THE PINHAS SAPIR CENTER FOR DEVELOPMENT
TEL AVIV UNIVERSITY

"Transfer Implementation in Congestion Games"

Itai Arieli[1]

Discussion Paper No. 9-14

October 2014

I would like to thank The Pinhas Sapir Center for Development at Tel Aviv University for their financial support.

We thank two anonymous referees for helpful comments.

[1] Itai Arieli – Faculty of Industrial Engineering and Management, Technion-Israel Institute of Technology. Email: iarieli@tx.technion.ac.il

Abstract

We study an implementation problem faced by a planner who can influence selfish behavior in a roadway network. It is commonly known that Nash equilibrium does not necessarily minimize the total latency on a network and that levying a tax on road users that is equal to the marginal congestion effect each user causes implements the optimal latency state. This holds however only under the assumption that taxes have no effect on the utility of the users. In this paper we consider taxes that satisfy the budget balance condition and that are therefore obtained using a money transfer among the network users. Hence at every stated the overall taxes imposed upon the users sums up to zero. We show that the optimal latency state can be guaranteed as a Nash equilibrium using a simple, easily computable transfer scheme that is obtained from a fixed matrix.

In addition, the resulting game remains a potential game and the levied tax on every edge is a function of its congestion.

1 Introduction

Roadway congestion is a source of enormous economic costs. The underlying assumption that users are selfish and their goal is to minimize their own latency time yields a Nash equilibrium that in general does not, and is even far from, minimizing the total latency. This inefficiency motivates the construction of economic incentives that improve efficiency in equilibrium and has therefore given rise to a large body of literature that studies the influence that taxing roads has on latency time.

As a first illustration of these ideas consider the classical Braess's paradox (see Figure 1). One unit of traffic commutes from the initial node s to the terminal node t. Each edge of the network in Figure 1 is labelled with its latency function, giving the delay incurred by traffic on the link as a function of the amount of traffic that uses the link. At Nash equilibrium, all traffic uses the route $s \to v \to w \to t$ and experiences two units of latency. On the other hand, if one unit of tax is levied on the edge (v, w), then in the Nash equilibrium of the resulting game half of the traffic uses each of the routes $s \to v \to t$ and $s \to w \to t$. In particular, the route $s \to v \to w \to t$ has a latency of 1 and a cost of 2 with respect to this flow, and hence does not offer an attractive alternative to the users. In this new flow at Nash equilibrium, everyone experiences a latency of $3/2$ and no taxes are paid. This outcome is clearly superior to the original flow at Nash equilibrium in the absence of taxes since $3/2$ is the minimal total latency for this network. The example highlights the known paradox that taxing some of the routes may improve efficiency in equilibrium and lead to a superior outcome in terms of the total latency time.

An old, related idea that was introduced informally by Pigou [7], and implemented formally by Beckman et al. [1] and more recently by Sandholm [11], is the principle of *marginal cost pricing.* To better understand this idea, consider an implementation problem faced by a social planner who would like to implement the optimal latency flow among a continuum population of network users. Under the marginal cost pricing principle each user pays an additional tax that is equal to the marginal delay he causes to the other users. If the users of the roadway network take into consideration both the imposed tax and the

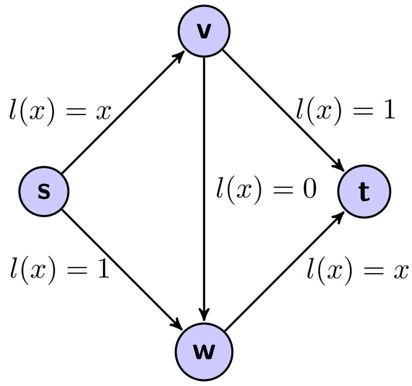

Figure 1: Braess's paradox.

delay on each route, then the proposed tax scheme yields a Nash equilibrium that minimizes the total latency time.

Sandholm [12] (see also [11]) provides a family of tax schemes that is based on marginal cost pricing. Each member of this family *evolutionarily implements* the optimal latency state under a wide range of evolutionarily based behavior adjustment processes, called *revision protocols,* employed by the users. These tax schemes alter the underlying game by pricing every edge at every given time as a function of the local congestion on that particular edge. Hence in practice the planner does not need to know the precise choice of routes by the users. Moreover, the pricing schemes introduced by Sandholm change the potential function of the game to be the total latency function. Therefore the optimal latency flow is globally stable under any reasonable adjustment process employed by the users.

A recent paper by Fleischer et al. [5] provides a general existence result of a tax scheme (or tolls) that implements efficient behavior in equilibrium for a large class of congestion games. Further developments on pricing networks can be found in [2], [4], [6], [13]. As pointed out by Cole et al. [3], these results ignore the disutility caused to the users due to the levied tax. Cole et al. [3] demonstrate that if one takes into consideration the negative effect of the tax in the marginal cost pricing, then the Nash equilibrium of the original game is always superior in terms of latency for every network with linear latency functions. In addition, Cole et al. show that if one takes into consideration the tax levied upon the users when calculating optimal latency, then finding

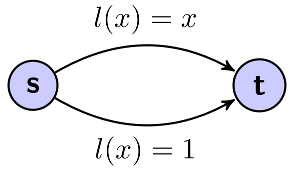

Figure 2: Transfer among users improves total latency.

optimal taxation is computationally hard (see Theorem 6.2 in [3]).

To highlight further the limitation of taxation, consider the example depicted in Figure 2. In the unique Nash equilibrium all traffic uses the upper route with a total latency of 1, and the optimal flow is obtained when half of the users use the upper route and the other half use the lower route. This yields a total latency of $1/4 + 1/2 = 3/4$. Under marginal cost pricing, an additional tax of $2x$ is levied on the upper route, which implements $(1/2, 1/2)$ as the unique equilibrium. If, however, one takes the tax levied into consideration when calculating the total latency time, then the resulting latency at equilibrium is $1/2 \cdot (3/2) + 1/2 \cdot 1 = 5/4$. Furthermore, it can easily be verified that any tax scheme that involves only negative payments does not improve latency in equilibrium.

As an alternative consider a (progressive) tax scheme that prices the upper route by a fixed amount of $1/4$ and benefits the lower route by a fixed amount of $1/4$. According to this tax scheme the optimal latency of $(1/2, 1/2)$ is obtained as a unique equilibrium. Moreover, this tax scheme satisfies the *budget balance condition* in equilibrium. That is, the overall tax paid by the users is 0; the tax that is levied on the users of the upper route subsidizes the benefit given to the users in the lower route, where the users of the upper route pay $1/4$ to the users of the lower route. Hence under the proposed tax scheme the budget balance condition implies that at equilibrium all taxes are obtained by a money transfer among the network users, and there is no additional payment from the planner. The budget balance condition is only guaranteed in equilibrium, whereas off equilibrium it might not hold.

Here we consider taxes that satisfy the budget balance condition and thus are obtained using a money transfer among the network users. Our goal is to study ways to implement the optimal latency flow as a Nash equilibrium using a simple *transfer scheme.* In particular we study transfer scheme that is obtained using a fixed transfer matrix that determines the amount of money transferred between every pair of edges. In addition, we are interested in implementing a transfer scheme for which the resulting game is a potential game. Thus, as in Sandholm [12], the optimal latency flow would be globally stable when the users apply any reasonable myopic adjustment learning rules.

Our main result demonstrates the construction of a transfer matrix for which the resulting game retains a potential function such that its unique equilibrium minimizes latency time. Moreover, the fixed transfer matrix that we construct has the property that the tax levied on every edge solely depends on its congestion and may be calculated in time that is polynomial in the number of edges.[1]

2 Model

Consider a directed graph $G = (V, E)$ with a source $s \in V$ and a sink $t \in V$. Denote the set of simple $s - t$ routes in G by $\mathcal{P}$, which is assumed to be nonempty. We allow parallel edges but we assume that G has no cycles. Consider one unit of traffic wishing to travel from s to t. A flow over the graph G is a probability distribution $x = \{x_i\}_{i \in \mathcal{P}} \in Y = \Delta(\mathcal{P})$ indexed by $s-t$ routes, with x_i representing the proportion of traffic using route i as the chosen route from s to t.

For a route $i \in \mathcal{P}$ let Φ_i be the set of edges that comprises i. A flow on routes induces a unique flow on edges, defined as a vector $\{x_e\}_{e \in E}$ where $x_e = \sum_{i: e \in \Phi_i} x_P$ represents the congestion of edge e. We note that a flow on edges may correspond to many different flows on routes. Let X be the set of all possible flows on edges.[2]

A congestion game $C = (G, (s, t), (l_e)_{e \in E})$ over G comprises a nonnegative,

[1] In contrast, Cole et al. [3] demonstrate that when only fixed positive taxes are allowed, calculating an approximately optimal tax scheme is NP hard.

[2] We note that X corresponds to the set of all non-negative weights $\{x_e\}_{e \in E}$ such that the outflow from node s and the inflow to node t are 1, and for any other node v the inflow to v equals the outflow from v.

continuous, nondecreasing latency function l_e for each edge e. l_e describes the delay incurred by traffic on edge e as a function of the congestion x_e. The latency of a route $i \in \mathcal{P}$ with respect to a flow $x \in X$ is then given by $F_i(x) = \sum_{e \in \Phi_i} l_e(x_e)$. We measure the quality of a flow by its total latency $L(x)$, defined by $L(x) = \sum_{i \in \mathcal{P}} F_i(x) x_i$, or, equivalently, by $L(x) = \sum_{e \in E} l_e(x_e) x_e$. We will call an edge flow that minimizes $L(\cdot)$ *optimal.* Such a flow always exists since X is a compact set and $L(\cdot)$ is a continuous function on X. We shall assume that $L(\cdot)$ is strictly convex; hence, it has a unique minimizer over X.[3]

A tax scheme $\tau = \{\tau_e\}_{e \in E}$ is a set of functions to be placed on the edges of the network G such that $\tau_e : X \to \mathbb{R}$ where $\tau_e(x)$ is the tax levied on the edge e when the flow is x. We further assume that all agents trade time and money equally, and that avoiding one unit of latency time equals one units of money. Denote the game obtained from the tax scheme τ by F^τ. Under the tax scheme τ, when the flow is induced by x a route i incurs a total cost of

$$F_i^\tau(x) = \sum_{e \in \Phi_i} l_e(x_e) + \tau_e(x).$$

We say that the tax scheme τ satisfies the *budget balance condition* at $x \in X$ if

$$\sum_{e \in E} \tau_e(x) x_e = 0.$$

Under the budget balance condition the taxes levied at state x are obtained by a money transfer among the network users, which means that overall tax money is neither wasted nor invested by the planner at state x. We shall focus on a tax scheme that satisfies the budget balance condition.

For the game in Figure 2 we show that one can find a tax scheme $\tau = \{\tau_e\}_{e \in E}$ such that τ_e are constant, the optimal latency x^* comprises a unique equilibrium of F^τ, and τ satisfies the budget balance condition at x^*. The first question that naturally arises is whether one can always find a tax scheme such that the optimal flow would be realized in equilibrium at which the balance budget condition holds. Our first observation demonstrates that this is in fact possible. We state this simple preliminary result.[4]

[3]This standard assumption is guaranteed, for example, whenever $x_e l_e''(x_e) + 2l_e'(x_e) > 0$ for every $e \in E$ and $x_e \in [0, 1]$.

[4]The proof of Lemma 1 is presented as part of the proof of our Main Theorem in Section 2.

Lemma 1. *Let $(G,(s,t),(l_e)_{e\in E})$ be a congestion game. There exists a real numbers $\tau=(\tau_e)_{e\in E}$ (the tax levied on every edge e is constant and equals τ_e) such that the unique equilibrium x^* of F^τ is the unique optimal latency and the tax scheme τ satisfies the budget balance condition at x^**

$$\sum_{i=1}^{n} x_e^* \tau_e = 0. \tag{1}$$

Lemma 1 shows that one can always have a tax scheme such that in equilibrium the tax comprises only money transfer among network users. The problem with this approach can best be understood by considering the framework of Sandholm [12]. In his model a myopic adjustment process is implemented by the users in light of new information on the congestion of an alternative route. When these dynamical processes are considered, Lemma 1 does not provide a satisfactory answer since throughout the learning process, in a non-equilibrium states the taxes that users pay do not necessarily satisfy the budget balance condition. In some states money will be wasted or, alternatively, the planner will have to invest money in order to implement the tax scheme.

We consider next a tax scheme that is based solely on a money transfer among network users. Formally, a *transfer scheme* is defined as follows.

Definition 1. Let $C=(G,(s,t),(l_e)_{e\in E})$ be a congestion game. Let K be a set of distinct pairs of edges, that is, $K=\{(e,f): e,f\in E \text{ and } e\neq f\}$. A *transfer scheme* is a Lipschitz continuous function $Q: X\to \mathbb{R}_+^K$ such that $Q_{ef}(x)$ represents the transfer of money from users who are using edge e to users who are using edge f at state x.

The new cost $F_e^Q(x)$ of using edge e at state x is determined as follows:

$$F_e^Q(x) = l_e(x_e) + \sum_{f\neq e}(Q_{ef}(x) - Q_{fe}(x)\cdot \frac{x_f}{x_e}). \tag{2}$$

The left-hand side of the sum $Q_{ef}(x)$ represents the transfer of money from the users of edge e to the users of edge f at state x. The right-hand side represents the transfer from the f users to the e users. Given a transfer value $Q_{fe}(x)$, the actual transfer from the f to the e users depends on x_f and x_e. The expression $x_f\cdot Q_{fe}(x)$ represents the actual amount of money collected from the f users to be transferred to the e users. This amount is evenly distributed among the e users, and hence the payoff that an e user obtains from the f users is $Q_{fe}(x) \cdot \frac{x_f}{x_e}$. We use the convention that if $Q_{fe}(x)$ is zero or if both x_e and x_f are zero there are no transfers from f to e. Note that the definition of transfer scheme above allows infinite payoffs for the cases $x_f > 0$, $x_e = 0$, and $Q_{fe}(x) > 0$. In the Appendix we shall consider the case of a bounded transfer (see also the Remark at the end of the proof of Theorem 1).

Let F^Q be the resulting game from using transfer scheme Q. That is, the payoff to a user of route i at state[5] $x \in \mathrm{int}X$ is determined as follows:

$$F_i^Q(x) = \sum_{e \in \Phi_i} F_e^Q(x).$$

It follows directly from the definition that at every state $x \in \mathrm{int}(X)$,

$$\sum_{e \in E} x_e F_e^Q(x) = \sum_{e \in E} x_e F_e(x). \tag{3}$$

This implies that the total latency time is not changed in the presence of the transfer scheme Q. Therefore, F^Q satisfies the budget balance condition at every state $x \in \mathrm{int}(X)$. When the transfer scheme of Definition 1 is independent of the current distribution of routes we shall call it *simple*. That is, a transfer scheme is simple when $Q_{ef}(x) = Q_{ef}$ is independent of $x \in X$ for every pair $(e, f) \in K$. We shall henceforth focus solely on simple transfer schemes.

By incorporating a transfer scheme we might lose some of the natural properties of the underlying congestion game. Most importantly, it might no longer be true that the resulting game F^Q admits a potential function. Having a potential function is particularly important in the dynamic setup considered by Sandholm [12], where the players update their choice of routes myopically in accordance with some evolutionary adaptive process. In this framework the existence of a potential function guarantees that for any reasonable adaptive process from any interior initial conditions the flow converges to an equilibrium. In addition, with the presence of a simple transfer scheme Q, it might hold that the actual cost of using an edge e depends on the entire flow x and not just on the congestion x_e. Therefore, another natural property to consider is that the cost of using edge e solely depends on x_e and not on the entire flow x.

[5] $\mathrm{int}(X)$ corresponds to the set of flows $x \in X$ with $x_e > 0$ for every edge e.

For a congestion game C, let $n = |E|$ be the number of distinct edges on G. We say that the game F^Q resulting from the transfer scheme Q admits a potential function if there exists a function $h : \mathbb{R}^n \to \mathbb{R}$ that is finite and differential on $\mathrm{int}(X)$ such that for every $x \in \mathrm{int}X$, and every route[6] i,

$$\frac{\partial h}{\partial x_i}(x) = F_i^Q(x).$$

Our main goal is therefore to define a transfer scheme Q that satisfies the following properties:

1. F^Q attains a unique equilibria x^* that corresponds to the optimal flow.
2. F^Q admits a potential function.
3. For every edge e and state $x \in \mathrm{int}(X)$, the payoff $F_e^Q(x)$ depends solely on x_e.
4. The matrix Q is computable in a time that is polynomial in the number of edges.

Our Main Theorem asserts that:

Theorem 1. *For every congestion game there exists a simple transfer scheme Q, that satisfies the above properties.*

Our result demonstrates the existence of a uniform transfer scheme such that the only equilibrium of the resulting game lies in the point that minimizes the latency time. The fact that we can have a transfer scheme such that the resulting game remains a potential game means that, as in Sandholm [12], for any reasonable adjustment processes of the users the unique optimal latency state is globally stable.

3 Proof of Theorem 1 and Lemma 1

Let $\{x_e^*\}_{e \in E}$ be the unique optimal flow in X. That is, it is the point that satisfies:

$$\sum_{e \in E} x_e^* l_e(x_e^*) \leq \sum_{e \in E} x_e l_e(x_e) \quad \forall x \in X.$$

[6] Note that $\frac{\partial h}{\partial x_i}(x) = \sum_{e \in \Phi_i} \frac{\partial h}{\partial x_e}(x)$.

There exists a constant tax scheme $\tau = \{\tau_e\}_{e\in E}$ such that $\tau_e \geq 0$ for every e, and the optimal latency x^* is the unique equilibrium of the game F^τ. As an example of such a tax scheme we can take $\tau_e = x_e^* l'(x_e^*)$, the marginal cost price. Let E' be the set of edges e for which $x_e^* > 0$, and assume that for every edge $e \in E \backslash E'$, it holds that $\tau_e = a$. This can be clearly obtained by increasing the tax on some unused edges.

Let $(e_1, \ldots, e_k)$ be the edges pointing out of the source node s. We note that every user must use a unique edge e_j for some $1 \leq j \leq k$. That is, for every route i there exists a unique j such that $e_j \in \Phi_i$. Since $\tau_e \geq 0$ we have that $\sum_{e\in E} \tau_e x_e^* \geq 0$. For every $h \geq 0$ define a the tax scheme $\{\tau_e^h\}_{e\in E}$ as follows:

$$\tau_e^h = \begin{cases} \tau_e - h & \text{if } e = e_j \text{ for some } 1 \leq j \leq k, \text{ and } e \in E', \\ \tau_e & \text{if } e \neq e_j \text{ for } j = 1, \ldots, k \text{ or } e \in E \setminus E'. \end{cases}$$

Note that x^* is still an equilibrium of F^{τ^h} for every $h \geq 0$. To see this note that the cost of routes that are used by a positive fraction of the population in x^* is smaller by exactly h in F^{τ^h} compared to F^τ. And, the cost of an unused route does not decrease by more than h in F^{τ^h}.

Let $h_0 \geq 0$ be such that[7]

$$\sum_{e\in E} \tau_e^{h_0} x_e^* = 0. \tag{4}$$

Therefore the tax scheme $\{\tau_e^{h_0}\}_{e\in E}$ satisfies the budget balance condition at x^*, and hence it is obtained by a money transfer among the network users. Based on this we shall define a simple transfer scheme that satisfies the required properties. For every $e \in E$, let $\kappa_e = \tau_e^{h_0}$. For every $e \in E'$, define a vector $v^e \in \mathbb{R}^n$ as follows:

$$v_f^e = \begin{cases} -\frac{1-x_e^*}{x_e^*} & \text{if } f = e, \\ 1 & \text{if } f \neq e. \end{cases}$$

Let $M \subset \mathbb{R}^n$ be the following subspace:

$$M = \{y \in \mathbb{R}^n | \sum_{e\in E} y_e x_e^* = 0 \text{ and } y_e = y_f \text{ for every } e, f \in E \setminus E'\}.$$

[7] Note that $h_0 = \sum_{e\in E} x_e^* \tau_e$.

Note that by definition the vector $\kappa = (\kappa_e)_{e \in E}$ and all the vectors v^e for $e \in E'$, lie in M. Note further that the set $\{v^e\}_{e \in E'}$ is a spanning set of M. It can be easily verified that one can write any vector $y \in M$ as a positive linear combination of the vectors $\{v^e\}_{e \in E'}$. Hence in particular there exists $\{q_e\}_{e \in E'} \subset \mathbb{R}^n_+$ such that $\sum_{e \in E'} q_e v^e = \kappa$. Let $q_e = 0$ for every $e \in E \setminus E'$. Define the matrix Q as follows: $Q_{ef} = q_f$ for every $(e,f) \in K$. We shall show that the matrix Q has the desired properties. By equation (2) the game F^Q is defined as follows:

$$\begin{aligned} F^Q_e(x) &= l_e(x_e) + \sum_{f \neq e} [q_f - (q_e \cdot \frac{x_f}{x_e})] = l_e(x_e) + \sum_{f \neq e} q_f - q_e \cdot \frac{1 - x_e}{x_e} \\ &= l_e(x_e) + \sum_f q_f - \frac{q_e}{x_e}. \end{aligned}$$

Hence $F^Q_e(x)$ is a function of x_e. By definition, F^Q has the following potential function:

$$h(x) = \sum_{e \in E} \int_0^{x_e} l_e(x_e) \mathrm{d}x + x_e(\sum_f q_f) - q_e \ln(x_e).$$

To see this note that for every route $i \in \mathcal{P}$,

$$\frac{\partial h}{\partial x_i}(x) = \sum_{e \in \Phi_i} [l_e(x_e) + \sum_f q_f - \frac{q_e}{x_e}] = \sum_{e \in \Phi_i} F^Q_e(x) = F^Q_i(x).$$

To see that x^* is an equilibrium for F^Q note that for every $e \in E$ such that $x^*_e > 0$ the payoff $F^Q_e(x^*_e)$ can be written as follows:

$$\begin{aligned} F^Q_e(x^*_e) &= l_e(x^*_e) + \sum_{f \neq e} q_f - \frac{q_e(1 - x^*_e)}{x^*_e} \\ &= l_e(x^*_e) + (\sum_f q_f v^f)_e \qquad (5) \\ &= l_e(x^*) + \kappa_e = F^\kappa_e(x^*_e). \qquad (6) \end{aligned}$$

Equality (5) follows from the definition of the vectors $\{v^f\}_{f \in E'}$. And, equality (6) follows from the definition of the non negative numbers $\{q_f\}_{f \in E}$. For every $e \in E \setminus E'$ such that $x^*_e = 0$ it follows from the definition of the game F^Q and from the fact that $q_e = 0$ that

$$F^Q_e(x^*_e) = l_e(0) + \sum_f q_f = l_e(0) + a = F^\kappa_e(x^*_e).$$

Hence the payoff $F^Q_i(x^*)$ to the i user is equal to the payoff $F^\kappa_i(x^*)$ for every $i \in \mathcal{P}$. By construction x^* is an equilibrium of F^κ and so it is an equilibrium of

F^Q. Note that the potential function $h(x)$ is convex and as such has a unique minimizer at X. Since any equilibrium of the game F^Q is a local minimizer of the potential function one can deduce that x^* is the unique equilibrium of F^Q.

It remains to show that the matrix Q can be calculated in polynomial time. First note that an optimal tax scheme τ such that F^τ has x^* as its unique equilibrium can be obtained as the solution of a linear programming problem. The number of constraints is polynomial in the size of the graph G and hence such τ can be computed in polynomial time (see Cole et al. [4] for details). The tax scheme $\kappa = \tau^{h_0}$ can obviously be obtained from τ in linear time. And finally the calculation of the transfer matrix Q that yields the appropriate transfer matrix can be obtained from κ by another linear programming problem with a number of constraints that are polynomial in G. This concludes the proof of our Main Theorem.

Remark 3.1. As noted, the payoffs in the game F^Q that results from the transfer scheme Q introduced in Theorem 1 may be unbounded on $\mathrm{int}(X)$ and infinite on the boundaries of X. In particular, when $q_e > 0$ the payoff from using edge e goes to $-\infty$ as x_e goes to 0. A natural question is whether one can redefine the game in such a way that makes both the payoffs and the potential function bounded over all X. We claim that to some extent this is indeed possible.

To do so, consider a bounded implementation mechanism for the same simple transfer matrix guaranteed by Theorem 1, under which the transfer value to every edge e is bounded by $M > 0$. That is, assume that every edge e cannot receive a money transfer that is greater than M from the other users at every state x. When such a bound is imposed, the resulting game has a bounded Lipschitz continuous payoff function that is well defined over all X, is still a potential game with a potential function that is differentiable over all X, and the tax levied on every edge e is still a function solely of the congestion on this edge. In addition, for every large enough M the optimal latency flow x^* is still a unique equilibrium of the resulting game. The only property that is violated at some states is the budget balance condition. I.e., for some states close to the boundary, money can be wasted and the tax may not be obtained by a money transfer among the network users. However, we claim that for every *reasonable* revision protocol implemented by the users the tax that is wasted during the learning process gets arbitrarily close to zero as M grows large. That is, controlling for the value of M can guarantee that the money being wasted throughout the learning process becomes negligible. We outline the above argument in the Appendix.

4 Conclusion

In this work we propose a tax scheme that is based on a money transfer among network users. We demonstrate the existence of an easily computable transfer matrix for which the resulting game has the optimal latency state as its unique equilibrium and it achieves a potential function. Moreover, similarly to marginal cost pricing, the levied tax on every edge e depends solely on the congestion x_e.

A Appendix: Bounded Implementation

Let Q be the matrix guaranteed by Theorem 1 and let $F^{Q,M}$ be the game obtained when the transfer to any edge e is bounded by $M > 0$. The new payoff from using edge e is then,

$$F_e^{Q,M}(x) = l_e(x_e) + \sum_{f \neq e} q_f - \min\{(q_e \frac{1 - x_e}{x_e}), M\}.$$

We note that the function $F_e^{Q,M}$ is Lipschitz continuous for every edge e. This in turn determines a Lipschitz continuous function for every route $i \in \mathcal{P}$.

For every edge e such that $q_e > 0$, let $x_e^M = \frac{q_e}{M+q_e}$ and let $c_e = x_e^M(q_e + M) - q_e \ln(x_e^M)$. For every such e define the function $\eta_e(x_e)$ as follows:

$$\eta_e(x_e) = \begin{cases} x_e(\sum_f q_f) - q_e \ln(x_e) & \text{if } x_e \geq x_e^M \\ x_e(\sum_{f \neq e} q_f) - x_e M + c_e & \text{if } x_e < x_e^M. \end{cases} \tag{7}$$

Note that by the choice c_e the function $\eta_e(x_e)$ is continuous and differentiable. For any edge e for which $q_e = 0$ let $\eta_e(x_e) = x_e(\sum_{f \neq e} q_f)$. The resulting differentiable potential function h^M of the game $F^{Q,M}$ may be defined by the functions $\{\eta_e(x_e)\}_{e\in E}$ as follows:

$$h^M(x) = \sum_{e\in E}[\int_0^{x_e} l_e(x)\mathrm{d}x + \eta_e(x_e)].$$

Let $X^M = \{x \in X : \forall e \in E,\ x_e \geq x_e^M\}$. Note that X^M approaches X (in the Hausdorff distance) as M grows. Since $\tau_e(x_e)$ are weakly convex it follows that the potential h^M is a weakly convex function and is strictly convex over X^M. Hence, as in Theorem 1, there exists $M_0 > 0$ such that the optimal latency x^* is the unique equilibrium of $F^{Q,M}$ for every $M \geq M_0$.

The game $F^{Q,M}$ no longer satisfies the budget balance condition. That is, if we let $\tau_e^M(x_e) = \sum_{f\neq e} q_f - \min\{(q_e\frac{1-x_e}{x_e}), M\}$ be the levied tax on edge e in the game $F^{Q,M}$, then for $x \notin X^M$ it holds that,[8]

$$V^M(x) = \sum_{e\in E} x_e\tau_e^M(x_e) > 0.$$

The budget balance condition does hold for any $x \in X^M$. Let $|\mathcal{P}| = m$. A *revision protocol* ρ is a Lipschitz function,

$$\rho : X \times \mathbb{R}^m \to \mathbb{R}_+^{m\times m}.$$

For each vector payoff $\pi \in \mathbb{R}^m$ over the routes, every state x, and all pairs of distinct routes $i, j \in \mathcal{P}$, the function $\rho_{ij}(x,\pi)$ determine the switching rate of revision from route i to route j.[9] Any revision protocol and initial state $y \in \Delta(\mathcal{P})$ determines a differential equation $z : \mathbb{R}_+ \to Y$ that describes the learning process as follows:

$$\forall i \in \mathcal{P},\ \dot{z}_i(t) = \sum_{j\in\mathcal{P}} z_j(t)\rho_{ji}(z(t), F^{Q,M}(z(t))) - z_i(t)\rho_{ij}(z(t), F^M(z(t))). \quad (8)$$

$\{z(t)\}_{t\geq 0}$ naturally defines a flow on edges $\{x(t)\}_{t\geq 0}$.

Consider,

$$\int_0^\infty V^M(x(t))\mathrm{d}t. \quad (9)$$

This expression represents the tax that is lost during the learning process. Under mild conditions on the revision protocol the flow $\{x(t)\}_{t\geq 0}$ defined by

[8]Note that $V^M(X)$ is uniformly bounded by $n\sum_{e\in E} q_e$ for every M and x.

[9]See Chapter 4.1.2 in [9] for details.

equation (8) converges to the equilibrium x^* of $F^{Q,M}$ for every $M \geq M_0$.[10] Therefore, the flow $\{x(t)\}_{t\geq 0}$ must enter X^M in a bounded time, independently of the initial conditions. Since X^M approaches X for large M, this bounded time must approach zero as M grows. Since $V^M(x)$ is bounded and zero on X^M, it must be the case that, (9) goes to zero as M increases.

References

[1] M. Beckmann, C.B. McGuire, C.B. Winsten, *Studies in the Economics of Transportation,* Yale University Press, 1956.

[2] P. Bergendorff, D.W. Hearn, M.V. Ramana, Congestion toll pricing of traffic networks, in: P.M. Pardalos, D.W. Hearn, W.W. Hager (Eds.), *Network Optimization,* Springer, 1997, pp. 51-71.

[3] R. Cole, Y. Dodis, T. Roughgarden, How much can taxes help selfish routing? *Journal of Computer and System Sciences* 72, 444-467, 2006.

[4] R. Cole, Y. Dodis, T. Roughgarden, Pricing network edges for heterogeneous selfish users. *ACM Symposium on Theory of Computing,* 521-530, 2003.

[5] L. Fleischer, K. Jain, M. Mahdian, Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games, in: *Proceedings of the 45th Symposium on Foundations of Computer Science,* 2004, pp. 277-285.

[6] D.W. Hearn, M.V. Ramana, Solving congestion toll pricing models, in: P. Marcotte, S. Nguyen (Eds.), *Equilibrium and Advanced Transportation Modeling,* Kluwer Academic Publishers, Dordrecht, 1998, pp. 109-124.

[7] A. C. Pigou, *The Economics of Welfare,* Macmillan, 1920.

[8] T. Roughgarden and E. Tardos. How bad is selfish routing? *Journal of the ACM* 49(2), 236-259, 2002.

[10] This indeed holds for many studied learning dynamics such as pairwise comparisons, projection dynamic, and better-reply dynamics, as well as whenever the potential function serves as a Lyaponov function for the learning process described in equation (7). Again, see Chapter 7.1 in [9].

[9] B. Sandholm, *Population Games and Evolutionary Dynamics,* MIT Press, 2010.

[10] B. Sandholm, Pigouvian pricing and stochastic evolutionary implementation. *Journal of Economic Theory* 132, 367-382, 2007.

[11] B. Sandholm, Negative externalities and evolutionary implementation. *Review of Economic Studies,* 72, 885-915, 2005.

[12] B. Sandholm, Evolutionary implementation and congestion pricing. *Review of Economic Studies,* 69, 667-689, 2002.

[13] M.J. Smith, The marginal cost taxation of a transportation network. *Transportation Res* Part B 13(3), 237-242, 1979.