1})^{k-1} x^d$. These resources are arranged in a tree with resources from group A_i at level i of the tree.

Player weights: Group players into groups P_1, \dots, P_k . For $i = 1, 2, \dots, k$, group P_i contains one player for each resources in A_i with player weight $w_i = (\alpha / \phi_d)^{k-i}$.

Optimal strategy: Players in group P_i play on resources in group A_i in the optimal strategy. Denote this outcome by \mathbf{s}^* .

Nash strategy: Players in group P_i play on resources in group A_{i-1} in the Nash strategy with α players on each resource. Denote this outcome by \mathbf{s} .

The example described above is symmetric — all players have access to all of the strategies. Theorem 2 then follows from the upper bound in [1] and the following lemma.

Lemma 4. *For the games in Lower Bound 2, outcome s is a pure Nash equilibrium and $\lim_{k \rightarrow \infty} \frac{C(s)}{C(s^*)} = \phi_d^{d+1}$.*

5 Unweighted Congestion Games

We show that for symmetric unweighted congestion games the worst-case POA is the same as that for general unweighted congestion games. A result in [21] gives an upper bound on the POA of unweighted congestion games. For a class of congestion cost functions \mathcal{C} , one uses the set of parameter pairs

$$\mathcal{A}(\mathcal{C}) = \{(\lambda, \mu) : \mu < 1; x^* c(x+1) \leq \lambda x^* c(x^*) + \mu x c(x)\}, \quad (6)$$

where the constraints range over cost functions $c \in \mathcal{C}$ and integers $x \geq 0, x^* > 0$. Then for each $(\lambda, \mu) \in \mathcal{A}(\mathcal{C})$, the pure POA is at most $\lambda/(1-\mu)$ in unweighted congestion games with cost functions in \mathcal{C} . We use $\gamma(\mathcal{C})$ to denote the best such upper bound: $\gamma(\mathcal{C}) = \inf \{\lambda/(1-\mu) : (\lambda, \mu) \in \mathcal{A}(\mathcal{C})\}$. Roughgarden [21] also showed how to construct an *asymmetric* unweighted congestion game that matches the upper bound $\gamma(\mathcal{C})$ (for each \mathcal{C}). Here we show that even symmetric games can be used to achieve this upper bound, in the limit as the number of players and resources tends to infinity.

We establish the following theorem.

Theorem 3. *For every set of cost functions \mathcal{C} , there exist symmetric congestion games with cost functions in \mathcal{C} and POA arbitrarily close to $\gamma(\mathcal{C})$.*

Along the lines of [21] we define $\gamma(\mathcal{C}, n)$ as the minimal value of $\lambda/(1-\mu)$ that can be obtained when the load on each edge is restricted to be at most n . Then $\gamma(\mathcal{C}, n)$ approaches $\gamma(\mathcal{C})$ as n approaches ∞ . For any finite n the feasible region for (λ, μ) is the intersection of a finite number of half planes, one for each value of x and x^* . An additional constraint on the feasible region is that $\mu < 1$. We now recall the following lemma from [21].

Lemma 5. *Fix finite n and a set of functions \mathcal{C} and suppose there exists $(\hat{\lambda}, \hat{\mu})$ such that, $\frac{\hat{\lambda}}{1-\hat{\mu}} = \gamma(\mathcal{C}, n)$. Then there exist $c_1, c_2 \in \mathcal{C}$, $x_1, x_2 \in \{0, 1, \dots, n\}$, $x_1^*, x_2^* \in \{1, 2, \dots, n\}$ such that, $c_j(x_j+1)x_j^* = \hat{\lambda}c_j(x_j^*)x_j^* + \hat{\mu}c_j(x_j)x_j$ for j in $\{1, 2\}$, $\beta_{c_1, x_1, x_1^*} < 1$, and $\beta_{c_2, x_2, x_2^*} \geq 1$, where β_{c, x, x^*} is defined as $\frac{xc(x)}{x^*c(x+x^*)}$.*

Note that the lemma as stated here differs from the one in [21] in the additional condition on $\beta_{c_1, x_1, x_1^*}, \beta_{c_2, x_2, x_2^*}$. However the modified version can be easily obtained by noting that we are always guaranteed a half plane with $\beta < 1$ and as long as the value of $\gamma(\mathcal{C}, n)$ is attained there is another half plane with $\beta \geq 1$.

We now describe the lower bound.

Lower Bound 3. Let N be an integer which will denote the number of players. Let c_1, x_1, x_1^* and c_2, x_2, x_2^* be defined as in the above lemma.

Resources and Cost functions: There are two groups of resources A_1 and A_2 . For $j = 1, 2$ group A_j contains $\binom{N}{x_j} \binom{N}{x_j^*}$ resources with the congestion cost function $\alpha_j c_j(x)$. We choose α_1, α_2 later.

Players: There are N players each with unit weight. Each player i has access to only two strategies, s_i and s_i^* .

Optimal strategy: For each $j \in \{1, 2\}, i \in \{1, 2, \dots, N\}$: define set $O(i, j)$ composed of $\Delta_j = \binom{N}{x_j} \binom{N-1}{x_j^*-1}$ resources of type A_j . Each resource is contained in x_j^* sets. Define the optimal strategy for player i , s_i^* , as the union of $O(i, 1)$ and $O(i, 2)$.

Equilibrium strategy: For each $j \in \{1, 2\}, i \in \{1, 2, \dots, N\}$: define set $E(i, j)$ as composed of $\frac{x_j}{N} \Delta_j$ resources from each $O(i', j)$ for all $i' \in \{1, 2, \dots, N\}$. The number of resources in $E(i, j)$ is $\frac{x_j}{x_j^*} \Delta_j$. Each resource is contained in x_j sets. Define the equilibrium strategy s_i for player i as the union of $E(i, 1)$ and $E(i, 2)$.

We choose α_1, α_2 s.t. each player is indifferent between his strategy in the equilibrium and optimal outcomes when all other players play their equilibrium strategies. The following lemma establishes that such non-negative α_1, α_2 exist.

Lemma 6. *If tuples c_1, x_1, x_1^* and c_2, x_2, x_2^* with $x_1, x_2 \geq 0, x_1^*, x_2^* > 0$ are s.t. $\beta_{c_1, x_1, x_1^*} < 1$ and $\beta_{c_2, x_2, x_2^*} \geq 1$, then for the game instance described in Lower Bound 3, there exist $\alpha_1, \alpha_2 \geq 0$ s.t. for each player i , $C_i(s_i, s_{-i}) = C_i(s_i^*, s_{-i})$.*

It now remains to prove that the Lower Bound 3 has the desired POA.

Lemma 7. *The POA of the game in Lower Bound 3 approaches $\hat{\lambda}/(1 - \hat{\mu})$ as $N \rightarrow \infty$.*

Even when $\gamma(\mathcal{C}, n)$ is not attained by any pair (λ, μ) , a similar construction yields an instance with POA arbitrarily close to $\gamma(\mathcal{C}, n)$.

Lemma 8. *For every class \mathcal{C} of functions and integer n such that the value $\gamma(\mathcal{C}, n)$ is not attained by any feasible pair (λ, μ) , there exist symmetric unweighted congestion games with cost functions in \mathcal{C} and POA arbitrarily close to $\gamma(\mathcal{C}, n)$.*

Lemmas 7 and 8 together establish Theorem 3.

References

1. S. Aland, D. Dumrauf, M. Gairing, B. Monien, and F. Schoppmann. Exact price of anarchy for polynomial congestion games. In *STACS*, pages 218–229, 2006.
2. E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.

3. B. Awerbuch, Y. Azar, and A. Epstein. Large the price of routing unsplitable flow. In *STOC*, pages 57–66, 2005.
4. D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Athena Scientific, 1997.
5. I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *ICALP (1)*, pages 311–322, 2006.
6. G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *ESA*, pages 59–70, 2005.
7. G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *STOC*, pages 67–73, 2005.
8. D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. In *SAGT*, pages 33–45, 2008.
9. D. Fotakis, S. C. Kontogiannis, and P. G. Spirakis. Atomic congestion games among coalitions. In *ICALP (1)*, pages 572–583, 2006.
10. M. Gairing and F. Schoppmann. Total latency in singleton congestion games. In *WINE*, pages 381–387, 2007.
11. T. Harks and M. Klimm. On the existence of pure nash equilibria in weighted congestion games. In *ICALP*, 2010.
12. A. Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *STOC*, pages 89–98, 2006.
13. R. Holzman and N. Law-Yone. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46:193–205, 2003.
14. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *STACS*, pages 404–413, 1999.
15. T. Lücking, M. Mavronicolas, B. Monien, and M. Rode. A new model for selfish routing. *STACS 2004, TCS 2008*, 406(3):187–206, 2008.
16. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, March 1996.
17. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.
18. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
19. T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
20. T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proceedings of the ICM*, volume III, pages 1071–1094, 2006.
21. T. Roughgarden. Intrinsic robustness of the price of anarchy. In *STOC*, pages 513–522, 2009.
22. L. S. Shapley. *Additive and Non-Additive Set Functions*. PhD thesis, Department of Mathematics, Princeton University, 1953.
23. S. Suri, C. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1):79–96, 2007.
24. H. P. Young. *Strategic Learning and its Limits*. Oxford University Press, London, U.K., 2005.