

Cost-Balancing Tolls for Atomic Network Congestion Games

Dimitris Fotakis and Paul G. Spirakis

Abstract. We investigate the existence of optimal tolls for atomic symmetric network congestion games with unsplittable traffic and arbitrary nondecreasing latency functions. We focus on pure Nash equilibria, and consider a natural toll mechanism, which we call *cost-balancing tolls*. A set of cost-balancing tolls turns every path with positive traffic on its edges into a minimum-cost path. Hence any given configuration is induced as a pure Nash equilibrium of the modified game with the corresponding cost-balancing tolls. We show how to compute in linear time a set of cost-balancing tolls for the optimal solution such that the total amount of tolls paid by any player in any pure Nash equilibrium of the modified game does not exceed the latency on the maximum-latency path in the optimal solution. Our main result is that for congestion games on series-parallel networks with strictly increasing latency functions, the optimal solution is induced as the *unique* pure Nash equilibrium of the game with the corresponding cost-balancing tolls. To the best of our knowledge, only linear congestion games on parallel links were known to admit optimal tolls prior to this work. To demonstrate the difficulty of computing a better set of optimal tolls, we show that even for two-player linear congestion games on series-parallel networks, it is NP-hard to decide whether the optimal solution is the unique pure Nash equilibrium or there is another pure Nash equilibrium of total cost at least $6/5$ times the optimal cost.

1. Introduction

Congestion games provide a natural model for noncooperative resource allocation in large-scale communication networks and have been the subject of intensive research in algorithmic game theory. In an (atomic) *congestion game*

[Rosenthal 73], a finite set of noncooperative players, each controlling an un-splittable unit of traffic, compete over a finite set of resources. All players using a resource experience a latency given by a nonnegative and nondecreasing function of the resource’s traffic (or congestion). Among a given set of resource subsets (or strategies), each player selects one selfishly, trying to minimize her *individual cost*, that is, the sum of the latencies on the resources in the chosen strategy. A natural solution concept is that of a *pure Nash equilibrium*, a configuration in which no player can decrease her individual cost by unilaterally changing her strategy.

At the other end, the network manager seeks to minimize the *social cost* measured by the total cost incurred by all players. It is well known that a Nash equilibrium does not need to optimize the social cost. To mitigate the performance degradation due to the players’ noncooperative and selfish behavior, the network manager can introduce economic incentives that influence the players’ selfish choices and act to induce an optimal network configuration.

Economic incentives can be naturally modeled by nonnegative per-unit-of-traffic *tolls* (also called taxes or prices) assigned to the resources. The tolls are levied by the network manager and constitute an additional cost factor that the players should take into account. In the modified congestion game with tolls, a player’s cost for using a resource is equal to the latency due to the resource’s congestion plus the toll for using the resource. The player’s individual cost for adopting a strategy is equal to the sum of the latencies and the tolls for the resources in the chosen strategy. Although tolls increase the players’ individual costs, they do not affect the social cost, because they are payments inside the system and can be feasibly refunded to the players. The goal is to find a set of moderate and efficiently computable *optimal tolls* that make the Nash equilibria of the modified game coincide with the optimal solution.

1.1. Related Work

In the nonatomic setting, where there is an infinite number of players each controlling an infinitesimal amount of traffic, the existence and the efficiency of optimal tolls has been investigated extensively (see, e.g., [Cole et al. 03b, Cole et al. 03a] and references therein). A classical result is that the optimal solution is realized as the Nash equilibrium of a nonatomic congestion game with *marginal cost tolls* [Beckmann et al. 56]. In simple words, the performance degradation due to the selfish and noncooperative behavior of nonatomic players can be eliminated by an appropriate set of tolls. Unfortunately, marginal cost tolls fail to

induce the optimal solution even for simple congestion games with unsplittable traffic.¹

Recent work on tolls for nonatomic congestion games was motivated by the limitations of marginal cost tolls. The first to consider *heterogeneous* players who may have a different valuation of time (latency) in terms of money (toll) was [Cole et al. 03b], in which the authors established the existence of optimal tolls for nonatomic symmetric network congestion games. Their proof was based on Brouwer’s fixed-point theorem and was nonconstructive. Furthermore, Cole et al. showed how to compute a set of optimal tolls efficiently if the number of player types is finite and the latency functions are convex. [Fleischer 05] extended the results of [Cole et al. 03b] and proved that the optimal toll on each edge need not exceed the latency of the maximum latency path in the optimal solution times the maximum valuation of time. For series-parallel networks, Fleischer showed how to compute a set of optimal tolls efficiently even if there are infinitely many player types. Subsequently, [Fleischer et al. 04] and [Karakostas and Kolliopoulos 04] independently proved that the existence of optimal tolls for nonatomic congestion games with heterogeneous players and arbitrary strategies follows from linear programming duality. Therefore, optimal tolls can be computed efficiently by solving a linear program.

For nonatomic congestion games, the Nash equilibrium is essentially unique (under mild assumptions on the latency functions; see, e.g., [Roughgarden and Tardos 02]). Hence the tolls of [Beckmann et al. 56, Cole et al. 03b, Fleischer 05, Fleischer et al. 04, Karakostas and Kolliopoulos 04] induce the optimal solution as the unique equilibrium of the game with tolls.² On the other hand, atomic congestion games (even with splittable traffic) may admit many different Nash equilibria. Therefore, when considering atomic games, one has to distinguish between the case in which a set of tolls *weakly enforces* the optimal solution, in the sense that the optimal solution is realized as some (pure) Nash equilibrium of the game with tolls, and that in which a set of tolls *strongly enforces* the optimal solution, in the sense that the optimal solution is realized as the unique (pure) Nash equilibrium of the game with tolls.

¹Let $d_e(x)$ be the (differentiable) latency function of a resource e , let $d'_e(x)$ denote the first derivative of $d_e(x)$, and let o_e be the traffic of e in the optimal solution. Then the marginal cost toll of e is $o_e d'_e(o_e)$. For a congestion game with unsplittable traffic where marginal cost tolls fail to induce the optimal solution, consider two players and two parallel links with latency functions $d_1(x) = x/2$ and $d_2(x) = (1+\varepsilon)x$, for some $\varepsilon \in (0, 1/2)$. In the optimal configuration, there is one player on every link, while in the unique pure Nash equilibrium, both players choose the first link. The latter configuration remains the unique pure Nash equilibrium of the modified game with marginal cost tolls.

²The uniqueness of the Nash equilibrium in nonatomic games is also exploited by the algorithm of [Kaporis and Spirakis 06], which computes the smallest fraction of coordinated players required by a Stackelberg routing strategy to induce the optimal solution.

For atomic congestion games with splittable traffic and heterogeneous players, [Swamy 07] proved that a set of tolls that weakly enforce the optimal solution can be computed efficiently by solving a convex program. For homogeneous players with splittable traffic, [Cominetti et al. 06] presented a set of tolls that reduces the price of anarchy,³ though it is not known whether it weakly enforces the optimal solution.

To the best of our knowledge, the only work prior to ours that investigates the efficiency of tolls for atomic congestion games with unsplittable traffic is [Caragiannis et al. 06], which considered games with linear latency functions and homogeneous players, and investigated how much tolls can improve the price of anarchy. On the negative side, they presented a simple asymmetric game for which the price of anarchy remains at least 1.2 under any set of tolls. Therefore, asymmetric congestion games do not necessarily admit tolls that strongly enforce the optimal solution. On the positive side, Caragiannis et al. presented a set of tolls that strongly enforces the optimal solution for linear congestion games on parallel links. In addition, they presented two efficiently computable toll mechanisms that improve the price of anarchy for linear games with arbitrary strategies. The first mechanism [Caragiannis et al. 06, Theorem 3] is simple and improves the pure price of anarchy to 2.155 (from 2.5 [Awerbuch et al. 05, Christodoulou and Koutsoupias 05]). The second mechanism [Caragiannis et al. 06, Theorem 5] applies to the more general setting of mixed equilibria and weighted players, and improves the price of anarchy to 2 (from 2.618 [Awerbuch et al. 05, Christodoulou and Koutsoupias 05]). However, the former mechanism may not weakly enforce the optimal solution even for linear games on parallel links, while the latter mechanism may not strongly enforce the optimal solution even for linear games on series-parallel networks.⁴

1.2. Contribution

Despite the considerable interest in optimal toll mechanisms for atomic and nonatomic congestion games, it is still unknown whether there is an optimal toll mechanism for symmetric games with unsplittable traffic. This is true even

³The *price of anarchy* [Koutsoupias and Papadimitriou 99] is a widely accepted measure of the performance degradation due to the players’ noncooperative and selfish behavior. The (pure) price of anarchy is the worst-case ratio of the total cost of a (pure) Nash equilibrium to the optimal total cost. For a survey on the price of anarchy for congestion games, see, e.g., [Gairing et al. 05].

⁴For a game in which the tolls of [Caragiannis et al. 06, Theorem 3] do not weakly enforce the optimal solution, consider two players and two parallel links with latency functions $d_1(x) = x/2$ and $d_2(x) = x$. For a game on a series-parallel network in which the tolls of [Caragiannis et al. 06, Theorem 5] do not strongly enforce the optimal solution, consider the instance of Figure 1.b with $k = 2$, $q_1 = q_2 = 2$, and a (slightly different) latency function $d_{(s,t)}(x) = (3 + \varepsilon)x$, where $\varepsilon \in (0, 1)$, for the direct (s, t) edge.

for relatively simple symmetric network congestion games, such as games on series-parallel networks and games on parallel links with nonlinear latencies.

In this work, we investigate the existence of optimal tolls for symmetric network congestion games with unsplittable traffic, homogeneous players, and arbitrary nondecreasing latency functions. We focus on pure Nash equilibria, and consider a natural toll mechanism, which we call *cost-balancing tolls*. Cost-balancing tolls are motivated by the optimal tolls for nonatomic games [Cole et al. 03b, Fleischer 05, Fleischer et al. 04, Karakostas and Kolliopoulos 04]. A set of cost-balancing tolls turns every path with positive traffic on its edges into a minimum-cost path (the optimal tolls for linear congestion games on parallel links [Caragiannis et al. 06] are also based on the same principle). Hence any given configuration is induced as a pure Nash equilibrium of the game with the corresponding cost-balancing tolls. We show how to compute in linear time a set of cost-balancing tolls for the optimal configuration such that the total amount of tolls paid by any player in any pure Nash equilibrium of the modified game does not exceed the latency on the maximum latency path in the optimal configuration. Roughly speaking, we prove that the optimal solution is weakly enforceable by a set of moderate cost-balancing tolls computable in linear time. Moreover, we present a simple example whereby the optimal solution cannot be weakly enforced by tolls substantially smaller than the cost-balancing tolls.

Motivated by the recent interest in analyzing toll mechanisms (see, e.g., [Fleischer 05]) and Stackelberg routing strategies (see, e.g., [Correa and Stier-Moses 07, Swamy 07]) for games on series-parallel networks, we study the efficiency of cost-balancing tolls for such games. Our main result is that for congestion games on series-parallel networks with increasing latencies, the optimal solution is *strongly enforceable* by the corresponding cost-balancing tolls. Therefore, congestion games on series-parallel networks with increasing latencies admit a set of moderate optimal tolls computable in linear time. To the best of our knowledge, only linear congestion games on parallel links were known to admit optimal tolls [Caragiannis et al. 06, Theorem 1] prior to this work. Our result is considerably stronger, since it applies to arbitrary increasing latency functions and to series-parallel networks, which are significantly more complex than parallel-link networks. On the negative side, we show that if the network is not series-parallel, the cost-balancing tolls may not strongly enforce the optimal solution even for linear latency functions.

Given the existence of an efficiently computable optimal toll mechanism for congestion games on series-parallel networks, it is natural to ask for a set of optimal tolls that minimizes some objective function (e.g., sum of tolls, maximum toll) on the amount of tolls assigned to the edges of the network. To demonstrate the difficulty of computing a better set of optimal tolls, we prove that even for

two-player linear congestion games on series-parallel networks, it is NP-hard to distinguish between the case in which the optimal solution is the unique pure Nash equilibrium (and thus any set of tolls serves only to increase the players’ disutility) and that in which there is another pure Nash equilibrium of total cost at least $6/5$ times the optimal cost (and hence a set of tolls is required to strongly enforce the optimal solution).

1.3. Organization

We start by formalizing the model and presenting some useful definitions and notation in Section 2. In Section 3, we formally define cost-balancing tolls, present a simple linear-time algorithm for computing a set of moderate cost-balancing tolls, and show that the optimal solution is weakly enforceable by them (Theorem 3.4). Our main result is established in Section 3.3, where we prove that for congestion games on series-parallel networks with increasing latencies, the optimal solution is strongly enforceable by cost-balancing tolls (Theorem 3.7). We conclude by showing that even for simple games on series-parallel networks, it is NP-hard to decide whether the use of tolls is really necessary to strongly enforce the optimal solution (Theorem 4.1 in Section 4).

2. Model, Definitions, and Notation

For a vector $x = (x_1, \dots, x_n)$, we set $x_{-i} = (x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ and $(x_{-i}, x'_i) = (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$. For any integer $k \geq 2$, we set $[k] = \{1, \dots, k\}$.

2.1. Congestion Games

A *congestion game* is a tuple $\Gamma(N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E})$, where N denotes the set of players, E denotes the set of resources, $\Sigma_i \subseteq 2^E \setminus \{\emptyset\}$ denotes the strategy space of each player i , and $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ is a nondecreasing latency function associated with each resource $e \in E$. A congestion game is *symmetric* if all players share the same strategy space. A congestion game is *linear* if the latency function of each resource e is $d_e(x) = a_e x + b_e$, where $a_e, b_e \geq 0$.

A *configuration* is a vector $\sigma = (\sigma_1, \dots, \sigma_n)$ consisting of a strategy $\sigma_i \in \Sigma_i$ for each player $i \in N$. For every resource e , $\sigma_e = |\{i \in N : e \in \sigma_i\}|$ denotes the congestion induced on e by σ . The *individual cost* of each player i in the configuration σ is $c_i(\sigma) = \sum_{e \in \sigma_i} d_e(\sigma_e)$. A configuration σ is a *pure Nash equilibrium* if no player can improve her individual cost by unilaterally changing

her strategy. Formally, σ is a Nash equilibrium if for every player i and every strategy $s_i \in \Sigma_i$, $c_i(\sigma) \leq c_i(\sigma_{-i}, s_i)$.

We say that a congestion game Γ admits a *unique* pure Nash equilibrium if all pure Nash equilibria of Γ induce the same congestion on every resource. It is proved in [Rosenthal 73] that the pure Nash equilibria of a congestion game correspond to the local optima of a natural potential function. Therefore, every congestion game admits a pure Nash equilibrium, which is not necessarily unique.

In the following, we let n denote the number of players, and m the number of resources. We restrict our attention to *symmetric network* congestion games, in which the players' strategies are determined by a directed network $G(V, E)$ with a distinguished source s and destination t . The network edges play the role of the resources, and the common strategy space of all players is the set of simple s - t paths in G , denoted by \mathcal{P} .

2.2. Flows and Configurations

Let $G(V, E)$ be a directed network with source s and destination t . An s - t flow f is a vector $(f_e)_{e \in E} \in \mathbb{R}_{\geq 0}^m$ that satisfies the flow conservation at all vertices other than s and t . Formally, for all $u \in V \setminus \{s, t\}$, $\sum_{v:(v,u) \in E} f(v,u) = \sum_{v:(u,v) \in E} f(u,v)$. The *volume* of an s - t flow, denoted by $|f|$, is the total flow leaving s ; formally, $|f| = \sum_{v:(s,v) \in E} f(s,v)$. A flow is *acyclic* if there is no cycle in G with positive flow on all its edges. For a flow f and a path $p \in \mathcal{P}$, we let $f_p^{\min} = \min_{e \in p} \{f_e\}$.

Given a configuration σ for a symmetric network congestion game, we refer to the congestion vector $(\sigma_e)_{e \in E}$ as the *flow* induced by σ . We note that many different configurations may induce the same flow. We say that a flow σ is *feasible* if there is a configuration inducing congestion σ_e on every edge e . We slightly abuse notation by letting the same symbol denote both a configuration and the feasible flow induced by it.

2.3. Social Cost

We evaluate configurations and the corresponding feasible flows using the objective of *total cost*. The total cost $C(\sigma)$ of a configuration σ is the sum of players' costs in σ . Formally,

$$C(\sigma) = \sum_{i=1}^n c_i(\sigma) = \sum_{e \in E} \sigma_e d_e(\sigma_e).$$

The optimal configuration, denoted by o , minimizes the total cost $C(o)$ among all configurations in \mathcal{P}^n . In the following, we let o denote both the optimal configuration and the optimal flow induced by it.

Every s - t network with nondecreasing latency functions admits an integral acyclic min-cost flow of volume n computable in polynomial time if the $xd_e(x)$ are convex. Therefore, if the $xd_e(x)$ are convex, an optimal configuration for a symmetric network congestion game can be computed in polynomial time by a min-cost flow computation followed by a flow decomposition in n s - t paths. For series-parallel networks with nondecreasing latency functions, an optimal configuration can be computed in $\mathcal{O}(m + n \log m)$ time by the greedy algorithm [Bein et al. 85].

2.4. Tolls

We consider a scenario in which the network manager levies tolls on the edges of a network trying to influence the players' selfish choices and induce an optimal configuration. Let $\Gamma(N, E, (\Sigma_i)_{i \in N}, (d_e)_{e \in E})$ be the original congestion game. A set of *tolls* is a function $\tau : E \mapsto \mathbb{R}_{\geq 0}$ that assigns a nonnegative per-unit-of-traffic toll τ_e to every edge e . The modified congestion game induced by τ is $\bar{\Gamma}_\tau(N, E, (\Sigma_i)_{i \in N}, (\bar{d}_e)_{e \in E})$, where $\bar{d}_e(x) = d_e(x) + \tau_e$ for each edge e . In $\bar{\Gamma}_\tau$, the players have the same strategy space as in the original game. The cost of each player i in a configuration σ increases by the total amount of tolls on the edges in σ_i and becomes

$$\bar{c}_i(\sigma) = \sum_{e \in \sigma_i} \bar{d}_e(\sigma_e) = \sum_{e \in \sigma_i} (d_e(\sigma_e) + \tau_e) = c_i(\sigma) + \sum_{e \in \sigma_i} \tau_e.$$

The selfish players reach a pure Nash equilibrium of $\bar{\Gamma}_\tau$. Since the tolls are payments inside the network, a common assumption is that the tolls can be feasibly refunded to the players and thus do not affect the social cost (see, e.g., [Cole et al. 03b, Fleischer et al. 04, Karakostas and Kolliopoulos 04], and the case of *refundable* tolls in [Caragiannis et al. 06]). Hence the social cost of a configuration σ remains $C(\sigma) = \sum_{e \in E} \sigma_e d_e(\sigma_e)$ as in the original congestion game Γ .

The goal is to compute a set of tolls τ that motivates the selfish players to induce a given flow f (in particular, the optimal flow). A first natural requirement is that every configuration corresponding to f should be a pure Nash equilibrium of $\bar{\Gamma}_\tau$. Namely, if the players take the tolls into account and adopt any configuration inducing congestion f_e on every edge e , they should not have an incentive to deviate. Formally, a feasible flow f is *weakly enforceable* by tolls τ if every configuration inducing congestion f_e on every edge e is a Nash equilibrium of $\bar{\Gamma}_\tau$. By definition, if a feasible flow f is weakly enforceable, there is at least one pure Nash equilibrium of $\bar{\Gamma}_\tau$ with congestion f_e on all $e \in E$.

Since $\bar{\Gamma}_\tau$ does not need to admit a unique pure Nash equilibrium, some equilibria of $\bar{\Gamma}_\tau$ may induce flows quite different from f . To exclude this possibility,

we require that f not be only weakly enforceable by τ , but also the unique pure Nash equilibrium of $\bar{\Gamma}_\tau$. Formally, a feasible flow f is *strongly enforceable* by tolls τ when a configuration σ is a Nash equilibrium of $\bar{\Gamma}_\tau$ if and only if $\sigma_e = f_e$ for all $e \in E$. A set of tolls τ is *optimal* if the optimal flow is strongly enforceable by τ .

2.5. Series-Parallel Networks

A directed s - t network is *series-parallel* if it either consists of a single edge (s, t) or can be obtained from two series-parallel graphs with terminals (s_1, t_1) and (s_2, t_2) composed either in series or in parallel. In a *series composition*, t_1 is identified with s_2 , s_1 becomes s , and t_2 becomes t . In a *parallel composition*, s_1 is identified with s_2 and becomes s , and t_1 is identified with t_2 and becomes t . A maximal set of contiguous series compositions is a *series component*, and a maximal set of parallel compositions is a *parallel component*.

A series-parallel network can be completely specified by its *decomposition tree*, which is a rooted tree with a leaf for each edge. Each internal node of the decomposition tree represents either a series or a parallel component obtained from series (respectively parallel) compositions of the networks represented by its subtrees. The root of the tree represents the entire network. The decomposition tree of a series-parallel network $G(V, E)$ can be computed in $\mathcal{O}(|V| + |E|)$ time (see, e.g., [Valdez et al. 82] for more details).

3. Cost-Balancing Tolls

We consider a natural toll mechanism, which we call cost-balancing tolls. A set of tolls τ is *cost-balancing* for a feasible flow f if for every path $p \in \mathcal{P}$ with $f_p^{\min} > 0$ and every path $p' \in \mathcal{P}$,

$$\sum_{e \in p} \bar{d}_e(f_e) = \sum_{e \in p} (d_e(f_e) + \tau_e) \leq \sum_{e \in p'} (d_e(f_e) + \tau_e) = \sum_{e \in p'} \bar{d}_e(f_e). \quad (3.1)$$

The following proposition shows that the cost-balancing property also holds for segments of s - t paths with common endpoints.

Proposition 3.1. *Let f be any acyclic feasible s - t flow, let τ be any set of cost-balancing tolls for f , and let u, v be two vertices connected by simple paths π and π' . If $f_\pi^{\min} > 0$, then*

$$\sum_{e \in \pi} \bar{d}_e(f_e) = \sum_{e \in \pi} (d_e(f_e) + \tau_e) \leq \sum_{e \in \pi'} (d_e(f_e) + \tau_e) = \sum_{e \in \pi'} \bar{d}_e(f_e).$$

Proof. Since f is an acyclic s - t flow with $f_{\pi}^{\min} > 0$, there are a path π_1 with $f_{\pi_1}^{\min} > 0$ connecting s to u (π_1 is empty if u is s) and a path π_2 with $f_{\pi_2}^{\min} > 0$ connecting v to t (π_2 is empty if v is t). Then $p = (\pi_1, \pi, \pi_2)$ is an s - t path with $f_p^{\min} > 0$ and $p' = (\pi_1, \pi', \pi_2)$ is an s - t path. Therefore,

$$\sum_{e \in p} \bar{d}_e(f_e) \leq \sum_{e \in p'} \bar{d}_e(f_e) \Rightarrow \sum_{e \in \pi} \bar{d}_e(f_e) \leq \sum_{e \in \pi'} \bar{d}_e(f_e),$$

where the first inequality follows from (3.1), because the tolls τ are cost-balancing tolls for f . \square

We observe that every feasible flow is weakly enforceable by any set of cost-balancing tolls for it.

Proposition 3.2. *Let f be a feasible flow that admits a set of cost-balancing tolls τ . Then f is weakly enforceable by τ .*

Proof. We prove that any configuration σ with $\sigma_e = f_e$ on every edge e is a pure Nash equilibrium of the congestion game $\bar{\Gamma}_{\tau}$. In particular, let i be any player, and let σ_i be i 's strategy in σ . Then, for every path $p \in \mathcal{P}$, the individual cost of i (in $\bar{\Gamma}_{\tau}$) does not decrease if i switches from σ_i to p :

$$\bar{c}_i(\sigma) = \sum_{e \in \sigma_i} \bar{d}_e(f_e) \leq \sum_{e \in p} \bar{d}_e(f_e) \leq \bar{c}_i(\sigma_{-i}, p).$$

The equality holds because $\sigma_e = f_e$ for all $e \in E$. The first inequality holds because τ is a set of cost-balancing tolls for f and $f_{\sigma_i}^{\min} \geq 1$. The last inequality follows from the definition of the players' individual costs. \square

3.1. Efficient Computation of Cost-Balancing Tolls

We present a simple linear-time algorithm, called BALANCE, that computes a set of cost-balancing tolls for any acyclic s - t flow f . The input of BALANCE consists of an s - t network $G(V, E)$ and an acyclic s - t flow f . BALANCE returns a set of cost-balancing tolls for f . The algorithm works as follows:

1. Let $E_f = \{e \in E : f_e > 0\}$, and let $G_f(V, E_f)$ be the spanning subgraph of G consisting of the edges with positive flow in f . Since f is acyclic, G_f is a directed acyclic graph (DAG).
2. BALANCE computes the longest paths from s to the vertices in G_f with respect to the edge lengths $\{d_e(f_e)\}_{e \in E_f}$. Let $\ell_s = 0$, and for every vertex $u \in V \setminus \{s\}$ reachable from s , let ℓ_u be the length of the longest s - u path in G_f .

3. The toll of every edge $e = (u, v) \in E_f$ is $\tau_e = \ell_v - (\ell_u + d_e(f_e))$. The toll of every edge $e \notin E_f$ is $\tau_e = \tau^{\max}$, where $\tau^{\max} = \delta + \max_{p \in \mathcal{P}} \sum_{e \in p} d_e(n)$, with $\delta > 0$ chosen arbitrarily small.⁵

Since G_f is a DAG, the longest paths from s can be computed in $\mathcal{O}(|V| + |E_f|)$ time by negating the edge lengths and solving the corresponding single-source shortest path problem (see, e.g., [Cormen et al. 01, Section 24.2]). Hence BALANCE can be implemented in time linear in the size of the network. The following lemma shows that the tolls τ computed by BALANCE are cost-balancing for f . In the following, we refer to the tolls computed by BALANCE as the *cost-balancing tolls for f* .

Lemma 3.3. *For any acyclic s - t flow f , the tolls τ computed by BALANCE are cost-balancing for f .*

Proof. We first show that $\tau_e \geq 0$ for all $e \in E$. The claim is trivial for the edges not in E_f . For the edges in E_f , we recall that a longest path in a DAG becomes a shortest path if we negate the edge lengths. Therefore, for every vertex v reachable from s , $-\ell_v$ is the length of the shortest s - v path in $G_f(V, E_f, (-d_e(f_e))_{e \in E_f})$. Hence, for every edge $e = (u, v) \in E_f$, we have $-\ell_v \leq -\ell_u - d_e(f_e)$, which implies that $\tau_e = \ell_v - (\ell_u + d_e(f_e)) \geq 0$. By the same reasoning, if an edge $e \in E_f$ is included in some longest path from s , then $\tau_e = 0$.

We have also to show that for every s - t path p with $f_p^{\min} > 0$, and every s - t path p' ,

$$\sum_{e \in p} (d_e(f_e) + \tau_e) \leq \sum_{e \in p'} (d_e(f_e) + \tau_e).$$

Let $p = (s = u_0, u_1, \dots, u_k = t)$ be an s - t path with positive flow on all edges. Since $\tau_e = \ell_{u_{i+1}} - (\ell_{u_i} + d_e(f_e))$ for every edge $e = (u_i, u_{i+1}) \in p$, it follows that

$$\sum_{e \in p} (d_e(f_e) + \tau_e) = \sum_{i=0}^{k-1} (\ell_{u_{i+1}} - \ell_{u_i}) = \ell_t - \ell_s = \ell_t.$$

Consequently, the total cost on every s - t path p with $f_p^{\min} > 0$ is precisely ℓ_t . On the other hand, the cost on every s - t path p' containing an edge e not used by f is at least $\tau^{\max} \geq \ell_t$ (recall that $\tau^{\max} > \max_{p \in \mathcal{P}} \sum_{e \in p} d_e(n)$ and that $\ell_t = \max_{p: f_p^{\min} > 0} \sum_{e \in p} d_e(f_e)$). \square

⁵In fact, any $\tau^{\max} \geq \ell_t$ suffices for the tolls τ to be cost-balancing for f (see also the proof of Lemma 3.3). We use a toll of $\tau^{\max} > \max_{p \in \mathcal{P}} \sum_{e \in p} d_e(n)$ to ensure that any edge with $f_e = 0$ (and thus with a toll greater than ℓ_t) remains unused in any pure Nash equilibrium of Γ_τ (see also Theorem 3.4).

3.2. Weakly Enforcing the Optimal Flow

Proposition 3.2 and Lemma 3.3 imply that the optimal flow is weakly enforceable by a set of moderate cost-balancing tolls computable in linear time.

Theorem 3.4. *For every symmetric network congestion game Γ , the optimal flow o is weakly enforceable by the cost-balancing tolls τ for o , which have the following properties:*

- (a) *Given the optimal flow o , τ is computed in time linear in the size of the network.*
- (b) *The maximum toll on any edge is at most $\tau^{\max} = \delta + \max_{p \in \mathcal{P}} \sum_{e \in p} d_e(n)$, for any $\delta > 0$. Every edge with toll τ^{\max} remains unused in any pure Nash equilibrium of $\bar{\Gamma}_\tau$.*
- (c) *The total amount of tolls paid by any player in any pure Nash equilibrium of $\bar{\Gamma}_\tau$ does not exceed $\ell_t = \max_{p: o_p^{\min} > 0} \sum_{e \in p} d_e(o_e)$.*

Proof. (a). Every symmetric network congestion game with nondecreasing latency functions admits an optimal configuration that corresponds to a feasible integral acyclic flow o . The linear-time algorithm BALANCE computes the cost-balancing tolls τ for o (Lemma 3.3) and o is weakly enforceable by τ (Proposition 3.2).

(b). Let $G(V, E)$ be the s - t network determining the strategy space of Γ . Since o is an acyclic s - t flow, $\ell_v \leq \ell_t$ for all vertices v reachable from s in G_o . Therefore, a toll no greater than ℓ_t is assigned to any edge with positive flow in o , and a toll equal to τ^{\max} is assigned to the remaining edges.

We then prove that no edge with toll τ^{\max} is used in any pure Nash equilibrium of $\bar{\Gamma}_\tau$. To reach a contradiction, we assume that there is a pure Nash equilibrium σ of $\bar{\Gamma}_\tau$ in which some player i uses an edge e with $o_e = 0$. Then $\bar{c}_i(\sigma) \geq \tau^{\max}$. Let p_t be the longest s - t path in G_o . In the proof of Lemma 3.3, we show that for every edge e lying on some longest path in G_o , $\tau_e = 0$. Therefore, if player i switches to p_t , her cost becomes at most $\max_{p: o_p^{\min} > 0} \sum_{e \in p} d_e(n) < \tau^{\max} \leq \bar{c}_i(\sigma)$, a contradiction.

(c). By (b), every Nash equilibrium σ of $\bar{\Gamma}_\tau$ induces congestion $\sigma_e = 0$ on every edge e with $o_e = 0$. Hence the strategy σ_i of any player i consists of edges with positive flow in o . In the proof of Lemma 3.3, we show that for every s - t path p with $o_p^{\min} > 0$, $\sum_{e \in p} (d_e(o_e) + \tau_e) = \ell_t$. Consequently, the total amount of tolls paid by any player i in any pure Nash equilibrium of $\bar{\Gamma}_\tau$ is at most $\ell_t = \max_{p: o_p^{\min} > 0} \sum_{e \in p} d_e(o_e)$. \square

Remark 3.5. There are symmetric network games for which the optimal solution cannot be weakly enforced by tolls substantially smaller than the cost-balancing tolls. For instance, let us consider a discrete version of Pigou’s example with an even number n of players and two parallel links with latencies $d_1(x) = x/n$ and $d_2(x) = 1$. In the optimal solution, there are $n/2$ players on every link. For this example, the cost-balancing tolls coincide with the marginal cost tolls and are $\tau_1 = 1/2$ and $\tau_2 = 0$. The optimal solution is weakly (but not strongly) enforced by τ . On the other hand, for any set of tolls τ' , either $\frac{1}{2} - \frac{1}{n} \leq \tau'_1 - \tau'_2 \leq \frac{1}{2}$, or the optimal solution is not weakly enforced by τ' .

3.3. Optimal Tolls for Series-Parallel Networks

If the latency functions are increasing, it is not difficult to establish the optimality of cost-balancing tolls for congestion games on *extension-parallel* networks, an interesting class of s - t networks including parallel links (see, e.g., [Milchtaich 06]). Intuitively, this result is based on the fact that (symmetric) congestion games on extension-parallel networks retain two important properties of congestion games on parallel links, namely that every feasible acyclic s - t flow corresponds to a unique configuration of the congestion game (uniqueness is up to players’ permutations), and that every pure Nash equilibrium minimizes Rosenthal’s potential function (see, e.g., [Holzman and Law-Yone 97, Section 6]).

In the following, we significantly generalize this result to series-parallel networks, an important class of s - t networks considerably more complex than extension-parallel networks. This generalization is particularly interesting because series-parallel networks do not have the above properties of parallel-link (and extension-parallel) networks. More precisely, in a series-parallel network, (i) a feasible acyclic s - t flow may be realized by many different configurations of the corresponding congestion game, and (ii) different pure Nash equilibria may have quite different potential values (see, e.g., the game in the proof of Theorem 4.1). Nevertheless, we are able to show that for any series-parallel network with increasing latencies and any feasible acyclic flow f , the game induced by the cost-balancing tolls τ for f admits a *unique* pure Nash equilibrium (which coincides with f , since f is weakly enforceable by τ).

Lemma 3.6. *Let Γ be a symmetric congestion game on a series-parallel network $G(V, E)$ with increasing latency functions, let f be any feasible acyclic flow, and let τ be any set of cost-balancing tolls for f . Then every pure Nash equilibrium σ of the congestion game $\bar{\Gamma}_\tau$ induces congestion $\sigma_e = f_e$ on every edge e .*

Proof. To reach a contradiction, we assume that there is a pure Nash equilibrium σ of $\bar{\Gamma}_\tau$ that induces congestion $\sigma_e \neq f_e$ on some edge e . The corresponding flow σ is acyclic because the network is series-parallel, the latency functions are nonnegative and increasing, and σ is a Nash equilibrium of $\bar{\Gamma}_\tau$. Hence both f and σ are integral acyclic s - t flows of volume n .

We refer to an edge e as a *forward* edge if $\sigma_e > f_e$, and as a *backward* edge if $\sigma_e < f_e$. Since the σ_e 's and f_e 's are integral, $\sigma_e \geq f_e + 1$ for all forward edges e , and $f_{e'} \geq \sigma_{e'} + 1$ for all backward edges e' . Since σ and f are different acyclic flows of the same volume, there is at least one forward edge and at least one backward edge.

Let H be a component in the series-parallel decomposition of G (H is the induced subgraph of G determined by the vertices in the corresponding component) such that H contains both forward and backward edges and none of H 's subcomponents H_1, \dots, H_k has this property. In other words, H corresponds to the entire subtree of an internal node u_H in G 's decomposition tree, and H_1, \dots, H_k correspond to the entire subtrees of u_H 's children (an H_i may consist of a single edge and correspond to a leaf of the decomposition tree). The component H is *minimal* in the series-parallel decomposition of G with respect to the property that it contains both forward and backward edges. Such a component exists and can be found by traversing G 's decomposition tree bottom-up, because every edge e with $\sigma_e \neq f_e$ is either a forward or a backward edge but not both.

Let H_1 be a component of H that contains at least one forward edge and no backward edges, and let H_2 be a component of H that contains at least one backward edge and no forward edges (their existence is guaranteed by the definition of H). Since σ is acyclic and H_1 does not contain any backward edges, by flow conservation, the number of players going through H_1 in σ is greater than the number of players going through H_1 in f (see also Proposition 5.1). Similarly, since f is acyclic and H_2 does not contain any forward edges, the number of players going through H_2 in σ is less than the number of players going through H_2 in f (see also Proposition 5.2). Hence, by flow conservation, H_1 and H_2 are not connected in series. Therefore, H is formed by a parallel composition of H_1, \dots, H_k . Let s_H and t_H be the common endpoints of H and its components H_1, \dots, H_k .

Let e^+ be a forward edge in H_1 , let i be any player with $e^+ \in \sigma_i$ (such a player exists because e^+ is a forward edge, and thus $\sigma_{e^+} \geq 1$), and let p^+ denote the restriction of σ_i to H_1 (i.e., p^+ is the path segment of σ_i between s_H and t_H). Since H_1 does not contain any backward edges, $\sigma_e \geq f_e$ and $\bar{d}_e(\sigma_e) \geq \bar{d}_e(f_e)$ for all $e \in p^+$. Moreover, for the forward edge e^+ , $\bar{d}_{e^+}(\sigma_e) > \bar{d}_{e^+}(f_e)$ because the latency functions $d_e(x)$ (and thus the functions $\bar{d}_e(x)$) are increasing. Therefore,

the individual cost of player i on p^+ is

$$\sum_{e \in p^+} \bar{d}_e(\sigma_e) > \sum_{e \in p^+} \bar{d}_e(f_e). \quad (3.2)$$

On the other hand, since f is acyclic and H_2 does not contain any forward edges, H_2 contains an s_H - t_H path p^- consisting entirely of backward edges (see also Proposition 5.2). Therefore $\sigma_e + 1 \leq f_e$ for all $e \in p^-$, which implies that

$$\sum_{e \in p^-} \bar{d}_e(\sigma_e + 1) \leq \sum_{e \in p^-} \bar{d}_e(f_e) \leq \sum_{e \in p^+} \bar{d}_e(f_e). \quad (3.3)$$

For the last inequality, we recall that tolls τ are cost-balancing for f , and apply Proposition 3.1 to p^- and p^+ , since $f_{p^-}^{\min} \geq 1$ and p^- and p^+ are path segments with common endpoints s_H and t_H .

Combining (3.2) and (3.3), we conclude that $\sum_{e \in p^+} \bar{d}_e(\sigma_e) > \sum_{e \in p^-} \bar{d}_e(\sigma_e + 1)$. Therefore player i can decrease her cost by changing her path between s_H and t_H from p^+ to p^- . This contradicts the hypothesis that σ is a pure Nash equilibrium of $\bar{\Gamma}_\tau$. \square

Theorem 3.4 and Lemma 3.6 immediately imply that for symmetric congestion games on series-parallel networks with increasing latencies, the optimal flow is *strongly enforceable* by the cost-balancing tolls computed by BALANCE.

Theorem 3.7. *Every symmetric congestion game on a series-parallel network with increasing latency functions admits a set of optimal tolls with the properties (a), (b), and (c) in Theorem 3.4.*

Remark 3.8. If the network is not series parallel, the optimal flow may not be strongly enforceable by the cost-balancing tolls even for (increasing) linear latencies $d_e(x) = a_e x$, $a_e > 0$. For example, let us consider the four-player game in Figure 1.a. The set of s - t paths consists of $p^u = (e_1, e_3, e_7, e_{11})$ (upper path), $p_1^m = (e_1, e_4, e_8, e_{11})$ (upper middle path), $p_2^m = (e_2, e_5, e_9, e_{12})$ (lower middle path), $p^l = (e_2, e_6, e_{10}, e_{12})$ (lower path), $p_1^c = (e_1, e_4, e_9, e_{12})$ (first cross path), and $p_2^c = (e_2, e_5, e_8, e_{11})$ (second cross path).

In the optimal flow, there are two players on e_1 , e_2 , e_{11} , and e_{12} , and one player on each of the remaining edges. An optimal configuration is $o = (p^u, p_1^m, p_2^m, p^l)$, with social cost 92. The longest s - t path in o is p_2^c , and it has length 24. The cost-balancing tolls assign a toll of 1 to each of e_4 , e_7 , and e_{12} , and no tolls to the remaining edges. The configuration $\sigma = (p_1^c, p_1^c, p_2^c, p_2^c)$ has social cost 104 and is a pure Nash equilibrium of the modified game with cost-balancing tolls.




Figure 1. (a) A symmetric network congestion game with (increasing) linear latency functions for which the cost-balancing tolls do not strongly enforce the optimal flow. Each edge is labeled with its latency function and its identifier. The edges with positive cost-balancing tolls are bold and the tolls appear in parentheses next to the original latency functions. (b) The series-parallel network used in the proof of Theorem 4.1.

4. Hardness of Deciding the Necessity of Tolls

Given an efficiently computable optimal toll mechanism for congestion games on series-parallel networks, it is natural to ask for a set of optimal tolls that minimizes some objective function (e.g., sum of tolls, maximum toll) on the amount of tolls assigned to the edges of the network. To demonstrate the difficulty of computing a better set of optimal tolls, we prove that even for very simple games on series-parallel networks, it is NP-hard to decide whether the use of tolls is really necessary to strongly enforce the optimal flow. In particular, we present a simple two-player game Γ on a series-parallel network with (increasing) linear latency functions $d_e(x) = a_e x$, $a_e > 0$ (see Figure 1.b). The optimal configuration o is a pure Nash equilibrium of Γ (and thus o is weakly enforceable by the trivial tolls $\tau_e = 0$ for all e). We prove that it is NP-hard to decide whether o is the unique pure Nash equilibrium of Γ , in which case tolls serve only to increase the players' disutility, or Γ admits another pure Nash equilibrium of social cost at least $\frac{6}{5}C(o)$, in which case a set of tolls is necessary to strongly enforce the optimal flow.

Theorem 4.1. *Given a two-player linear congestion game Γ on a series-parallel network for which the optimal configuration o is a Nash equilibrium, it is NP-hard to distinguish between the case in which o is the unique pure Nash equilibrium of Γ and that in which Γ admits another pure Nash equilibrium of social cost at least $\frac{6}{5}C(o)$.*

Proof. We use a reduction from PARTITION, which is weakly NP-complete [Garey and Johnson 79, Problem SP12]. Let q_1, \dots, q_k be positive integers such that $\sum_{i=1}^k q_i = B$, with B even. Then PARTITION asks whether there is a set $A^* \subseteq [k]$ such that $\sum_{i \in A^*} q_i = \sum_{i \notin A^*} q_i = B/2$.

Given q_1, \dots, q_k , we construct a two-player game Γ on the series-parallel network in Figure 1.b. The latency functions are $d_{e_{i,1}}(x) = q_i x$, $d_{e_{i,2}}(x) = (q_i/2)x$, $i \in [k]$, and $d_{(s,t)}(x) = (3B/4)x$. In the optimal configuration o , one player uses path $p_2 = (e_{i,2})_{i \in [k]}$ and has individual cost $B/2$, and the other player uses the direct edge from s to t and has individual cost $3B/4$. The optimal configuration has social cost $C(o) = 5B/4$ and is a Nash equilibrium of Γ (thus o is weakly enforceable by the trivial tolls $\tau_e = 0$ for all $e \in E$).

We prove that Γ admits a pure Nash equilibrium of social cost $3B/2$ iff there exists a set $A^* \subseteq [k]$ such that $\sum_{i \in A^*} q_i = B/2$. For every set $A \subseteq [k]$, we define a configuration $\sigma_A = (p_A, p_{\bar{A}})$ consisting of a pair of complementary paths $p_A = (e_{i,j_i})_{i \in [k]}$, and $p_{\bar{A}} = (e_{i,3-j_i})_{i \in [k]}$, where $j_i = 1$ if $i \in A$, and $j_i = 2$ if $i \notin A$. Let $S_A = \sum_{i \in A} q_i$. The cost of the player on p_A is $(B + S_A)/2$, and the cost of the player on $p_{\bar{A}}$ is $B - S_A/2$. The total cost of σ_A is $C(\sigma_A) = 3B/2$.

The configuration σ_A is a Nash equilibrium iff $S_A = B/2$. More specifically, if $S_A = B/2$, the individual cost of both players is $3B/4$, and none of them has an incentive to deviate to the direct edge (s, t) . Hence the configuration σ_A is a Nash equilibrium. If $S_A > B/2$ (respectively $S_A < B/2$), the player on p_A (respectively $p_{\bar{A}}$) has an individual cost greater than $3B/4$ and can decrease her cost by switching to the direct edge (s, t) .

To conclude the proof, we show that Γ does not admit any pure Nash equilibrium other than the optimal configuration and the configurations σ_A corresponding to sets $A \subseteq [k]$ with $S_A = B/2$. First we observe that in any pure Nash equilibrium of Γ , at most one player uses the direct edge (s, t) . If one player uses (s, t) , the other player uses p_2 , and we have the optimal configuration. If no player uses (s, t) , there is one player on every edge $e_{i,j}$, $i \in [k]$, $j \in \{1, 2\}$. If there are two players on some edge $e_{i,2}$, the total cost is greater than $3B/2$. Hence, some player has cost greater than $3B/4$ and can decrease her cost by switching to (s, t) . If both players use an edge $e_{i,1}$, one of them can decrease her cost by switching to $e_{i,2}$. Every configuration with one player on every edge

$e_{i,j}$, $i \in [k]$, $j \in \{1,2\}$, corresponds to a configuration σ_A for an appropriate set $A \subseteq [k]$. □

5. Appendix: A Simple Fact about Acyclic Flows

Let $G(V, E)$ be a directed network with source s and destination t . For a pair of s - t flows f and g , the network $G_{f-g}(V, E_{f-g})$ of $f - g$ has the same vertex set V as G and contains at most one capacitated edge for each edge (u, w) . In particular, if $f_{uw} = g_{uw}$, there is no edge in E_{f-g} corresponding to $(u, w) \in E$; if $f_{(u,w)} > g_{(u,w)}$, E_{f-g} contains a *forward* edge (u, w) with flow $f_{(u,w)} - g_{(u,w)}$; and if $f_{(u,w)} < g_{(u,w)}$, E_{f-g} contains a *backward* edge (w, u) with flow $g_{(u,w)} - f_{(u,w)}$. If f and g are identical, E_{f-g} is empty. If f and g have the same volume, a flow decomposition of $f - g$ yields only cycles and no s - t paths. If f and g are nonidentical acyclic flows of the same volume, G_{f-g} contains at least one cycle with at least one forward edge and at least one backward edge.

Proposition 5.1. *Let $G(V, E)$ be a network with terminals s and t , let f and g be s - t flows, and let $G_{f-g}(V, E_{f-g})$ be the network of $f - g$. If f is acyclic and E_{f-g} is nonempty and contains only forward edges, then $|f| > |g|$ and every forward edge belongs to an s - t path in G_{f-g} .*

Proof. Since f is acyclic and G_{f-g} does not contain any backward edges, G_{f-g} is acyclic too. Let e be any forward edge in G_{f-g} , and let π be a *maximal* path in G_{f-g} that contains e . Let u and w be the vertices connected by π . The path π is maximal in the sense that u does not have any incoming edges and w does not have any outgoing edges in G_{f-g} (recall that G_{f-g} contains only forward edges). Since G_{f-g} is acyclic, u is different from w .

Since u has some outgoing but no incoming edges in G_{f-g} , and since G_{f-g} does not contain any backward edges,

$$\sum_{v:(u,v) \in E_{f-g}} f_{(u,v)} > \sum_{v:(u,v) \in E_{f-g}} g_{(u,v)}$$

(i.e., the volume of f -flow leaving u is greater than the volume of g -flow leaving u). By flow conservation, $\sum_{v:(v,u) \in E_{f-g}} f_{(v,u)} > \sum_{v:(v,u) \in E_{f-g}} g_{(v,u)}$ (i.e., the volume of f -flow entering u is greater than the volume of g -flow entering u). If we assume that u is different from s , we obtain that u has some incoming forward edge in G_{f-g} , which contradicts the assumption that π is maximal. Hence, u is s and $|f| > |g|$.

By a symmetric argument, since w has some incoming but no outgoing edges in G_{f-g} , and since G_{f-g} does not contain any backward edges,

$$\sum_{v:(v,w) \in E_{f-g}} f(v,w) > \sum_{v:(v,w) \in E_{f-g}} g(v,w)$$

(i.e., the volume of f -flow entering w is greater than the volume of g -flow entering w). By flow conservation, $\sum_{v:(w,v) \in E_{f-g}} f(w,v) > \sum_{v:(w,v) \in E_{f-g}} g(w,v)$ (i.e., the volume of f -flow leaving w is greater than the volume of g -flow leaving w). If we assume that w is different from t , we obtain that w has some outgoing forward edge in G_{f-g} , which contradicts the assumption that π is maximal. \square

Proposition 5.2. *Let $G(V, E)$ be a network with terminals s and t , let f and g be s - t flows, and let $G_{f-g}(V, E_{f-g})$ be the network of the flow $f-g$. If g is acyclic and E_{f-g} is nonempty and contains only backward edges, then $|g| > |f|$ and every backward edge belongs to an s - t path in G_{f-g} .*

Proof. We apply Proposition 5.1 to $G_{g-f}(V, E_{g-f})$ and observe that forward edges in G_{g-f} correspond to backward edges in G_{f-g} . \square

Acknowledgments. An extended abstract of this work appeared in *Internet and Network Economics: Third International Workshop, WINE 2007, San Diego, CA, USA, December 12–14, 2007, Proceedings*, Lecture Notes in Computer Science 4858, edited by X. Deng and F. C. Graham, pp. 179–190, Berlin, Springer, 2007. This work was done while the first author was with the Department of Information and Communication Systems Engineering, University of the Aegean, Greece. The second author was partially supported by EU/6th Framework Programme, contract 001907 (DELIS).

References

- [Awerbuch et al. 05] B. Awerbuch, Y. Azar, and A. Epstein. “The Price of Routing Unsplittable Flow.” In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pp. 57–66. New York: ACM Press, 2005.
- [Beckmann et al. 56] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. New Haven: Yale University Press, 1956.
- [Bein et al. 85] W. Bein, P. Brucker, and A. Tamir. “Minimum Cost Flow Algorithms for Series-Parallel Networks.” *Discrete Applied Mathematics* 10 (1985), 117–124.
- [Caragiannis et al. 06] I. Caragiannis, C. Kaklamani, and P. Kanellopoulos. “Taxes for Linear Atomic Congestion Games.” In *Algorithms—ESA 2006: 14th Annual European Symposium, Zurich, Switzerland, September 11–13, 2006, Proceedings*, Lecture Notes in Computer Science 4168, pp. 184–195. Berlin: Springer, 2006.

- [Christodoulou and Koutsoupias 05] G. Christodoulou and E. Koutsoupias. “The Price of Anarchy of Finite Congestion Games.” In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pp. 67–73. New York: ACM Press, 2005.
- [Cole et al. 03a] R. Cole, Y. Dodis, and T. Roughgarden. “How Much Can Taxes Help Selfish Routing.” In *Proceedings of the 4th ACM Conference on Electronic Commerce*, pp. 98–107. New York: ACM Press, 2003.
- [Cole et al. 03b] R. Cole, Y. Dodis, and T. Roughgarden. “Pricing Network Edges for Heterogeneous Selfish Users.” In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pp. 521–530. New York: ACM Press, 2003.
- [Cominetti et al. 06] R. Cominetti, J. R. Correa, and N. E. Stier-Moses. “The Impact of Oligopolistic Competition in Networks.” Working Paper DRO-2006-03, Columbia Business School, New York, 2006.
- [Cormen et al. 01] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*, Second Edition. New York: McGraw-Hill, 2001.
- [Correa and Stier-Moses 07] J. R. Correa and N. E. Stier-Moses. “Stackelberg Routing in Atomic Network Games.” Working Paper DRO-2007-03, Columbia Business School, New York, 2007.
- [Fleischer 05] L. Fleischer. “Linear Tolls Suffice: New Bounds and Algorithms for Tolls in Single Source Networks.” *Theoretical Computer Science* 348 (2005), 217–225.
- [Fleischer et al. 04] L. Fleischer, K. Jain, and M. Mahdian. “Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games.” In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 277–285. Los Alamitos, CA: IEEE Press, 2004.
- [Gairing et al. 05] M. Gairing, T. Lücking, B. Monien, and K. Tiemann. “Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture.” In *Automata, Languages and Programming: 32nd International Colloquium, ICALP 2005, Lisbon, Portugal, July 11–15, 2005, Proceedings*, Lecture Notes in Computer Science 3580, pp. 51–65. Berlin: Springer, 2005.
- [Garey and Johnson 79] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. New York: W. H. Freeman, 1979.
- [Holzman and Law-Yone 97] R. Holzman and N. Law-Yone. “Strong Equilibrium in Congestion Games.” *Games and Economic Behaviour* 21 (1997) 85–101.
- [Kaporis and Spirakis 06] A. C. Kaporis and P. G. Spirakis. “The Price of Optimum in Stackelberg Games on Arbitrary Single Commodity Networks and Latency Functions.” In *Proceedings 18th Annual ACM Symposium on Parallelism in Algorithms and Architectures*, pp. 19–28. New York: ACM Press, 2006.
- [Karakostas and Kolliopoulos 04] G. Karakostas and S. Kolliopoulos. “Edge Pricing of Multicommodity Networks for Heterogeneous Users.” In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 268–276. Los Alamitos, CA: IEEE Press, 2004.
- [Koutsoupias and Papadimitriou 99] E. Koutsoupias and C. Papadimitriou. “Worst-Case Equilibria.” In *STACS ’99: 16th Annual Symposium on Theoretical Aspects of Computer Science, Trier, Germany, March 4–6, 1999 Proceedings*, Lecture Notes in Computer Science 1563, pp. 404–413. Berlin: Springer, 1999.

- [Milchtaich 06] I. Milchtaich. “Network Topology and the Efficiency of Equilibrium.” *Games and Economic Behaviour* 57 (2006), 321–346.
- [Rosenthal 73] R. W. Rosenthal. “A Class of Games Possessing Pure-Strategy Nash Equilibria.” *International Journal of Game Theory* 2 (1973), 65–67.
- [Roughgarden and Tardos 02] T. Roughgarden and É. Tardos. “How Bad is Selfish Routing?” *Journal of the Association for Computing Machinery* 49:2 (2002), 236–259.
- [Swamy 07] C. Swamy. “The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games.” In *Proceedings of the 18th Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 1133–1142. Philadelphia: SIAM, 2007.
- [Valdez et al. 82] J. Valdez, R. E. Tarjan, and E. L. Lawler. “The Recognition of Series-Parallel Digraphs.” *SIAM Journal on Computing*, 11:2 (1982), 298–313.

Dimitris Fotakis, School of Electrical and Computer Engineering, National Technical University of Athens, Iroon Polytechniou Str. 9, 15780 Athens, Greece (fotakis@cs.ntua.gr)

Paul G. Spirakis, Research Academic Computer Technology Institute, N. Kazantzaki Str., Rion, 26500 Patras, Greece (spirakis@cti.gr)

Received March 18, 2008; accepted January 27, 2009.