Cost-Balancing Tolls for Atomic Network Congestion Games*

Dimitris Fotakis[1] and Paul G. Spirakis[2]

[1] Dept. of Information and Communication Systems Engineering,
University of the Aegean, 83200 Samos, Greece. fotakis@aegean.gr
[2] Research Academic Computer Technology Institute,
P.O. Box 1382, N. Kazantzaki Str., Rion, 26500 Patras, Greece. spirakis@cti.gr

Abstract. We investigate the existence of optimal tolls for atomic symmetric network congestion games with unsplittable traffic and arbitrary non-negative and non-decreasing latency functions. We focus on pure Nash equilibria and 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 increasing latencies, 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 2-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 equilibrium of total cost at least $6/5$ times the optimal cost.

1 Introduction

Congestion games provide a natural model for non-cooperative resource allocation in large-scale communication networks and have been the subject of intensive research in algorithmic game theory. In an (atomic) *congestion game* [19], a finite set of non-cooperative players, each controlling an unsplittable unit of traffic, compete over a finite set of resources. All players using a resource experience a latency given by a non-negative and non-decreasing 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 where no player can decrease her individual cost by unilaterally changing her strategy.

* Partially supported by EU / 6th Framework Programme, contract 001907 (DELIS).

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' non-cooperative and selfish behaviour, the network manager can introduce economic incentives that influence the players' selfish choices and hopefully induce an optimal network configuration.

Economic incentives can be naturally modelled by non-negative per-unit-of-traffic *tolls* (aka taxes or prices) assigned to the resources. The tolls are levied by the network manager and comprise an additional cost factor which 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 cost, 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*, which make the Nash equilibria of the modified game coincide with the optimal solution.

Related Work. In the non-atomic 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. [7, 6] and references therein). A classical result is that the optimal solution is realized as the Nash equilibrium of a non-atomic congestion game with *marginal cost tolls* [2]. In simple words, the performance degradation due to the selfish and non-cooperative behaviour of non-atomic 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 non-atomic congestion games was motivated by the limitations of marginal cost tolls. Cole *et al.* [7] were the first to consider *heterogeneous* players, who may have a different valuation of time (latency) in terms of money (toll), and established the existence of optimal tolls for non-atomic symmetric network congestion games. Their proof was based on Brouwer's fixed point theorem and was non-constructive. In addition, 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 [11] extended the results of [7] 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.* [12] and Karakostas and Kolliopoulos [16] independently

[1] Let $d_e(x)$ 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 $d_2(x) = (1+\varepsilon)x$, $\varepsilon > 0$. 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.

proved that the existence of optimal tolls for non-atomic 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 non-atomic congestion games, the Nash equilibrium is essentially unique (under mild assumptions on the latency functions, see e.g. [20]). Hence the tolls of [2, 7, 11, 12, 16] 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 where a set of tolls *weakly enforces* the optimal solution, in the sense that the optimal solution is realized as some equilibrium of the game with tolls, and the case where a set of tolls *strongly enforces* the optimal solution, in the sense that the optimal solution is realized as the unique equilibrium of the game with tolls.

For atomic congestion games with splittable traffic and heterogeneous players, Swamy [21] 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.* [8] presented a toll mechanism 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 toll mechanisms for atomic congestion games with unsplittable traffic is [4]. Caragiannis *et al.* 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 non-symmetric game for which the price of anarchy remains at least 1.2 under any toll mechanism. Therefore, non-symmetric 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 strongly enforcing 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 of linear games with arbitrary strategies. The first mechanism [4, Theorem 3] is simple and improves the pure price of anarchy to 2.155 (from 2.5 [1, 5]). The second mechanism [4, 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 [1, 5]). 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].

[2] The uniqueness of Nash equilibrium in non-atomic games is also exploited by the algorithm of [15], which computes the smallest fraction of coordinated players required by a Stackelberg routing strategy to induce the optimal solution.

[3] The *price of anarchy* [17] is a widely accepted measure of the performance degradation due to the players' non-cooperative and selfish behaviour. 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 of congestion games, see e.g. [13].

[4] For a game where the tolls of [4, Theorem 3] do not weakly enforce the optimal solution, consider 2 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 where the tolls of [4, Theorem 5] do not strongly enforce the optimal solution, consider the instance of Fig. 1.b with $k = 2$, $q_1 = q_2 = 2$, and latency function $d_{(s,t)}(x) = (3 + \varepsilon)x$, $\varepsilon \in (0, 1)$, for the direct (s, t) edge.

Contribution. Despite the considerable interest in optimal toll mechanisms for atomic and non-atomic congestion games, it is still unknown whether there is an optimal toll mechanism for symmetric games with unsplittable traffic. This is true even for relatively simple symmetric network congestion games, such as games on series-parallel networks and games on parallel links with non-linear latencies.

In this work, we investigate the existence of optimal tolls for symmetric network congestion games with unsplittable traffic, homogeneous players, and arbitrary non-negative and non-decreasing 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 toll mechanisms for non-atomic games [7, 11, 12, 16]. 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 [4] 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 give a simple example where 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 (e.g. [11]) and Stackelberg routing strategies (e.g. [10, 21]) 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 strictly 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 [4, 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.

To highlight the difficulty of computing a better set of optimal tolls, we prove that even for 2-player linear congestion games on series-parallel networks, it is NP-hard to distinguish between the case where the optimal solution is the unique pure Nash equilibrium (thus any positive tolls only serve to increase the players' disutility) and the case where there is another equilibrium of total cost at least $6/5$ times the optimal cost (and hence some positive tolls are required to strongly enforce the optimal solution).

2 Model, Definitions, and Notation

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, $\Sigma_i \subseteq 2^E$ denotes the strategy space of each player i, and $d_e : \mathbb{N} \mapsto \mathbb{R}_{\geq 0}$ is a non-negative and non-decreasing 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 $d_e(x) = a_e x + b_e$, $a_e, b_e \geq 0$.

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 each resource e, $\sigma_e = |\{i \in N : e \in \sigma_i\}|$ denotes the congestion induced on e by σ. The individual cost of 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 his individual cost by unilaterally changing his strategy. Formally, σ is a Nash equilibrium if for every player i and every strategy $s_i \in \Sigma_i$, $c_i(\sigma) \leq c_i(\sigma_{-i}, s_i)$ [5].

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. Rosenthal [19] proved 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 denote the number of resources. We restrict our attention to *symmetric network* congestion games, where 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 $\mathcal{P}$.

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 $(f_e)_{e \in E} \in \mathbb{R}^m_{\geq 0}$ that satisfies the flow conservation at all vertices other than s and t. The *volume* of an $s - t$ flow is the total flow leaving s. A flow is *acyclic* if there is no directed 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 say that a flow σ is *feasible* if there is a configuration inducing congestion σ_e on every edge e. We slightly abuse the notation by letting the same symbol denote both a configuration and the feasible flow induced by it.

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, usually denoted 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 non-negative and non-decreasing latency functions admits an integral acyclic min-cost flow of volume n computable in polynomial time if $x d_e(x)$ are convex. Therefore, if $x d_e(x)$ 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 arbitrary non-negative and non-decreasing latencies, an optimal configuration can be computed in $\mathrm{O}(m + n \log m)$ time by the greedy algorithm [3].

[5] For a vector $x = (x_1, \dots, x_n)$, $(x_{-i}, x'_i) \equiv (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n)$.

Tolls. We consider a scenario where the network manager levies tolls on the edges of the network trying to influence the players' selfish choices and induce an optimal configuration. A set of *tolls* is a function $\tau : E \mapsto \mathbb{R}_{\geq 0}$ that assigns a non-negative per-unit-of-traffic toll τ_e to each resource e. The modified congestion game induced by τ is $\bar{\Gamma}_\tau(N, E, (\Sigma_i)_{i \in N}, (\overline{d}_e)_{e \in E})$, where $\overline{d}_e(x) = d_e(x) + \tau_e$ for all $e \in E$. 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 resources in σ_i and becomes:

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

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. [7, 12, 16], and the case of *refundable* tolls in [4]). 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 motivate the selfish players to induce a given feasible 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 an arbitrary configuration inducing congestion f_e on each 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 each 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 should be not 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.

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 (resp. 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 $\mathrm{O}(|V| + |E|)$ time (see e.g. [22] for more details).

3 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 $p' \in \mathcal{P}$,

$$\sum_{e \in p} \overline{d}_e(f_e) = \sum_{e \in p} (d_e(f_e) + \tau_e) \leq \sum_{e \in p'} (d_e(f_e) + \tau_e) = \sum_{e \in p'} \overline{d}_e(f_e) \tag{1}$$

Proposition 1. *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 each edge e is a pure Nash equilibrium of the congestion game $\bar{\Gamma}_\tau$. In particular, for every player i and 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:

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

where the first inequality follows from the definition 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 a $s - t$ network $G(V, E)$ and an acyclic $s - t$ flow f. BALANCE 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.
2. BALANCE assigns a non-negative length $d_e(f_e)$ to each edge $e \in E_f$ and computes the longest path from s to any vertex in G_f reachable from s. 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.

If f is acyclic, G_f is a directed acyclic graph (DAG). Therefore, the longest paths from s can be computed in $\mathrm{O}(|V| + |E_f|)$ time by negating the edge lengths and solving the corresponding single-source shortest path problem (see e.g. [9, Section 24.2]). Hence BALANCE can be implemented in time linear in the size of the network.

Lemma 1. *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 $e = (u, v) \in E_f$, $-\ell_v \leq -\ell_u - d_e(f_e)$, which implies that $\tau_e = \ell_v - (\ell_u + d_e(f_e)) \geq 0$. By the same reasoning, if an edge $e \in E_f$ lies on some longest path, 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 $p = (s = u_0, u_1, \ldots u_k = t)$ be an $s-t$ path with positive flow on all edges. Since for every edge $e = (u_i, u_{i+1}) \in p$, $\tau_e = \ell_{u_{i+1}} - (\ell_{u_i} + d_e(f_e))$, $\sum_{e\in p}(d_e(f_e)+\tau_e) = \ell_t$. On the other hand, the cost on every $s-t$ path p' containing an edge e with $f_e = 0$ is at least $\tau^{\max} \geq \ell_t$. □

In the following, we refer to the tolls computed by BALANCE for an acyclic flow f as *the cost-balancing tolls for f.*

Theorem 1. *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 $\delta > 0$. No edge with toll $\tau^{\max}$ is used 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 non-negative and non-decreasing 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 1) and o is weakly enforceable by τ (Proposition 1).

(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, every edge with positive flow in o is assigned a toll no greater than ℓ_t. The remaining edges are assigned a toll equal to $\tau^{\max}$.

We then prove that all edges with toll $\tau^{\max}$ remain unused 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 1, 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 each player i entirely consists of edges with positive flow in o. In the proof of Lemma 1, 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 player i in σ is at most $\ell_t = \max_{p: o_p^{\min}>0} \sum_{e\in p} d_e(o_e)$. □

Remark 1. 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 the 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 each 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 enforced by τ (note that it is not strongly enforced by τ because the latency of the second link is constant). 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'.

3.1 Optimality of Cost-Balancing Tolls for Series-Parallel Networks

It is not hard to prove that if the latency functions are increasing, the optimal solution is strongly enforceable by the cost-balancing tolls for *networks with linearly independent routes* (e.g. [18]), an interesting class of $s-t$ networks including parallel-link networks. Next we generalize this result to series-parallel networks, an important class of networks considerably more complex than networks with linearly independent routes. This generalization is particularly interesting because in a series-parallel network, an $s-t$ flow may be realized by many different configurations (see e.g. the proof of Theorem 3).

Lemma 2. *Let Γ be a symmetric congestion game on a series-parallel network $G(V,E)$ with increasing latency functions, and let f be any feasible acyclic flow that admits a set of cost-balancing tolls τ. Then every pure Nash equilibrium σ of the congestion game $\bar{\Gamma}_\tau$ induces congestion $\sigma_e = f_e$ on each 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 non-negative 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 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 $H_1, \ldots, H_k$ has this property. In other words, H corresponds to the entire subtree of an internal node u_H in G's decomposition tree, and $H_1, \ldots, H_k$ 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). H is a *minimal* component in the series-parallel decomposition of G wrt 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. 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. Hence, by flow conservation, H_1 and H_2 are not connected in series. Therefore, H is formed by a parallel composition of $H_1, \ldots, H_k$. Let s_H and t_H be the common endpoints of H and its components $H_1, \ldots, H_k$.

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 part of σ_i between s_H and t_H). Since H_1 does not contain any backward edges, $\sigma_e \geq f_e$ and $\overline{d}_e(\sigma_e) \geq \overline{d}_e(f_e)$ for all $e \in p^+$. Moreover, for the forward edge e^+, $\overline{d}_{e^+}(\sigma_e) > \overline{d}_{e^+}(f_e)$ because the latency functions $d_e(x)$ (and thus the functions $\overline{d}_e(x)$) are increasing. Therefore, the individual cost of player i on p^+ is:

$$\sum_{e \in p^+} \overline{d}_e(\sigma_e) > \sum_{e \in p^+} \overline{d}_e(f_e) \tag{2}$$

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^- entirely consisting of backward edges. Therefore $\sigma_e + 1 \leq f_e$ for all $e \in p^-$, which implies that:

$$\sum_{e \in p^-} \overline{d}_e(\sigma_e + 1) \leq \sum_{e \in p^-} \overline{d}_e(f_e) \leq \sum_{e \in p^+} \overline{d}_e(f_e) \tag{3}$$

For the last inequality, we observe that $f_{p^-}^{\min} \geq 1$ and that (1) holds for p^- and p^+, since they are segments of $s - t$ paths with common endpoints s_H and t_H.

Combining (2) and (3), we conclude that $\sum_{e \in p^+} \overline{d}_e(\sigma_e) > \sum_{e \in p^-} \overline{d}_e(\sigma_e + 1)$. Therefore player i can decrease her cost by changing her path between s_H and t_H from p^+ to p^-. This contradicts the assumption that σ is a Nash equilibrium of $\bar{\Gamma}_\tau$. □

Theorem 1 and Lemma 2 immediately imply the following:

Theorem 2. *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 1.*

Remark 2. If the network is not series parallel, the optimal flow may not be strongly enforceable by the cost-balancing tolls even for linear latencies $d_e(x) = a_e x$, $a_e > 0$. For example, consider the 4-player game in Fig. 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).

An optimal configuration is $o = (p^u, p_1^m, p_2^m, p^l)$ and has cost 92. The longest $s - t$ path in o is p_2^c and has length 24. The cost-balancing tolls assign a toll of 1 to 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 cost 104 and is a pure Nash equilibrium of the modified game with cost-balancing tolls.

4 Hardness of Deciding the Necessity of Tolls

To highlight the difficulty of computing a better set of optimal tolls, we prove 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 flow.

Theorem 3. *Given a 2-player linear congestion game Γ on a series-parallel network for which the optimal configuration o is Nash equilibrium, it is* NP*-hard to distinguish between the case where o is the unique pure Nash equilibrium of Γ, and the case where Γ admits another pure Nash equilibrium of cost at least $\frac{6}{5}C(o)$.*

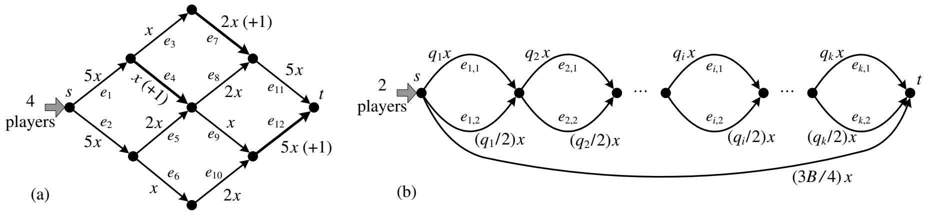

Fig. 1. (a). A symmetric network congestion game with (increasing) linear latency functions for which the cost-balancing tolls do not strongly enforce the optimal flow. Each edge is labeled with its latency function and its identifier. The edges with positive cost-balancing tolls are bold and the tolls appear in parenthesis next to the original latency functions. (b). The series-parallel network used in the proof of Theorem 3.

Proof. We use a reduction from Partition, which is weakly NP-complete [14, Problem SP12]. Let $q_1, \dots, q_k$ be positive integers such that $\sum_{i=1}^{k} q_i = B$, with B even. 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 2-player game Γ on the series-parallel network in Fig. 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 cost $B/2$ and the other player uses the direct edge from s to t and has cost $3B/4$. The optimal configuration has cost $C(o) = 5B/4$ and is a Nash equilibrium of Γ (in fact, 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 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 $j_i = 2$ otherwise. 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 $C(\sigma_A) = 3B/2$. We prove that σ_A is a Nash equilibrium iff $S_A = B/2$. If $S_A = B/2$, the 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 Nash equilibrium. If $S_A > B/2$ (resp. $S_A < B/2$), the player on p_A (resp. $p_{\bar{A}}$) has 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 some 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 $A \subseteq [k]$. □

References

1. B. Awerbuch, Y. Azar, and A. Epstein. The Price of Routing Unsplittable Flow. *Proc. of STOC '05*, pp. 57–66, 2005.
2. M. Beckmann, C.B. McGuire, and C.B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. W. Bein, P. Brucker, and A. Tamir. Minimum Cost Flow Algorithms for Series-Parallel Networks. *Discrete Applied Mathematics*, 10:117–124, 1985.
4. I. Caragiannis, C. Kaklamanis, and P. Kanellopoulos. Taxes for Linear Atomic Congestion Games. *Proc. of ESA '06*, LNCS 4168, pp. 184–195, 2006.
5. G. Christodoulou and E. Koutsoupias. The Price of Anarchy of Finite Congestion Games. *Proc. of STOC '05*, pp. 67–73, 2005.
6. R. Cole, Y. Dodis, and T. Roughgarden. How Much Can Taxes Help Selfish Routing. *Proc. of EC '03*, pp. 98–107, 2003.
7. R. Cole, Y. Dodis, and T. Roughgarden. Pricing Network Edges for Heterogeneous Selfish Users. In *Proc. of STOC '03*, pp. 521–530, 2003.
8. R. Cominetti, J.R. Correa, and N.E. Stier-Moses. The Impact of Oligopolistic Competition in Networks. *DRO-2006-03*, Columbia Business School, 2006.
9. T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms*, (2nd Edition). McGraw-Hill, 2001.
10. J.R. Correa and N.E. Stier-Moses. Stackelberg Routing in Atomic Network Games. *DRO-2007-03*, Columbia Business School, 2007.
11. L. Fleischer. Linear Tolls Suffice: New Bounds and Algorithms for Tolls in Single Source Networks. *Theoretical Computer Science*, 348:217–225, 2005.
12. L. Fleischer, K. Jain, and M. Mahdian. Tolls for Heterogeneous Selfish Users in Multicommodity Networks and Generalized Congestion Games. *Proc. of FOCS '04*, pp. 277–285, 2004.
13. M. Gairing, T. Lücking, B. Monien, and K. Tiemann. Nash Equilibria, the Price of Anarchy and the Fully Mixed Nash Equilibrium Conjecture. *Proc. of ICALP '05*, LNCS 3580, pp. 51–65, 2005.
14. M.R. Garey and D.S. Johnson. *Computers and Intractability: A Guide to the Theory of* NP*-Completeness*. W.H. Freeman, 1979.
15. A.C. Kaporis and P.G. Spirakis. The Price of Optimum in Stackelberg Games on Arbitrary Single Commodity Networks and Latency Functions. *Proc. of SPAA '06*, pp. 19–28, 2006.
16. G. Karakostas and S. Kolliopoulos. Edge Pricing of Multicommodity Networks for Heterogeneous Users. *Proc. of FOCS '04*, pp. 268–276, 2004.
17. E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. *Proc. of STACS '99*, LNCS 1563, pp. 404–413, 1999.
18. I. Milchtaich. Network Topology and the Efficiency of Equilibrium. *Games and Economic Behaviour*, 57:321–346, 2006.
19. R.W. Rosenthal. A Class of Games Possessing Pure-Strategy Nash Equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
20. T. Roughdarden and É. Tardos. How Bad is Selfish Routing? *Journal of the ACM*, 49(2):236–259, 2002.
21. C. Swamy. The Effectiveness of Stackelberg Strategies and Tolls for Network Congestion Games. *Proc. of SODA '07*, pp. 1133–1142, 2007.
22. J. Valdez, R.E. Tarjan, and E.L. Lawler. The Recognition of Series-Parallel Digraphs. *SIAM Journal on Computing*, 11(2):298–313, 1982.