Efficiency and Stability of Nash Equilibria in Resource Allocation Games

Tobias Harks Konstantin Miller

***Abstract*— We study resource allocation games, where users send data along paths and links in the network charge a price equal to marginal cost. When users are price taking, it is known that there exist distributed dynamics that converge towards a fully efficient Nash equilibrium. When users are price anticipating, however, a Nash equilibrium does not maximize total utility in general. In this paper, we explore the inefficiency of Nash equilibria for general networks and semi-convex marginal cost functions. While it is known that for $m \geq 2$ users and convex marginal cost functions, no efficiency guarantee is possible, we prove that an additional differentiability assumption on marginal cost functions implies a bounded efficiency loss of $2/(2m+1)$. For polynomial marginal cost functions with nonnegative coefficients, we precisely characterize the price of anarchy. We also prove that the efficiency of Nash equilibria significantly improves if all users have the same strategy space and the same utility function.**

We propose a class of distributed dynamics and prove that whenever a game admits a potential function, these dynamics globally converge to a Nash equilibrium. Finally, we show that in general the *only* class of marginal cost functions that guarantees the existence of a potential function are affine linear functions.

I. Introduction

We study a resource allocation problem in communication networks, where users want to route flow from their source node to some sink node in the network and may strategically vary their flow rates. It is assumed that each user has a utility function mapping the total flow rate to a nonnegative value measuring the received utility. Each link maintains a flow dependent cost function modeling congestion on that link. An efficient allocation maximizes *total* utility defined as aggregate utility less aggregate cost. Kelly et al. [16] proposed a pricing mechanism (proportionally-fair pricing), in which every link charges a price per unit resource equal to the *marginal cost* on that link. Despite the simplicity and scalability of this mechanism, Kelly et al. showed that an optimal solution can be achieved as a competitive equilibrium if users are *price takers*, that is, they do not anticipate the consequence of price change in response to a change of their communication rates. Over the past years, the concept of congestion pricing , such as marginal cost pricing, has been successfully applied to understand and develop transport protocols, see for instance Kelly and Voice [15], Kunniyur and Srikant [18], or the book by Srikant [28].

Tobias Harks is with the Technische Universität Berlin, Institut für Mathematik, Straße des 17. Juni 136, 10623 Berlin, Germany. Email: harks@math.tu-berlin.de.

Konstantin Miller is with the Technische Universität Berlin, Telecommunication Networks Group, Einsteinufer 25, 10587 Berlin, Germany. Email: miller@tkn.tu-berlin.de.

Most of the remarkable results that have been obtained so far require that users behave as price takers. If users anticipate the price change caused by changing their communication rates, the structure and efficiency of Nash equilibria are substantially different. Even for simple single link networks, there are instances in which the efficiency loss of a Nash equilibrium is unbounded, see Johari [12] and Yang and Hajek [30].

Johari and Tsitsiklis [14] considered resource allocation games, where users choose rates along each available path in the network, and optimize their payoff only based on the aggregate price of each path. If users are price anticipating and prices are set equal to marginal costs, Johari and Tsitsiklis [13] showed that no non-trivial performance guarantee is possible provided at least two users participate in the game. It is worth noting that their worst case instance uses a non-differentiable marginal cost function. For the special case of linear marginal cost functions, Johari and Tsitsiklis [14] showed that the efficiency loss is bounded by $2/3$. Remarkably, this result holds for an arbitrary collection of concave utility functions (asymmetric games). For symmetric games (equal utility functions) with m users on a single link, Johari and Tsitsiklis [13] proved a bound of $2m/(2m+1)$ for arbitrary convex marginal cost functions.

Despite these positive results, several important questions are still open:

1) How efficient are Nash equilibria when realistic, that is, *differentiable and nonlinear* marginal cost functions are considered?
2) Can we design distributed dynamics that provably converge to a Nash equilibrium?

The first question is particularly relevant in practice, since for example link delays grow super-linearly with link flows in close-to-capacity regions, e.g., $M/M/1$ functions. In road networks, for instance, the most frequently used functions modeling delay are polynomials whose degrees and coefficients are determined from real-world data through statistical evaluation methods, see Patrikkson [22], Branston [3], or the Bureau of Public Roads [4]. The second question concerns the design of stable transport protocols that are capable of anticipating their influence of rate changes on prices.

A. Our Results

We study resource allocation games with price anticipating users and marginal cost pricing. Our contributions to the above two questions are the following (for an overview, see Table I):

TABLE I

KNOWN AND NEW LOWER BOUNDS ON THE WORST CASE EFFICIENCY OF NASH EQUILIBRIA DEPENDING ON THE CLASS $\mathcal{C}$ OF ALLOWABLE MARGINAL COST FUNCTIONS. THE CLASS $\mathcal{C}^0_{conv}$ DENOTES CONTINUOUS, NONDECREASING AND CONVEX FUNCTIONS. THE CLASS $\mathcal{C}^1_{conv} \subset \mathcal{C}^0_{conv}$ ADDITIONALLY REQUIRES DIFFERENTIABILITY. THE CLASS $\mathcal{C}^1_{conc}$ CONTAINS NONDECREASING, DIFFERENTIABLE, CONCAVE, AND SEMI-CONVEX FUNCTIONS, E.G., $\log(x+1)$. RESULTS THAT ONLY HOLD FOR COURNOT GAMES (GAMES ON A SINGLE LINK) ARE MARKED WITH $(^*)$. THE CLASS $\mathcal{C}_j, j \in \mathbb{N} \cup \{\infty\}$ DENOTES POLYNOMIAL FUNCTIONS WITH NONNEGATIVE COEFFICIENTS AND MAXIMUM DEGREE j.

	Asymmetric Games					Symmetric Games			
$\lvert\mathcal{S}\rvert$	$\mathcal{C}^0_{conv}$	$\mathcal{C}^1_{conv}$	$\mathcal{C}^1_{conc}$	$\mathcal{C}_1$	$\mathcal{C}_d$	$\mathcal{C}^0_{conv}$	$\mathcal{C}^1_{conv}$	$\mathcal{C}^1_{conc}$	$\mathcal{C}_\infty$
1	$2/3$ [13]	$2/3$ [13]	$1/2^*$	$3/4$ [5]	$3/4$	$2/3$ [13]	$2/3$	$1/2^*$	$3/4$
2	0 [13]	$2/5$	$1/2^*$	$8/11$ [5]	$\Omega(1/\sqrt{d})$	$4/5^*$ [13]	$4/5$	$2/3^*$	$8/9$
3	0 [13]	$2/7$	$1/2^*$	$5/7$ [5]	$\Omega(1/\sqrt{d})$	$6/7^*$ [13]	$6/7$	$3/4^*$	$15/16$
m	0 [13]	$\frac{2}{2m+1}$	$1/2^*$	$\frac{2m+4}{3m+5}$ [5]	$\Omega(1/\sqrt{d})$	$\frac{2m}{2m+1}^*$ [13]	$\frac{2m}{2m+1}$	$\frac{m}{m+1}^*$	$\frac{m(m+2)}{(m+1)^2}$
∞	0 [13]	0	$1/2^*$	$2/3$ [14]	$\Omega(1/\sqrt{d})$	1^* [13]	1	1^*	1

1) For asymmetric games on general networks (users have arbitrary differentiable, nondecreasing and concave utility functions) with $m \geq 2$ users, we prove a bound of $2/(2m+1)$ on the worst case efficiency for differentiable, nondecreasing and convex marginal cost functions. In particular, this bound carries over to practically relevant M/M/1 functions that model queuing delays with arc-capacities. Moreover, we characterize the price of anarchy for polynomial marginal cost functions with nonnegative coefficients (previous results, e.g. [14], only covered linear marginal costs).
2) For symmetric games (users have equal utility functions and equal strategy space), we present a series of results showing that the efficiency of Nash equilibria significantly improves. In particular, we prove that the worst case efficiency for polynomial marginal costs is exactly $3/4$.
3) For Cournot games on a single link with differentiable, nondecreasing, semi-convex and concave marginal cost functions, we prove that the price of anarchy is at most $1/2$. This result holds for an arbitrary number of users. If additionally users have the same utility function, we prove a bound of $m/(m+1)$.
4) We define a class $\mathcal{D}$ of distributed dynamics that can be implemented by end users. We show that this class contains, among others, the gradient method and certain replicator dynamics known from evolutionary game theory, see Wardrop [29] and Fischer et al. [7]. We prove that dynamics from $\mathcal{D}$ converge to a Nash equilibrium from any initial value if the game admits a potential function. We show that a potential function always exists if (i) marginal cost functions are linear, or (ii) all users have the same utility function and share a common set of paths. We also show that without restrictions on utility functions and the underlying network, the only marginal cost functions that guarantee the existence of a potential are affine linear functions.

B. *Significance and Techniques*

Our first results generalize the result of Johari and Tsitsiklis [14] for linear marginal cost functions. We prove a tight characterization of the price of anarchy for polynomial marginal cost functions. It is worth noting that our proof technique is significantly simpler than that of [14]. In [14] the authors explicitely identify the worst possible game by analytically solving a sequence of quadratic optimization problems. Hence, this approach becomes increasingly complicated if such optimization problems involve polynomial cost functions of higher degree. Our approach hinges on variational inequalities, which are used to relate the total utility of a Nash flow to that of an optimal flow. As a consequence, this technique can be applied to derive bounds on the price of anarchy for arbitrary subclasses of semi-convex marginal cost functions, see for instance our results for concave marginal cost functions. Additionally, our proof technique does not make use of the combinatorial structure of networks. In fact, most of our results for bounding the price of anarchy carry over to general congestion games with fractional assignments and elastic demands.

While our focus is on the marginal pricing scheme, Chen and Zhang [5] recently presented a class of pricing mechanisms satisfying certain axioms for which they proved improved bounds on the price of anarchy, if users are price anticipating. Their results hold for quadratic total cost, which correspond to linear marginal cost. For proving bounds on the price of anarchy for more general cost functions and general pricing schemes the technique presented in this paper can be applied to the setting of Chen and Zhang.

Finally, we study distributed dynamics that can be implemented by users. Using potential theory, we derive conditions under which these dynamics globally converge to a Nash equilibrium. As a byproduct of our analysis, we establish (under mild differentiability assumptions) a characterization of the existence of potentials. Since the initiating paper of Rosenthal [24] about congestion games and potential functions, a central topic of game theory is to determine classes of games that admit a potential. Thus, we believe that our result is of independent interest as it precisely describes, which classes of resource allocation games (depending on the class of marginal cost functions) admit a potential function.

C. Related Work

It is well known that Nash equilibria can be *inefficient* in the sense that they need not achieve socially desirable objectives. Koutsoupias and Papadimitriou [17] 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 social welfare of a system optimum and that of a Nash equilibrium.

Moulin [20] studied the price of anarchy for resource allocation games with three different pricing mechanisms that are based on cost sharing principles. The used social welfare function, however, differs from our setting and, thus, the derived bounds are not transferable. Chen and Zhang [5] defined certain axioms for a feasible pricing mechanism and derived for quadratic cost functions (which corresponds to linear marginal cost functions) a slightly better efficiency guarantee (0.686) than the bound $(2/3)$ proved for the marginal cost pricing, see Johari and Tsitsiklis [14].

Related to network resource allocation games are network routing games. In a seminal work, Roughgarden and Tardos [27] showed that the price of anarchy for network routing games with nonatomic players and linear latency functions is $4/3$. The case of more general families of latency functions can be found in the book by Roughgarden [26] and the survey by Altman et al. [2].

Even closer to the model considered in this paper are network routing games with a finite number of players who can split the flow along available paths, see Altman et al. [1], Haurie and Marcotte [10], Harks [8], Hayrapetyan et al. [11] and Cominetti et al. [6]. Haurie and Marcotte presented a general framework for studying atomic splittable network games with elastic demands. This class of games implicitly contains the resource allocation games with price anticipating users if latency functions are interpreted as marginal cost functions and the elastic demand functions model the equilibrium demand functions for the resource allocation game involving concave utility functions. Haurie and Marcotte, however, do not study the efficiency of Nash equilibria with respect to an optimal solution. Altman et al. [1] and Cominetti et al. [6] studied the atomic splittable selfish routing model. Altman et al. bounded the price of anarchy for monomial latency functions. They also derived conditions under which a Nash equilibrium is unique. Uniqueness of Nash equilibria has been further studied by Orda et al. [21]. The main difference between our model and that of Harks [8], Hayrapetyan et al. [11] and Cominetti et al. [6] is that our model involves elastic demands that are strategically set by players according to their utility functions.

II. The Model

In a network resource allocation game we are given a directed network $\mathcal{G} = (\mathcal{V}, \mathcal{A})$ and a set $\mathcal{S}$ of users with cardinality $|\mathcal{S}| = m$. Each user $s \in \mathcal{S}$ is characterized by an origin-destination pair of nodes and a collection of paths $\mathcal{P}_s$ available through the network. We will assume that sets of paths $\mathcal{P}_s$ are disjoint, so that for each path P there exists a unique user s such that $P \in \mathcal{P}_s$. This can be done without loss of generality because if two users share the same physical path, we capture it in our model by creating two paths which use the same subset of physical links. Let $\mathcal{P} = \cup_s \mathcal{P}_s$ be the set of all available paths of all users. A *flow* is then a function $x : \mathcal{P} \to \mathbb{R}^{|\mathcal{P}|}_{\geq 0}$. For every user $s \in \mathcal{S}$, the demand $d_s = \sum_{P \in \mathcal{P}_s} x_P$ specifies the amount of flow over the available paths. For a given flow x, we define the aggregate flow on an arc $a \in \mathcal{A}$ as $x_a = \sum_{P \in \mathcal{P} : a \in P} x_P$. The flow of user s is denoted by $x^s : \mathcal{P}_s \to \mathbb{R}^{|\mathcal{P}_s|}_{\geq 0}$ and the aggregate flow of user s on arc a is given as $x^s_a = \sum_{P \in \mathcal{P}_s : a \in P} x_P$. The benefit of sending at rate d_s is measured by the utility function $U_s(d_s)$.

Assumption 2.1: The utility functions $U_s(d_s), s \in \mathcal{S}$, are differentiable, strictly increasing, concave, and satisfy $U_s(0) \geq 0$ for all $s \in \mathcal{S}$.

Furthermore, each arc $a \in \mathcal{A}$ has an associated flow dependent cost function, which is specified as follows.

Assumption 2.2: The total cost function for a flow x is defined as $C(x) = \sum_{a \in \mathcal{A}} \int_0^{x_a} c_a(z)\, dz$, where every function $c_a(x_a)$, $a \in \mathcal{A}$, is a differentiable, semi-convex, nondecreasing function over $x_a \geq 0$, with $c_a(0) \geq 0$ and $\lim_{x_a \to \infty} c_a(x_a) = \infty$. Note that a function is semi-convex if $x_a\, c_a(x)$ is a convex function of x_a. Such functions are also called *standard* [25].

Let us briefly explain the differences of our assumption on marginal cost functions compared to Johari and Tsitsiklis [14]. In contrast to [14], we do not assume convexity of marginal cost functions. In fact, semi-convexity is already enough to ensure the existence of Nash equilibria. Semi-convex functions may also be concave, thus, are capable of modeling the effect of economies of scale. We do assume (in contrast to [14]) that marginal cost functions are differentiable. Even though this assumption is slightly more restrictive, we are rewarded with a series of non-trivial bounds on the price of anarchy, see Section III.

Remark 2.3: The total cost function $C(x)$ is thus a nondecreasing and convex function over $x \geq 0$. We call the functions $c_a(z)$, $a \in \mathcal{A}$, marginal cost functions or simply cost functions.

The *total utility* of a flow x satisfying demand $(d_s, s \in \mathcal{S})$ is defined as

$$\mathcal{U}(x) := \sum_{s \in \mathcal{S}} U_s(d_s) - C(x)\,. \tag{1}$$

A flow of maximum total utility is called *optimal*.

Suppose each user s chooses desired rates x_P for every path $P \in \mathcal{P}_s$. Given the resulting flow x, the links choose prices according to marginal cost $c_a(x_a)$, for every $a \in \mathcal{A}$.

If the flow of user s is x^s supporting the demand d_s, then this user receives a utility of $U_s(d_s)$ and pays $\sum_{P\in\mathcal{P}_s} x_P \cdot \big(\sum_{a\in P} c_a(x_a)\big)$. It is well known that if users are price taking, that is, they treat $c_a(x_a)$ as a constant, every Nash equilibrium of the marginal pricing mechanism maximizes total utility, see Kelly et al [16], Johari and Tsitsiklis [14]. When users anticipate the cost function $c_a(x_a)$ on the links $a \in \mathcal{A}$, the payoff for every user $s \in \mathcal{S}$ is given by

$$Q_s(x^s; x^{-s}) := U_s\,(d_s) - \sum_{P\in\mathcal{P}_s}\sum_{a\in P} c_a(x_a)\, x_P, \qquad (2)$$

where x^{-s} denotes the flow of all users except s. Unless stated otherwise, the strategy space of source s is $R_s = \mathbb{R}^{|\mathcal{P}_s|}_{\geq 0}$. We denote by $R = \underset{s\in\mathcal{S}}{\times} R_s$ the set of strategy profiles of all sources. Then, a resource allocation game I is completely described by the tuple $I = \big(\mathcal{G}, \mathcal{S}, R, U, c\big)$.

Remark 2.4: Using Assumptions 2.1 and 2.2, we obtain $\lim\limits_{|x^s|\to\infty} Q_s\,(x^s; x^{-s}) = -\infty$, hence, we can effectively restrict the strategy space for every user to a compact flow polytope. As the payoff functions are concave, a Nash equilibrium exists, see the result of Rosen [23].

III. Price of Anarchy

In the following, we will study the price of anarchy with respect to a class of marginal cost functions $\mathcal{C}$ satisfying Assumption 2.2. Throughout the paper we assume that utility functions for every instance satisfy Assumption 2.1.

Definition 3.1: Let $\mathcal{C}$ be a class of marginal cost functions. Let $\mathcal{I}\big(\mathcal{C}\big)$ be the set of all resource allocation games with marginal cost functions in $\mathcal{C}$. For a game $I \in \mathcal{I}\big(\mathcal{C}\big)$, let y_I be an optimal flow and let $\mathcal{X}_I$ be the set of pure Nash equilibria, respectively. Then, the *worst case efficiency* of Nash equilibria is defined by

$$\rho\big(\mathcal{C}\big) = \inf_{I\in\mathcal{I}(\mathcal{C})}\ \inf_{x_I\in\mathcal{X}_I} \tfrac{\mathcal{U}_I(x)}{\mathcal{U}_I(y)},$$

where $\mathcal{U}_I$ is the total utility function for instance I. Conversely, $\rho\big(\mathcal{C}\big)^{-1}$ defines the *price of anarchy*.

With this definition, the result of Johari and Tsitsiklis [14] reads as $\rho\big(\mathcal{C}_1\big) = 2/3$, where $\mathcal{C}_1 := \big\{c(z) = a_1\, z + a_0,\ a_0 \geq 0,\ a_1 \geq 0\big\}$.

In the following, we define the *user corrected marginal path costs* $\hat{c}_P(x)$ for every $s \in \mathcal{S}$ by

$$\hat{c}_P(x) = \sum_{a\in P} \big(c_a(x_a) + c'_a(x_a)\, x^s_a\big), \quad P \in \mathcal{P}_s\,.$$

Lemma 3.2: Consider a game I. Let x be a Nash flow and let y be an optimal flow with flow value d and $\bar{d}$, respectively. Then, the following conditions hold for all $s \in \mathcal{S}$:

$$\nabla Q_s(x^s; x^{-s})\,(g^s - x^s) \leq 0, \quad \forall g^s \geq 0\,, \qquad (3)$$

$$\begin{aligned} U'_s\,(d_s) &= \hat{c}_P(x), \quad \text{for } x_P > 0,\ P \in \mathcal{P}_s\,,\\ U'_s\,(d_s) &\leq \hat{c}_P(x), \quad \text{for } x_P = 0,\ P \in \mathcal{P}_s\,, \end{aligned} \qquad (4)$$

$$\begin{aligned} U'_s\,(\bar{d}_s) &= \sum_{a\in P} c_a(y_a),\ \text{for } y_P > 0,\ P \in \mathcal{P}_s\,,\\ U'_s\,(\bar{d}_s) &\leq \sum_{a\in P} c_a(y_a),\ \text{for } y_P = 0,\ P \in \mathcal{P}_s\,. \end{aligned} \qquad (5)$$

Proof: The function Q_s is differentiable and concave with respect to x^s. Furthermore, the set of flows is convex. Since x is a Nash equilibrium, the flow x^s is a maximizer of $Q_s(x^s; x^{-s})$. Thus, we can invoke the variational inequality as a necessary (and sufficient) optimality condition giving (3). Note that the derivative with respect to $x_P,\ P \in \mathcal{P}_s$, is given by $\frac{\partial Q_s}{\partial x_P}(x^s; x^{-s}) = U'_s\,(d_s) - \hat{c}_P(x)$. The second and third conditions follow directly from the Karush-Kuhn-Tucker conditions for maximizing $Q_s(x^s; x^{-s})$ and $\mathcal{U}(y)$, respectively. ∎

The next lemma, which can be found in Johari and Tsitsiklis [14], Moulin [20], and Chen and Zhang [5], shows that for bounding the price of anarchy it is sufficient to bound the price of anarchy for linear utility functions and single link networks. The idea for proving the lemma is to observe that every Nash flow x for a game I can be transformed to a Nash flow $\bar{x}$ for a modified instance $\bar{I}$, where linear utility functions of the form $\bar{U}_s(\bar{d}_s) = U'_s(d_s)\,\bar{d}_s$ are used, where d_s is the equilibrium demand of source s in the game I.

Lemma 3.3: [[5],[14],[20]] For bounding the price of anarchy, it is enough to consider instances in which utility functions are linear.

In the following, we represent the total utility of a Nash flow and that of an optimal flow in terms of the involved cost functions.

Lemma 3.4: Consider a game I in which utility functions are linear, that is, $U_s(d_s) = u_s\, d_s$, $u_s \geq 0$, $s \in \mathcal{S}$. Let y be an optimal flow and x be a Nash flow. Then, y and x generate total utility of $\mathcal{U}(y) = \sum_{a\in\mathcal{A}} \big(c_a(y_a)\, y_a - \int_0^{y_a} c_a(z)\, dz\big)$ and $\mathcal{U}(x) = \sum_{a\in\mathcal{A}} \big(c(x_a)\, x_a + \sum_{s\in\mathcal{S}} c'_a(x_a)\,(x^s_a)^2 - \int_0^{x_a} c_a(z)\, dz\big)$, respectively.

Proof: Using the optimality condition (5) in Lemma 3.2 we get $u_s = \sum_{a\in P} c_a(y_a)$ for all $s \in \mathcal{S}, P \in \mathcal{P}_s$ with $y_P > 0$. Thus, we obtain

$$\begin{aligned} \mathcal{U}(y) &= \sum_{s\in S} u_s \,\big(\sum_{P\in\mathcal{P}_s} y_P\big) - \sum_{a\in\mathcal{A}} \int_0^{y_a} c_a(z)\, dz\\ &= \sum_{a\in\mathcal{A}} \big(c_a(y_a)\, y_a - \int_0^{y_a} c_a(z)\, dz\big) \end{aligned}$$

proving the first part of the lemma. For proving the second equation, we use the optimality condition (4), which implies $u_s = \sum_{a\in P} c_a(x_a) + c'_a(x_a)\, x^s_a$, for all $s \in \mathcal{S}, P \in \mathcal{P}_s$ with $x_P > 0$. ∎

Lemma 3.5: Consider a game I with $|\mathcal{S}| = m$ in which utility functions are linear. Let $\lambda > 0$, x be a Nash flow, and y be an optimal flow. For every $a \in \mathcal{A}$, we define $\mu_a := \max_{j\in\mathcal{S}} \big\{\frac{x^j_a}{x_a}\big\} \in [\frac{1}{m}, 1]$, if $x_a > 0$ and $\mu_a := 0$, otherwise. Then, the following inequality is valid:

$$\begin{aligned} \mathcal{U}(y) \leq \lambda\,\mathcal{U}(x) + \sum_{a\in\mathcal{A}} \big[&c_a(x_a)\, y_a + c'(x_a)\,\mu_a\, x_a\, y_a\\ &- \lambda\,\tau_a(c_a, x_a, \mu_a) - \int_0^{y_a} c_a(z)\, dz\,\big] \end{aligned}$$

where $\tau_a(c_a, x_a, \mu_a) := c_a(x_a)\,x_a + \mu_a^2\,c_a'(x_a)x_a^2 - \int_0^{x_a} c_a(z)\,dz$.
The proof can be found in Section VIII.

Lemma 3.5 provides an inequality of the form $\mathcal{U}(y) \leq \lambda\,\mathcal{U}(x) + \sum_{a\in\mathcal{A}} \gamma_a(c_a, x_a, y_a, \mu_a)$. The main idea for proving bounds on the price of anarchy is to bound γ from above in terms of $\omega\,\mathcal{U}(y)$ for some $\omega < 1$. This would imply the inequality $\mathcal{U}(y) \leq \lambda\,\mathcal{U}(x) + \omega\,\mathcal{U}(y)$, which yields a bound on the worst case efficiency of $\frac{1-\omega}{\lambda}$. As a consequence, we will then optimize over λ so as to derive the best possible bound. This technique (λ-technique) has been previously applied to bounding the price of anarchy in congestion games, see Harks [8].

To this end, we define for a cost function c and parameters $\lambda > 0$ and $\mu \in \{0\} \cup [\frac{1}{m}, 1]$ the following value:

$$\omega_m(c;\lambda) := \sup_{\substack{\mu\in\{0\}\cup[\frac{1}{m},1]\\ (x,y)\in\mathbb{R}_+^2}} \frac{\gamma(c,x,y,\mu)}{c(y)\,y - \int_0^y c(z)\,dz}. \qquad (6)$$

For a given class of functions $\mathcal{C}$, we further define $\omega_m(\mathcal{C};\lambda) := \sup_{c\in\mathcal{C}} \omega_m(c;\lambda)$. Moreover, for a class $\mathcal{C}$ of marginal cost functions that satisfies Assumption 2.2, we define the feasible λ-region as $\Lambda_m(\mathcal{C}) := \{\lambda > 0 |\, \omega_m(\mathcal{C};\lambda) < 1\}$.

Theorem 3.6: Let $\mathcal{C}$ be a class of marginal cost functions. Consider the set $\mathcal{I}(\mathcal{C})$ of games with at most $m \in \mathbb{N}$ users. Then, the worst case efficiency is at least

$$\rho(\mathcal{C}) \geq \sup_{\lambda\in\Lambda_m(\mathcal{C})} \left[\frac{1-\omega_m(\mathcal{C};\lambda)}{\lambda}\right].$$

Proof: The proof follows directly from Lemma 3.3, Lemma 3.5 and the definition of $\omega_m(\mathcal{C};\lambda)$. ■

Notice that Theorem 3.6 can now be used to derive bounds on the price of anarchy for arbitrary classes of cost functions (satisfying Assumption 2.2). The challenge is to calculate the function $\omega_m(\mathcal{C};\lambda)$ for a given class $\mathcal{C}$.

A. Convex Marginal Cost Functions

We start with applying Theorem 3.6 for arbitrary convex marginal cost functions.

Theorem 3.7: Let $\mathcal{C}^{conv}$ be a class of convex marginal cost function. Consider the set $\mathcal{I}(\mathcal{C}^{conv})$ of games with at most $m \in \mathbb{N}$ users. Then, $\rho(\mathcal{C}^{conv}) \geq \frac{2}{2\,m+1}$.
The proof can be found in Section VIII.

The above result gives a finite price of anarchy for convex marginal cost functions provided a finite number users participate in the game. Compared to the negative result of Johari and Tsitsiklis [13], this result shows that a mild differentiability assumption on cost functions is enough to obtain a bounded efficiency loss.

B. Polynomial Marginal Cost Functions

In practice, the most frequently used functions modeling delay are polynomials whose degrees and coefficients are determined from real-world data through statistical evaluation methods, see Patrikkson [22] and Branston [3]. Thus, we will explicitly calculate the price of anarchy for the class $\mathcal{C}_d := \{c(z) = \sum_{j=1}^d a_j\,z^j,\ a_j \geq 0,\ j = 0,\ldots,d\}$ of polynomials with nonnegative coefficients. To simplify the analysis, we focus on the general case $|\mathcal{S}| \in \mathbb{N}\cup\{\infty\}$. Let us define $\omega_\infty(c;\lambda) := \lim_{m\to\infty} \omega_m(c;\lambda)$. Then, it is easy to see that $\omega_\infty(c;\lambda) \geq \omega_m(c;\lambda)$ for any $m \in \mathbb{N}$.

Remark 3.8: We observe that for polynomial marginal cost functions the total cost function $C(x)$ is linear in each of the marginal cost functions $c_a(\cdot)$. We can therefore reduce the analysis to monomial price functions. For this, we subdivide each arc a into d arcs $a_0,\ldots,a_d$ with monomial price functions $c_{a_s}(x) = c_s\,x^s$ for any $s = 0,\ldots,d$.

Lemma 3.9: Consider the class $\mathcal{M}_j := \{c(z) = a_j\,z^j,\ a_j \geq 0,\ j \in \mathbb{N}\}$. Then, $\omega_\infty(\mathcal{M}_j;\lambda)$ is at most $\Big[\big(\frac{1+\mu(j)\,j}{\lambda\,(1+\mu(j)^2\,j+\mu(j)^2)}\big)^j\,(\mu(j)\,j+1) - 1\Big]/j$, where $\mu(j) = \frac{1}{\sqrt{j}+1}$.

Proof: Using (6) for $c \in \mathcal{M}_j$, $\omega_\infty(c;\lambda)$ is at most $\sup_{\mu\in[0,1],\beta\geq 0} \big(\beta^j\,(1+j\,\mu) - \lambda\,(1+\mu^2\,j - \frac{1}{j+1})\,\beta^{j+1} - \frac{1}{j+1}\big)/\big(1-\frac{1}{j+1}\big)$, where $\beta := \frac{x}{y}$ (the case $y = 0$ can be excluded since then the expression becomes negative). The supremum with respect to β is a strictly convex program with the unique global maximizer $\beta^* = \frac{1+\mu\,j}{\lambda\,(1+\mu^2\,j+\mu^2)}$. Thus, since $c \in \mathcal{M}_j$ was arbitrary, $\omega_\infty(\mathcal{M}_j;\lambda)$ is bounded from above by

$$\sup_{\mu\in[0,1]} \Big[\big(\tfrac{1+\mu\,j}{\lambda\,(1+\mu^2\,j+\mu^2)}\big)^j\,(\mu\,j+1) - 1\Big]/j.$$

The unique maximizer for this supremum is given by $\mu(j) = \frac{1}{\sqrt{j}+1}$. ■

Theorem 3.10: Let $\mathcal{C}_d$ be the class of polynomial cost functions with nonnegative coefficients and maximum degree $d \in \mathbb{N}$. Then,

$$\rho(\mathcal{C}_d) \geq \Big[\big(1+\mu(d)\,d\big)^{1+\frac{1}{d}}\Big]\Big/\Big[1+\mu(d)^2\,d+\mu(d)^2\Big],$$

where $\mu(d) = \frac{1}{\sqrt{d}+1}$. Furthermore, this bound is tight.
The proof of Theorem 3.10 can be found in Section VIII.

Remark 3.11: The worst case efficiency for marginal cost functions in $\mathcal{C}_d$ is asymptotically bounded from below by $\Omega(1/\sqrt{d})$.

IV. Symmetric Games

In this section, we consider *symmetric games* in which all users have the same utility function $U(\cdot)$ and the same strategy space, that is, $\mathcal{P}_i = \mathcal{P}_j$ for all $i, j \in \mathcal{S}$. In this case, we get improved bounds on the price of anarchy.

Consider a symmetric game with $|\mathcal{S}| = m$ users. Then, there exists a symmetric optimal flow y such that $y^s = \frac{y}{m}$ for all $s \in \mathcal{S}$. Using an adapted version of Lemma 3.5, we get the following variational inequality relating any Nash equilibrium flow to a symmetric optimal flow.

$$\sum_{a\in\mathcal{A}} c_a(y_a)\,y_a \leq \sum_{a\in\mathcal{A}} \big[c_a(x_a)\,y_a + c'(x_a)\,x_a\,\frac{y_a}{m}\big]. \qquad (7)$$

Furthermore, Lemma 3.4 implies that $\mathcal{U}(x)$ is greater or equal than $\sum_{a\in\mathcal{A}} \big(c(x_a)\,x_a + \sum_{s\in\mathcal{S}} c_a'(x_a)\,\frac{x_a^2}{m} - \int_0^{x_a} c_a(z)\,dz\big)\,dz$.

In the following, we evaluate the efficiency of Nash equilibria for symmetric games and several classes of marginal cost functions using a similar technique as in the asymmetric case. For a cost function c and parameters $\lambda > 0$ and $m \in \mathbb{N}$ we define the value $\delta_m(c;\lambda)$ as in (6) except that

$$\gamma(c,x,y,\mu) = c(x)\,y + c'(x)\,\frac{x\,y}{m} - \int_0^y c(z)\,dz - \lambda\left(c(x)\,x + c'(x)\,\frac{x^2}{m} - \int_0^x c(z)\,dz\right).$$

For a given class of functions $\mathcal{C}$, we further define $\delta_m(\mathcal{C};\lambda) := \sup_{c\in\mathcal{C}} \delta_m(c;\lambda)$. Moreover, we define $\Delta_m(\mathcal{C}) := \{\lambda > 0 |\, \delta_m(\mathcal{C};\lambda) < 1\}$.

Theorem 4.1: Let $\mathcal{C}$ be a class of marginal cost functions. Consider a set $\mathcal{I}(\mathcal{C})$ of symmetric games with at most $m \in \mathbb{N}$ users. Then, the worst case efficiency is at least

$$\rho(\mathcal{C}) \geq \sup_{\lambda\in\Delta_m(\mathcal{C})} \left[\frac{1-\delta_m(\mathcal{C};\lambda)}{\lambda}\right].$$

Proof: The proof follows directly from Lemma 3.5, the representation of $\mathcal{U}(x)$ and the definition of $\delta_m(\mathcal{C};\lambda)$. ∎

A. Convex Marginal Cost Functions

The following result for convex marginal cost functions has been previously obtained by Johari and Tsitsiklis [13] for the special case of single link networks. We present here a more general result (arbitrary symmetric strategy space) with a simpler proof.

Proposition 4.2: Let $\mathcal{C}^{conv}$ be the class of convex marginal cost function. Consider the set $\mathcal{I}(\mathcal{C}^{conv})$ of games with symmetric utility functions and strategy space and at most $m \in \mathbb{N}$ users. Then, $\rho(\mathcal{C}^{conv}) \geq \frac{2\,m}{2\,m+1}$.

Proof: The proof proceeds along the lines of the proof of Theorem 3.7, except that $\lambda = \frac{1+2\,m}{2\,m}$ and the values μ and μ^2 are replaced by $\frac{1}{m}$. ∎

B. Polynomial Marginal Cost Functions

For polynomials with nonnegative coefficients and arbitrary degree $d \in \mathbb{N}\cup\{\infty\}$ we prove the following.

Theorem 4.3: Let $\mathcal{C}_\infty$ be the class of polynomial marginal cost function with nonnegative coefficients and arbitrary degree $d \in \mathbb{N}\cup\{\infty\}$. Consider the set $\mathcal{I}(\mathcal{C}_\infty)$ of games with symmetric utility functions and strategy space. Then, $\rho(\mathcal{C}_\infty) \geq \frac{3}{4}$. Furthermore, this bound is tight.
The proof proceeds along similar lines as in the asymmetric case and is omitted.

V. Concave Cournot Games

In the following, we analyze the efficiency loss of Nash equilibria for Cournot games. In Cournot games, it is assumed that the network consists of a single link only. We will establish bounds on the price of anarchy for Cournot games involving concave marginal cost functions. These functions are of particular interest as they model the effect of economy of scale.

Theorem 5.1: Let $\mathcal{C}^{conc}$ be a class of concave marginal cost functions. Consider the set $\mathcal{I}(\mathcal{C}^{conc})$ of resource allocation games on a single link. Then, $\rho(\mathcal{C}^{conc}) \geq \frac{1}{2}$.

Proof: Let x and y be Nash and optimal flows, respectively. Let $x := \beta\,y$. Then, on the one hand side, the variational inequality gives

$$c(y) \leq c(\beta\,y) + c'(\beta\,y)\,\beta\,\mu\,y.$$

On the other hand, since c is concave, we have

$$c(y) \leq c(\beta\,y) + c'(\beta\,y)\,(1-\beta)\,y.$$

Thus, it follows that $\beta \geq \frac{1}{1+\mu}$ must hold. Then, since $\mu \leq 1$ and $\mathcal{U}(z)$ is concave, we have $\beta \geq 1/2$ and $\mathcal{U}(\beta\,y) \geq \beta\,\mathcal{U}(y)$. ∎

Similar to the previous section, we will also provide a bound that holds for symmetric Cournot games with concave marginal cost functions.

Proposition 5.2: Let $\mathcal{C}^{conc}$ be the class of concave marginal cost functions. Consider the set $\mathcal{I}(\mathcal{C})$ of games with symmetric utility functions and at most $m \in \mathbb{N}$ users. Then, $\rho(\mathcal{C}^{conc}) \geq \frac{m}{m+1}$.
For the proof, we refer to the full version [9]. *Proof:* For concave marginal cost functions, we have the following relationship

$$c(x) + c'(x)\,(y-x) \geq c(y), \text{ for all } y \geq 0.$$

Let x and y be Nash and optimal flows, respectively. Let $x := \beta\,y$. Then, on the one hand side, the variational inequality gives

$$c(y)\,y \leq c(\beta\,y)\,y + c'(\beta\,y)\,\frac{\beta\,y^2}{m},$$

which implies $(y > 0)$

$$c(y) \leq c(\beta\,y) + c'(\beta\,y)\,\frac{\beta\,y}{m}.$$

On the other hand, we have

$$c(y) \leq c(\beta\,y) + c'(\beta\,y)\,(1-\beta)\,y.$$

Thus, it follows that $\beta \geq \frac{m}{m+1}$ must hold. Then, since $\mathcal{U}(z)$ is concave, we have $\mathcal{U}(\beta\,y) \geq \beta\,\mathcal{U}(y)$. ∎

VI. Stability of Distributed Dynamics and Potential Functions

In this section, we define a class of distributed dynamics $\mathcal{D}_I$ for the resource allocation game I and show that all dynamics from this class converge to a Nash equilibrium from any initial state, provided that I admits a potential function. One representative of this class is the well known gradient descent method. Another representative is a combination of the gradient method with *atomic splittable* routing principles, inspired by replicator dynamics known from evolutionary game theory, see Wardrop [29] and Fischer et al. [7]. For a detailed description of this dynamic, we refer to the full version [9].

As we link the stability of a class of dynamics with a potential function argument, we consequently study necessary and sufficient conditions for a game to possess a potential function. We show that a game I with affine linear marginal costs admits a potential function. We also show that affine linear cost functions are the *only* functions that always guarantee the existence of a potential. For nonlinear marginal costs, we show that if we restrict the set of joint strategies R to symmetric flows, then a potential function exists.

A. Stability of Distributed Dynamics

We will now define a class $\mathcal{D}_I$ of dynamics that are stable if game I admits a potential function.

Definition 6.1: Given a game I, we say that a dynamic described by a differential equation $\dot{x}_P = f_P^I(x),\ P \in \mathcal{P}$, belongs to the class $\mathcal{D}_I$ if

1) x being a Nash equilibrium of I implies $f_P^I(x) = 0,\ \forall P \in \mathcal{P}$,
2) $\sum_{P\in\mathcal{P}} \frac{\partial Q_s(x)}{\partial x_P} \cdot f_P^I(x) \geq 0$, for all flows $x \in R$ and $s \in \mathcal{S}$ such that $P \in \mathcal{P}_s$,
3) $\sum_{P\in\mathcal{P}} \frac{\partial Q_s(x)}{\partial x_P} \cdot f_P^I(x) = 0$ if and only if $x \in R$ is a Nash equilibrium.

Before we study stability of dynamics in $\mathcal{D}_I$, we define the notion of an (exact) potential function for the game I and present two necessary and sufficient conditions for a game to admit a potential. The following definition is due to Monderer and Shapley [19].

Definition 6.2 (Monderer and Shapley [19]): A function $\Phi : \mathbb{R}_{\geq 0}^{|\mathcal{P}|} \rightarrow \mathbb{R}$ is a potential function for the game I if and only if

$$\Phi(x) - \Phi(y^s, x^{-s}) = Q_s(x) - Q_s(y^s, x^{-s}),$$

forall $x, y \in \mathbb{R}_{\geq 0}^{|\mathcal{P}|},\ \forall s \in \mathcal{S}$.

In other words, a potential function for game I is a real-valued function on the strategy space, which exactly tracks the difference in the payoff that occurs if one player unilaterally deviates. Similar to [19], we observe that if the payoffs are continuously differentiable, we obtain a characterization of potential functions in terms of the first derivatives of the payoffs. Different to [19], where strategies are scalars, in our case the strategies are vectors from $\mathbb{R}_{\geq 0}^{|\mathcal{P}|}$.

Lemma 6.3: Given that payoff functions are continuously differentiable, a function $\Phi : \mathbb{R}_{\geq 0}^{|\mathcal{P}|} \rightarrow \mathbb{R}$ is a potential function for the game I if and only if

$$\frac{\partial \Phi}{\partial x_P} = \frac{\partial Q_s}{\partial x_P},\ \forall s \in \mathcal{S},\ \forall P \in \mathcal{P}_s .$$

Proof: Obvious. ■

Again, similar to [19], given that the payoffs are twice continuously differentiable, we obtain the following characterization of games admitting a potential function.

Lemma 6.4: A game I, where all payoffs are twice continuously differentiable, admits a potential function if and only if

$$\frac{\partial^2 Q_s}{\partial x_P \partial x_Q} = \frac{\partial^2 Q_t}{\partial x_P \partial x_Q},\ \forall s, t \in \mathcal{S},\ \forall P \in \mathcal{P}_s,\ Q \in \mathcal{P}_t.$$

Proof: Obvious. ■

Now, the following theorem establishes a stability result for dynamics in $\mathcal{D}_I$.

Theorem 6.5: If a game I admits a potential function Φ_I, all dynamics in $\mathcal{D}_I$ converge to a Nash equilibrium from any initial value $x \in R$.

Proof: We will show that Φ_I is a Lyapunov function for an arbitrary dynamic in $\mathcal{D}_I$ defined by $\dot{x}_P = f_P^I(x),\ P \in \mathcal{P}$. According to the definition of a Lyapunov function, we need to show that

$$\sum_{P\in\mathcal{P}} \frac{\partial \Phi_I(x)}{\partial x_P} \cdot f_P^I(x) \geq 0,\ \forall x \in R,$$

and that $\sum_{P\in\mathcal{P}} \frac{\partial \Phi_I(x)}{\partial x_P} \cdot f_P^I(x) = 0$ if and only if $x \in R$ is a Nash equilibrium. However, both conditions follow from the definition of $\mathcal{D}_I$ and Lemma 6.3. ■

Note that the first condition in Definition 6.1 is not required for the proof. Indeed, what we show in Theorem 6.5 is that any trajectory converges to the set of Nash equilibria. However, Theorem 6.5 does not exclude the case, where the trajectory continues to oscillate within this set. Condition 1 is a sufficient condition to preclude such oscillations. From Theorem 6.5 it immediately follows that the gradient method is asymptotically stable for all games I that admit a potential function.

Corollary 6.6: Let I admit a potential function. Then, the gradient method

$$\dot{x}_P = \kappa_s\big(d_s\big)\big[U_s'\big(d_s\big) - \hat{c}_P(x)\big]_{x_P}^+, \quad \forall P \in \mathcal{P}_s,\ s \in \mathcal{S}, \quad (8)$$

where $\kappa_s\big(d_s\big)$ is a parameter determining the step size along the gradient, and

$$[a]_b^+ = \begin{cases} 0 & \text{if } b = 0 \text{ and } a < 0, \\ a & \text{otherwise}, \end{cases}$$

converges to a Nash equilibrium of game I from any initial value $x \in R$.

Proof: All we need to show is that the gradient method is in $\mathcal{D}_I$. First, note that condition 1 in Definition 6.1 is implied by equation (4) in Lemma 3.2. Next, observe that $\sum_{P\in\mathcal{P}} \frac{\partial Q_s(x)}{\partial x_P} \cdot f_P^I(x)$ is equal to $\kappa_s(d_s) \cdot \sum_{P\in\mathcal{P}} \Big(\big[U_s'(d_s) - \hat{c}_P(x)\big]_{x_P}^+\Big)^2 \geq 0$, and that due to (4) in Lemma 3.2 equality holds if and only if x is a Nash equilibrium. ■

In the following, we propose a distributed dynamic, which is a combination of the gradient method with *atomic splittable* routing, inspired by replicator dynamics known from evolutionary game theory, see Wardrop [29] and Fischer et al. [7]. We show that this dynamic is in $\mathcal{D}_I$ so that it is asymptotically stable whenever I admits a potential function.

The interpretation of atomic splittable routing is that the flow of each source is seen as a population of agents each controlling a very small fraction of flow. An agent continuously samples an alternative path in the network and switches to the sampled path with a probability depending on the difference of costs of the own and the sampled path. We apply this model in the fluid limit so that these two logical steps result in a single dynamic expressed by a differential equation. Our setting generalizes the model of [7] as we allow players to strategically vary their demands. Further, instead of summarized marginal link costs, we let the agents minimize user corrected path costs in order to count for the ability of the sources to anticipate the costs, which makes the stability analysis more challenging.

The routing update consists of two *logical* steps: sampling and migration. During the sampling step, each agent using path $P \in \mathcal{P}_s$ samples a path $Q \in \mathcal{P}_s$ with probability σ_{PQ}. During the migration step, each agent switches to the sampled path Q with probability $\mu_{PQ}(\hat{c}_P - \hat{c}_Q)$ depending on the difference of path costs. An example for a sampling policy is uniform sampling with $\sigma_{PQ} = \frac{1}{|\mathcal{P}_s|}$ for $P, Q \in \mathcal{P}_s$, where each path is sampled with an equal probability. An example for a migration policy is linear migration policy $\mu_{PQ}(\hat{c}_P - \hat{c}_Q) = \max\{\frac{\hat{c}_P - \hat{c}_Q}{\hat{c}_{\max}}, 0\}$. We restrict the class of considered sampling policies and migration probabilities by the following Assumption.

Assumption 6.7:

1) Sampling policies are assumed to be strictly positive: $\sigma_{PQ} > 0, \; \forall P, Q \in \mathcal{P}_s, \; \forall s \in \mathcal{S}$.
2) Migration probability functions $\mu_{PQ}(\cdot), \quad P, Q \in \mathcal{P}_s, \; s \in \mathcal{S}$, are assumed to be continuous and strictly increasing with $\mu_{PQ}(0) = 0$. For brevity, we will write μ_{PQ} instead of $\mu_{PQ}(\hat{c}_P - \hat{c}_Q)$.

Let us denote by γ_P the fraction of demand of user s that is routed over the path $P \in \mathcal{P}_s$ so that $\sum_{P \in \mathcal{P}_s} \gamma_P = 1, \;\; \forall s \in \mathcal{S}$. Then, the sample and migration probabilities induce a migration rate $r_{PQ} = \gamma_P \cdot \sigma_{PQ} \cdot \mu_{PQ}$. The growth rate of the fraction of flow on path P is then $\dot{\gamma}_P = \sum_{Q \in \mathcal{P}_s} (r_{PQ} - r_{QP})$. Note that this dynamic is a pure rerouting, it does not change the total demand of user s since $\sum_{P \in \mathcal{P}_s} \dot{\gamma}_P = 0$. Combining this with a gradient method to update the demand of a source, we obtain the following dynamics:

$$\dot{d}_s = \kappa_s(d_s) \left[U_s'(d_s) - \hat{c}_s(x)\right]^+_{d_s}, \quad \forall s \in \mathcal{S}, \tag{9}$$

$$\dot{\gamma}_P = \sum_{Q \in \mathcal{P}_s} \left(r_{QP} - r_{PQ}\right), \; \forall P \in \mathcal{P}_s, \; s \in \mathcal{S}. \tag{10}$$

Here, $\hat{c}_s(x) = \sum_{P \in \mathcal{P}_s} \gamma_P \cdot \hat{c}_P(x)$ are the average user-corrected path costs of the user s. Note that this dynamic can also be expressed in terms of the path flows as follows.

$$\dot{x}_P = \dot{\gamma}_P \, d_s + \gamma_P \, \dot{d}_s \,, \quad \forall P \in \mathcal{P}_s, \; s \in \mathcal{S} \,,$$

where $\dot{\gamma}$ and $\dot{d}$ are as in (9), (10).

Similar to a feasible flow $x \geq 0$, we want to define a feasible tuple (γ, d).

Definition 6.8: A tuple (d, γ) is *feasible* if and only if $d, \gamma \geq 0$ and $\sum_{P \in \mathcal{P}_s} \gamma_P = 1, \; \forall s \in \mathcal{S}$.

Obviously, each feasible flow $x \geq 0$ corresponds to a unique feasible tuple (d, γ) and vice versa. Now we are ready to prove the following Theorem.

Theorem 6.9: For a game I, the dynamic defined by (9), (10) is contained in $\mathcal{D}_I$ and thus converges to a Nash equilibrium whenever I admits a potential.

The proof is contained in the full version [9].

B. Potential Functions

So far, we defined a class of dynamics that are stable whenever a game I admits a potential function. In the following, we study necessary and sufficient conditions for existence of such a function. The next theorem shows that without any restriction of the network topology and the class of utility functions, the *only* class of marginal cost functions that always admits a potential is the class of affine linear functions.

Theorem 6.10: Let $\mathcal{I}(\bar{\mathcal{C}})$ be the set of games with marginal cost functions in $\bar{\mathcal{C}}$ such that payoff functions are twice continuously differentiable. Then, the following statements are equivalent:

1) Every $I \in \mathcal{I}(\bar{\mathcal{C}})$ is a potential game
2) $\bar{\mathcal{C}}$ contains only affine linear functions on $\mathbb{R}_{\geq 0}$

Proof: Calculating the corresponding second derivatives, Lemma 6.4 implies that the game I possess a potential if and only if for all flows $x \in R$

$$\sum_{a \in P \cap Q} c_a''(x_a) \left(x_a^s - x_a^t\right) = 0 \tag{11}$$

for all $s, t \in \mathcal{S}$, $P \in \mathcal{P}_s$, $Q \in \mathcal{P}_t$. The direction 2. $\Rightarrow$ 1. is proved by simply verifying that affine linear marginal cost functions satisfy the above condition. For proving 1. $\Rightarrow$ 2., we will assume that marginal cost functions are not linear and then construct a game that violates condition (11). First, observe that c is affine linear on $\mathbb{R}_{\geq 0}$ if and only if $c''(z) = 0$ for all $z \in \mathbb{R}_{\geq 0}$. Assume by contradiction that $\bar{\mathcal{C}}$ contains a function c that is not affine linear on $\mathbb{R}_{\geq 0}$. Then, there exists a $z_0 \in \mathbb{R}_{\geq 0}$ such that $c''(z_0) \neq 0$. W.l.o.g., we assume $c''(z_0) > 0$. We further assume w.l.o.g. that $z_0 > 0$, since if $z_0 = 0$, by continuity of $c''(x)$, there exists $\tilde{z}_0 > 0$ with $c''(\tilde{z}_0) > 0$. Now consider a game I, where two sources share a single link. Let the flow of source 1 be $x^1 = 1/3 \, z_0$ and the flow of source 2, $x^2 = 2/3 \, z_0$. It is easy to see that for $x = z_0$ condition (11) fails as $c''(z_0) \left(x^2 - x^1\right) = c''(x) \, 1/3 \, z_0 > 0$. ■

A potential function for linear marginal cost functions is given in the full version [9].

Theorem 6.10 implies that if all marginal cost functions are affine linear then all dynamics in $\mathcal{D}_I$ converge to a Nash equilibrium from any initial state $x \in R$. Further, it implies that in order to find a potential function for games with nonlinear marginal costs, we have to restrict the topology, the utility functions, or the set of common strategies R. Condition (11) implies another sufficient condition for an instance I to possess a potential, which is based on the notion of a symmetric flow.

Definition 6.11: A flow $x \geq 0$ is called symmetric, if $x_a^s = x_a^t, \; \forall a \in \mathcal{P}_s \cap \mathcal{P}_t, \; \forall s, t, \in \mathcal{S}$.

Theorem 6.12: Let I be a game with twice continuously payoff functions. Let the set of common strategies R be restricted to symmetric flows. Then, I admits a potential function.

Proof: Follows from condition (11) in Lemma 6.4 and the definition of a symmetric flow. ■

The following corollary follows straight-forwardly.

Corollary 6.13: Let I be a game, where all sources share the same set of physical paths, have the same utility functions, and all marginal cost functions are twice continuously differentiable. Then, the gradient method and the combination of gradient method and atomic splittable routing converge to a Nash equilibrium from any symmetric initial state.

A potential function for this setting is given in the full version [9].

VII. Extensions and Future Work

In this work, we studied the efficiency and stability of Nash equilibria in resource allocation games with price anticipating users. We considered the marginal cost pricing mechanism and derived various results about the price of anarchy depending on the structure of allowable marginal cost functions. In particular, we were able to provide tight bounds for the price of anarchy for polynomial marginal link costs. As this class of functions is quite rich and widely used for modeling for instance queuing delays at links, we see our results as an important contribution towards the applicability of the marginal pricing mechanism in practice.

The second contribution of this paper concerns the design of a class of distributed dynamics that converge towards a Nash equilibrium. We identified conditions under which global stability of the proposed dynamics can be proved. An open issue is the stability of the proposed class of algorithms if communication delays are considered. While the stability of delayed differential equations for resource allocation games with price taking users has been well studied, this issue is largely open for games with price-anticipating users.

VIII. Proofs

Proof of Lemma 3.5

Proof: We first sum the variational inequality (3) in Lemma 3.2 with $g^s = y^s$ over all $s \in \mathcal{S}$.

$$\sum_{s\in\mathcal{S}} U_s'(d_s)\,(\bar{d}_s - d_s) - \sum_{a\in\mathcal{A}} \big[c_a(x_a)\,(y_a - x_a) + \sum_{s\in\mathcal{S}} c'(x_a)\, x_a^s\,(y_a^s - x_a^s)\big] \leq 0.$$

Using that utility functions are linear and rewriting yields

$$\sum_{s\in\mathcal{S}} u_s\,\bar{d}_s \leq \sum_{s\in\mathcal{S}} u_s\, d_s + \sum_{a\in\mathcal{A}} \big[c_a(x_a)\,(y_a - x_a) + \sum_{s\in\mathcal{S}} c_a'(x_a)\, x_a^s\,(y_a^s - x_a^s)\big].$$

Using Lemma 3.4 and the definition of μ_a we get

$$\sum_{a\in\mathcal{A}} c_a(y_a)\, y_a \leq \sum_{s\in\mathcal{S}} u_s\, d_s + \sum_{a\in\mathcal{A}} \big[c_a(x_a)\,(y_a - x_a) + c'(x_a)\,\mu_a\, x_a\, y_a - \sum_{s\in\mathcal{S}} c'(x_a)\,(x^s)^2\big] = \sum_{a\in\mathcal{A}} \big[c_a(x_a)\, y_a + c'(x_a)\,\mu_a\, x_a\, y_a\big].$$

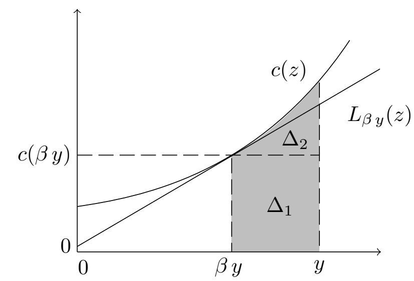

Fig. 1. The gray-shaded area illustrates the term $\int_0^y c(z)\,dz - \int_0^{\beta\, y} c(z)\,dz = \Delta_1 + \Delta_2$. The linear approximation $L_{\beta\, y}(\cdot)$ of the convex function $c(\cdot)$ bounds $c(z)$ from below, i.e., $L_{\beta\, y}(z) \leq c(z)$. Then, we have $\Delta_1 = (y - \beta\, y)\, c(\beta\, y)$ and $\Delta_2 \geq \frac{(y-\beta\, y)^2}{2}\, c'(\beta\, y)$.

Subtracting $\sum_{a\in\mathcal{A}} \int_0^{y_a} c_a(z)\,dz$ on both sides gives

$$\mathcal{U}(y) \leq \sum_{a\in\mathcal{A}} \big[c_a(x_a)\, y_a + c'(x_a)\,\mu_a\, x_a\, y_a - \int_0^{y_a} c_a(z)\,dz\big].$$

Now we add and subtract $\lambda\,\mathcal{U}(x)$ for $\lambda \geq 0$ on the right-hand side.

Finally, we observe that $\lambda\,\mathcal{U}(x) \geq \lambda\,\big(\sum_{a\in\mathcal{A}} \big[c_a(x_a)\, x_a + \mu_a^2\, c'(x_a) x_a^2 - \int_0^{x_a} c_a(z)\,dz\big]\big)$. Thus, the claim is proved. ■

Proof of Theorem 3.7

Proof: We define $\lambda = \frac{1}{2} + m$ and prove the claim by showing $\omega_m(c;\lambda) \leq 0$ for $c \in \mathcal{C}^{conv}$. We proceed by a case distinction. First, we assume $x \geq y$. Then, the nominator of (6) can be bounded from above as follows.

$$c(x)\, y + c'(x)\,\mu\, x\, y - \lambda\,\tau(c, x, \mu) - \int_0^y c(z)\,dz \leq c'(x)\,\big(\mu\, x\, y - \lambda\,\mu^2\, x^2\big) \leq c'(x)\, x^2\big(\mu - \lambda\,\mu^2\big).$$

For the first inequality, we used that

$$c(x)\, y - \lambda\, c(x)\, x + \lambda\, \int_0^x c(z)\,dz - \int_0^y c(z)\,dz \leq 0,$$

since $y \leq x$ and $\lambda \geq 1$. The second inequality follows from $y \leq x$ and $c'(x) \geq 0$. Then, $\lambda = \frac{1}{2} + m$ yields $\omega_m(c;\lambda) \leq 0$ as $\max_{\mu\in\{0\}\cup[\frac{1}{m},1]} \mu - (\frac{1}{2} + m)\,\mu^2 \leq 0$.

Now, we consider the case $x < y$. We define $\beta := \frac{x}{y} \in [0, 1)$. We now observe that

$$\int_0^y c(z)\,dz - \lambda\, \int_0^{\beta\, y} c(z)\,dz = \int_0^y c(z)\,dz - \int_0^{\beta\, y} c(z)\,dz - (\lambda - 1)\, \int_0^{\beta\, y} c(z)\,dz.$$

Then, we use the following inequality, which is illustrated in Fig. 1.

$$\int_0^y c(z)\,dz - \int_0^{\beta\, y} c(z)\,dz \geq (y - \beta\, y)\, c(\beta\, y) + \frac{(y - \beta\, y)^2}{2}\, c'(\beta\, y).$$

Together with

$$(\lambda - 1) \int_0^{\beta y} c(z)\, dz \le (\lambda - 1)\, c(\beta\, y) \beta\, y,$$

we obtain that $\omega_m(c; \lambda)$ is at most

$$\sup_{\substack{\beta \in [0,1) \\ \mu \in \{0\} \cup [\frac{1}{m}, 1] \\ y \in \mathbb{R}_+}} \frac{c'(\beta\, y)\, y^2 \big(\beta\, \mu - \lambda\, \mu^2\, \beta^2 - \frac{(1-\beta)^2}{2}\big)}{c(y)\, y - \int_0^y c(z)\, dz}.$$

We now arrive at

$$\max_{\beta \in [0,1)} \big(\beta\, \mu - \lambda\, \mu^2\, \beta^2 - \tfrac{(1-\beta)^2}{2}\big) \le \tfrac{\mu\,(2+\mu-2\,\lambda\,\mu)}{2\,(2\,\lambda\,\mu^2+1)},$$

where $\beta^* = \frac{\mu+1}{2\,\lambda\,\mu^2+1}$ is the unique maximizer. Thus, since $\lambda = \frac{1}{2} + m$ and using that $\mu \ge \frac{1}{m}$ we obtain $\omega_m(c; \lambda) \le 0$. Notice that also $\mu = 0$ implies $\omega_m(c; \lambda) \le 0$.

Applying Theorem 3.6 for both cases proves the claim. ■

Proof of Theorem 3.10

Proof: We define $\lambda = \Big[\big(1 + \mu(d)\, d\big)^{1+\frac{1}{d}}\Big] / \Big[1 + \mu(d)^2\, d + \mu(d)^2\Big]$. Then, invoking Lemma 3.9 implies $\omega_\infty(\mathcal{M}_j; \lambda) \le 0$ for all $j < d$ and $\omega_\infty(\mathcal{M}_j; \lambda) = 0$ for $j = d$. Thus, using Theorem 3.6, we have $\lambda\, \mathcal{U}(x) \ge \mathcal{U}(y)$.

Now we prove the upper bound. Consider a single link with marginal cost function $c(x) = x^d$ for some $d \in \mathbb{N}$. Assume we have N users, where user 1 has the utility function $U_1(x_1) = x_1$, while the remaining $N-1$ users have utility functions $U_k(x_k) = bx_k$ for some $b \in [0, 1]$ specified later. We denote the total rate on the link by $x(N)$. Then, the Nash condition for user 1 yields $1 = x(N)^d + d\, x(N)^{d-1}\, x_1$. Thus, we have $x_1(N) = \frac{1-x^d}{d\, x(N)^{d-1}}$. The conditions for users k, $k = 2, \ldots, N$ yield $x_k(N) = \frac{b-x(N)^d}{d\, x(N)^{d-1}}$. Summing all rates we get

$$\begin{aligned} x(N) &= \tfrac{1-x(N)^d}{d\, x(N)^{d-1}} + \big(N - 1\big)\, \tfrac{b-x(N)^d}{d\, x(N)^{d-1}} \\ \Leftrightarrow x(N) &= \big(\tfrac{1+(N-1)\, b}{d+N}\big)^{\frac{1}{d}}. \end{aligned}$$

We get $\lim_{N\to\infty} x(N) = b^{\frac{1}{d}}$, $\lim_{N\to\infty} x_1(N) = \frac{b^{\frac{1}{d}}\,(1-b)}{d\, b}$, $\lim_{N\to\infty} b\,(N-1)\, x_k(N) = \frac{b^{\frac{1}{d}}\,(b\, d-1+b)}{d}$. Thus, we get in the limit for the total utility of the Nash flow x

$$\lim_{N\to\infty} \mathcal{U}(x(N)) = \tfrac{b^{\frac{1}{d}}\,(1-b)}{d\, b} + \tfrac{b^{\frac{1}{d}}\,(b\, d-1+b)}{d} - \tfrac{b^{\frac{1}{d}}\, b}{d+1}.$$

The optimal solution is given by $y = (1, 0, \ldots, 0)$ with total utility of $\mathcal{U}(y) = 1 - \frac{1}{d+1}$. Now choosing $b = \frac{1+d^{\frac{3}{2}}}{d^2+d+1}$ one can check that the ratio $\frac{\mathcal{U}(x)}{\mathcal{U}(y)}$ coincides with the lower bound of the theorem. ■

References

[1] E. Altman, T. Basar, T. Jimnez, and N. Shimkin. Competitive routing in networks with polynomial costs. *IEEE Trans. Automatic Control*, 47:92–96, 2002.

[2] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter. A survey on networking games in telecommunications. *Comput. Oper. Res.*, 33(2):286–311, 2006.

[3] D. Branston. Link capacity functions. *Transportation Research*, 10:223–236, 1976.

[4] Bureau of Public Roads. Traffic assignment manual. *U.S. Department of Commerce, Urban Planning Division, Washington, DC*, 1964.

[5] Y.-J. Chen and J. Zhang. Design of price mechanisms for network resource allocation via price of anarchy. *Operations Research, forthcoming*, 2008.

[6] R. Cominetti, J. R. Correa, and N. E. Stier-Moses. Network games with atomic players. In *Proc. ICALP*, 2006.

[7] S. Fischer and B. Vöcking. Adaptive routing with stale information. In *Proc. ACM PODC*, 2005.

[8] T. Harks. Stackelberg strategies and collusion in network games with splittable flow. In *Proc. WAOA*, 2008.

[9] T. Harks and K. Miller. Efficiency and stability of nash equilibria in resource allocation games. Technical report, Berlin Institute of Technology, COGA Preprint 032-2008, 2008.

[10] A. Haurie and P. Marcotte. On the relationship between nash-cournot and wardrop equilibria. *Networks*, 15:295–308, 1985.

[11] A. Hayrapetyan, Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proc. ACM STOC*, 2006.

[12] R. Johari. *Efficiency loss in market mechanisms for resource allocation*. PhD thesis, Mass. Inst. Technol., 2004.

[13] R. Johari and J. N. Tsitsiklis. Efficiency loss in cournot games. Technical report, LIDS-P-2639, Laboratory for Information and Decision Systems, MIT,, 2005.

[14] R. Johari and J. N. Tsitsiklis. A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE Journal on Selected Areas in Communications*, 24(5):992–999, 2006.

[15] F. Kelly and T. Voice. Stability of end-to-end algorithms for joint routing and rate control. *ACM SIGCOMM Computer Communication Review*, 35:5–12, 2005.

[16] F. P. Kelly, A. K. Maulloo, and D. K. H. Tan. Rate Control in Communication Networks: Shadow Prices, Proportional Fairness, and Stability. *J. Oper. Res. Soc.*, 49:237–52, 1998.

[17] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proc. STACS*, 1999.

[18] S. Kunniyur and R. Srikant. Stable, scalable, fair congestion control and aqm schemes that achieve high utilization in the internet. *IEEE Trans. on Automatic Control*, 48:2024–2029, 2003.

[19] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[20] H. Moulin. The price of anarchy of serial, average and incremental cost sharing. *Economic Theory*, 36(3):379–405, 2008.

[21] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

[22] M. Patriksson. *The Traffic Assignment Problem—Models and Methods*. Topics in Transportation, VSP, 1994.

[23] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33:520–534, 1965.

[24] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. J. of Game Theory*, 2:65–67, 1973.

[25] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67:341–364, 2002.

[26] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. The MIT Press, 2005.

[27] T. Roughgarden and E. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.

[28] R. Srikant. *The Mathematics of Internet Congestion Control*. Birkhäuser Boston, 2003.

[29] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institute of Civil Engineers*, 1(Part II):325–378, 1952.

[30] S. Yang and B.Hajek. Revenue and stability of a mechanism for efficient allocation of a divisible good. Submitted, Oct. 2006.