Competitive Routing in Multi-User Communication Networks

Ariel Orda *Raphael Rom* *Nahum Shimkin*

Faculty of Electrical Engineering
Technion
Haifa, Israel

Published in the IEEE/ACM Transactions on Networking, Vol. 1, pp. 510-521, October 1993.

Abstract

We consider a communication network shared by several selfish users. Each user seeks to optimize its own performance by controlling the routing of its given flow demand, giving rise to a non-cooperative game. We investigate the Nash equilibrium of such systems. For a two-node multiple-links system, uniqueness of the Nash equilibrium is proved under reasonable convexity conditions. It is shown that this Nash equilibrium point possesses interesting monotonicity properties. For general networks, these convexity conditions are not sufficient for guaranteeing uniqueness, and a counter example is presented. Nonetheless, uniqueness of the Nash equilibrium for general topologies is established under various assumptions.

1 Introduction

Traditional computer networks were designed with a single administrative domain in mind. That is, the network is designed and operated as a single entity with a single control objective. A single control objective does not mean that control is centralized but it means that users are essentially passive and would quite often reduce their own performance for the good of the entire network. For example, many advanced network routing protocols attempt to optimize average *network* delay.

In modern networking, a single administration is no longer a valid assumption. Internetworking, for example [1, 2], is a coalition of networks each belonging to a different administration sharing gateways and sometimes internal network resources. Another example is a set of different companies in the same neighborhood using wireless local area networks and sharing the same portion of the spectrum. It is evident therefore that single control objectives cannot provide solutions in more modern environments.

An alternative approach is to view the network as a resource shared by a group of active users. Users may have completely different measures of performance and satisfaction and completely different demands which at times may be contradictory. One possible way of managing such a network is to let the individual users compete with one another in a way that allows each of them to reach its (subjective) optimal working state. In such an environment users change their behavior based on the state of the network. The change in behavior of one user is likely to cause changes in other users' behavior resulting in a dynamic system. Several questions need be asked in this context such as whether there exists an equilibrium point of operation such that no user would find it beneficial to change its working parameters (i.e., a Nash equilibrium), whether such an equilibrium point is unique, and whether the dynamic system actually converges to the equilibrium point. These are fundamental questions in game theory [3, 4].

Existing networks have avoided dealing with the above mentioned issues. The Internet[5, 6], for example, uses a routing protocol that is based on topological considerations alone without regard to other optimization criteria. Recently, policy-based routing has been deliberated[7], but this approach does not accommodate general dynamics or individual user characteristics. This paper addresses the most basic networking problem–the routing problem– from a game theoretical standpoint, contributing to the understanding of the dynamics of modern networks.

Most previous work in applying game-theory and economic techniques to computer networks deals with flow control, whereas routing is assumed to be given or centrally managed. Kurose et al. [8] and Ferguson et al. [9] use economic pricing tools in order to deal with network resource-allocation problems. More related to the present paper are the works of Bovopoulos and Lazar[10] and Hsiao and Lazar[11], where the Nash equilibrium is examined in the context of multi-controller network flow control problems. In particular, in [10] uniqueness of the Nash equilibrium is established for a BCMP-type queuing network, under a power-based criterion. A thorough study of competitive multi-class flow control to a single node (server) may be found in [12], [13] and the references cited therein. Yet another flow control analysis is given by Shenker[14] who considers an internetwork gateway problem. The approach there is to assume that users operate selfishly and it is the task of the designer to set the gateway parameters such that overall network resources are used as efficiently as possible. In another paper Shenker [15] discusses at some length game theoretical issues related to networking problems.

The routing standpoint has been scarcely considered in the literature. Lee and Cohen[16] consider a set of parallel M/M/c queues (which, in our context, could represent parallel links connecting two nodes). The users control the amount of flow through each queue (making this a routing problem) and consider a linear combination of the average queue length and customer delay as performance criteria. They show that in such a setting at most one Nash equilibrium exists with the property that each user ships positive flow through each queue. This result is employed in order to establish the uniqueness of the Nash equilibrium for the case of identical queues. They fall short, however, of establishing uniqueness in the general case. Indeed, for non-identical queues (or communications links) there do exist in general equilibrium points where some users find it optimal to use only a subset of all available links (a simple example is given in Section 2.3). This uniqueness problem is completely resolved in the present work. Another game theoretic treatment of a routing problem was considered in [17], where existence and uniqueness of the Nash equilibrium was shown for the special case of two exponential servers working in parallel, and two users, each employing a different cost function (namely, average delay and blocking probability). Noncooperative games in the context of routing were studied also in the area of transportation networks. A fundamental result due to Dafermos and Sparrow [18] shows that the game problem can be solved via a standard network optimization problem, by a simple transform of the cost function. Nonetheless, the "user" considered in the context of transportation networks is one that controls just an infinitesimally small portion of the network flow (e.g., a car on the road), whereas we are concerned with users that control non-negligible portions of flow.

In general, uniqueness (or even existence) of the Nash equilibrium is not guaranteed. A result due to Rosen[19] defines conditions for existence, uniqueness, and stability of Nash equilibria in convex games. This result is the basis of several subsequent works, such as that of Bovopoulos and Lazar[10]. Unfortunately, the specific requirements of Rosen's uniqueness result (diagonal strict convexity) are not generally satisfied in the problem posed in this paper (see Section 3.2). Another result is that of Li and Basar[20] who describe a distributed (uncoordinated) environment in which the users play the game. They use contraction conditions to guarantee uniqueness and stability, but these conditions are too complicated to be verified in the problem we are interested in.

The work presented here deals with routing, meaning that the network topology is known, each user knows its individual throughput demands, each user can measure the load on the network links, and routing is selected by each user so as to optimize a certain selfish criterion such as its own average delay. This paper addresses mainly the uniqueness problem, and also investigates properties of the flow traffic at the Nash equilibrium point. Also, the stability issue (i.e., convergence to the Nash equilibrium point) is addressed briefly. We note that the routing problem in a network with a single common (convex) objective can be solved in a fairly standard way using convex programming techniques. Centralized and distributed algorithms of that type have been described in the literature (e.g., [21, 22]). When the objective function is convex but not common to all users the setting becomes that of a convex game. As mentioned above and shown in the sequel, uniqueness of the Nash point cannot be derived directly from available results of convex game theory (such as [19]). This leads us to exploit the specific structure of our problem in order to prove uniqueness.

We analyze the routing problem in two phases. First we consider a case of two nodes connected by a set of parallel links, similar to the setting considered by Lee and Cohen[16], except that ours allows more general functions and not just those resulting from a queuing framework. This model is presented in Section 2. Existence and uniqueness of the Nash equilibrium is established under fairly weak convexity properties, which are satisfied by standard network cost functions. We also derive characteristics of the equilibrium. In Section 3 we extend the discussion to a general network. We show through an example that the above weak convexity conditions are not enough in order to guarantee the uniqueness of the Nash equilibrium. Nonetheless, we prove uniqueness of the Nash equilibrium in several cases. First, sufficient conditions for uniqueness are derived based on [19], and their applicability is discussed. A second result establishes uniqueness for the case of symmetric users. Finally, we obtain a result similar to that of [16], but for the general network environment. Several conclusions and open problems are discussed in Section 4.

2 A Network of Parallel Links

2.1 Model and Problem Formulation

We are given a set $\mathcal{I} = \{1, 2, \ldots, I\}$ of users, which share a set of parallel communication links $\mathcal{L} = \{1, 2, \ldots, L\}$ interconnecting a common source node to a common destination node. We assume that users are selfish and do not cooperate in managing the communication links. Each user $i \in \mathcal{I}$ has a throughput demand, which is some ergodic process with average rate of r^i. Without loss of generality, we assume that $r^1 \leq r^2 \leq \cdots \leq r^I$. A user ships its demand by splitting it through the communication links $\mathcal{L}$. A user is able to decide (at any time) how its demand is split among the links, i.e., user i decides what fraction of r^i should be sent through each link. We denote by f_l^i the expected flow of user $i \in \mathcal{I}$

on link $l \in \mathcal{L}$. Thus, user i can fix any value for f_l^i, as long as $f_l^i \geq 0$ (nonnegativity constraint) and $\sum_{l \in \mathcal{L}} f_l^i = r^i$ (demand constraint). Turning our attention to a link $l \in \mathcal{L}$, let f_l be the total flow on that link i.e., $f_l = \sum_{i \in \mathcal{I}} f_l^i$; also, denote by $\mathbf{f_l}$ the vector of all user flows on link $l \in \mathcal{L}$, i.e., $\mathbf{f_l} = (f_l^1, f_l^2, \ldots, f_l^I)$. The *flow configuration* $\mathbf{f^i}$ of user i is the vector $\mathbf{f^i} = (f_1^i, f_2^i, \ldots, f_L^i)$. The *system flow configuration* $\mathbf{f}$ is the vector of all user flow configurations, $\mathbf{f} = (\mathbf{f^1}, \mathbf{f^2}, \ldots, \mathbf{f^I})$. We say that a user flow configuration is *feasible* if its components obey the nonnegativity and demand constraints and we denote by $\mathbf{F^i}$ the set of all feasible $\mathbf{f^i}$'s. Similarly, a system configuration is feasible if it is composed of feasible user flow configurations and we denote by $\mathbf{F}$ the set of all feasible $\mathbf{f}$'s.

The performance measure of a user $i \in \mathcal{I}$ is given by a cost function $J^i(\mathbf{f})$. The aim of each user is to minimize its cost. Since the cost functions depend on the flow configuration of all users, it turns out that the optimal decision of each user depends on the decisions made by other users, and since users are selfish, we are faced with a noncooperative game[3, 4]. Thus, we are interested in the Nash solution of the game. In other words, we seek a system flow configuration such that no user finds it beneficial to change its flow on any link. Formally, a feasible system flow configuration $\tilde{\mathbf{f}} = (\tilde{\mathbf{f}}^\mathbf{1}, \tilde{\mathbf{f}}^\mathbf{2}, \ldots, \tilde{\mathbf{f}}^\mathbf{I})$ is a Nash Equilibrium Point (NEP) if, for all $i \in \mathcal{I}$, the following condition holds:

$$J^i(\tilde{\mathbf{f}}) = J^i(\tilde{\mathbf{f}}^\mathbf{1}, \ldots, \tilde{\mathbf{f}}^\mathbf{i-1}, \tilde{\mathbf{f}}^\mathbf{i}, \tilde{\mathbf{f}}^\mathbf{i+1}, \ldots, \tilde{\mathbf{f}}^\mathbf{I}) = \min_{\mathbf{f^i} \in \mathbf{F^i}} J^i(\tilde{\mathbf{f}}^\mathbf{1}, \ldots, \tilde{\mathbf{f}}^\mathbf{i-1}, \mathbf{f^i}, \tilde{\mathbf{f}}^\mathbf{i+1}, \ldots, \tilde{\mathbf{f}}^\mathbf{I}) \quad (1)$$

We remark that the NEP concept is of special importance from a dynamic standpoint: in a practical scenario, a user changes its flow repeatedly, in response to the varying load conditions. The stability points of such systems are exactly those in which no user finds it beneficial to change its flow, i.e., the NEPs. An interesting question is whether the system indeed converges to an NEP.

The following general assumptions on the cost function J^i of each user are imposed throughout the paper (some additional structural assumptions will be considered in the sequel):

G1 J^i is the sum of link cost functions i.e., $J^i(\mathbf{f}) = \sum_{l \in \mathcal{L}} J_l^i(\mathbf{f_l})$. Each J_l^i satisfies:

G2 $J_l^i : [0, \infty)^I \rightarrow [0, \infty]$, a continuous function.

G3 J_l^i is convex in f_l^i.

G4 Wherever finite, J_l^i is continuously differentiable in f_l^i. We denote: $K_l^i = \frac{\partial J_l^i}{\partial f_l^i}$.

Note the inclusion of $+\infty$ in the range of J_l^i, which is useful to incorporate implicitly and compactly additional constraints such as link capacities (as in the type-C functions below). We emphasize that only "gradual" constraints, where the cost function increases continuously to infinity, may be incorporated; the addition of other ("abrupt") constraints would involve modification of the Kuhn-Tucker conditions below, and is not included in our analysis.

An additional assumption concerning the entire model data is:

G5 For every system flow configuration $\mathbf{f}$, if not all costs are finite then at least one user with infinite cost ($J^i(\mathbf{f}) = \infty$) can change its own flow configuration to make its cost finite.

The last assumption immediately implies that in any NEP the costs of all users must be finite. This assumption can be simply re-stated (in the present two-node network) in the typical case where cost functions take infinite values only due to link capacity constraints (i.e., the cost for a user is infinite if and only if it uses a link on which the total flow exceeds the link's capacity). Assumption G5 is then equivalent to the sum of link capacities being greater than the sum of the users' demands.

Under the above assumptions, the routing game is equivalent to a convex game in the sense of [19], and thus the existence of an NEP is guaranteed (Theorem 1 in [19]). Since some semantic differences do exist, we briefly outline the proof. Consider the point-to-set mapping $\mathbf{f} \in \mathbf{F} \to \Gamma(\mathbf{f}) \subset \mathbf{F}$, defined by

$$\Gamma(\mathbf{f}) = \{\tilde{\mathbf{f}} \in \mathbf{F} : \tilde{\mathbf{f}}^{\mathbf{i}} \in \arg\min_{\mathbf{g}^{\mathbf{i}} \in \mathbf{F}^{\mathbf{i}}} J^i(\mathbf{f}^{\mathbf{1}}, \ldots, \mathbf{g}^{\mathbf{i}}, \ldots, \mathbf{f}^{\mathbf{I}})\} .$$

Then Γ is an upper semicontinuous mapping (by the continuity assumption G2) which maps each point of the convex compact set $\mathbf{F}$ into a closed (by G2) convex (by G3) subset of $\mathbf{F}$. Then by the Kakutani fixed point theorem there exists a fixed point $\mathbf{f} \in \Gamma(\mathbf{f})$, and such a point is easily seen to be a Nash equilibrium.

It also follows from our assumptions that the minimization in (1) is equivalent to the following Kuhn-Tucker conditions: for every $i \in \mathcal{I}$ there exist a (Lagrange multiplier) λ^i such that, for every link $l \in \mathcal{L}$,

$$f_l^i > 0 \to K_l^i(\mathbf{f_l}) = \lambda^i \tag{2}$$

$$f_l^i = 0 \to K_l^i(\mathbf{f_l}) \geq \lambda^i \tag{3}$$

In other words, the Kuhn-Tucker conditions as stated above constitute necessary and sufficient conditions for a feasible system flow configuration to be an NEP.

Given the existence of an NEP, we investigate its uniqueness as well as other interesting properties. We also discuss a simple dynamic system and prove its convergence.

We shall mainly consider cost functions that comply with the following assumptions:

A1 J_l^i is a function of two arguments, namely user i's flow on link l and the total flow on that link. In other words: $J_l^i(\mathbf{f_l}) = \bar{J}_l^i(f_l^i, f_l)$.

A2 $\bar{J}_l^i$ is increasing in each of its two arguments.

A3 Note that $K_l^i = K_l^i(f_l^i, f_l)$ is now a function of two arguments. We assume that, wherever J_l^i is finite, $K_l^i(f_l^i, f_l)$ is strictly increasing in each of its two arguments.

Functions that comply with the above assumptions shall be referred to as *type-A* functions. We point out that the above assumptions encompass a large family of interesting cost functions, some of which are described in the sequel. In particular, we note that the first assumption relates the performance of a user on a link to both its amount of flow on that link, which measures its "investment" on that link, and to the total amount of flow through the link, which determines the link performance. In order to facilitate the presentation, and by an (harmless) abuse of notation, whenever referring to cost functions of type-A we shall denote them as $J_l^i(f_l^i, f_l)$ (instead of $\bar{J}_l^i(f_l^i, f_l)$).

Typically, the performance of a link l is manifested through some function $T_l(f_l)$, which measures the cost per unit of flow on the link, and depends on the link's total flow. Thus, it is of interest to consider cost functions of the following form (see also [11, 16]):

B1 $J_l^i(f_l^i, f_l) = f_l^i \cdot T_l(f_l)$.

B2 $T_l : [0, \infty) \to (0, \infty]$.

B3 $T_l(f_l)$ is positive, strictly increasing and convex.

B4 $T_l(f_l)$ is continuously differentiable.

Functions that comply with the above assumptions shall be referred to as *type-B* functions. Note that a type-B function is a special case of type-A. Note also that if $T_l(f_l)$ is the average delay per unit of flow, then the corresponding type-B function is the widely used average delay function (in our case, per user). We note that for type-B functions we have

$$K_l^i = f_l^i \cdot T_l' + T_l$$

where $T_l' = \frac{dT_l}{df_l}$.

A special kind of type-B cost functions is that which corresponds to an M/M/1 link model. In other words, suppose that:

C1 $J_l^i(f_l^i, f_l) = f_l^i \cdot T_l(f_l)$ is a type-B cost function.

C2 $T_l = \begin{cases} \frac{1}{C_l - f_l} & f_l < C_l \\ \infty & f_l \geq C_l \end{cases}$, where C_l is the capacity of link l.

Functions that comply with the above assumptions shall be referred to as *type-C* functions. Such delay functions are broadly used in modeling the behavior of links in computer communication networks [23, 24].

2.2 Uniqueness of the Nash Equilibrium

The following result establishes the uniqueness of the NEP for the parallel lines network.

Theorem 2.1 *In a network of parallel links where the cost function of each user is of type-A the NEP* $\tilde{\mathbf{f}}$ *is unique.*

Proof: Let $\mathbf{f} \in F$ and $\hat{\mathbf{f}} \in F$ be two NEPs. As observed above, $\mathbf{f}$ and $\hat{\mathbf{f}}$ satisfy the Kuhn-Tucker conditions (2-3), which may be written as

$$K_l^i(f_l^i, f_l) \geq \lambda^i \; ; \;\; K_l^i(f_l^i, f_l) = \lambda^i \text{ if } f_l^i > 0 \quad \forall i, l \, . \tag{4}$$

$$K_l^i(\hat{f_l}^i, \hat{f_l}) \geq \hat{\lambda}^i \; ; \;\; K_l^i(\hat{f_l}^i, \hat{f_l}) = \hat{\lambda}^i \text{ if } \hat{f_l}^i > 0 \quad \forall i, l \, . \tag{5}$$

These relations, and the fact that $K_l^i(\cdot, \cdot)$ is increasing in each of its arguments, will now be employed to establish that $\mathbf{f} = \hat{\mathbf{f}}$, i.e. $f_l^i = \hat{f_l}^i$ for every l, i.

The first step is to establish that $f_l = \hat{f_l}$ for each line l. To this end, we prove that for each l and i, the following relations hold:

$$\{\hat{\lambda}^i \leq \lambda^i, \hat{f_l} \geq f_l\} \text{ implies that } \hat{f_l}^i \leq f_l^i \, , \tag{6}$$

$$\{\hat{\lambda}^i \geq \lambda^i, \hat{f_l} \leq f_l\} \text{ implies that } \hat{f_l}^i \geq f_l^i \, . \tag{7}$$

We shall only prove (6), since (7) is symmetric. Assume that $\hat{\lambda}^i \leq \lambda^i$ and $\hat{f_l} \geq f_l$ for some l and i. Note that (6) holds trivially if $\hat{f_l}^i = 0$. Otherwise, if $\hat{f_l}^i > 0$, then (4)-(5) together with our assumption imply that

$$K_l^i(\hat{f_l}^i, \hat{f_l}) = \hat{\lambda}^i \leq \lambda^i \leq K_l^i(f_l^i, f_l) \leq K_l^i(f_l^i, \hat{f_l}) \, , \tag{8}$$

where the last inequality follows from the monotonicity of K_l^i in its second argument. Now, since K_l^i is non-decreasing in its first argument, this implies that $\hat{f_l}^i \leq f_l^i$, and (6) is established.

Let $\mathcal{L}_1 = \{l \, : \, \hat{f_l} > f_l\}$. Also denote $\mathcal{I}_a = \{i \, : \, \hat{\lambda}^i > \lambda^i\}$, $\mathcal{L}_2 = \mathcal{L} - \mathcal{L}_1 = \{l \, : \, \hat{f_l} \leq f_l\}$. Assume that $\mathcal{L}_1$ is not empty. Recalling that $\sum_l \hat{f_l}^i = \sum_l f_l^i = r^i$, it follows by (7) that for every i in $\mathcal{I}_a$,

$$\sum_{l \in \mathcal{L}_1} \hat{f_l}^i = r^i - \sum_{l \in \mathcal{L}_2} \hat{f_l}^i \leq r^i - \sum_{l \in \mathcal{L}_2} f_l^i = \sum_{l \in \mathcal{L}_1} f_l^i \, , \quad i \in \mathcal{I}_a \, . \tag{9}$$

Noting that (6) implies that $\hat{f_l}^i \leq f_l^i$ for $l \in \mathcal{L}_1$ and $i \notin \mathcal{I}_a$, it follows that

$$\sum_{l \in \mathcal{L}_1} \hat{f_l} = \sum_{l \in \mathcal{L}_1} \sum_{i \in \mathcal{I}} \hat{f_l}^i \leq \sum_{l \in \mathcal{L}_1} \sum_{i \in \mathcal{I}} f_l^i = \sum_{l \in \mathcal{L}_1} f_l \, . \tag{10}$$

This inequality obviously contradicts our definition of $\mathcal{L}_1$, which implies that $\mathcal{L}_1$ is an empty set. By symmetry it may also be concluded that the set $\{l \, : \, \hat{f_l} < f_l\}$ is empty. Thus, it has been established that

$$\hat{f_l} = f_l \;\; \text{for every} \;\; l \in \mathcal{L} \, . \tag{11}$$

We now proceed to show that $\hat{\lambda}^i = \lambda^i$ for each user i. To this end, note that (6) may be strengthened as follows:

$$\{\hat{\lambda}^i < \lambda^i, \hat{f}_l = f_l\} \text{ implies that either } \hat{f}_l^i < f_l^i \text{ or } \hat{f}_l^i = f_l^i = 0\,. \tag{12}$$

Indeed, if $\hat{f}_l^i = 0$ then the implication is trivial. Otherwise, if $\hat{f}_l^i > 0$, it follow similarly to (8) that $K_l^i(\hat{f}_l^i, \hat{f}_l) < K_l^i(f_l^i, \hat{f}_l)$, so that $\hat{f}_l^i < f_l^i$ as required.

Assume now that $\hat{\lambda}^i < \lambda^i$ for some $i \in \mathcal{I}$. Since $\sum_{l\in\mathcal{L}} \hat{f}_l^i = r^i > 0$, then $\hat{f}_l^i > 0$ for at least one link l, and (12) implies that

$$\sum_{l\in\mathcal{L}} f_l^i > \sum_{l\in\mathcal{L}} \hat{f}_l^i = r^i\,,$$

which contradicts the demand constraint for user i. We therefore conclude that $\hat{\lambda}^i < \lambda^i$ does not hold for any user i. A symmetric argument may be used to show that $\hat{\lambda}^i > \lambda^i$ cannot hold as well. Thus, $\hat{\lambda}^i = \lambda^i$ for every $i \in \mathcal{I}$. Combined with (11), this implies by (6)-(7) that $\hat{f}_l^i = f_l^i$ for every l, i, and uniqueness of the NEP is thus proved. □

2.3 Properties of the Nash Equilibrium

In this subsection we derive several properties of the (unique) NEP for type-A cost functions that are *identical for all users,* i.e., all users use the same function: for all $i \in \mathcal{I}$ and $l \in \mathcal{L}$ $J_l^i(f_l^i, f_l) = J_l(f_l^i, f_l)$ (to which correspond $K_l(f_l^i, f_l)$). Note that type-B functions belong to this class. It should be observed that, even though users use the same cost functions, they still have different objectives since they use different arguments (i.e., f_l^i) in these functions. All references to flow values are to those at the NEP.

Lemma 2.1 *Suppose that* $f_{\hat{l}}^i > f_{\hat{l}}^j$ *holds for some link* $\hat{l}$ *and users* i *and* j*. Then* $f_l^i \geq f_l^j$ *for all* $l \in \mathcal{L}$*; moreover, the last inequality is strict if* $f_l^j > 0$.

Proof: Since cost functions are identical then so are their derivatives, namely $K_l^i(\cdot,\cdot) = K_l^j(\cdot,\cdot)$. Recall that the latter are *strictly* increasing in their first argument, by our definition of class-A functions.

Choose an arbitrary link l. The claim holds trivially for $f_l^j = 0$. Assume, then, that $f_l^j > 0$. From the Kuhn-Tucker conditions we have that

$$K_l^j(f_l^j, f_l) \leq K_{\hat{l}}^j(f_{\hat{l}}^j, f_{\hat{l}})$$

Also, since $f_{\hat{l}}^i > f_{\hat{l}}^j$ implies $f_{\hat{l}}^i > 0$, we have

$$K_{\hat{l}}^i(f_{\hat{l}}^i, f_{\hat{l}}) \leq K_l^i(f_l^i, f_l)$$

Thus, we have

$$K_l^j(f_l^j, f_l) \leq K_{\hat{l}}^j(f_{\hat{l}}^j, f_{\hat{l}}) = K_{\hat{l}}^i(f_{\hat{l}}^j, f_{\hat{l}}) < K_{\hat{l}}^i(f_{\hat{l}}^i, f_{\hat{l}}) \leq K_l^i(f_l^i, f_l) = K_l^j(f_l^i, f_l)$$

i.e., $K_l^j(f_l^j, f_l) < K_l^j(f_l^i, f_l)$ which implies $f_l^j < f_l^i$. □

Theorem 2.2 *Consider a network of parallel links with identical type-A cost functions. For any pair of users i and j, $r^i \geq r^j$ implies that $f_l^i \geq f_l^j$ for all $l \in \mathcal{L}$. Moreover, if $r^i > r^j$ then equality holds only for $f_l^i = f_l^j = 0$.*

Proof: Assume $r^i \geq r^j$. If $r^i > r^j$ then there must be at least one line $\hat{l}$ for which $f_{\hat{l}}^i > f_{\hat{l}}^j$, and the required conclusions follow directly from the last lemma. Consider now the case $r^i = r^j$ and assume by contradiction that $f_{\hat{l}}^i < f_{\hat{l}}^j$ for some $\hat{l}$. Then, by the last lemma we have $f_l^i \leq f_l^j$ on all other lines, which upon summation yields $r^i < r^j$, contradicting $r^i = r^j$. □

The above theorem shows that, for identical type-A cost functions, there is a monotonicity among users in their use of links: a user with a higher demand uses more of each and every link. We conclude that

Corollary 2.1 *For users i and j such that $r^i = r^j$, holds $f_l^i = f_l^j$ for all $l \in \mathcal{L}$.* □

In particular, if all users have the same demand i.e., $r^i \equiv r$ for all $i \in \mathcal{I}$, then, for all $l \in \mathcal{L}$ and for all $i \in \mathcal{I}$ we have $f_l^i = f_l / I$.

Consider two users, say i and j, such that $r^i \geq r^j$ (i.e., $i > j$). Suppose that at the NEP user i refrains from using link l. It follows from Theorem 2.2 that so does user j i.e., $f_l^i = 0$ implies $f_l^j = 0$. Thus, at equilibrium we can partition the set of links into a sequence of sets $\mathcal{L}_1, \mathcal{L}_2, ..., \mathcal{L}_I$, such that $\mathcal{L}_n \subseteq \mathcal{L}$ for $1 \leq n \leq I$ and $\mathcal{L}_n$ is the set of links that is used exclusively by users $n, n+1, ..., I$. We have that $\mathcal{L}_n \supset \mathcal{L}_{n+1}$; also, since each user, including the one with the smallest demand, should use some link, we have that $\mathcal{L}_1 \neq \emptyset$ (other sets may be empty).

We observe a nice monotonic partition of users among links: a user with a higher demand uses more links, and uses more of each link. There is another monotonic property that can be derived, regarding the order of preference of links as seen by each user: suppose that a user , say i, prefers link l over link $\hat{l}$, i.e., $f_l^i \geq f_{\hat{l}}^i$; does this relation between links l and $\hat{l}$ hold for all users? The following lemma shows that this property holds for some types of Type-B cost functions.

Lemma 2.2 *Assume that for links $l, \hat{l} \in \mathcal{L}$ the following condition holds:*

$$T_l(f_l) > T_{\hat{l}}(f_{\hat{l}}) \Leftrightarrow T_l'(f_l) > T_{\hat{l}}'(f_{\hat{l}})$$

Then, $f_{\hat{l}}^i > f_l^i$ implies $f_{\hat{l}}^j \geq f_l^j$ for all $j \in \mathcal{I}$ and if $f_l^j > 0$ then $f_{\hat{l}}^j > f_l^j$.

Proof: Assume $f_{\hat{l}}^i > f_l^i$. Since the claims hold trivially for $f_l^j = 0$, we may further assume that $f_l^j > 0$. By our first assumption $f_{\hat{l}}^i > 0$, so that the Kuhn-Tucker conditions imply

$$T_{\hat{l}}(f_{\hat{l}}) + f_{\hat{l}}^i \cdot T_{\hat{l}}'(f_{\hat{l}}) \leq T_l(f_l) + f_l^i \cdot T_l'(f_l)$$

Since $f_{\hat{l}}^i > f_l^i$, it follows that either $T_{\hat{l}}(f_{\hat{l}}) < T_l(f_l)$ or $T_{\hat{l}}'(f_{\hat{l}}) < T_l'(f_l)$. $f_{\hat{l}}^i > f_l^i$, we have that either $T_{\hat{l}}(f_{\hat{l}}) < T_l(f_l)$ or $T_{\hat{l}}'(f_{\hat{l}}) < T_l'(f_l)$. However, the assumption made in the lemma implies that both of the last two inequalities hold (since each implies the other). Now, since $f_l^j > 0$, we similarly have

$$T_{\hat{l}}(f_{\hat{l}}) + f_{\hat{l}}^j \cdot T_{\hat{l}}'(f_{\hat{l}}) \geq T_l(f_l) + f_l^j \cdot T_l'(f_l)$$

Since it has just been established that $T_{\hat{l}}(f_{\hat{l}}) < T_l(f_l)$ and $T_{\hat{l}}'(f_{\hat{l}}) < T_l'(f_l)$, we conclude that $f_{\hat{l}}^j > f_l^j$. □

We note that the type-C function complies with the condition of the last lemma. Thus, we have:

Theorem 2.3 *In a network of parallel links where the cost function of each user is of type-C, $C_{\hat{l}} > C_l$ implies $f_{\hat{l}}^i \geq f_l^i$ for all $i \in \mathcal{I}$ and if $f_l^i > 0$ then $f_{\hat{l}}^i > f_l^i$.*

Proof: The claim holds trivially for $f_l^i = 0$. Assume, then, that $f_l^i > 0$. From the Kuhn-Tucker conditions we have that

$$T_l(f_l) + f_l^i \cdot T_l'(f_l) \leq T_{\hat{l}}(f_{\hat{l}}) + f_{\hat{l}}^i \cdot T_{\hat{l}}'(f_{\hat{l}}) \tag{13}$$

By contradiction, assume that $f_{\hat{l}}^i \leq f_l^i$. Since the type-C function complies with the condition of Lemma 2.2, it follows that $f_{\hat{l}}^j \leq f_l^j$ for all $j \in \mathcal{I}$, and thus $f_{\hat{l}} \leq f_l$. Since $C_{\hat{l}} > C_l$, and since $T_{l'}(f_{l'}) = \frac{1}{C_{l'} - f_{l'}}$ for $l' \in \{l, \hat{l}\}$, we have that $T_{\hat{l}}(f_{\hat{l}}) < T_l(f_l)$ and $T_{\hat{l}}'(f_{\hat{l}}) \leq T_l'(f_l)$. Since, by assumption, $f_{\hat{l}}^i \leq f_l^i$, we have a contradiction to inequality (13). □

The last theorem, together with Theorem 2.2, show that with type-C cost functions we get a partition of users among links at NEP: starting with a link with minimal capacity and moving towards links with higher capacities, we observe more and more users joining the links, and each user increasing its usage on the next link.

The property described in Lemma 2.2 does not hold for general average-delay cost functions, as the following example shows. Consider two users ($\mathcal{I} = \{1, 2\}$) sharing two links ($\mathcal{L} = \{1, 2\}$). Suppose that their cost functions are average-delay ones, and that, for $l \in \{1, 2\}$, $T_l = f_l + \frac{1}{C_l - f_l}$, where C_l is a parameter of link l (its capacity). Assume that $C_1 = 201$, $C_2 = 100$, $r^1 = 9.53$ and $r^2 = 290$. It can be verified (e.g., by substituting in the Kuhn-Tucker conditions) that the NEP is $f_1^1 = 0$, $f_2^1 = 9.53$, $f_1^2 = 200.0$ and $f_2^2 = 90.0$, which contradicts the property described in Lemma 2.2. This is because the property relies deeply on the assumption made in Lemma 2.2, and the above counter-example shows why that property may not hold for general average-delay cost functions: a user with a low demand seeks links with small delay, whereas one with a high demand seeks links with small delay derivatives.

2.4 A Simple Convergence Result

In this subsection we consider briefly the stability of the NEP. A reasonable dynamic model for users' behavior in non-equilibrium is suggested, and convergence of the flow configuration to the NEP is demonstrated. Only the special case of two users and a two link network is considered. We point out that this convergence result is not readily extendible to more general cases, however it may shed some light onto the important stability issue.

Consider a system of two users ($\mathcal{I} = \{1,2\}$) sharing two links ($\mathcal{L} = \{1,2\}$). The system starts with some (non-equilibrium) flow configuration $\mathbf{f}(0)$. From time to time, each user measures the current load on each link, and (after performing the necessary calculations) adjusts its own flows to minimize its cost function. We assume that exact minimization is achieved at each stage, and that all the above sequence of operations (measuring, calculating and adjusting) are done instantly. We shall refer to the above as the *Elementary Stepwise System (ESS)*. Essentially, the system can be modeled as a sequence of steps in each of which a user updates its routing decisions; we use the notation $f_l^i(n)$ to denote user's i flow on link l at the completion of step n. A similar dynamic scheme has been considered, e.g., in [20].

Let an ESS be initialized with a system configuration denoted by

$$\mathbf{f}(0) = (\mathbf{f}^1(0), \mathbf{f}^2(0)) = ((f_1^1(0), f_2^1(0)), (f_1^2(0), f_2^2(0)))$$

Relabel the users so that user 1 is the first to adjust its flow. The resulting flow configuration after this first step is denoted by $\mathbf{f}(1) = (\mathbf{f}^1(1), \mathbf{f}^2(1))$, where $\mathbf{f}^2(1) = \mathbf{f}^2(0)$ and $\mathbf{f}^1(1)$ is the optimal flow for user 1 against $\mathbf{f}^2(0)$. Since exact optimization is performed at each step, then the two users alternate in updating their flows. Thus, user 2 next updates its flow to yield $\mathbf{f}(2)$. Proceeding in this manner, user 1 updates its flow at each odd step n and user 2 updates its flow at each even step n, with the resulting system flow denoted by $\mathbf{f}(n)$ in each case. It is our purpose to show that $\mathbf{f}(n)$ converges to the NEP. The proof will be based on the following simple observation:

Lemma 2.3 *Let $\mathbf{f}^1$ and $\hat{\mathbf{f}}^1$ be two feasible flows for user 1, and let $\mathbf{f}^2$ (resp. $\hat{\mathbf{f}}^2$) be an optimal feasible flow of user 2 against $\mathbf{f}^1$ (resp. $\hat{\mathbf{f}}^1$). For $l \in \{1,2\}$, if $f_l^1 \geq \hat{f}_l^1$, then $f_l^2 \leq \hat{f}_l^2$.*

Proof: Assume, to the contrary, that $f_l^1 \geq \hat{f}_l^1$ and $f_l^2 > \hat{f}_l^2$ hold for some link l. Denoting the other link by m, this implies that $f_m^1 \leq \hat{f}_m^1$, $f_m^2 < \hat{f}_m^2$, $f_l > \hat{f}_l$ and $f_m < \hat{f}_m$. The Kuhn-Tucker conditions (2)-(3) for user 2 together with the monotonicity properties of the marginal costs K_l^i now lead to the following contradiction

$$K_m^2(\hat{f}_m^2, \hat{f}_m) = \hat{\lambda}^2 \leq K_l^2(\hat{f}_l^2, \hat{f}_l) < K_l^2(f_l^2, f_l) = \lambda^2 \leq K_m^2(f_m^2, f_m) < K_m^2(\hat{f}_m^2, \hat{f}_m),$$

and the conclusion follows. □

Proposition 2.1 *Assume that the ESS is initialized with a feasible system configuration* $\mathbf{f}(0)$. *Then the system configuration converges over time to the (unique) NEP* $\hat{\mathbf{f}}$, *i.e.:*

$$\lim_{n\to\infty} \mathbf{f}(n) = \hat{\mathbf{f}}$$

Proof: We first establish that each component of $\mathbf{f}(n)$ increases or decreases monotonically in n. Let l be a link for which $f_l^1(1) \leq f_l^1(3)$. Noting that by its definition $\mathbf{f}^2(2)$ is optimal for user 2 against $\mathbf{f}^1(1)$ and similarly that $\mathbf{f}^2(4)$ is optimal against $\mathbf{f}^1(3)$, it follows from the above lemma that $f_l^2(2) \geq f_l^2(4)$. A symmetric argument can now be employed to establish that $f_l^1(3) \leq f_l^1(5)$. Proceeding inductively, and recalling that user 1's flows remain fixed at each even step, it follows that for each odd n

$$f_l^1(n) = f_l^1(n+1) \leq f_l^1(n+2)\,,$$

and similarly for each even n

$$f_l^2(n) = f_l^2(n+1) \geq f_l^2(n+2)\,.$$

Since the flows are bounded, this implies that $f_l^1(n)$ and $f_l^2(n)$ converge as $n \to \infty$. Since the sum of flows on both links is constant, this obviously implies similar convergence for the flows on the second link, so that $\mathbf{f}(n)$ converges to some flow vector $\mathbf{f} = (\mathbf{f}^1, \mathbf{f}^2)$. Due to the continuity of the cost functions it follows that $\mathbf{f}^1$ is optimal for user 1 against $\mathbf{f}^2$ and $\mathbf{f}^2$ is optimal for user 2 against $\mathbf{f}^1$, so that $\mathbf{f}$ is the NEP. □

3 General Networks

3.1 Model and Problem Formulation

We consider now a network $\mathcal{G}(\mathcal{V}, \mathcal{L})$, where $\mathcal{V}$ is a finite set of nodes and $\mathcal{L} \subseteq \mathcal{V} \times \mathcal{V}$ is a set of directed links. For simplicity of notation, we assume that at most one link exists between each pair of nodes (in each direction). This assumption involves no loss of generality, since any network can be reduced to this form by introducing fictitious nodes. For a link $l \in \mathcal{L}$, we denote by $S(l)$ the identity of the node at the starting point of l and $D(l)$ is the node at the ending point. We shall at times denote a link l as (u, v), where $u = S(l)$, $v = D(l)$. Considering now a node $v \in \mathcal{V}$, we denote by $In(v)$ (correspondingly, $Out(v)$) the set of node v's in-going links (correspondingly, out-going links) i.e.: $In(v) = \{l | D(l) = v\}$ $(Out(v) = \{l | S(l) = v\})$.

As before, we are given a set $\mathcal{I} = \{1, 2, \ldots, I\}$ of selfish users, which now share the network $\mathcal{G}$. With each user i we associate a unique pair of source node $s(i)$ and destination node $t(i)$, and a throughput demand, which is some ergodic process with average rate of r^i. (In Sections 3.3 and 3.4 we shall specialize to the case where all users have the same sources and destinations.)

A user ships its demand (from $s(i)$ to $t(i)$) by splitting it through the various paths connecting the source to the destination. This is the essence of the routing operation performed by users. A user is able to decide (at any time) how to route its demand, i.e., user i decides what fraction of r^i should be sent through each path. We note that routing decisions can be interpreted also at the nodal level i.e., a user decides (perhaps distributedly) what amount of the flow entering into each node should be sent through each out-going link. We denote by $f_l^i \equiv f_{uv}^i$ the expected flow of user $i \in \mathcal{I}$ on link $l = (u, v) \in \mathcal{L}$. User i can fix any value for the f_l^i's, as long as

F1 For all $l \in \mathcal{L}$, $f_l^i \geq 0$ (nonnegativity constraint).

F2 For all $v \in \mathcal{V}$, $\sum_{l \in Out(v)} f_l^i = \sum_{l \in In(v)} f_l^i + r_v^i$, where $r_{s(i)}^i = r^i$, $r_{t(i)}^i = -r^i$ and $r_v^i = 0$ for $v \neq s(i), t(i)$ (conservation constraint).

Turning our attention to a link $l = (u, v) \in \mathcal{L}$, let $f_l = f_{uv}$ be the total flow on that link, i.e., $f_l = \sum_{i \in \mathcal{I}} f_l^i$. As before, the *flow configuration* $\mathbf{f^i}$ of user i is the vector of all f_l^i's (i.e., for all $l \in \mathcal{L}$), and the *system flow configuration* $\mathbf{f}$ is the vector of all user flow configurations. Feasible user- and system- flow configurations are defined as before (but considering the above constraints F1-F2), and denoted correspondingly as $\mathbf{F^i}$ and $\mathbf{F}$.

Assumptions G from Section 2 are now assumed to hold for the general network model. In particular, the cost function J^i of user i is the sum (over all network links) of link cost functions i.e., $J^i = \sum_{l \in \mathcal{L}} J_l^i$. The definition of K_l^i is as before.

We face a noncooperative game played by users on a *network*, and investigate the properties of the corresponding NEP(s). We note that the existence of an NEP in the network environment is guaranteed due to the same reasons as in the parallel-links environment. The Kuhn-Tucker conditions take now the following form: for every $i \in \mathcal{I}$ there exists a set of (Lagrange multipliers) $\{\lambda_u^i\}_{u \in \mathcal{V}}$, such that, for every link $(u, v) \in \mathcal{L}$:

$$f_{uv}^i > 0 \rightarrow \lambda_u^i = K_{uv}^i(\mathbf{f}_{uv}) + \lambda_v^i. \tag{14}$$

$$f_{uv}^i = 0 \rightarrow \lambda_u^i \leq K_{uv}^i(\mathbf{f}_{uv}) + \lambda_v^i. \tag{15}$$

The investigation of NEPs in network environments proves to be a much harder task than in the previous environment. In the Appendix we present an example of a four-node network with two users and type-A functions, for which the NEP is not unique. Nonetheless, in the following we prove uniqueness of the NEP under further conditions. We note that the properties of NEPs, derived in the previous section for certain classes of cost functions, may not hold in the network environment. In particular, in the Appendix we present an example of a four-node network with two users and type-B functions, in which the properties indicated in Lemma 2.1 and Corollary 2.1 do not hold. These indeed point at the complexity of the network environment.

3.2 Uniqueness of the NEP under Diagonal Strict Convexity

Uniqueness of the NEP for convex games has been established in [19], under certain convexity-like conditions on the cost functions which were termed *diagonal strict convexity* (DSC) conditions. In the present sub-section this result is applied to the network flow game. We first introduce the DSC conditions for our problem. Since these conditions may be hard to verify directly, we derive sufficient conditions on the *line* cost functions which guarantee DSC, and hence the uniqueness of the NEP. Some examples will be presented to illustrate the applicability of these results; these examples indicate that the derived conditions are quite useful for lightly loaded networks, but may fail to hold otherwise.

Following [19], the following notation will be used. Let $\nabla_i J^i(\mathbf{f})$ denote the gradient of $J^i(\mathbf{f})$ with respect to $\mathbf{f}^\mathbf{i}$, i.e. the column vector

$$\nabla_i J^i(\mathbf{f}) = \frac{\partial J^i}{\partial \mathbf{f}^\mathbf{i}}(\mathbf{f}) = \left(\frac{\partial J^i_l}{\partial f^i_l}(\mathbf{f_l}) \right)_{l \in \mathcal{L}} .$$

Let $\rho \in \Re^I$ be a fixed positive vector ($\rho_i > 0$, $i \in \mathcal{I}$). Define the weighted sum

$$\sigma(\mathbf{f}, \rho) = \sum_{i \in \mathcal{I}} \rho_i J^i(\mathbf{f}) \,,$$

and the associated *pseudogradient* vector

$$g(\mathbf{f}, \rho) = \begin{bmatrix} \rho_1 \nabla_1 J^1(\mathbf{f}) \\ \vdots \\ \rho_I \nabla_I J^I(\mathbf{f}) \end{bmatrix} .$$

(Note that ρ_i represents a positive scaling of the cost function for user i, and that such scaling does not affect the NEPs.) The function $\sigma(\cdot, \rho) : \mathbf{F} \to \Re$ is called *diagonally strictly convex* (DSC) if for every $\mathbf{f}, \hat{\mathbf{f}} \in \mathbf{F}$ holds

$$(\hat{\mathbf{f}} - \mathbf{f}) \, (g(\hat{\mathbf{f}}, \rho) - g(\mathbf{f}, \rho)) > 0 \,. \tag{16}$$

Theorem 3.1 *(Rosen [19]): If $\sigma(\mathbf{f}, \rho)$ is DSC for some $\rho > 0$, then the NEP system flow configuration is unique.*

Note that constraints F on the flows of each user are independent of the others' flows, as required in Theorem 2 of [19]. We also note that allowing infinite values for the cost functions causes no problem, since by our Assumption G5 the costs are finite at each NEP, and the uniqueness proof of [19] goes through without modification.

We proceed to formulate sufficient conditions for DSC which may be verified for each link separately. For every link l, let $\mathbf{F_l}$ denote the set of feasible flow vectors $\mathbf{f_l} = (f^1_l, \ldots, f^I_l)$ on that link for which the link costs $J^i_l(\mathbf{f_l})$ are all finite. Note that $\mathbf{F_l}$ has the simple form

$\mathbf{F}_l = \{\mathbf{f}_l : m_l^i \leq f_l^i \leq M_l^i,\, J_l^i(\mathbf{f}_l) < \infty,\, i \in \mathcal{I}\}$, where m_l^i and M_l^i are, respectively, the minimal and maximal values that f_l^i may take in any feasible flow configuration of user i. For each $0 < \rho \in \Re^I$, define the *pseudo-Jacobian* for link l as the following $I \times I$ matrix

$$G_l(\mathbf{f}_l, \rho) = \left\{ \rho_i \frac{\partial^2 J_l^i}{\partial f_l^i \partial f_l^j}(\mathbf{f}_l) \right\}_{i,j \in \mathcal{I}} \quad , \quad \mathbf{f}_l \in \mathbf{F}_l \,.$$

The following terminology shall be used: a square matrix M is said to be *positive definite* (denoted $M > 0$) if the symmetric matrix $(M + M')$ is positive definite, namely all eigenvalues of the latter are strictly positive. Here M' denotes the transpose of M.

Corollary 3.1 *Assume that for some positive $\rho \in \Re^I$, the matrix $G_l(\mathbf{f}_l, \rho)$ is positive definite for every $\mathbf{f}_l \in \mathbf{F}_l$ and $l \in \mathcal{L}$. Then the NEP is unique.*

Proof: As shown in [19, Theorem 6], a sufficient condition for $\sigma(\mathbf{f}, \rho)$ to be DSC over $\mathbf{F}$ is that $G(\mathbf{f}, \rho) > 0$ for every $\mathbf{f} \in \mathbf{F}$, where $G(\mathbf{f}, \rho)$ is the $(LI) \times (LI)$ Jacobian matrix of $g(\mathbf{f}, \rho)$ with respect to $\mathbf{f}$. It may be easily seen that, up to re-indexing of rows and columns, $G(\mathbf{f}, \rho)$ equals $\mathrm{diag}\{G_l(\mathbf{f}_l, \rho),\, l \in \mathcal{L}\}$, and the required conclusion follows. $\square$

Remark: In the conditions of the last corollary, the system flow constraints are manifested only through their projection on each of the links. While this makes them easier to verify, it also implies that these conditions are stronger (more demanding) than DSC condition (16), where the system flow constraints are fully taken into account.

We now present a few examples and observations to illustrate the conditions of Corollary 3.1. For simplicity, only the case of two users is considered, namely $\mathcal{I} = \{1, 2\}$.

Consider a link $l \in \mathcal{L}$ with type-C cost functions, namely

$$J_l^i(\mathbf{f}_l) = \frac{f_l^i}{C_l - f_l}\,, \quad i = 1, 2\,, \tag{17}$$

where $\mathbf{f}_l = (f_l^1, f_l^2)$ and $f_l = f_l^1 + f_l^2$. Then

$$G_l(\mathbf{f}_l, \rho) = \frac{1}{(C_l - f_l)^3} \begin{bmatrix} 2\rho_1(C_l - f_l^2) & \rho_1(C_l + f_l^1 - f_l^2) \\ \rho_2(C_l + f_l^2 - f_l^1) & 2\rho_2(C_l - f_l^1) \end{bmatrix}$$

for $f_l < C_l$. The following facts may now be easily verified:

(*i*) Let $\mathbf{F}_l = \{\mathbf{f}_l : 0 \leq f_l^i \leq x^i,\, i = 1, 2\}$, where (x^1, x^2) are constants which satisfy $x^1 + x^2 < C_l$. Then, for $\rho = (x^2, x^1)$, $G_l(\mathbf{f}_l, \rho) > 0$ over $\mathbf{F}_l$.

(*ii*) Let $\mathbf{F}_l = \{\mathbf{f}_l : f_l^i \geq 0,\, f_l^1 + f_l^2 < C_l\}$. Then no vector $\rho > 0$ exists for which $G_l(\mathbf{f}_l, \rho) > 0$ over $\mathbf{F}_l$. Indeed, for any fixed ρ the matrix $G_l(\mathbf{f}_l, \rho)$ will not be positive definite if (say) f_l^1 is close enough to C_l.

As a consequence of fact (i) above, the following result is evident. Consider a network with two users, cost functions (17) on each link, and flow requirements (r^1, r^2) such that $r^1 + r^2 \leq C_l$ for every $l \in \mathcal{L}$. Then the NEP is unique.

It is evident from fact (ii) above that a similar result cannot be deduced from Corollary 3.1 if $r^1 + r^2 > C_l$ for some link l. Thus, the usefulness of Corollary 3.1 is limited in this case to lightly loaded networks.

To further illustrate this point, consider cost functions of the form

$$J_l^i(\mathbf{f_l}) = f_l^i P_m(f_l)\,, \quad i = 1, 2 \tag{18}$$

where P_m is a monic polynomial of degree $m \geq 1$. Let $\mathbf{F_l}$ be the positive quadrant. It may then be verified that $G_l(\mathbf{f_l}) > 0$ over $\mathbf{f_l} \in \mathbf{F_l}$ if $m \leq 7$, but not if $m \geq 8$. Thus, if the cost functions on each line of a two-user network are of the form (18) with $m \leq 7$, then uniqueness of the NEP is guaranteed without any restrictions on the flow requirements of the users. Both this and the "lightly loaded network" condition alluded to above can be interpreted as requiring that the cost functions will not increase "too steeply" as the load on the line increases.

3.3 Symmetrical Users

Suppose that all users have the same demands (in particular the same source node and same destination node) and use the same *type-A* cost functions, i.e.: for all $i, j \in \mathcal{I}$ $r^i \equiv r^j$, and for all $i \in \mathcal{I}$ and $l \in L$ $J_l^i(f_l^i, f_l) = J_l(f_l^i, f_l)$. We call such users *symmetrical*. We recall that even though symmetrical users use the same cost functions, they still have different objectives. Note that for symmetrical users we also have $K_l^i(\cdot,\cdot) = K_l^j(\cdot,\cdot)$ for all $i, j \in \mathcal{I}$.

Lemma 3.1 *In a network with symmetrical users, the flow values at an NEP are such that*

$$f_l^i \equiv \frac{f_l}{I}$$

for all $i \in \mathcal{I}$ and $l \in \mathcal{L}$.

Proof: Assume, by contradiction, that there is a link $\hat{l} \in \mathcal{L}$ and a user $i \in \mathcal{I}$ such that $f_{\hat{l}}^i \neq \frac{f_{\hat{l}}}{I}$. It follows that there is another user $j \in \mathcal{I}$ such that $f_{\hat{l}}^i \neq f_{\hat{l}}^j$, and without loss of generality assume that $f_{\hat{l}}^i > f_{\hat{l}}^j$. We construct a directed network $\mathcal{G}'(\mathcal{V}', \mathcal{L}')$, whose set of nodes is identical to that of $\mathcal{G}$ (i.e., $\mathcal{V}' = \mathcal{V}$) and the set of links $\mathcal{L}'$ is constructed as follows:

- for each link $l = (u, v) \in \mathcal{L}$ such that $f_l^i \geq f_l^j$ we have a link $l' = (u, v) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = f_l^i - f_l^j$.
- for each link $l = (u, v) \in \mathcal{L}$ such that $f_l^i < f_l^j$ we have a link $l' = (v, u) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = f_l^j - f_l^i$.

In other words, we redirect links according to the relation between f_l^i and f_l^j. It is easy to verify that the values $x_{l'}$ constitute a nonnegative, directed flow in the network. Since symmetrical users have the same demand (i.e., $r^i - r^j = 0$), this flow has no sources (the total flow into each node equals the total flow out of that node, i.e. it is a circulation). Thus, either $x_{l'} \equiv 0$ or else there is a cycle C of links in $\mathcal{G}'$ such that $x_{l'} > 0$ for all $l' \in \mathcal{L}'$. Since for the link $\hat{l}$, $f_{\hat{l}}^i > f_{\hat{l}}^j$, we have that $x_{\hat{l}'} > 0$, we conclude that a cycle C as described above exists.

Consider now a link $l' = (u, v) \in \mathcal{L}'$ for which $x_{l'} > 0$. Clearly, therefore, either $f_{uv}^i > f_{uv}^j$ or else $f_{vu}^j > f_{vu}^i$. In the case where $f_{uv}^i > f_{uv}^j \geq 0$ we have that

$$\lambda_u^i - \lambda_v^i = K_{uv}^i(f_{uv}^i, f_{uv}) = K_{uv}^j(f_{uv}^i, f_{uv}) > K_{uv}^j(f_{uv}^j, f_{uv}) \geq \lambda_u^j - \lambda_v^j \tag{19}$$

where the first transition follows from the Kuhn-Tucker conditions for $f_{uv}^i > 0$, the second is because all users have the same cost functions, the third is due to the assumption $f_{uv}^i > f_{uv}^j$, and the fourth is again due to the Kuhn-Tucker conditions. In the second case (i.e., $f_{vu}^j > f_{vu}^i \geq 0$) we have by symmetry that

$$\lambda_v^j - \lambda_u^j > \lambda_v^i - \lambda_u^i \tag{20}$$

Note that the results of equations 19 and 20 are in fact identical.

Denote $\Delta\lambda_w = \lambda_w^i - \lambda_w^j$ (for all $w \in \mathcal{V}$). From (19) and (20) we conclude that, for $l' = (u, v)$, $x_{l'} > 0$ implies that $\Delta\lambda_u > \Delta\lambda_v$. This means that along the cycle C described above we would have a monotonically increasing sequence of $\Delta\lambda$'s, which is a contradiction.

We conclude that, for all $i, j \in \mathcal{I}$ and all $l \in \mathcal{L}$, we have $f_l^i = f_l^j$. This means that $f_l^i = \frac{f_l}{\mathcal{I}}$. □

Theorem 3.2 *A network with symmetrical users has a unique NEP.*

Proof: Suppose by contradiction that there are two NEPs, and denote by f_l^i, f_l the flow values of one NEP and by $\hat{f}_l^i, \hat{f}_l$ those of the other NEP. Also, λ_u, $\hat{\lambda}_u$ are, respectively, the Lagrange multipliers at a node $u \in \mathcal{V}$ in the two NEPs. Since the two NEPs are different, there are some $i \in \mathcal{I}$ and $\bar{l} \in \mathcal{L}$ such that $f_{\bar{l}}^i \neq \hat{f}_{\bar{l}}^i$, and without loss of generality assume that$f_{\bar{l}}^i > \hat{f}_{\bar{l}}^i$. We construct a directed network $\mathcal{G}'$ in the same way as in the proof of Lemma 3.1, only that now we consider the relation between f_l^i and $\hat{f}_l^i$. In other words, the set of links $\mathcal{L}'$ is constructed as follows:

- for each link $l = (u, v) \in \mathcal{L}$ such that $f_l^i \geq \hat{f}_l^i$ we have a link $l' = (u, v) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = f_l^i - \hat{f}_l^i$.
- for each link $l = (u, v) \in \mathcal{L}$ such that $f_l^i < \hat{f}_l^i$ we have a link $l' = (v, u) \in \mathcal{L}'$; to such a link l' we assign a (flow) value $x_{l'} = \hat{f}_l^i - f_l^i$.

Again, it is easy to verify that the values $x_{l'}$ constitute a nonnegative, directed flow in the network. Since user i has the same demand r^i at both NEPs, this flow has no sources. This, together with $f_{\bar{l}}^i > \hat{f}_{\bar{l}}^i$, mean that there is a cycle C of links in $\mathcal{G}'$ such that $x_{l'} > 0$ for all $l' \in \mathcal{L}'$.

From the proof of Lemma 3.1 we have that, for all $j \in \mathcal{I}$ and $l \in \mathcal{L}$, $f_l^j = \frac{f_l}{I}$ and $\hat{f}_l^j = \frac{\hat{f}_l}{I}$. Thus, $f_l^j > \hat{f}_l^j$ implies $f_l > \hat{f}_l$.

Consider now any link $l' = (u, v) \in \mathcal{L}'$ for which $x_{l'} > 0$. We have that either $f_{uv}^i > \hat{f}_{uv}^i$ or else $\hat{f}_{vu}^i > f_{vu}^i$. In the first case (i.e., $f_{uv}^i > \hat{f}_{uv}^i \geq 0$) we have that

$$\lambda_u^i - \lambda_v^i = K_{uv}^i(f_{uv}^i, f_{uv}) > K_{uv}^i(\hat{f}_{uv}^i, \hat{f}_{uv}) \geq \hat{\lambda}_u^i - \hat{\lambda}_v^i$$

where the first transition is due to the Kuhn-Tucker conditions (for the first NEP and for $f_{uv}^i > 0$), the second is due to the the assumption $f_l^i > \hat{f}_l^i$ (which implies $f_l > \hat{f}_l$), and the third is again due to the Kuhn-Tucker conditions (for the second NEP). In the second case (i.e., $\hat{f}_{vu}^i > f_{vu}^i \geq 0$) we have by symmetry that

$$\hat{\lambda}_v^i - \hat{\lambda}_u^i > \lambda_v^i - \lambda_u^i$$

Denote $\Delta\lambda_w = \lambda_w^i - \hat{\lambda}_w^i$ (for all $w \in \mathcal{V}$). The contradiction follows as in the proof of Lemma 3.1. □

Corollary 3.2 *If all users have the same demand ($r^i = r^j$) and the cost functions are of type B, then there is a unique NEP. The flow values at the NEP are such that $f_l^i \equiv \frac{f_l}{I}$ for all $i \in \mathcal{I}$ and $l \in \mathcal{L}$.*

Proof: Follows directly from the last lemma and theorem, since cost functions of type B are identical for all users. □

3.4 Uniqueness of NEPs with All-Positive Flows

In this sub-section we prove a result, which implies, in particular, that at most one NEP exists with the property that all users have strictly positive flows on each link of the network. Note that this all-positive assumption makes sense only when all users have the same source and destination nodes, which will be assumed in this sub-section. A similar result has been established in [16], for the special case of the two-node parallel lines network.

Theorem 3.3 *Assume type-B cost functions. Let $\mathbf{f}$ and $\hat{\mathbf{f}}$ be two NEPs such that: there exists a set of links $\mathcal{L}_1 \subset \mathcal{L}$ such that $\{f_l^i > 0$ and $\hat{f}_l^i > 0,\ i \in \mathcal{I}\}$ for $l \in \mathcal{L}_1$, and $\{f_l^i = \hat{f}_l^i = 0,\ i \in \mathcal{I}\}$ for $l \notin \mathcal{L}_1$. Then $\mathbf{f} = \hat{\mathbf{f}}$.*

The theorem defines certain classes of flow configurations, each characterized by a set of links such that every user ships flow only through each of these links. The theorem states that within each class at most one NEP may exist. Obviously, to deduce global uniqueness of the NEP, one needs to verify independently that an NEP may exist within exactly one of these classes. This latter fact is not true in general, e.g., see the end of subsection 2.3 for an example in which the unique NEP is not within any such class (one user uses a link not used by the other).

Proof: We start by re-writing the Kuhn-Tucker conditions (14)-(15) for $\mathbf{f}$ and $\hat{\mathbf{f}}$ over the set of links $\mathcal{L}_1$. Note that, since all flows are strictly positive on these links, then the equality condition (14) holds. Recall also that for type-B cost functions, i.e. $J_l^i = f_l^i T_l(f_l)$, we have $K_l^i = f_l^i T_l' + T_l$. Thus, there exist constants $\{\lambda_u^i\}$ and $\{\hat{\lambda}_u^i\}$ such that, for every link $l = (u,v) \in \mathcal{L}_1$ and $i \in \mathcal{I}$,

$$\begin{aligned} \hat{f}_{uv}^i T_{uv}'(\hat{f}_{uv}) + T_{uv}(\hat{f}_{uv}) &= \hat{\lambda}_u^i - \hat{\lambda}_v^i \,, \\ f_{uv}^i T_{uv}'(f_{uv}) + T_{uv}(f_{uv}) &= \lambda_u^i - \lambda_v^i \,. \end{aligned}$$

Summing each of these equations over i, we get

$$S_{uv}(\hat{f}_l) \triangleq \hat{f}_{uv} T_{uv}'(\hat{f}_{uv}) + I \cdot T_{uv}(\hat{f}_{uv}) = \hat{\lambda}_u - \hat{\lambda}_v \tag{21}$$

$$S_{uv}(f_l) \triangleq f_{uv} T_{uv}'(f_{uv}) + I \cdot T_{uv}(f_{uv}) = \lambda_u - \lambda_v \,, \tag{22}$$

where $\hat{\lambda}_u \triangleq \sum_i \hat{\lambda}_u^i$ and $\lambda_u \triangleq \sum_i \lambda_u^i$. From the last two equations, combined with the conservation constraints F2, it will now be deduced that $\hat{f}_l = f_l$ for every $l \in \mathcal{L}_1$. (Note that these equations are very similar to the Kuhn-Tucker conditions for a single-user optimization problem of link flows, with respect to a modified (convex) link cost function with derivatives $S_l(f_l)$. Thus, the uniqueness of their solution is actually a consequence of standard convex programming results. However, a direct proof is provided below.) Subtracting (22) from (21), multiplying by $(\hat{f}_l - f_l)$, summing over $(u,v) \in \mathcal{L}_1$ and noting that $f_l = \hat{f}_l = 0$ for $l \notin \mathcal{L}_1$, we get

$$\sum_{(u,v)\in\mathcal{L}} (\hat{f}_{uv} - f_{uv})(S_{uv}(\hat{f}_{uv}) - S_{uv}(f_{uv})) = \sum_{(u,v)\in\mathcal{L}} (\hat{f}_{uv} - f_{uv}) \left[(\hat{\lambda}_u - \lambda_u) - (\hat{\lambda}_v - \lambda_v) \right] \,. \tag{23}$$

By the properties of type-B cost functions, it follows that each function $S_{uv}(\cdot)$ is strictly increasing. This implies that (each term in the sum on) the left-hand side of (23) is non-negative, and equals zero only if $\hat{f}_l = f_l$ for every l. However, the right-hand side of this equation sums up to 0. Indeed, changing summation variables and applying the flow conservation constraints to $\mathbf{f}$ and $\hat{\mathbf{f}}$ separately yields

$$\begin{aligned} \sum_{(u,v)\in\mathcal{L}} (\hat{f}_{uv} - f_{uv})(\hat{\lambda}_u - \lambda_u) &- \sum_{(w,u)\in\mathcal{L}} (\hat{f}_{wu} - f_{wu})(\hat{\lambda}_u - \lambda_u) \\ &= \sum_{u\in\mathcal{V}} (\hat{\lambda}_u - \lambda_u) \left[\sum_{l\in Out(u)} (\hat{f}_l - f_l) - \sum_{l\in In(u)} (\hat{f}_l - f_l) \right] \\ &= 0 \,. \end{aligned}$$

It has thus been established that $\hat{f}_l = f_l$ for every $l \in \mathcal{L}$. Proceeding exactly as in the proof of Theorem 2.1 (starting from equation (11)) it may now be inferred that $\hat{f}_l^i = f_l^i$ for every i and l, i.e. $\hat{\mathbf{f}} = \mathbf{f}$. □

4 Conclusion

This paper considered the fundamental problem of routing in an environment composed of several selfish users. The problem was posed as a non-cooperative game, for which the Nash equilibrium was investigated.

The main thrust of the analysis was to establish the uniqueness of the NEP under conditions that comply with reasonable network environment, and in particular encompass performance criteria commonly employed for routing. The investigation of NEPs in network routing problems seems to be complicated by the fact that each user is faced with a multi-variable decision. As was shown, standard theory on convex games fails to yield satisfactory results even for simple networks.

Problem-specific analysis yielded a complete uniqueness result for the parallel links case. Moreover, we presented monotonicity properties of the Nash flows that characterize the NEP in an intuitively appealing way. General network topologies prove to be considerably harder to tackle. This is indicated by the fact that the type-A assumptions are not sufficient in order to guarantee uniqueness of the NEP and that "intuitive" properties that hold in the parallel-links case fail to hold in general. Nonetheless, we presented several uniqueness results for various network conditions.

This paper presented initial steps towards the understanding of multi-user routing games. Several questions for further research stem from the present work. At this point it is not clear if and how the strong uniqueness result, obtained for parallel links, can be extended to a general network. The counter-example presented in this paper shows that such an extension is not possible for type-A functions, however more restricted classes (e.g. type-B functions) might guarantee the uniqueness of the NEP. Another major issue is the stability of the NEP. A preliminary result in this vein was presented. The uniqueness results obtained in this paper encourage further investigation on the stability issue.

Several other open questions of practical value deserve attention. For example, in many networks users are restricted to route their flow along a single path (with strict rules of changing them). Under such circumstances an NEP may not exist at all and complicated oscillatory behavior is likely to arise. Another example is that of the delay encountered by measuring and adjusting network flows which will affect the convergence rate and might affect convergence altogether.

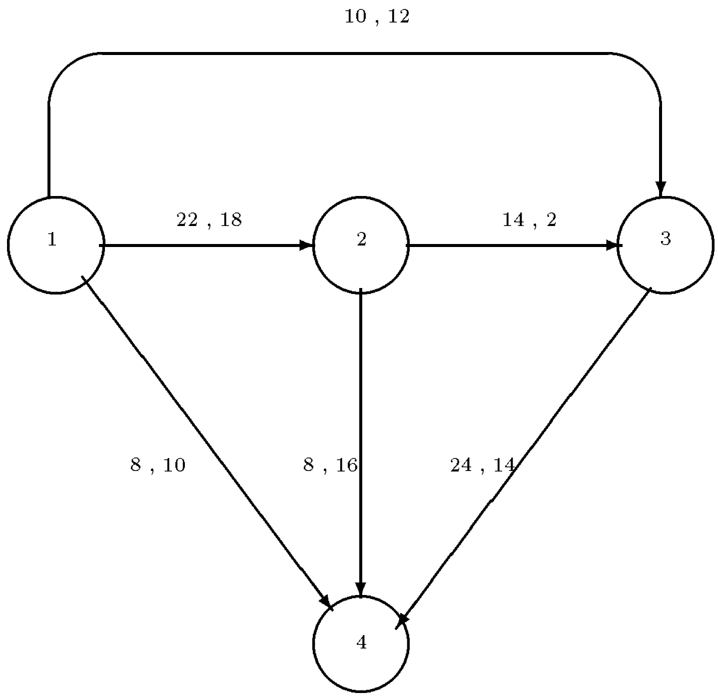

Figure 1: First NEP

Acknowledgment

We would like to thank Aurel Lazar and Yannis Korilis for insightful comments on an earlier version of this work.

5 Appendix

Nonunique NEP Example for General Topologies

Consider the network of Figures 1 and 2, having four nodes and two users with demands $r^1 = r^2 = 40$ between source node 1 and destination node 4. Figures 1 and 2 describe the flow values at two different NEP's; each pair of numbers adjacent to a link are the flow values of the two users on that link at the corresponding NEP (the left value corresponds to user 1 and the right value to user 2). Let the values of the K_l^i functions be as follows:

For the first user and the first NEP:

$K_{12}^1(22,40) = 102$; $K_{13}^1(10,22) = 109$; $K_{14}^1(8,18) = 201$;

$K_{23}^1(14,16) = 7$; $K_{24}^1(8,24) = 99$;

$K_{34}^1(24,38) = 92$.

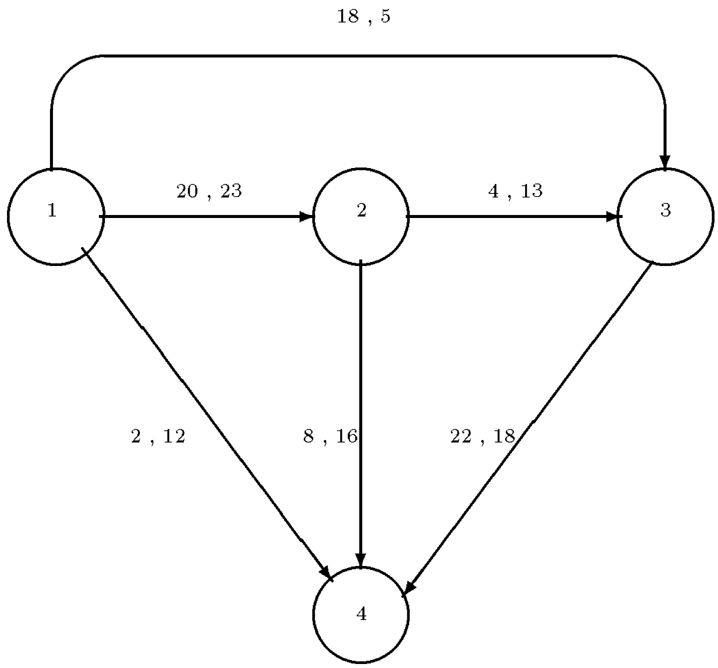

Figure 2: Second NEP

For the first user and the second NEP:

$K^1_{12}(20, 43) = 100$; $K^1_{13}(18, 23) = 110$; $K^1_{14}(2, 14) = 200$;

$K^1_{23}(4, 17) = 10$; $K^1_{24}(16, 26) = 100$;

$K^1_{34}(22, 40) = 90$.

For the second user and the first NEP:

$K^2_{12}(18, 40) = 20$; $K^2_{13}(12, 22) = 30$; $K^2_{14}(10, 18) = 120$;

$K^2_{23}(2, 16) = 10$; $K^2_{24}(16, 24) = 100$;

$K^2_{34}(14, 38) = 90$.

For the second user and the second NEP:

$K^2_{12}(23, 43) = 30$; $K^2_{13}(5, 23) = 50$; $K^2_{14}(12, 14) = 150$;

$K^2_{23}(13, 17) = 20$; $K^2_{24}(10, 26) = 120$;

$K^2_{34}(18, 40) = 100$.

The above values comply with the Kuhn-Tucker conditions 14-15, as well as with the monotonicity property of the K^i_l functions.

We now present a scheme for constructing proper K^i_l functions to which the above values correspond. We denote by $\tilde{f}_l^{\,i}, \tilde{f}_l$ the flow values of the first NEP and by $\hat{f}_l^{\,i}, \hat{f}_l$ those of the second NEP. For each user i and link l we have two values of the function K_l^i, namely $K_l^i(\tilde{f_l}^i, \tilde{f_l})$ and $K_l^i(\hat{f_l}^i, \hat{f_l})$. Consider some i and l, and suppose that $K_l^i(\tilde{f_l}^i, \tilde{f_l}) > K_l^i(\hat{f_l}^i, \hat{f_l})$. Suppose first that $\tilde{f_l}^i > \hat{f_l}^i$. We define K_l^i as $K_l^i(f_l^i, f_l) = a_l^i \cdot (f_l^i)^n + b_l^i \cdot (f_l)^{1/n}$. For a large enough n we have $(\tilde{f_l})^{1/n} \approx (\hat{f_l})^{1/n} \approx 1$ thus

$$a_l^i \approx \frac{K_l^i(\tilde{f_l}^i, \tilde{f_l}) - K_l^i(\hat{f_l}^i, \hat{f_l})}{(\tilde{f_l}^i)^n - (\hat{f_l}^i)^n} > 0$$

and

$$b_l^i \approx (\hat{f_l}^i)^n \cdot K_l^i(\tilde{f_l}^i, \tilde{f_l}) \cdot \frac{(\frac{\tilde{f_l}^i}{\hat{f_l}^i})^n \cdot \frac{K_l^i(\hat{f_l}^i, \hat{f_l})}{K_l^i(\tilde{f_l}^i, \tilde{f_l})} - 1}{(\tilde{f_l}^i)^n - (\hat{f_l}^i)^n}$$

For a large enough n we have $(\frac{\tilde{f_l}^i}{\hat{f_l}^i})^n > \frac{K_l^i(\tilde{f_l}^i, \tilde{f_l})}{K_l^i(\hat{f_l}^i, \hat{f_l})}$, thus $b_l^i > 0$.

Suppose now that $\tilde{f_l}^i \leq \hat{f_l}^i$. Then we must have $\tilde{f_l} > \hat{f_l}$. We repeat the above only that now we choose

$$K_l^i(f_l^i, f_l) = a_l^i \cdot (f_l^i)^{1/n} + b_l^i \cdot (f_l)^n$$

and we get a symmetrical solution for a_l^i and b_l^i. In each of the two cases we have a function K_l^i that is continuous and strictly increasing in each of its two arguments. The corresponding cost function is

$$J_l^i(f_l^i, f_l) = \frac{a_l^i}{n+1} \cdot (f_l^i)^{n+1} + \frac{b_l^i}{1 + 1/n} \cdot (f_l)^{1+1/n}$$

for the first case and

$$J_l^i(f_l^i, f_l) = \frac{b_l^i}{n+1} \cdot (f_l)^{n+1} + \frac{a_l^i}{1 + 1/n} \cdot (f_l^i)^{1+1/n}$$

for the second case. In either case, we have a proper type-A cost function.

Nonmonotonous NEP Example for General Topologies

Consider the network of Figure 3, having four nodes and two users with type-B cost functions and with demands $r^1 = 7$, $r^2 = 4$ between source node 1 and destination node 4. Each pair of numbers adjacent to a link are the flow values of the two users on that link at the NEP (the left value corresponds to user 1 and the right value to user 2). Let the values of the T_l and T_l' functions be as follows:

$T'_{ab}(7) = 5$, $T_{ab}(7) = 4$,

$T'_{ac}(4) = 2$, $T_{ac}(4) = 20$,

$T'_{bc}(3) = 1$, $T_{bc}(3) = 1$,

$T'_{bd}(4) = 2$, $T_{bd}(4) = 21$,

$T'_{cd}(7) = 5$, $T_{bd}(7) = 5$.

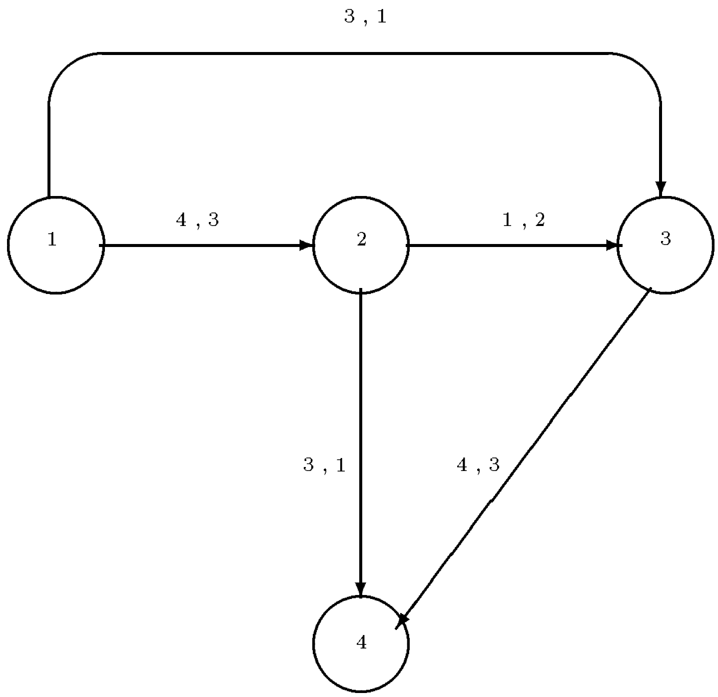

Figure 3: Nonmonotonous NEP

The above values comply with the Kuhn-Tucker conditions 14-15, as well as with the monotonicity property of the T_l and $T_l{}'$ functions. Note that $f^1_{bc} < f^2_{bc}$ although $r^1 > r^2$.

It can be verified that, for every link, there are positive values C_l and d_l such that $T_l(f_l) = \frac{1}{C_l - f_l} + d_l$ satisfies the above NEP values of T_l and $T_l{}'$. Clearly, such functions comply with the requirements of type-B cost functions.

References

[1] C. Partridge, *Innovations in Internetworking.* Norwood, Ma: Artech House, 1988.

[2] C. Sunshine, "Network interconnection and gateways," *IEEE Journal on Selected Areas in Communications*, vol. 8, pp. 4–11, 1990.

[3] R. Luce and H. Raiffa, *Games and Decisions.* New-York: Wiley, 1957.

[4] T. Basar and G. Olsder, *Dynamic Noncooperative Game Theory.* London: Academic Press, 1982.

[5] J. Postel, C. Sunshine, and D. Cohen, "The ARPA Internet protocol," *Computer Networks*, vol. 5, pp. 261–271, 1981.

[6] R. Hinden, J. Haverty, and A. Sheltzer, "The DARPA Internet: Interconnecting heterogeneous computer networks with gateways," *Computer*, vol. 16, pp. 38–48, September 1983.

[7] D. Clark, "Policy routing in internet protocols," Tech. Rep. RFC 1102, DDN Network Information Center, Menlo Park, California, May 1989.

[8] M. S. J. Kurose and Y. Yemini, "A microeconomic approach to optimization of channel access policies in multiaccess networks," in *Proceedings of the Fifth International Conference on Distributed Computer Systems*, (Denver, CO), pp. 70–80, May 1985.

[9] D. Ferguson, C. Nikolaou, and Y. Yemini, "An economy for flow control in computer networks," in *Proceedings of INFOCOM'89*, pp. 110–118, 1989.

[10] A. Bovopoulos and A. Lazar, "Synchronous and asynchronous iterative algorithms for optimal load balancing," in *Proceedings of the 22nd Conference on Information Sciences and Systems*, (Princetion, NJ), pp. 1051–1057, March 1988.

[11] M. Hsiao and A. Lazar, "Optimal decentralized flow control of Markovian queueing networks with multiple controllers," *Performance Evaluation*, vol. 13, pp. 181–204, 1991.

[12] Z. Zhang and C. Douligeris, "Convergence of synchronous and asynchronous greedy algorithms in a multiclass telecommunications environment," *IEEE Transactions on Communications*, vol. 40, pp. 1277–1281, August 1992.

[13] R. Mazumdar, L. Mason, and C. Douligeris, "Fairness in network optimal flow control: Optimality of product forms," *IEEE Transactions on Communications*, vol. 39, pp. 775–782, May 1991.

[14] S. Shenker, "Efficient network allocations with selfish users," in *Proccedings of Performance '90*, (Edinburgh, Scotland), pp. 279–285, North Holland, September 1990.

[15] S. Shenker, "Making greed work on networks: A game-theoretic analysis of gateway service disciplines," in *Proc. of ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, pp. 241–242, May 1990.

[16] H. Lee and M. Cohen, "Multi-agent customer allocation in a stochastic service system," *Management Science*, vol. 31, pp. 752–763, June 1985.

[17] A. A. Economides and J. A. Silvester, "Multi-objective routing in integrated services networks: A game theory approach," in *Proceedings of INFOCOM'91*, pp. 1220–1225, 1991.

[18] S. C. Dafermos and F. T. Sparrow, "The traffic assignment problem for a general network," *Journal of Research of the NBS (B)*, vol. 73B, pp. 91–118, April-June 1969.

[19] J. Rosen, "Existence and uniqueness of equilibrium points for concave n-person games," *Econometrica*, vol. 33, pp. 520–534, July 1965.

[20] S. Li and T. Basar, "Distributed algorithms for the computation of noncooperative equilibria," *Automatica*, vol. 23, no. 4, pp. 523–533, 1987.

[21] M. G. L. Fratta and L. Kleinrock, "The flow deviation method: An approach to store and forward communication network design," *Networks*, vol. 3, pp. 97–133, 1973.

[22] R. G. Gallager, "A minimum delay routing algorithm using distributed computation," *IEEE Transactions on Communications*, vol. COM-25, pp. 73–85, 1977.

[23] D. Bertsekas and R. Gallager, *Data Networks*. New-Jersey: Prentice Hall, Inc., 1987.

[24] L. Kleinrock, *Queueing Systems*. New York: J. Wiley, 1975.