Edge Pricing of Multicommodity Networks for Selfish Users with Elastic Demands

George Karakostas[1,⋆] and Stavros G. Kolliopoulos[2,⋆⋆]

[1] Department of Computing and Software, McMaster University
[2] Department of Informatics and Telecommunications, University of Athens

Abstract. We examine how to induce selfish heterogeneous users in a multicommodity network to reach an equilibrium that minimizes the social cost. In the absence of centralized coordination, we use the classical method of imposing appropriate taxes (tolls) on the edges of the network. We significantly generalize previous work [20,13,9] by allowing user demands to be *elastic*. In this setting the demand of a user is not fixed a priori but it is a function of the routing cost experienced, a most natural assumption in traffic and data networks.

1 Introduction

We examine a network environment where uncoordinated users, each with a specified origin-destination pair, select a path to route an amount of their respective commodity. Let f be a flow vector defined on the paths of the network, which describes a given routing according to the standard multicommodity flow conventions. The users are selfish: each wants to choose a path P that minimizes the cost $T_P(f)$. The quantity $T_P(f)$ depends typically on the latency induced on P by the aggregated flow of all users using some edge of the path.

We model the interaction of the selfish users by studying the system in the steady state captured by the classic notion of a *Wardrop equilibrium* [19]. This state is characterized by the following principle: in equilibrium, for every origin-destination pair (s_i, t_i), the cost on every used $s_i - t_i$, path is equal and less than or equal to the cost on any unused path between s_i and t_i. The Wardrop principle states that in equilibrium the users have no incentive to change their chosen route; under some minor technical assumptions the Wardrop equilibrium concept is equivalent to the Nash equilibrium in the underlying game. The literature on traffic equilibria is very large (see, e.g., [2,6,5,1]). The framework is in principle applicable both to transportation and decentralized data networks. In recent years, starting with the work of Roughgarden and Tardos [17], the latter area motivated a fruitful treatment of the topic from a computer science perspective.

⋆ (www.cas.mcmaster.ca/~gk). Research supported by an NSERC Discovery grant and MITACS.

⋆⋆ (www.di.uoa.gr/~sgk). This work was partially supported by the EU-GME (Greek Ministry of Education) EPEAEK II project no. 2.2.3 Archimedes 2: Enhancing Research Groups in Technological Institutes.

D.Z. Chen and D.T. Lee (Eds.): COCOON 2006, LNCS 4112, pp. 23–32, 2006.

The behavior of uncoordinated selfish users can incur undesirable consequences from the point of view of the system as a whole. The *social cost* function, usually defined as the total user latency, expresses this societal point of view. Since for several function families [17] one cannot hope that the uncoordinated users will reach a traffic pattern which minimizes the social cost, the system designer looks for ways to induce them to do so. A classic approach, which we follow in this paper, is to impose economic disincentives, namely put nonnegative per-unit-of-flow *taxes* (tolls) on the network edges [2,12]. The tax-related monetary cost will be, together with the load-dependent latency, a component of the cost function $T_P(f)$ experienced by the users. As in [3,20] we consider the users to be *heterogeneous*, i.e., belonging to classes that have different sensitivities towards the monetary cost. This is expressed by multiplying the monetary cost with a factor $a(i)$ for user class i. We call *optimal* the taxes inducing a user equilibrium flow which minimizes the social cost.

The existence of a vector of optimal edge taxes for heterogeneous users in multicommodity networks is not a priori obvious. It has been established for fixed demands in [20,13,9]. In this paper we significantly generalize this previous work by allowing user demands to be *elastic.* Elastic demands have been studied extensively in the traffic community (see, e.g., [10,1,12]). In this setting the demand d_i of a user class i is not fixed a priori but it is a function $D_i(u)$ of the vector u of routing costs experienced by the various user classes. Demand elasticity is natural in traffic and data networks. People may decide whether to travel based on traffic conditions. Users requesting data from a web server may stop doing so if the server is slow. Even more elaborate scenarios, such as multi-modal traffic, can be implemented via a judicious choice of the demand functions. E.g., suppose that origin-destination pairs 1 and 2 correspond to the same physical origin and destination points but to different modes of transit, such as subway and bus. There is a total amount d of traffic to be split among the two modes. The modeler could prescribe the modal split by following, e.g., the well-studied logit model [1]:

$$D_1(u) = d\frac{e^{\theta u_1 + A_1}}{e^{\theta u_1 + A_1} + e^{\theta u_2 + A_2}}, \quad D_2(u) = d - D_1(u)$$

for given negative constant θ and nonnegative constants A_1 and A_2. Here u_1 (resp. u_2) denotes the routing cost on all used paths of mode 1 (resp. 2).

For the elastic demand setting we show in Section 3 the existence of taxes that induce the selfish users to reach an equilibrium that minimizes the total latency. Note that for this result we only require that the vector $D(u)$ of the demand functions is monotone according to Definition 1. The functions $D_i(u)$ do not have to be strictly monotone (and therefore invertible) individually, and for some $i \neq j$, $D_i(u)$ can be increasing while $D_j(u)$ can be decreasing on a particular variable (as for example in the logit model mentioned above). The result is stated in Theorem 1 and constitutes the main contribution of this paper. The existence results for fixed demands in [20,13,9] follow as corollaries. Our proof is developed over several steps but its overall structure is explained at the the beginning of Section 3.1.

We emphasize that the equilibrium flow in the elastic demand setting satisfies the demand values that materialize in the same equilibrium, values that are not known a priori. This indeterminacy makes the analysis particularly challenging. On the other hand, one might argue that with high taxes, which increase the routing cost, the actual demand routed (which being elastic depends also on the taxes) will be unnaturally low. This argument does not take fully into account the generality of the demand functions $D_i(u)$ which do not even have to be decreasing; even if they do they do not have to vanish as u increases. Still it is true that the model is indifferent to potential lost benefit due to users who do not participate. Nevertheless, there are settings where users may decide not to participate without incurring any loss to either the system or themselves and these are settings we model in Section 3. Moreover in many cases the system designer chooses explicitly to regulate the effective use of a resource instead of heeding the individual welfare of selfish users. Charging drivers in order to discourage them from entering historic city cores is an example, among many others, of a social policy of this type.

A more user-friendly agenda is served by the study of a different social cost function which sums total latency and the lost benefit due to the user demand that was not routed [10,11]. This setting was recently considered in [4] from a price of anarchy [14] perspective. In this case the elasticity of the demands is specified implicitly through a function $\Gamma_i(x)$ (which is assumed nonincreasing in [4]) for every user class i. $\Gamma_i(d_i)$ determines the minimum per-user benefit extracted if d_i users from the class decide to make the trip. Hence $\Gamma_i(d_i)$ also denotes the maximum travel cost that each of the first d_i users (sorted in order of nonincreasing benefit) from class i is willing to tolerate, in order to travel. In the full version of the paper we show the existence of optimal taxes for this model. We demonstrate however that for these optimal taxes to exist, participating users must tolerate, in the worst-case, higher travel costs than those specified by their $\Gamma(\cdot)$ function.

In this extended abstract we omit many technical details. A full version of the paper is available as AdvOL-Report 2006/02 at `http://optlab.mcmaster.ca/`

2 Preliminaries

The model: Let $G = (V, E)$ be a directed network (possibly with parallel edges but with no self-loops), and a set of *users*, each with an infinitesimal amount of traffic (flow) to be routed from an origin node to a destination node of G. Moreover, each user α has a positive *tax-sensitivity* factor $a(\alpha) > 0$. We will assume that the tax-sensitivity factors for all users come from a finite set of possible positive values. We can bunch together into a single *user class* all the users with the same origin-destination pair and with the same tax-sensitivity factor; let k be the number of different such classes. We denote by $\mathcal{P}_i, a(i)$ the the flow paths that can be used by class i, and the tax-sensitivity of class i, for all $i = 1, \ldots, k$ respectively. We will also use the term 'commodity i' for class i. Set $\mathcal{P} \doteq \cup_{i=1,\ldots,k}\mathcal{P}_i$. Each edge $e \in E$ is assigned a *latency function* $l_e(f_e)$ which gives the latency experienced by any user that uses e due to congestion caused by the total flow f_e that passes through e. In other words, as in [3], we assume the additive model in which for any path $P \in \mathcal{P}$ the latency is $l_P(f) = \sum_{e \in P} l_e(f_e)$, where $f_e = \sum_{e \ni P} f_P$ and f_P is the flow through path P. If every edge is assigned a per-unit-of-flow tax $b_e \geq 0$, a selfish user in class i that uses a path $P \in \mathcal{P}_i$ experiences total cost $T_P(f)$ equal to $\sum_{e \in P} l_e(f_e) + a(i) \sum_{e \in P} b_e$ hence the name 'tax-sensitivity' for the $a(i)$'s: they quantify the importance each user assigns to the taxation of a path.

A function $g : \mathbb{R}^n \to \mathbb{R}^m$ is *positive* if $g(x) > 0$ when $x > 0$. We assume that the functions l_e are strictly increasing, i.e., $x > y \geq 0$ implies $l_e(x) > l_e(y)$, and that $l_e(0) \geq 0$. This implies that $l_e(f_e) > 0$ when $f_e > 0$, i.e., the function l_e is positive.

Definition 1. *Let $f : K \to \mathbb{R}^n$, $K \subseteq \mathbb{R}^n$. The function f is* monotone *on K if $(x - y)^T (f(x) - f(y)) \geq 0$, $\forall x \in K, y \in K$. The function f is* strictly monotone *if the previous inequality is strict when $x \neq y$.*

In what follows we will use heavily the notion of a nonlinear complementarity problem. Let $F(x) = (F_1(x), F_2(x), \ldots, F_n(x))$ be a vector-valued function from the n-dimensional space $\mathbb{R}^n$ into itself. Then the nonlinear complementarity problem of mathematical programming is to find a vector x that satisfies the following system:

$$x^T F(x) = 0, \quad x \geq 0, \quad F(x) \geq 0.$$

3 The Elastic Demand Problem

In this section the social cost function is defined as the total latency $\sum_e f_e l_e(f_e)$. We set up the problem in the appropriate mathematical programming framework and formulate the main result for this model in Theorem 1.

The traffic (or Wardrop) equilibria for a network can be described as the solutions of the following mathematical program (see [1] p. 216):

$$\begin{aligned} (T_P(f) - u_i) f_P = 0 \quad & \forall P \in \mathcal{P}_i, i = 1 \ldots k \\ T_P(f) - u_i \geq 0 \quad & \forall P \in \mathcal{P}_i, i = 1 \ldots k \\ \sum_{P \in \mathcal{P}_i} f_P - D_i(u) = 0 \quad & \forall i = 1 \ldots k \\ f, u \geq 0 \quad & \end{aligned}$$

where T_P is the cost of a user that uses path P, f_P is the flow through path P, and $u = (u_1, \ldots, u_k)$ is the vector of shortest travel times (or generalized costs) for the commodities. The first two equations model Wardrop's principle by requiring that for any origin-destination pair i the travel cost for all paths in $\mathcal{P}_i$ with nonzero flow is the same and equal to u_i. The remaining equations ensure that the demands are met and that the variables are nonnegative. Note that the formulation above is very general: every path $P \in \mathcal{P}_i$ for every commodity i has its own T_P (even if two commodities share the same path P, each may have its own T_P).

If the path cost functions T_P are positive and the $D_i(\cdot)$ functions take non-negative values, [1] shows that the system above is equivalent to the following nonlinear complementarity problem (Proposition 4.1 in [1]):

$$\begin{aligned} (T_P(f) - u_i)f_P &= 0 \quad \forall i,\ \forall P \in \mathcal{P}_i \\ T_P(f) - u_i &\geq 0 \quad \forall i,\ \forall P \in \mathcal{P}_i \\ u_i(\sum_{P\in\mathcal{P}_i} f_P - D_i(u)) &= 0 \quad \forall i \\ \sum_{P\in\mathcal{P}_i} f_P - D_i(u) &\geq 0 \quad \forall i \\ f, u &\geq 0 \end{aligned} \tag{CPE}$$

In our case the costs T_P are defined as $\sum_{e\in P} l_e(f_e) + a(i)\sum_{e\in P} b_e$, $\forall i,\ \forall P \in \mathcal{P}_i$, where b_e is the per-unit-of-flow tax for edge e, and $a(i)$ is the tax sensitivity of commodity i. In fact, it will be more convenient for us to define T_P slightly differently:

$$T_P(f) := \frac{l_P(f)}{a(i)} + \sum_{e\in P} b_e, \quad \forall i,\ \forall P \in \mathcal{P}_i.$$

The special case where $D_i(u)$ is constant for all i, was treated in [20,13,9]. The main complication in the general setting is that the minimum-latency flow $\hat{f}$ cannot be considered a priori given before some selfish routing game starts. At an equilibrium the u_i achieve some concrete value which in turn fixes the demands. These demands will then determine the corresponding minimum-latency flow $\hat{f}$. At the same time, the corresponding minimum-latency flow affects the taxes we impose and this, in turn, affects the demands. The outlined sequence of events serves only to ease the description. In fact the equilibrium parameters materialize simultaneously. We should not model the two flows (optimal and equilibrium) as a two-level mathematical program, since there is no the notion of leader-follower here, but as a complementarity problem as done in [1].

Suppose that we are given a vector u^* of generalized costs. Then the social optimum $\hat{f}^*$ for the particular demands $D_i(u^*)$ is the solution of the following mathematical program:

$$\begin{aligned} &\min \sum_{e\in E} l_e(\hat{f}_e)\hat{f}_e \qquad \text{s.t.} \\ &\sum_{P\in\mathcal{P}_i} \hat{f}_P \geq D_i(u^*) \qquad \forall i \\ &\hat{f}_e = \sum_{P\in\mathcal{P}:e\in P} \hat{f}_P \ \forall e \in E \\ &\hat{f}_P \geq 0 \qquad \forall P \end{aligned} \tag{MP}$$

Under the assumption that the functions $xl_e(x)$ are continuously differentiable and convex, it is well-known that $\hat{f}^*$ solves (MP) iff $(\hat{f}^*, \mu^*)$ solves the following pair of primal-dual linear programs (see, e.g., [8, pp. 9–13]):

$$\min \sum_{e\in E} \left(l_e(\hat{f}_e^*) + \hat{f}_e^* \frac{\partial l_e}{\partial f_e}(\hat{f}_e^*) \right) \hat{f}_e \text{ s.t.} \tag{LP2}$$

$$\sum_{P\in\mathcal{P}_i} \hat{f}_P \geq D_i(u^*), \quad \forall i$$

$$\hat{f}_e = \sum_{P\in\mathcal{P}:e\in P} \hat{f}_P, \quad \forall e \in E$$

$$\hat{f}_P \geq 0, \quad \forall P$$

$$\max \sum_i D_i(u^*)\mu_i \quad \text{s.t.} \tag{DP2}$$

$$\mu_i \leq \sum_{e\in P} \left(l_e(\hat{f}_e^*) + \hat{f}_e^* \frac{\partial l_e}{\partial f_e}(\hat{f}_e^*) \right) \forall i, P\in\mathcal{P}_i$$

$$\mu_i \geq 0 \quad \forall i$$

Let the functions $D_i(u)$ be bounded and set $K_1 := \max_i \max_{u\geq 0}\{D_i(u)\} + 1$. Then if n denotes $|V|$ the solutions $\hat{f}^*$, μ^* of (LP2), (DP2) are upper bounded as follows $\hat{f}_P^* \leq D_i(u^*) < K_1,\ \forall P \in \mathcal{P}_i$ $\mu_i \leq \sum_{e\in P} \left(l_e(\hat{f}_e^*) + \hat{f}_e^* \frac{\partial l_e}{\partial f_e}(\hat{f}_e^*) \right) < n \cdot \max_{e\in E} \max_{0\leq x\leq k\cdot K_1} \{l_e(x) + x\frac{\partial l_e}{\partial f_e}(x)\},\ \forall i$. It is important to note that these upper bounds are *independent of* u^*.

We wish to find a tax vector b that will steer the edge flow solution of (CPE) towards $\hat{f}$. Similarly to [13] we add this requirement as a constraint to (CPE): for every edge e we require that $f_e \leq \hat{f}_e$. By adding also the Karush-Kuhn-Tucker conditions for (MP) we obtain the following complementarity problem:

$$\begin{aligned}
f_P(T_P(f) - u_i) &= 0,\ \forall i, P & T_P(f) &\geq u_i,\ \forall i, P \\
u_i(\sum_{P\in\mathcal{P}_i} f_P - D_i(u)) &= 0,\ \forall i & \sum_{P\in\mathcal{P}_i} f_P &\geq D_i(u),\ \forall i \\
b_e(f_e - \hat{f}_e) &= 0,\ \forall e & f_e &\leq \hat{f}_e,\ \forall e \\
(\sum_{e\in P}(l_e(\hat{f}_e) + \hat{f}_e\frac{\partial l_e}{\partial f_e}(\hat{f}_e)) - \mu_i)\hat{f}_P &= 0, \forall i, P & \sum_{e\in P}(l_e(\hat{f}_e) + \hat{f}_e\frac{\partial l_e}{\partial f_e}(\hat{f}_e)) &\geq \mu_i, \forall i, P \\
\mu_i(\sum_{P\in\mathcal{P}_i} \hat{f}_P - D_i(u)) &= 0,\ \forall i & \sum_{P\in\mathcal{P}_i} \hat{f}_P &\geq D_i(u),\ \forall i \\
& & f_P, b_e, u_i, \hat{f}_P, \mu_i &\geq 0,\ \forall P, e, i
\end{aligned} \tag{GENERAL CP}$$

where $f_e = \sum_{P\ni e} f_P, \hat{f}_e = \sum_{P\ni e} \hat{f}_P$.

The users should be steered towards $\hat{f}$ without being conscious of the constraints $f_e \leq \hat{f}_e$; the latter should be felt only implicitly, i.e., through the corresponding tax b_e. Our main result is expressed in the following theorem. For convenience, we view $D_i(u)$ as the ith coordinate of a vector-valued function $D : \mathbb{R}^k \to \mathbb{R}^k$.

Theorem 1. *Consider the selfish routing game with the latency function seen by the users in class i being $T_P(f) := \sum_{e\in P} l_e(f_e) + a(i)\sum_{e\in P} b_e,\ \forall i,\ \forall P \in \mathcal{P}_i$. If (i) for every edge $e \in E$, $l_e(\cdot)$ is a strictly increasing continuous function with $l_e(0) \geq 0$ such that $xl_e(x)$ is convex and continuously differentiable and (ii) D_i*

are continuous functions bounded from above for all i such that $D(\cdot)$ is positive and $-D(\cdot)$ is monotone then there is a vector of per-unit taxes $b \in \mathbb{R}_+^{|E|}$ such that, if $\bar{f}$ is a traffic equilibrium for this game, $\bar{f}_e = \hat{f}_e$, $\forall e \in E$. Therefore $\bar{f}$ minimizes the social cost $\sum_{e \in E} f_e l_e(f_e)$.

3.1 Proof of the Main Theorem

The structure of our proof for Theorem 1 is as follows. First we give two basic Lemmata 1 and 2. We then argue that the two lemmata together with a proof that a solution to (GENERAL CP) exists imply Theorem 1. We establish that such a solution for (GENERAL CP) exists in Theorem 2. The proof of the latter theorem uses the fixed-point method of [18] and arguments from linear programming duality.

The following result of [1], can be easily extended to our case:

Lemma 1 (Theorem 6.2 in [1]). *Assume that the $l_e(\cdot)$ functions are strictly increasing for all $e \in E$, $D(\cdot)$ is positive and $-D(\cdot)$ is monotone. Then if more than one solutions (f, u) exist for* (CPE)*, u is unique and f induces a unique edge flow.*

Lemma 2. *Let $(f^*, b^*, u^*, \hat{f}^*, \mu^*)$ be any solution of* (GENERAL CP)*. Then $\sum_{P \in \mathcal{P}_i} f_P^* = D_i(u^*)$, $\forall i$ and $f_e^* = \hat{f}_e^*$, $\forall e \in E$.*

Let $(f^*, b^*, u^*, \hat{f}^*, \mu^*)$ be a hypothetical solution to (GENERAL CP). Then $\hat{f}^*$ is a minimum latency flow solution for the demand vector $D(u^*)$. Moreover $f_e^* \leq \hat{f}_e^*$, $\forall e \in E$. After setting $b = b^*$ in (CPE), Lemma 1 implies that any solution $(\bar{f}, \bar{u})$ to (CPE) would satisfy $\bar{f}_e = f_e^*$ and $\bar{u} = u^*$. Therefore $\bar{f}_e \leq \hat{f}_e^*$, $\forall e \in E$. Under the existing assumptions on $l_e(\cdot)$, We can show (proof omitted) that any *equilibrium flow* $\bar{f}$ for the selfish routing game where the users are conscious of the modified latency $T_P(f) := \frac{l_P(f)}{a(i)} + \sum_{e \in P} b_e^*$, $\forall i$, $\forall P \in \mathcal{P}_i$, is a minimum-latency solution for *the demand vector reached in the same equilibrium.* Therefore the b^* vector would be the vector of the optimal taxes. To complete the proof of Theorem 1 we will now show the existence of (at least) one solution to (GENERAL CP):

Theorem 2. *If $f_e l_e(f_e)$ are continuous, convex, strictly monotone functions for all $e \in E$, and $D_i(\cdot)$ are nonnegative continuous functions bounded from above for all i, then* (GENERAL CP) *has a solution.*

Proof. We provide only a sketch of the proof. See the full paper for details. (GENERAL CP) is equivalent in terms of solutions to the complementarity problem (GENERAL CP′) (proof omitted). The only difference between (GENERAL CP) and (GENERAL CP′) is that $T_P(f) = \sum_{e \in P}(\frac{l_e(f_e)}{a(i)} + b_e)$ is replaced by $T_P(\hat{f}) = \sum_{e \in P}(\frac{l_e(\hat{f}_e)}{a(i)} + b_e)$ in the first two constraints.

To show that (GENERAL CP′) has a solution, we will follow a classic proof method by Todd [18] that reduces the solution of a complementarity problem to a Brouwer fixed-point problem. In what follows, let $[x]^+ := \max\{0, x\}$. If $\phi : \mathbb{R}^n \to \mathbb{R}^n$ with $\phi(x) = (\phi_1(x), \phi_2(x), \ldots, \phi_n(x))$ is a function with components $\phi_1, \ldots, \phi_n$ defined as

$$\phi_i(x) = [x_i - F_i(x)]^+,$$

then $\hat{x}$ is a fixed point to ϕ iff $\hat{x}$ solves the complementarity problem $x^T F(x) = 0, F(x) \geq 0, x \geq 0$. Following [1], we will restrict ϕ to a large cube with an artificial boundary, and show that the fixed points of this restricted version of ϕ are fixed points of the original ϕ by showing that no such fixed point falls on the boundary of the cube.

Note that for (GENERAL CP) $x = (f, u, b, \hat{f}, \mu)$. We start by defining the cube which will contain x. Let $K_{\hat{f}} := \max_i \max_{u \geq 0}\{D_i(u)\}+1$, $K_f := K_{\hat{f}}$, $K_\mu := n \cdot \max_{e\in E} \max_{0\leq x \leq k\cdot K_{\hat{f}}}\{l_e(x) + x\frac{\partial l_e}{\partial f_e}(x)\}$. Let S be the maximum possible entry of the inverse of any ± 1 matrix of dimension at most $(k+m)\times(k+m)$, where m denotes $|E|$ (note that S depends only on $(k+m)$.) Also, let $a_{max} = \max_i\{1/a(i)\}$ and $l_{max} = \max_e\{l_e(k \cdot K_f)\}$. Then define $K_b := (k+m)Sma_{max}l_{max} + 1$, $K_u := n \cdot \left(\max_{e\in E, i\in\{1,\ldots,k\}}\left\{\frac{l_e(k\cdot K_f)}{a(i)}\right\} + K_b\right) + 1$.

We allow x to take values from the cube $\{0 \leq f_P \leq K_f, P \in \mathcal{P}\}, \{0 \leq u_i \leq K_u, i = 1, \ldots k\}, \{0 \leq b_e \leq K_b, e \in E\}, \{0 \leq \hat{f}_P \leq K_{\hat{f}}, P \in \mathcal{P}\}, \{0 \leq \mu_i \leq K_\mu, i = 1, \ldots k\}$. We define $\phi = (\{\phi_P : P \in \mathcal{P}\}, \{\phi_i : i = 1, \ldots, k\}, \{\phi_e : e \in E\}, \{\phi_{\hat{P}} : P \in \mathcal{P}\}, \{\phi_{\hat{i}} : i = 1, \ldots k\})$ with $|\mathcal{P}| + k + m + |\mathcal{P}| + k$ components as follows:

$$\phi_P(f, u, b, \hat{f}, \mu) = \min\{K_f, [f_P + u_i - T_P(\hat{f})]^+\} \qquad \forall i, \forall P \in \mathcal{P}_i$$
$$\phi_i(f, u, b, \hat{f}, \mu) = \min\{K_u, [u_i + D_i(u) - \sum_{P\in\mathcal{P}_i} f_P]^+\} \qquad i = 1, \ldots, k$$
$$\phi_e(f, u, b, \hat{f}, \mu) = \min\{K_b, [b_e + f_e - \hat{f}_e]^+\} \qquad \forall e \in E$$
$$\phi_{\hat{P}}(f, u, b, \hat{f}, \mu) = \min\{K_{\hat{f}}, [\hat{f}_P + \mu_i - \sum_{e\in P}\frac{\partial l_e}{\partial f_e}(\hat{f}_e)]^+\} \qquad \forall i, \forall P \in \mathcal{P}_i$$
$$\phi_{\hat{i}}(f, u, b, \hat{f}, \mu) = \min\{K_{\hat{i}}, [\mu_i + D_i(u) - \sum_{P\in\mathcal{P}_i} \hat{f}_P]^+\} \qquad i = 1, \ldots, k$$

where $f_e = \sum_{P\ni e} f_P$, $\hat{f}_e = \sum_{P\ni e} \hat{f}_P$. By Brouwer's fixed-point theorem, there is a fixed point x^* in the cube defined above, i.e., $x^* = \phi(x^*)$. In particular we have that $f^*_P = \phi_P(x^*), u^*_i = \phi_i(x^*), b^*_e = \phi_e(x^*), \hat{f}^*_P = \phi_{\hat{P}}(x^*), \mu^*_i = \phi_{\hat{i}}(x^*)$ for all $P, \hat{P} \in \mathcal{P}, i = 1, \ldots, k$, $e \in E$.

Following the proof of Theorem 5.3 of [1] we can show that

$$\hat{f}^*_P = [\hat{f}^*_P + \mu^*_i - \sum_{e\in P}(l_e(\hat{f}^*_e) + \hat{f}^*_e\frac{\partial l_e}{\partial f_e}(\hat{f}^*_e))]^+, \ \forall P \quad \mu^*_i = [\mu^*_i + D_i(u^*) - \sum_{P\in\mathcal{P}_i}\hat{f}^*_P]^+, \ \forall i$$
$$f^*_P = [f^*_P + u^*_i - T_P(\hat{f}^*)]^+, \ \forall P. \qquad (1)$$

Note that this implies that $(\hat{f}^*, \mu^*)$ satisfy the KKT conditions of (MP) for u^*. Here we prove only (1) (the other two are proven in a similar way). Let $f^*_P = K_f$

for some $i, P \in \mathcal{P}_i$ (if $f_P^* < K_f$ then (1) holds). Then $\sum_{P\in\mathcal{P}_i} f_P^* > D_i(u^*)$, which implies that $u_i^* + D_i(u^*) - \sum_{P\in\mathcal{P}_i} f_P^* < u_i^*$, and therefore by the definition of ϕ_i we have that $u_i^* = 0$. Since $T_P(\hat{f}^*) \geq 0$, this implies that $f_P^* \geq f_P^* + u_i^* - T_P(\hat{f}^*)$. If $T_P(\hat{f}^*) > 0$, the definition of ϕ_P implies that $f_P^* = 0$, a contradiction. Hence it must be the case that $T_P(\hat{f}^*) = 0$, which in turn implies (1).

If there are $i, P \in \mathcal{P}_i$ such that $f_P^* > 0$, then (1) implies that $u_i^* = T_P(\hat{f}^*) = \sum_{e\in P} \frac{l_e(\hat{f}_e^*)}{a(i)} + \sum_{e\in P} b_e^*$. In this case we have that $u_i^* < K_u$, because $u_i^* = K_u \Rightarrow \sum_{e\in P} \frac{l_e(\hat{f}_e^*)}{a(i)} + \sum_{e\in P} b_e^* = n \cdot \left(\max_{e\in E, i\in\{1,\ldots,k\}} \left\{\frac{l_e(K_f)}{a(i)}\right\} + K_b\right) + 1$ which is a contradiction since $b_e^* \leq K_b$. On the other hand, if there are $i, P \in \mathcal{P}_i$ such that $f_P^* = 0$, then (1) implies that $u_i^* \leq T_P(\hat{f}^*)$. Again $u_i^* < K_u$, because if $u_i^* = K_u$ we arrive at the same contradiction. Hence we have that

$$u_i^* = [u_i^* + D_i(u^*) - \sum_{P\in\mathcal{P}_i} f_P^*]^+, \ \forall i. \tag{2}$$

Next, we consider the following primal-dual pair of linear programs:

$$\begin{array}{lll|lll}
\min \sum_i \sum_{P\in\mathcal{P}_i} f_P \frac{l_P(\hat{f}^*)}{a(i)} & \text{s.t.} & (\text{LP*}) & \max \sum_i D_i(u^*) u_i - \sum_{e\in E} \hat{f}_e^* b_e \ \text{s.t.} & & (\text{DP*}) \\
\sum_{P\in\mathcal{P}_i} f_P \geq D_i(u^*) & i = 1,\ldots,k & & u_i \leq \frac{l_P(\hat{f}^*)}{a(i)} + \sum_{e\in P} b_e & \forall i, \forall P \in \mathcal{P}_i & \\
f_e = \sum_{P\in\mathcal{P}: e\in P} f_P & \forall e \in E & & b_e, u_i \geq 0 & \forall e \in E, \forall i & \\
f_e \leq \hat{f}_e^* & \forall e \in E & & & & \\
f_P \geq 0 & \forall P & & & &
\end{array}$$

From the above, it is clear that $\hat{f}^*$ is a feasible solution for (LP*), and (u^*, b^*) is a feasible solution for (DP*). Moreover, since the objective function of (LP*) is bounded from below by 0, (DP*) has at least one bounded optimal solution as well. There is an optimal solution $(\hat{u}, \hat{b})$ of (DP*) such that all the $\hat{b}_e$'s are suitably upper bounded:

Lemma 3 (folklore). *There is an optimal solution* $(\hat{u}, \hat{b})$ *of* (DP*) *such that* $\hat{b}_e \leq K_b - 1, \ \forall e \in E$.

Let $\hat{f}$ be the optimal primal solution of (LP*) that corresponds to the optimal dual solution $(\hat{u}, \hat{b})$ of (DP*). Exploiting the fact that $(\hat{f}, \hat{u}, \hat{b})$ is a *saddle point* for the Lagrangian (see e.g. [16]) of (LP*)-(DP*) we can show (derivation omitted) that

$$b_e^* = [b_e^* + f_e^* - \hat{f}_e^*]^+, \ \forall e \in E. \tag{3}$$

Equations (1),(2),(3) imply that $(f^*, u^*, b^*, \hat{f}^*, \mu^*)$ is indeed a solution of (GENERAL CP′), and therefore a solution to (GENERAL CP). The proof of Theorem 2 is complete.

References

1. H. Z. Aashtiani and T. L. Magnanti. Equilibria on a congested transportation network. *SIAM J. Algebraic and Discrete Methods*, 2:213–226, 1981.
2. M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
3. R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 35th ACM STOC*, pp. 521–530, 2003.
4. R. Cole, Y. Dodis, and T. Roughgarden. Bottleneck links, variable demand and the tragedy of the commons. In *Proc. 17th ACM-SIAM SODA*, 668 - 677, 2006.
5. S. C. Dafermos. Traffic equilibria and variational inequalities. *Transportation Science* 14, pp. 42–54, 1980.
6. S. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *J. Research National Bureau of Standards, Series B*, 73B:91–118, 1969.
7. S. Dafermos. Toll patterns for multiclass-user transportation networks. *Transportation Science*, 7:211–223, 1973.
8. F. Facchinei and J.-S. Pang. *Finite-Dimensional Variational Inequalities and Complementarity Problems, Vol 1.* Springer-Verlag, Berlin, 2003.
9. L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proc. 45th IEEE Symposium on Foundations of Computer Science*, 277–285, 2004.
10. N. H. Gartner. Optimal traffic assignment with elastic demands: a review. Part I: analysis framework. *Transportation Science*, 14:174–191, 1980.
11. N. H. Gartner. Optimal traffic assignment with elastic demands: a review. Part II: algorithmic approaches. *Transportation Science*, 14:192–208, 1980.
12. D. W. Hearn and M. B. Yildirim. A toll pricing framework for traffic assignment problems with elastic demand. In M. Gendreau and P. Marcotte, editors, *Transportation and Network Analysis: Current Trends. Miscellanea in honor of Michael Florian.* Kluwer Academic Publishers, 2002.
13. G. Karakostas and S. G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proc. 45th IEEE Symposium on Foundations of Computer Science*, 268–276, 2004.
14. E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proc. 16th Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
15. A. Nagurney. A multiclass, multicriteria traffic network equilibrium model. *Mathematical and Computer Modelling*, 32:393–411, 2000.
16. R. T. Rockafellar. *Convex Analysis.* Princeton University Press, 1970.
17. T. Roughgarden and É. Tardos. How bad is selfish routing? *J. ACM*, 49:236–259, 2002.
18. M. J. Todd. The computation of fixed points and applications. *Lecture Notes in Economics and Mathematical Systems*, 124, Springer-Verlag, 1976.
19. J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers, Part II*, 1:325–378, 1952.
20. H. Yang and H.-J. Huang. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Research Part B*, 38:1–15, 2004.