

Available online at www.sciencedirect.com ScienceDirect

Games and Economic Behavior ••• (••••) ••••••••

**GAMES and
Economic
Behavior**

www.elsevier.com/locate/geb

A geometric approach to the price of anarchy in nonatomic congestion games

José R. Correa ^a, Andreas S. Schulz ^b, Nicolás E. Stier-Moses ^{c,*}^a School of Business, Universidad Adolfo Ibáñez, Av. Presidente Errázuriz 3485, Las Condes, Santiago, Chile^b Sloan School of Management, Massachusetts Institute of Technology, Office E53-361, 77 Massachusetts Avenue, Cambridge, MA 02139, USA^c Graduate School of Business, Columbia University, Uris Hall, Room 418, 3022 Broadway Avenue, New York, NY 10027, USA

Received 30 January 2006

Abstract

We present a short, geometric proof for the price-of-anarchy results that have recently been established in a series of papers on selfish routing in multicommodity flow networks and on nonatomic congestion games. This novel proof also facilitates two new types of theoretical results: On the one hand, we give pseudo-approximation results that depend on the class of allowable cost functions. On the other hand, we derive stronger bounds on the inefficiency of equilibria for situations in which the equilibrium costs are within reasonable limits of the fixed costs. These tighter bounds help to explain empirical observations in vehicular traffic networks. Our analysis holds in the more general context of nonatomic congestion games, which provide the framework in which we describe this work.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C02; C61; C72*Keywords:* Noncooperative games; Nonatomic games; Congestion games; Wardrop equilibrium; Price of anarchy

* Corresponding author.

E-mail addresses: correa@uai.cl (J.R. Correa), schulz@mit.edu (A.S. Schulz), stier@gsb.columbia.edu (N.E. Stier-Moses).

1. Introduction

Congestion games (Rosenthal, 1973) are noncooperative games in which players' strategies consist of subsets of resources, and the utility of a player depends only on the number of players choosing the same or some overlapping strategy. We consider *nonatomic* congestion games (Schmeidler, 1973), which model interactions involving a continuous number of players, each having a negligible affect on other players. Nonatomic congestion games have been studied, among others, by Milchtaich (2000, 2004), Chau and Sim (2003), and Roughgarden and Tardos (2004).

The most prominent example of a nonatomic congestion game is the traffic routing model of Wardrop (1952). The arcs in a given network represent the resources, the different origin–destination pairs correspond to the player types, and the strategies available to a particular player type are the paths in the network between its origin–destination pair. The cost of an arc describes the delay experienced by traffic traversing that arc as a function of the flow on that arc. A social optimum corresponds to a multicommodity flow of minimum total delay, whereas in a Wardrop equilibrium, every player is traveling on a shortest path.

Nash equilibria in general and Wardrop equilibria in particular are typically inefficient: they generally do not minimize the social cost. Koutsoupias and Papadimitriou (1999) proposed to analyze the inefficiency of equilibria from a worst-case perspective; this led to the notion of “price of anarchy” (Papadimitriou, 2001), which is the ratio of the worst social cost of a Nash equilibrium to the cost of an optimal solution. In the context of selfish routing (i.e., the traffic model described in the previous paragraph), the price of anarchy was analyzed in a series of papers for increasingly more general classes of cost functions and other model features; see, among others, Roughgarden and Tardos (2002, 2004), Roughgarden (2003), Chau and Sim (2003), Correa et al. (2004), and Perakis (2007).

In this article, we give alternative proofs for the price-of-anarchy results in the above-mentioned papers. Our proofs simplify and unify previous arguments, and they provide insights that enable us to extend these results to more general settings. The paper is organized as follows. Nonatomic congestion games are formally defined in Section 2. In Section 3, we present new proofs for two known bounds on the inefficiency of equilibria in nonatomic congestion games with affine, separable cost functions. These proofs rely on a new interpretation of the parameter β , originally introduced by Correa et al. (2004) in the context of traffic routing. This interpretation sets the stage for various generalizations. In Section 4, we consider nonseparable, nonlinear cost functions. We also discuss the special cases of separable cost functions (Section 4.1) and of situations in which the variable costs of resources do not exceed their fixed costs by too much (Section 4.2). Section 5 contains our concluding remarks.

2. Nonatomic congestion games

A nonatomic congestion game consists of a finite set A of resources and k different types of players. Players are infinitesimal agents, and the continuum of players of type i is represented by the interval $[0, n_i]$, for some $n_i > 0$, $i = 1, 2, \dots, k$. Each player type i possesses a set \mathcal{S}_i of strategies, and each strategy consists of a subset of the resources. For notational convenience, we assume that the sets \mathcal{S}_i , $i = 1, 2, \dots, k$, are disjoint. We denote their union by \mathcal{S} . The rate of consumption of a resource $a \in S$ by a strategy $S \in \mathcal{S}_i$ is given by $r_{a,S} \geq 0$. Each player selects a strategy, which leads to a strategy distribution $x = (x_S)_{S \in \mathcal{S}}$ with $\sum_{S \in \mathcal{S}_i} x_S = n_i$ for each player type i , and $x_S \geq 0$ for all $S \in \mathcal{S}$. A strategy distribution generates a utilization rate

$x_a := \sum_{i=1}^k \sum_{S \in \mathcal{S}_i} r_{a,S} x_{iS}$ for each resource $a \in A$. In the sequel, we use x interchangeably to denote a strategy distribution or its associated vector of utilization rates. Moreover, we let X represent both the space of feasible strategy distributions and that of feasible utilization rate vectors. This abuse of notation is harmless because we are interested only in the total cost of different strategy distributions, and the total cost of any strategy distribution is fully specified by its corresponding vector of utilization rates.

The cost of a resource $a \in A$ is given by a continuous function $c_a : \mathbb{R}_{\geq 0}^A \rightarrow \mathbb{R}_{\geq 0}$. The cost functions c_a considered here include separable functions, for which $c_a(x) = c_a(x_a)$, and affine functions, for which $c_a(x) = \sum_{a' \in A} c_{a,a'} x_{a'} + b_a$ for some coefficients $c_{a,a'}$ and b_a . The cost experienced by a player that selects a strategy $S \in \mathcal{S}$ is given by $c_S(x) := \sum_{a \in S} r_{a,S} c_a(x)$. We define the social cost $C(x)$ of a strategy distribution x as the total disutility experienced by all players:

$$C(x) := \sum_{i=1}^k \sum_{S \in \mathcal{S}_i} c_S(x) x_S = \sum_{a \in A} c_a(x) x_a = \langle c(x), x \rangle.$$

Here, $c(x) := (c_a(x))_{a \in A}$, and $\langle c(x), x \rangle$ denotes the inner product of the two vectors $c(x)$ and x .

A social optimum x^{OPT} is a feasible strategy distribution of minimum social cost; i.e., $C(x^{\text{OPT}}) \leq C(x)$ for all strategy distributions $x \in X$. Because c is continuous and X is compact and convex, social optima exist and are well-defined.

Extending the notion of Wardrop equilibrium (1952) to nonatomic congestion games, a strategy distribution x^{EQ} is called an equilibrium if all players of the same type experience the same cost, and there is no strategy of smaller cost; i.e., $c_S(x^{\text{EQ}}) \leq c_{S'}(x^{\text{EQ}})$ for any two strategies $S, S' \in \mathcal{S}_i$ with $x_S^{\text{EQ}} > 0$, for each $i = 1, 2, \dots, k$. It is well known (e.g., de Palma and Nesterov, 1998) that if a strategy distribution x^{EQ} is an equilibrium, then it satisfies

$$\langle c(x^{\text{EQ}}), x^{\text{EQ}} - x \rangle \leq 0 \quad \text{for all strategy distributions } x \in X. \tag{1}$$

Hartman and Stampacchia (1966) proved that for c continuous, (1) always has a solution (if X is compact and convex). In the following, we assume that the operator $c : \mathbb{R}_{\geq 0}^A \rightarrow \mathbb{R}_{\geq 0}^A$ is monotone; i.e., $\langle c(x) - c(y), x - y \rangle \geq 0$ for all $x, y \in \mathbb{R}_{\geq 0}^A$. Monotonicity guarantees the existence of equilibria, and every solution to the variational inequality (1) is, in fact, an equilibrium (de Palma and Nesterov, 1998).

For a given instance of a nonatomic congestion game, we measure the inefficiency of an equilibrium with the help of the ratio $C(x^{\text{EQ}})/C(x^{\text{OPT}})$, where x^{EQ} and x^{OPT} denote a (worst) equilibrium and a social optimum, respectively. The price of anarchy for a class of games is defined as the supremum of this ratio over all instances belonging to that class.

3. A graphical proof for affine and separable cost functions

Let us begin by studying nonatomic congestion games with affine and separable cost functions. We present a new proof of an existing bound on the worst-case inefficiency of equilibria, and a strengthening of a pseudo-approximation bound. The key insight into both results comes from a graphical representation of the situation.

We assume for the rest of this section that the cost of a resource $a \in A$ under the utilization rate vector x is $c_a(x) = c_a(x_a) = c_a x_a + b_a$ for some nonnegative coefficients c_a and b_a . Under this assumption, it is known not only that an equilibrium x^{EQ} exists, but also that different equilibria have the same social cost (Beckmann et al., 1956).

The following result shows that, in equilibrium, the social cost cannot increase by more than 33%, compared to a social optimum. This means that the lack of central coordination does not cause a significant loss of efficiency. The most compact proof of this fact was presented in the context of selfish routing by Correa et al. (2004); our proof below can be viewed as a geometric variant of it. The result itself is due to Roughgarden and Tardos (2002, 2004).

Theorem 3.1. *Let x^{EQ} be an equilibrium of a nonatomic congestion game with separable, affine cost functions, and let x^{OPT} be a social optimum. Then, $C(x^{\text{EQ}}) \leq 4/3C(x^{\text{OPT}})$.*

Proof. Let x be a feasible strategy distribution. Because of (1), we have

$$\begin{aligned} C(x^{\text{EQ}}) &= \sum_{a \in A} c_a(x_a^{\text{EQ}})x_a^{\text{EQ}} \leq \sum_{a \in A} c_a(x_a^{\text{EQ}})x_a \\ &= \sum_{a \in A} c_a(x_a)x_a + \sum_{a \in A} (c_a(x_a^{\text{EQ}}) - c_a(x_a))x_a. \end{aligned} \tag{2}$$

Since the functions c_a are nondecreasing, we need to focus only on the expressions $(c_a(x_a^{\text{EQ}}) - c_a(x_a))x_a$ for which $x_a < x_a^{\text{EQ}}$ to bound the last term in (2) from above. In this case, $(c_a(x_a^{\text{EQ}}) - c_a(x_a))x_a$ is equal to the area of the shaded rectangle in Fig. 1. Note that the area of any rectangle whose upper-left corner point is $(0, c_a(x_a^{\text{EQ}}))$ and whose lower-right corner point lies on the line representing $c_a(y_a) = c_a y_a + b_a$, is at most half that of the triangle defined by the three points $(0, c_a(x_a^{\text{EQ}}))$, $(0, b_a)$, and $(x_a^{\text{EQ}}, c_a(x_a^{\text{EQ}}))$. In turn, the area of the triangle is at most half that of the rectangle defined by the two points $(0, 0)$ and $(x_a^{\text{EQ}}, c_a(x_a^{\text{EQ}}))$. In particular,

$$(c_a(x_a^{\text{EQ}}) - c_a(x_a))x_a \leq \frac{1}{4}c_a(x_a^{\text{EQ}})x_a^{\text{EQ}}. \tag{3}$$

Setting $x = x^{\text{OPT}}$ proves the claim. \square

An immediate consequence of this proof is a pseudo-approximation result, which upper bounds the social cost of an equilibrium by that of an optimal strategy distribution for the same




Fig. 1. Illustration of the proof of Theorem 3.1.

game with more players of each type. More precisely, to obtain the following result, one needs only the following inequality derived in the preceding proof,

$$\sum_{a \in A} c_a(x_a^{\text{EQ}})x_a \leq C(x) + \frac{1}{4}C(x^{\text{EQ}}), \quad (4)$$

which holds for *any* nonnegative vector x (i.e., x need not be a feasible strategy distribution).

Corollary 3.2. *If x^{EQ} is an equilibrium of a nonatomic congestion game with separable, affine cost functions, and y^{OPT} is a social optimum for the same game with $5/4$ times as many players of each type,¹ then $C(x^{\text{EQ}}) \leq C(y^{\text{OPT}})$.*

Proof. Using (1) together with the feasibility of the vector $(4/5)y^{\text{OPT}}$ for the original game, and (4), we obtain

$$C(x^{\text{EQ}}) = \frac{5}{4} \sum_{a \in A} c_a(x_a^{\text{EQ}})x_a^{\text{EQ}} - \frac{1}{4}C(x^{\text{EQ}}) \leq \sum_{a \in A} c_a(x_a^{\text{EQ}})y_a^{\text{OPT}} - \frac{1}{4}C(x^{\text{EQ}}) \leq C(y^{\text{OPT}}). \quad \square$$

The smallest value for which this corollary remains true is, in fact, $5/4$. A tight example is presented in a more general context in Section 4.2. The first result of this kind was given by Roughgarden and Tardos (2002), who showed that an equilibrium traffic assignment causes a total travel time of at most that of a social optimum routing twice as much traffic. This result and its subsequent extension to general nonatomic congestion games (Roughgarden and Tardos, 2004) hold for arbitrary (albeit separable) cost functions. The selfish routing version of Corollary 3.2 is due to Chakrabarty (2004) and inspired us to qualify the pseudo-approximation bounds according to the class of cost functions considered; see the next section for details.

4. Nonseparable and nonlinear cost functions

In this section, we generalize the results of the previous section to nonatomic congestion games with general cost functions, as defined in Section 2. Afterwards, we discuss some important special cases for which we can provide additional insights.

In several practical situations, the cost of using one resource may depend on the rate of consumption of other resources, and the relations between utilization rates and costs may be nonlinear. For instance, the time a vehicle needs to cross through a stop sign clearly depends on the amount of traffic on the perpendicular street; the waiting time of passengers at a given bus stop depends on the number of passengers boarding the bus at previous stops; or, to give an example in the context of wireless communication networks, transmission delays might depend on the load of neighboring cells, because of interference. In transportation science, congestion effects are usually modeled with the help of degree-4 polynomials (Bureau of Public Roads, 1964), and in telecommunication engineering, delays usually arise from queuing effects; in either case, the functions are nonlinear.

In the previous section, we used the linearity of the cost functions c_a in one place only, namely when we proved (3). An appropriate generalization of (3) is, in fact, the key for extending The-

¹ Formally, the continuum of players of type i in the new game is represented by the interval $[0, \frac{5}{4}n_i]$, $i = 1, 2, \dots, k$.

orem 3.1 and Corollary 3.2 to more general classes of cost functions. To that effect, we define, for a cost function c and a vector $v \in X$,

$$\beta(c, v) := \max_{x \in \mathbb{R}_{\geq 0}^A} \frac{\langle c(v) - c(x), x \rangle}{\langle c(v), v \rangle},$$

where $0/0 = 0$, by convention. For a given class of cost functions \mathcal{C} , we let

$$\beta(\mathcal{C}) := \sup_{c \in \mathcal{C}, v \in X} \beta(c, v).$$

Note that, because of the monotonicity of $c \in \mathcal{C}$, $\beta(\mathcal{C}) \leq 1$. This parameter was originally introduced by Correa et al. (2004) in the context of traffic routing and separable cost functions. It leads directly to the following lemma, which is a generalization of (4).

Lemma 4.1. *Let x^{EQ} be an equilibrium of a nonatomic congestion game with cost functions drawn from a class \mathcal{C} of nonseparable cost functions, and let x be a nonnegative vector. Then, $\langle c(x^{\text{EQ}}), x \rangle \leq C(x) + \beta(\mathcal{C})C(x^{\text{EQ}})$.*

Proof. Using the definition of $\beta(\mathcal{C})$ and the linearity of the inner product, we get

$$\begin{aligned} \langle c(x^{\text{EQ}}), x \rangle &= \langle c(x), x \rangle + \langle c(x^{\text{EQ}}) - c(x), x \rangle \\ &\leq \langle c(x), x \rangle + \beta(c, x^{\text{EQ}}) \langle c(x^{\text{EQ}}), x^{\text{EQ}} \rangle \\ &\leq C(x) + \beta(\mathcal{C})C(x^{\text{EQ}}). \quad \square \end{aligned}$$

The following theorem generalizes the main results of Section 3, yielding price-of-anarchy and pseudo-approximation results for nonatomic congestion games with nonseparable and nonlinear cost functions. Actually, it also strengthens previously known price-of-anarchy bounds for the case of nonseparable cost functions. These prior bounds require stronger assumptions on the cost functions, such as convexity and differentiability. Chau and Sim (2003) proved that the price of anarchy for nonseparable and symmetric cost functions is bounded by a natural extension of the parameter $\alpha(\mathcal{C})$ of Roughgarden and Tardos (see the discussion in Section 4.1 below). Perakis (2007) considered general nonseparable cost functions and proved upper bounds on the price of anarchy using variational inequalities as well. Her bounds depend on two parameters that measure the asymmetry and the nonlinearity of the cost functions considered.

Theorem 4.2. *Let x^{EQ} be an equilibrium of a nonatomic congestion game with cost functions drawn from a class \mathcal{C} of nonseparable cost functions.*

- (a) *If x^{OPT} is a social optimum for this game, then $C(x^{\text{EQ}}) \leq (1 - \beta(\mathcal{C}))^{-1}C(x^{\text{OPT}})$.*
- (b) *If y^{OPT} is a social optimum for the same game with $1 + \beta(\mathcal{C})$ times as many players of each type, then $C(x^{\text{EQ}}) \leq C(y^{\text{OPT}})$.*

Proof. For (a), it suffices to use (1) and Lemma 4.1:

$$C(x^{\text{EQ}}) = \langle c(x^{\text{EQ}}), x^{\text{EQ}} \rangle \leq \langle c(x^{\text{EQ}}), x^{\text{OPT}} \rangle \leq C(x^{\text{OPT}}) + \beta(\mathcal{C})C(x^{\text{EQ}}).$$

Let us now prove part (b). Because of the feasibility of $(1 + \beta(\mathcal{C}))^{-1}y^{\text{OPT}}$ for the original game, we have that $\langle c(x^{\text{EQ}}), x^{\text{EQ}} \rangle \leq \langle c(x^{\text{EQ}}), (1 + \beta(\mathcal{C}))^{-1}y^{\text{OPT}} \rangle$. Therefore,

$$\begin{aligned}
 C(x^{EQ}) &= (1 + \beta(C))\langle c(x^{EQ}), x^{EQ} \rangle - \beta(C)\langle c(x^{EQ}), x^{EQ} \rangle \\
 &\leq (1 + \beta(C))\langle c(x^{EQ}), (1 + \beta(C))^{-1}y^{OPT} \rangle - \beta(C)\langle c(x^{EQ}), x^{EQ} \rangle \\
 &\leq C(y^{OPT}) + \beta(C)C(x^{EQ}) - \beta(C)C(x^{EQ}) \\
 &= C(y^{OPT}).
 \end{aligned}$$

The last inequality follows from Lemma 4.1. \square

Both parts of this theorem are tight for particular classes of cost functions, as we will illustrate in Section 4.2.

It remains to determine the value of $\beta(C)$ for certain classes of cost functions. We start off by considering affine cost functions of the form $c(x) = Ax + b$, with $b \geq 0$ and A symmetric and positive semidefinite. In this setting, Theorem 4.2 provides a simple proof of a result by Chau and Sim (2003), who established that the price of anarchy for this class of cost functions is at most $4/3$. Indeed, in this case

$$\beta(c, v) = \max_{x \in \mathbb{R}_{\geq 0}^A} \frac{\langle c(v) - c(x), x \rangle}{\langle c(v), v \rangle} = \frac{\max_{x \in \mathbb{R}_{\geq 0}^A} \langle A(v - x), x \rangle}{\langle Av, v \rangle + \langle b, v \rangle}.$$

As A is symmetric and positive semidefinite, the numerator amounts to a convex minimization problem, and the optimum is attained at $x = v/2$, leading to a β -value of $1/4$. Theorem 4.2 yields $C(x^{EQ}) \leq 4/3C(x^{OPT})$, where x^{OPT} is a social optimum for this game. Moreover, $C(x^{EQ}) \leq C(y^{OPT})$, for a social optimum y^{OPT} of the same game with $5/4$ times as many players of each type.

4.1. Separable cost functions

An important special case is that of nonatomic congestion games with separable, nonlinear cost functions. In practice, this class of cost functions is often used as a simplified model of reality when there are insufficient data to estimate the interdependencies among different resources. Moreover, several prior results were derived in this setting. A separable, continuous, and nondecreasing cost function $c_a : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ is “standard” if it is differentiable, and its product with the identity function is convex.

Theorem 4.3. (See Roughgarden and Tardos, 2004, Theorem 4.7.) For a nonzero cost function c_a in a given class \mathcal{C} of standard cost functions, define

$$\alpha(c_a) := \sup_{x_a > 0 : c_a(x_a) > 0} (\lambda\mu + (1 - \lambda))^{-1},$$

where $\lambda \in [0, 1]$ satisfies $c_a(\lambda x_a) + (\lambda x_a)c'_a(\lambda x_a) = c_a(x_a)$ and $\mu = c_a(\lambda x_a)/c_a(x_a) \in [0, 1]$. Let x^{EQ} be an equilibrium of a nonatomic congestion game with cost functions drawn from \mathcal{C} . If x^{OPT} is a social optimum of this game, then $C(x^{EQ}) \leq \alpha(C)C(x^{OPT})$, where $\alpha(C) := \sup_{0 \neq c_a \in \mathcal{C}} \alpha(c_a)$.

Under the assumptions of Theorem 4.3, one can show that $\alpha(C) = (1 - \beta(C))^{-1}$ (Correa et al., 2004). In particular, Theorem 4.3 is, in fact, a special case of Theorem 4.2 (a). Indeed, for such cost functions, $\beta(C) = \sup_{c_a \in \mathcal{C}, v_a \geq 0} \beta(c_a, v_a)$, where

$$\beta(c_a, v_a) := \max_{x_a \geq 0} \frac{(c_a(v_a) - c_a(x_a))x_a}{c_a(v_a)v_a} = \max_{0 \leq x_a \leq v_a} \frac{(c_a(v_a) - c_a(x_a))x_a}{c_a(v_a)v_a}.$$




Fig. 2. Illustration of the definition of β .

While this coincides with the original definition of $\beta(C)$ by Correa et al. (2004), the authors did not provide the interpretation offered by Fig. 2. If we go back to the proof of Theorem 3.1 and argue along similar lines, we encounter the situation depicted in this figure. It remains to observe that $\beta(C)$ is an upper bound on the ratio of the area of the shaded rectangle to that of the large rectangle.

Concrete values of $\beta(C)$ can be readily calculated for several classes of cost functions. For example, the maximum inefficiency of equilibria in games with quadratic or cubic cost functions is 1.626 and 1.896, respectively; for polynomials of degree at most 4, it is 2.151. More generally, when C contains polynomials of degree at most d with nonnegative coefficients, then $\beta(C) = d/((d+1)^{1+1/d})$, which implies that the price of anarchy grows asymptotically as $\Theta(d/\ln d)$. For details on these calculations, the reader is referred to Roughgarden (2003) and Correa et al. (2004).

Let us also point out that Theorem 4.2 (b) dominates the following result.

Theorem 4.4. (See Roughgarden and Tardos, 2004, Theorem 5.1.) *Let x^{EQ} be an equilibrium of a nonatomic congestion game with nondecreasing, continuous, and separable cost functions. If y^{OPT} is a social optimum of the same game with twice as many players of each type, then $C(x^{\text{EQ}}) \leq C(y^{\text{OPT}})$.*

4.2. Cost functions with limited congestion effects

To motivate our next set of results, it is helpful to consider selfish routing in traffic networks. The empirically observed ratio of the total travel time of a Wardrop equilibrium to that of a system optimum is typically significantly smaller than predicted by the price-of-anarchy results of Theorem 4.2. For instance, in the computational studies of Jahn et al. (2005), the largest ratio of Wardrop-equilibrium cost to system-optimum cost over several realistic instances was 1.15 (instead of the theoretical worst case of 2.151 for polynomials of degree 4). Qiu et al. (2006) made similar observations in the context of telecommunication networks. Indeed, for a given class of cost functions, the corresponding price of anarchy is a worst-case measure, taken over all possible instances. However, if one compares the time needed to drive to work during rush hour with the duration of the same trip at night, then the *free-flow travel time* is usually not a negligible fraction of the rush-hour experience. We therefore consider the following scenario:

the cost of any given resource at utilization rate zero is at least a constant fraction of that of the same resource at the utilization rate in equilibrium.²

The ratio between the cost at equilibrium and that under no utilization is usually referred to as the *travel time index* in studies of transportation systems. It measures the congestion level in a given urban area. According to the Texas Transportation Institute (2005), the most congested city in the United States is Los Angeles, which has a travel time index of 1.75. The average index among very large urban areas in the United States is 1.48.

A different instance of this concept may be found in production processes, where the fixed cost of any one resource accounts for a non-negligible fraction of the total cost (fixed plus variable costs) in equilibrium.

The next result makes the technical assumption that the family \mathcal{C} of allowable cost functions is closed under the subtraction of costs at utilization rate zero; i.e., $c - c(0) \in \mathcal{C}$ for all $c \in \mathcal{C}$. Without this assumption, the improvement resulting from explicitly considering fixed costs would remain hidden in the value of $\beta(\mathcal{C})$.

Theorem 4.5. *Let \mathcal{C} be a family of cost functions that is closed under the subtraction of costs at utilization rate zero. Let x^{EQ} be an equilibrium of a nonatomic congestion game with cost functions drawn from \mathcal{C} that satisfy $c(0) \geq \eta c(x^{\text{EQ}})$, for some $0 \leq \eta < 1$.*

- (a) *If x^{OPT} is a social optimum for this game, then $C(x^{\text{EQ}}) \leq (1 - (1 - \eta)\beta(\mathcal{C}))^{-1} C(x^{\text{OPT}})$.*
- (b) *If y^{OPT} is a social optimum of the same game with $1 + (1 - \eta)\beta(\mathcal{C})$ as many players of each type, then $C(x^{\text{EQ}}) \leq C(y^{\text{OPT}})$.*

Proof. Consider a game with cost functions as specified in the hypothesis. Let us write $c(x) = M(x) + c(0)$. Note that $M \in \mathcal{C}$, and $\langle c(x^{\text{EQ}}), x^{\text{EQ}} \rangle \geq \langle M(x^{\text{EQ}}), x^{\text{EQ}} \rangle / (1 - \eta)$. Thus,

$$\begin{aligned} \beta(c, x^{\text{EQ}}) &= \max_{x \in \mathbb{R}_{\geq 0}^A} \frac{\langle c(x^{\text{EQ}}) - c(x), x \rangle}{\langle c(x^{\text{EQ}}), x^{\text{EQ}} \rangle} \\ &\leq (1 - \eta) \max_{x \in \mathbb{R}_{\geq 0}^A} \frac{\langle M(x^{\text{EQ}}) - M(x), x \rangle}{\langle M(x^{\text{EQ}}), x^{\text{EQ}} \rangle} \\ &\leq (1 - \eta)\beta(\mathcal{C}). \end{aligned}$$

We can therefore replace Lemma 4.1 with $\langle c(x^{\text{EQ}}), x \rangle \leq C(x) + (1 - \eta)\beta(\mathcal{C})C(x^{\text{EQ}})$. The rest of the proof is then identical to that of Theorem 4.2. \square

We illustrate this result by computing the new bounds for the class of separable polynomials that have nonnegative coefficients and maximum degree d . The interesting aspect of this particular case is that we can use again the picture introduced in Section 3. Indeed, because $c_a(0) \geq \eta c_a(x_a^{\text{EQ}})$, the area $(c_a(x_a^{\text{EQ}}) - c_a(x_a))x_a$ of the small shaded rectangle in Fig. 3 is at most $\beta(\mathcal{C})$ times that of the rectangle with upper-left corner point $(0, c_a(x_a^{\text{EQ}}))$ and lower-right corner point $(x_a^{\text{EQ}}, c_a(0))$, which itself is of size at most $(1 - \eta)c_a(x_a^{\text{EQ}})x_a^{\text{EQ}}$.

Fig. 4 (a) displays the relationship between η and the refined bound on the price of anarchy established by Theorem 4.5 for polynomials of various degrees. It is worth mentioning that the

² To be formally correct, this assumption requires that the equilibrium utilization rates of all resources are unique, as it is the case when c is strictly monotone (de Palma and Nesterov, 1998). To circumvent this problem, we formulate the following results in terms of a specific equilibrium.




Fig. 3. Illustration of the proof of Theorem 4.5 for separable cost functions.

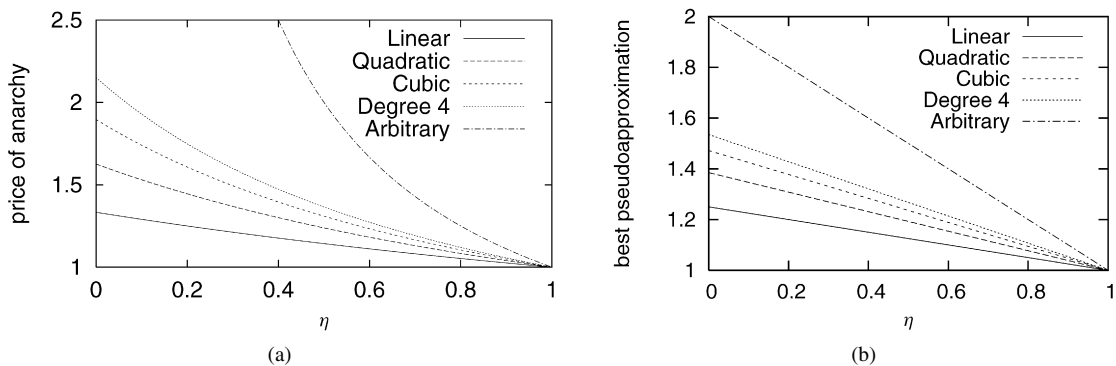


Fig. 4. (a) Price of anarchy as a function of fixed to total cost. (b) Minimum values for which the pseudo-approximation results hold. Each curve refers to a family of polynomials of a certain degree with nonnegative coefficients.

price of anarchy is at most $1/\eta$, even if we do not place any additional restriction on \mathcal{C} . This observation qualifies the discussion by Roughgarden and Tardos (2002, 2004) on the unboundness of the price of anarchy for instances with general cost functions.

As another example, consider a vehicular network in which users travel at most twice as long on each arc when the network is congested compared to the situation when it is not. Following the recommendation of the Bureau of Public Roads (1964), we let cost functions be polynomials of degree 4. Hence, Theorem 4.2 (a) gives an upper bound of 2.151 on the price of anarchy, if we do not take the congestion level into account. In contrast, Theorem 4.5 (a) gives a more accurate bound of 1.365. More generally, the bounds provided by Theorem 4.5 may offer a good explanation of the satisfactory performance of equilibria in many practical situations. Fig. 4 (b) illustrates Theorem 4.5 (b).

Let us also point out that the bounds in Theorem 4.5 (and, therefore, Theorem 4.2) are tight, for separable cost functions. Consider the traffic assignment instance in Fig. 5, where v units of flow must be routed from one node to the other over two parallel arcs a and a' . The cost functions are $c_a(x_a) = c(v)$ (a constant) and $c_{a'}(x_{a'}) = \eta c(v) + (1 - \eta)c(x_{a'})$, respectively. Here, the function $c : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ and the scalar v are chosen such that $\beta(\mathcal{C}) = \beta(c, v)$. Both bounds given in Theorem 4.5 are simultaneously tight for this instance.




Fig. 5. Instance for which the bounds provided by Theorem 4.5 are tight.

5. Discussion and concluding remarks

Most of the preceding results and proof techniques can almost effortlessly be carried forward to a variety of different settings. As an example to illustrate this, we sketch two extensions: a different type of social cost function, and general side constraints on the utilization rates. Let us also mention that the proof-by-picture idea put forward in Section 3 has recently been used by others to derive bounds on the price of anarchy in a variety of noncooperative games (e.g., Yang and Huang, 2005; Farzad et al., 2006; Acemoglu et al., 2007; Harks, 2007).

Another social cost function that is relevant in a variety of contextual areas such as evacuation planning, telecommunication networks, and supply chain management is the maximum cost of a used strategy (Roughgarden, 2004; Correa et al., 2007). For a given strategy distribution x , let us denote by $L(x)$ the maximum cost $c_S(x)$ over all strategies $S \in \mathcal{S}$ for which $x_S > 0$. The results of Correa et al. (2007) and Theorem 4.5 imply that the price of anarchy does not change.

Corollary 5.1. *Let x^{EQ} be an equilibrium of a nonatomic congestion game with a single player type and cost functions drawn from a class \mathcal{C} of nonseparable cost functions such that $c_a(0) \geq \eta c_a(x^{EQ})$, for some $0 \leq \eta < 1$.*

- (a) *If x is the vector of utilization rates of a feasible strategy distribution, then $L(x^{EQ}) \leq (1 - (1 - \eta)\beta(\mathcal{C}))^{-1}L(x)$.*
- (b) *If y is a feasible utilization rate vector of the same game with $1 + (1 - \eta)\beta(\mathcal{C})$ times as many players, then $L(x^{EQ}) \leq L(y)$.*

For the original social cost function, all results presented here can be extended to the setting with arbitrary side constraints in the space of resources. (The same is true if one were to consider side constraints in the space of strategies.) From a practical standpoint, capacity constraints are the simplest and perhaps most useful type of side constraints. In the context of traffic assignment, Hearn (1980), Larsson and Patriksson (1994, 1995, 1999), and Marcotte et al. (2004), among others, have advocated the explicit inclusion of resource capacities as an obvious way of improving the quality of models. An equilibrium with side constraints is an equilibrium in the same game without side constraints, but where players experience infinite disutilities when their actions would result in infeasible solutions. For formal definitions, details, and references on these equilibria, we refer the reader to Correa et al. (2004) and Stier-Moses (2004).

In the conference version of this article (Correa et al., 2005), we also considered instances for which the cost of each resource at utilization rate zero is equal to zero. This assumption helps

to capture situations in which variable costs dominate fixed costs, or where fixed costs can be neglected altogether. Subsequently, Dumrauf and Gairing (2006) computed the exact price of anarchy for polynomial cost functions of this kind. It turns out that when variable costs dominate, the price of anarchy is smaller than in the general case. Together with the results of Section 4.2, this suggests that the price of anarchy is largest in instances with intermediate levels of congestion: when congestion is low, players are aligned with the system, and when congestion is high, the system does poorly anyway because the problem is not the alignment of incentives, but the lack of capacity.

Acknowledgments

The authors are grateful to an anonymous referee for several suggestions that helped to improve the presentation of this paper. A preliminary version of this article appeared in the Proceedings of the 11th International Conference on Integer Programming and Combinatorial Optimization (Correa et al., 2005). The research of the first author was partially supported by the Chilean National Commission for Scientific and Technological Research (CONICYT), Anillo en Redes ACT08. The research of the second author was partially supported by NSF grant CMMI-0426686, and by ONR grant N00014-98-1-0317. The research of the last author was partially supported by the Center for International Business Education and Research (CIBER), and by the Center for Excellence in E-Business (CEBIZ), both at Columbia University.

References

- Acemoglu, D., Johari, R., Ozdaglar, A., 2007. Partially optimal routing. *IEEE J. Selected Areas Communications* 25, 1148–1160.
- Beckmann, M.J., McGuire, C.B., Winsten, C.B., 1956. *Studies in the Economics of Transportation*. Yale Univ. Press, New Haven, CT.
- Bureau of Public Roads, 1964. *Traffic Assignment Manual*. US Department of Commerce, Urban Planning Division, Washington, DC.
- Chakrabarty, D., 2004. Improved bicriteria results for the selfish routing problem. Unpublished manuscript.
- Chau, C.K., Sim, K.M., 2003. The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. *Operations Res. Lett.* 31, 327–334.
- Correa, J.R., Schulz, A.S., Stier-Moses, N.E., 2004. Selfish routing in capacitated networks. *Math. Operations Res.* 29, 961–976.
- Correa, J.R., Schulz, A.S., Stier-Moses, N.E., 2005. On the inefficiency of equilibria in congestion games. In: Jungert, M., Kaibel, V. (Eds.), *Proceedings of the 11th International Integer Programming and Combinatorial Optimization Conference (IPCO)*. Berlin, Germany. In: *Lecture Notes in Computer Science*, vol. 3509. Springer, Heidelberg, pp. 167–181.
- Correa, J.R., Schulz, A.S., Stier-Moses, N.E., 2007. Fast, fair, and efficient flows in networks. *Operations Research* 55, 215–225.
- de Palma, A., Nesterov, Y., 1998. Optimization formulations and static equilibrium in congested transportation networks. Discussion paper 9861. CORE, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.
- Dumrauf, D., Gairing, M., 2006. Price of anarchy for polynomial Wardrop games. In: Spirakis, P.G., Mavronicolas, M., Kontogiannis, S.C. (Eds.), *Second International Workshop on Internet and Network Economics (WINE)*. Patras, Greece. In: *Lecture Notes in Computer Science*, vol. 4286. Springer, Heidelberg, pp. 319–330.
- Farzad, B., Olver, N., Vetta, A., 2006. Flow-free selfish routing. Manuscript.
- Harks, T., 2007. On the price of anarchy for network games with nonatomic and atomic players. Manuscript.
- Hartman, G., Stampacchia, G., 1966. On some nonlinear elliptic differential equations. *Acta Math.* 115, 271–310.
- Hearn, D.W., 1980. Bounding flows in traffic assignment models. Technical report 80-4. Department of Industrial and Systems Engineering, University of Florida, Gainesville, FL.
- Jahn, O., Möhring, R.H., Schulz, A.S., Stier-Moses, N.E., 2005. System-optimal routing of traffic flows with user constraints in networks with congestion. *Oper. Res.* 53, 600–616.

- Koutsoupias, E., Papadimitriou, C.H., 1999. Worst-case equilibria. In: Meinel, C., Tison, S. (Eds.), Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS). Trier, Germany. In: Lecture Notes in Comput. Sci., vol. 1563. Springer, Heidelberg, pp. 404–413.
- Larsson, T., Patriksson, M., 1994. Equilibrium characterizations of solutions to side constrained asymmetric traffic assignment models. *Le Matematiche* 49, 249–280.
- Larsson, T., Patriksson, M., 1995. An augmented Lagrangean dual algorithm for link capacity side constrained traffic assignment problems. *Transportation Res.* 29B, 433–455.
- Larsson, T., Patriksson, M., 1999. Side constrained traffic equilibrium models—Analysis, computation and applications. *Transportation Res.* 33B, 233–264.
- Marcotte, P., Nguyen, S., Schoeb, A., 2004. A strategic flow model of traffic assignment in static capacitated networks. *Math. Operations Res.* 52, 191–212.
- Milchtaich, I., 2000. Generic uniqueness of equilibrium in large crowding games. *Math. Operations Res.* 25, 349–364.
- Milchtaich, I., 2004. Social optimality and cooperation in nonatomic congestion games. *J. Econ. Theory* 114, 56–87.
- Papadimitriou, C.H., 2001. Algorithms, games, and the Internet. In: Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC). Hersonissos, Greece. ACM Press, New York, NY, pp. 749–753.
- Perakis, G., 2007. The “price of anarchy” under nonlinear and asymmetric costs. *Math. Operations Res.* 32, 614–628.
- Qiu, L., Yang, Y.R., Zhang, Y., Shenker, S., 2006. On selfish routing in Internet-like environments. *Trans. Networking* 14, 725–738.
- Rosenthal, R.W., 1973. A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* 2, 65–67.
- Roughgarden, T., 2003. The price of anarchy is independent of the network topology. *J. Computer System Sci.* 67, 341–364.
- Roughgarden, T., 2004. The maximum latency of selfish routing. In: Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). New Orleans, LA. SIAM, Philadelphia, PA, pp. 973–974.
- Roughgarden, T., Tardos, É., 2002. How bad is selfish routing? *J. ACM* 49, 236–259.
- Roughgarden, T., Tardos, É., 2004. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games Econ. Behav.* 47, 389–403.
- Schmeidler, D., 1973. Equilibrium points of nonatomic games. *J. Stat. Phys.* 7, 295–300.
- Stier-Moses, N.E., 2004. Selfish versus coordinated routing in network games. PhD thesis. Operations Research Center, Massachusetts Institute of Technology, Cambridge, MA.
- Texas Transportation Institute, J., 2005. Urban Mobility Study. Available at <http://mobility.tamu.edu/ums>.
- Wardrop, J.G., 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers* II (1), 325–378.
- Yang, H., Huang, H.-J., 2005. *Mathematical and Economic Theory of Road Pricing*. Elsevier, Oxford, UK.