Proceedings of the
44th IEEE Conference on Decision and Control, and
the European Control Conference 2005
Seville, Spain, December 12-15, 2005

Uniqueness of the Nash Equilibrium in Convex Routing Games: Topological Conditions

Oran Richman and Nahum Shimkin

Abstract— **We consider the problem of non-cooperative routing in a congested network with atomic users, where each user wishes to minimize the cost of its own flow. Cost functions may differ among users, and are required to satisfy standard convexity and monotonicity conditions. A characterization is given of the class of two-terminal network topologies for which the Nash equilibrium is always unique. We further address the uniqueness issue for the mixed Nash-Wardrop equilibrium problem.**

I. INTRODUCTION

Congested networks have been an object of interest in engineering and operations research for over five decades, motivated by major applications to traffic engineering and communication networks. In these and other application domains, the network is often not centrally controlled but rather shared by a number of users who pursue their own objectives. This has led to extensive work on the analysis of multi-user networks within the framework of game theory, and to the investigation of equilibrium concepts for these models. For a recent survey on these issues from the telecommunications perspective see [2].

We consider here the problem of *competitive routing*, where each user needs to deliver a given amount of flow over the network from its designated origin node to its destination. A user can choose how to divide its flow between the available routes. On each link the user incurs a certain cost per unit flow, which in general will depend on the link congestion, namely the total flow over that link. In the context of computer networks the per-unit cost is often synonymous to the link *latency*, a terminology that we adopt here for simplicity. The latency of a path is simply the sum of the latencies along its links.

The fundamental notion of equilibrium in transportation networks has been proposed by Wardrop in [27]. Essentially, it requires all traffic to occupy paths with minimal latency. While this solution concept has been addressed by different names, including the Nash equilibrium for infinitesimal users, user equilibrium, or traffic equilibrium, we shall use the term *Wardrop equilibrium* to distinguish it from the finite-user Nash equilibrium that is the main topic of this paper. The Wardrop equilibrium arises naturally when the flow is considered to be composed of infinitesimal users, so that the effect of each user on link congestion is negligible. This equilibrium concept is also relevant in the context of computer networks, as many of the current dynamic routing protocols focus on shortest path routing. For recent overviews of the extensive literature that has been devoted to the Wardrop equilibrium and its variants see, for example, [23], [21], [2].

When cost-optimizing users control non-negligible amount of flow, we are led to consider the standard Nash equilibrium for finitely-many users. A flow pattern is considered to be a Nash equilibrium point (NEP) if no user can reduce the total cost incurred by its traffic by a unilateral change of its own flow pattern. While this problem has more recent roots than the Wardrop equilibrium, it has attracted considerable attention in recent literature. Existence, uniqueness and some basic properties of the Nash equilibrium are studied in [22], [1], [3]. The notion of a mixed Nash-Wardrop equilibrium, which combines infinitesimal users with finitely-sized ones, is considered in [12], [6]. Efficient network design and management are considered in [15], [16], [18], [17], [11], while [26] bounds the performance degradation relative to centralized routing (along with similar results for the Wardrop equilibrium). The convergence of some dynamic schemes to the Nash equilibrium is considered in [14], while [19] considers a repeated game version of the routing problem, and [4] considers the effect of adding side-constraints on the flows.

We focus here on the issue of uniqueness of the Nash equilibrium in noncooperative routing with finitely-manu users. For a two-node network with parallel links, uniqueness of the Nash equilibrium has been established under mild convexity assumptions on the link costs [22]. This result does not hold for networks of general topology, as demonstrated there by a counter-example. However, the question of whether there exist other network topologies for which uniqueness of the Nash equilibrium is guaranteed (under similar convexity assumptions) remained open until now.

For networks of general topology, the uniqueness of the Nash equilibrium has been established under various additional conditions on the cost functions. A general set of conditions is implied by the requirement of diagonal strict convexity, which is a well known sufficient condition for uniqueness of the Nash equilibrium in convex (or concave) games [24]. These conditions have been applied to the Nash routing problem in [13], [22]. Unfortunately, those conditions do not hold in many cases of interest – for example, they are violated by popular M/M/1 latency function under significant congestion. Other special cases are presented in [1] and [3]: the first considers link latencies that are polynomial with

O. Richman is a graduate student in the Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel roran@tx.technion.ac.il

N. Shimkin is with the Department of Electrical Engineering, Technion – Israel Institute of Technology, Haifa 32000, Israel shimkin@ee.technion.ac.il

0-7803-9568-9/05/$20.00 ©2005 IEEE

a low enough order, while the latter establishes uniqueness under some specific symmetry conditions.

For the Wardrop equilibrium, a corresponding line of uniqueness results exists, with the requirement of link cost convexity replaced by monotonicity of the link latency. This is sufficient to guarantee uniqueness in the single-class case, but not for the multi-class problem [9], [8]. Additional conditions on the costs that ensure uniqueness are considered in [7], [8], [3]. In a recent paper, Milchtaich [20] provides a complete characterization of all two-terminal network topologies (called *nearly parallel networks*) for which uniqueness is guaranteed under the basic monotonicity requirement. We shall further elaborate on these notions below.

Our goal here is to characterize those network topologies for which the *Nash equilibrium* is unique, for any number and size of users, as long as their link cost functions satisfy some mild convexity conditions. Our main results establish that the class of networks that satisfy this property coincides with the set of nearly-parallel networks.

II. MODEL AND PRELIMINARIES

A. The Network Model

Let the network topology be specified by an undirected graph $\mathcal{G} = \mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V}$ is a finite set of vertices (or *nodes*) and $\mathcal{E}$ is a finite set of edges. Each edge joins two distinct vertices. Thus, single-edge loops are not allowed, but more than one edge can join two vertices. Two of the vertices in this graph will be designated as the *terminal* vertices, O (for origin) and D (for destination). We refer to such a graph as a *two-terminal network*.

The actual network we consider is directional, and is obtained from the undirected graph by replacing each edge with two directional links. Thus, an edge e between vertices u and v is split into two directional links, one from u to v and the other from v to u. (We shall comment about that particular way of defining the directional network at the end of section 3. For now, note that by imposing large enough costs on some of the links we may be effectively obtain any subnetwork thereof.) The set of links that connect u to v is denoted by $\mathcal{L}(u, v)$ and the set of all links in the directional network by $\mathcal{L}$.

We are given a set $\mathcal{I} = \{1, 2, \ldots, I\}$ of users, which share the network. Each user i needs to deliver a positive amount d^i of flow from node O to node D, and should decide how to divide its flow between the different routes that connect these two nodes. Denote by f_l^i the flow of user i through link l, and let $f_l = \sum_{i\in\mathcal{I}} f_l^i$ denote the total flow on link l. User i's flow pattern is the vector $f^i = (f_l^i,\ l \in \mathcal{L})$. Finally, the system flow pattern f is the vector of all user flow patterns.

A feasible flow pattern must obey the following flow conservation and positivity constraints for each user:

$$\sum_{l\in Out(v)} f_l^i = \sum_{l\in In(v)} f_l^i + d_v^i,\ \ v \in \mathcal{V}; \quad f_l^i \geq 0\,. \tag{1}$$

Here $In(v)$ and $Out(v)$ are the set of input and output links for node v, and d_v^i is the external flow: $d_v^i = d^i$ for $v = O$, $d_v^i = -d^i$ for $v = D$, and $d_v^i = 0$ otherwise. We denote the set of feasible flow patterns f^i for user i by F^i. This is clearly a convex polyhedron. A system flow pattern f is feasible if $f^i \in F^i$ for all $i \in \mathcal{I}$.

B. Convex Network Games

The performance measure to be minimized by user $i \in \mathcal{I}$ is specified by a cost function $J^i(f)$. We shall consider additive cost functions of the form

$$J^i(f) = \sum_{l\in\mathcal{L}} J_l^i(f_l^i, f_l)\,. \tag{2}$$

Thus, the cost incurred by a user on a link l depends only on its own flow f_l^i, as well as on the total link flow f_l which measures the link congestion. Link costs are often taken to be in the form $J_l^i(f_l^i, f_l) = f_l^i T_l^i(f_l)$, where $T_l^i(f_l)$ represents the cost per unit flow (or *latency*). Note that the link cost functions J_l^i may depend on the user i, and similarly for the latency T_l^i.

Definition 1: A flow vector f_* is a Nash equilibrium point (NEP) if, for each $i \in \mathcal{I}$,

$$J^i(f_*) = \min_{f^i\in F^i} J^i(f_*^1, \ldots, f_*^{i-1}, f^i, f_*^{i+1}, \ldots, f_*^I). \tag{3}$$

Denote by $\mathcal{J}$ the vector of link costs functions $(J_l^i,\ l \in \mathcal{L}, i \in \mathcal{I})$. Also denote by d the vector $(d^1, \ldots, d^I)$ which specifies the demand of all users. Let

$$K_l^i(f_l^i, f_l) \triangleq \frac{\partial J_l^i(f_l^i, f_l)}{\partial f_l^i} + \frac{\partial J_l^i(f_l^i, f_l)}{\partial f_l} \tag{4}$$

denote the marginal costs of user i on link l with respect to its flow f_l^i (where the second term is required since $f_l = \sum_i f_l^i$). We shall assume that the link cost functions satisfy the following properties.

Assumption A1: J_l^i is a continuous, non-negative function on its domain $\{(x, y) \in \Re_+^2 : x \leq y\}$.

Assumption A2: J_l^i is strictly increasing in each of its arguments (except possibly when $f_l^i = 0$).

Assumption A3: J_l^i is continuously differentiable, and the marginal cost $K_l^i(f_l^i, f_l)$ is strictly increasing in each of its arguments.

Functions that satisfy these assumptions were termed *type-A cost functions* in [22]. We note that in this reference the costs were allowed to take infinite values, in order to accommodate popular cost functions such as the M/M/1 latency function: $J_l^i(f_l^i, f_l) = \frac{f_l^i}{c_l - f_l}$, with $J_l^i = \infty$ for $f_l \geq c_l$. This extension is easily accommodated in the present paper, with all uniqueness results directly applicable to finite NEPs (namely, NEPs where all users incur finite costs). For details see [25].

Definition 2: A *convex network game* over a two-terminal network $\mathcal{G}$ is a triple $(\mathcal{I}, d, \mathcal{J})$ over $\mathcal{G}$, with cost functions that satisfy Assumptions A1-A3 for each link l and user i.

For a convex network game, the minimization problem faced by each user (with the flow pattern of the others held fixed) is a convex optimization problem, as convexity of the link costs in each user's flow is implied by our Assumption A3. Necessary and sufficient conditions for a flow pattern to be a Nash equilibrium are therefore provided by the Karush-Kuhn-Tucker conditions, applied to each user in turn [22]: A flow pattern f is a Nash equilibrium if there exists a set of constants $\{\lambda_u^i : \; i \in \mathcal{I}, \; u \in \mathcal{V}\}$ (the marginal cost parameters) so that for every link $l \in \mathcal{L}(u,v)$, and for every user i:

$$\lambda_u^i = K_l^i(f_l^i, f_l) + \lambda_v^i \qquad \text{if } \; f_l^i > 0\,, \tag{5}$$

$$\lambda_u^i \leq K_l^i(f_l^i, f_l) + \lambda_v^i \qquad \text{if } \; f_l^i = 0\,. \tag{6}$$

Conditions (5) and (6) can also be expressed in the following path-oriented manner: For any two nodes in the network u and v, and any path p that connects u and v, if $f_l^i > 0$ for every $l \in p$ then

$$\lambda_{uv} \triangleq \lambda_u - \lambda_v = \sum_{l \in p} K_l^i(f_l^i, f_l) \leq \sum_{l \in p'} K_l^i(f_l^i, f_l) \tag{7}$$

where p' is any other path connecting u and v.

In general, the NEP of a convex network game need not be unique. Still, uniqueness is guaranteed for certain network topologies. We shall refer to this property as *topological uniqueness*. More precisely:

Definition 3: A network $\mathcal{G}$ has the *topological uniqueness property* if the NEP is unique for any convex network game over $\mathcal{G}$.

C. Nearly Parallel Networks

We briefly repeat here some definitions and results from [20] concerning nearly-parallel networks. These results show that network topologies can be divided into two classes. The first essentially contains the networks shown in Figure 1, and networks which are serial connection of those networks. The second class contains all networks in which one of the basic networks shown in Figure 2 is embedded, in the following sense.

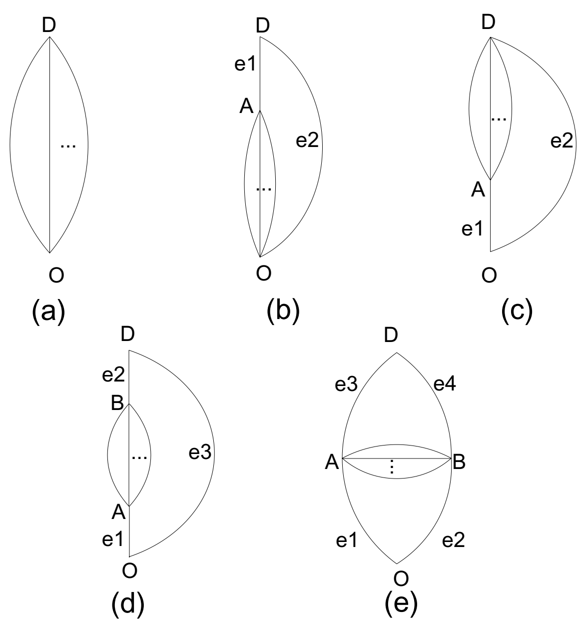

Fig. 1. Basic networks that define the class of nearly-parallel networks

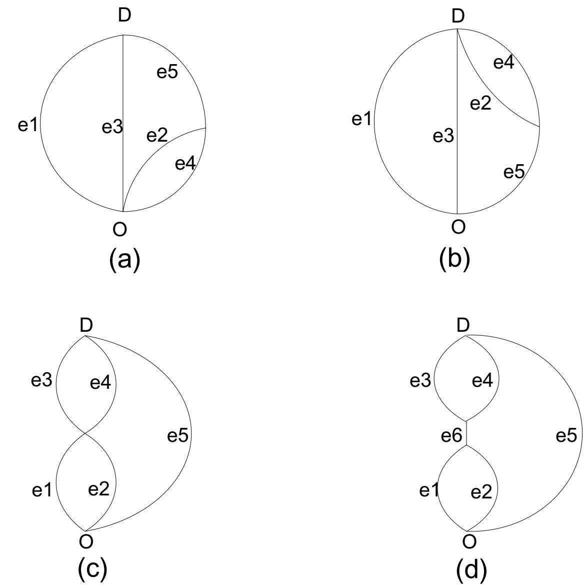

Fig. 2. Basic networks that are not nearly-parallel

Definition 4: A network $\mathcal{G}'$ is said to be embedded in the wide sense in network $\mathcal{G}''$ if the latter can be obtained from the former by some sequence of the following three operations: (1) *Edge subdivision:* An edge is replaced by two edges with a single common end vertex. (2) *Edge addition:* The addition of a new edge joining two existing vertices. (3) *Terminal vertex subdivision:* The addition of a new edge, joining the terminal vertex O or D with a new vertex v, such that a nonempty subset of the edges originally incident with the terminal vertex are incident with v instead.

Definition 5: A two-terminal network $\mathcal{G}$ is called *nearly parallel* if it is one of the networks in Figure 1, or can be constructed from one of the networks in Figure 1 by a series of edge subdivisions.

Of the five networks in Figure 1, network (e) is the most interesting, as the other four may be considered a special case of this network for routing purposes. However, the formal definition of nearly parallel networks does require to consider all these basic networks. Note also that only network (e) supports meaningful bi-directional traffic (between nodes A and B) when routing traffic between the prescribed source and destination nodes.

Proposition 1 ([20]): For every two-terminal network $\mathcal{G}$, one, and only one, of the following conditions holds:

(i) $\mathcal{G}$ is nearly parallel or is a serial connection of two or more nearly parallel networks.
(ii) One (or more) of the networks in Figure 2 is embedded in the wide sense in $\mathcal{G}$.

To simplify terminology, from here on we shall use the term "nearly parallel network" to refer to any network that meets condition (i) of the last proposition, namely both to nearly parallel networks in the sense of Definition 5 *and* to serial connections thereof.

III. UNIQUENESS OF THE NASH EQUILIBRIUM

It has been established in [22] that parallel-link networks posses the topological uniqueness property, namely that uniqueness of the NEP is guaranteed under Assumptions A1-A3. Our main result states that topological uniqueness (in the sense of Definition 3) holds for the larger class of nearly parallel networks, and *only* for that class.

Theorem 1: A two-terminal network $\mathcal{G}$ has the topological uniqueness property if, and only if, $\mathcal{G}$ is a nearly parallel network.

The proof is presented in the following subsections.

A. Uniqueness for Nearly Parallel Networks

The following proposition repeats the sufficiency part of Theorem 1.

Proposition 2: If $\mathcal{G}$ is nearly parallel, then the NEP is unique for every convex network game $(I, d, \mathcal{J})$ over $\mathcal{G}$.
The proof, which is omitted here due to space limitations, may be found in [25].

B. Counter-examples to Uniqueness

This subsection establishes the necessity part of Theorem 1. We will show that a counter-example (namely, a convex network game with multiple Nash equilibria) may be found for any network which is not nearly parallel. Recall from Proposition 1 that in any network which is not nearly parallel, one of the networks in Figure 2 is embedded in the wide sense. We start with two basic examples:

Example 1: Consider the network in Figure 2(a). Let the cost functions be of the form $J_l^i(f_l^i, f_l) = f_l^i T_l^i(f_l)$, with the latency functions $T_l^i(x)$ and the demand d^i for each user given in the following table

user	d^i	$e1$	$e2$	$e3$	$e4$	$e5$
1	6	"∞"	"∞"	$7x$	x	$f_1(x)$
2	4	"∞"	x	$f_2(x)$	"∞"	$2x$
3	4	$x+21$	"∞"	"∞"	$f_2(x)$	x

where

$$f_1(x) = \begin{cases} x & \text{if } x < 6 \\ \frac{1}{3}\left(e^{3(x-6)} + 17\right) & \text{if } x \geq 6 \end{cases} \tag{8}$$

and

$$f_2(x) = \begin{cases} x & \text{if } x < 4 \\ \frac{1}{3}\left(e^{3(x-4)} + 11\right) & \text{if } x \geq 4 \end{cases}. \tag{9}$$

Note that this functions are continuously differentiable. The infinite values in the table indicate large enough functions, so that the user always prefers not to use this link.
Each user can thus choose to divide its flow between two different routes. User 1 can choose between $e3$ and $e4$–$e5$, user 2 can chose between $e2$–$e5$ and $e3$, and user 3 can choose between $e4$–$e5$ and $e1$. It is easily verified that one Nash equilibrium is formed if each user diverts all of its flow to its first option, and another Nash equilibrium is formed if each user will divert all of its flow to the second option. For example, in the first NEP, (7) may be verified by noting that for user 1,

$$K_{e3}^1(6,6) < K_{e4}^1(0,4) + K_{e5}^1(0,8) \tag{10}$$

and in the second NEP:

$$K_{e4}^1(6,6) + K_{e5}^1(6,6) < K_{e3}^1(0,4). \tag{11}$$

Similar inequalities may be verified for users 2 and 3.

Example 2: A similar example can be constructed for the network in Figure 2(c). The demand and $T_l^i(x)$ functions for each user are:

user	d^i	$e1$	$e2$	$e3$	$e4$	$e5$
1	6	$2x$	"∞"	$2x$	"∞"	$f_1(x)$
2	4	$f_2(x)$	"∞"	"∞"	x	$6x$
3	4	"∞"	x	$f_2(x)$	"∞"	$6x$

where $f_1(x)$ and $f_2(x)$ are as defined in Example 1. Each user can choose how to divide its flow between link $e5$ and some other route, $e1$–$e3$ for user 1, $e1$–$e4$ for user 2 and $e2$–$e3$ for user 3. It may be verified as above that one Nash equilibrium is obtained when user 1 ships all its flow through $e5$, while users 2 and 3 avoid $e5$. Another Nash equilibrium is obtained when users 2 and 3 ship all their demand on $e5$, while user 1 selects the $e1$–$e3$ path.

These two examples show that multiple equilibria exist in the networks of Figure 2(a) and 2(c). We need now to extend the examples to the other networks in Figure 2, and then to any network in which these basic networks are embedded. To this end, we will require the considered equilibrium point to be stable with respect to small perturbations, so that the series addition of links with small enough cost does not alter the equilibrium. We use the following definition.

Definition 6: A Nash equilibrium of the network game is called *strong* if for any path p from O to D that is used by user i, namely $f_l^i > 0$ for every $l \in p$, it holds that

$$\sum_{l \in p} K_l^i(f_l^i, f_l) < \sum_{l \in p'} K_l^i(f_l^i, f_l) \tag{12}$$

for any other path p' that connects O and D.
Note that in a strong NEP each user employs a unique path from origin to destination. It may be verified that the NEPs in examples 1 and 2 are strong.

Lemma 1: Let $\mathcal{G}$ be a network over which there exists a convex network game with two distinct strong NEPs. Then for any network $\mathcal{G}'$ in which $\mathcal{G}$ is embedded in the wide sense, there exists a convex network game with two different strong NEPs.

Proof: The proof is similar to that of a corresponding claim in [20]. Let f and $\hat{f}$ be the two strong NEPs in $\mathcal{G}$ and denote by p^i and $\hat{p}^i$ the unique paths of user i in f and $\hat{f}$ respectively. Since the definition of embedding in the wide sense is recursive we need only consider the case where $\mathcal{G}'$ was obtained from $\mathcal{G}$ by one of the following operations: (1) the subdivision of an edge, (2) the addition of an edge, or (3) the subdivision of a terminal vertex. In case (1) the cost function of each direction of the edge that was subdivided is equally split between its two parts. It is trivially seen that the new game over $\mathcal{G}'$ remains a convex network game, which supports the two distinct equilibria of $\mathcal{G}$. In case (2) we may set the cost functions of each user on the added edge so that $K_l^i(0,0)$ is higher then

$\max(\sum_{l\in p^i} K_l^i(f_l^i, f_l), \sum_{l\in \hat{p}^i} K_l^i(\hat{f}_l^i, \hat{f}_l))$. In that case no user has an incentive to use the new edge and the equilibrium points do not change. In case (3) a new node is added to $\mathcal{G}$, with an edge that connects it either to D or to O. If we would set the cost on that link to zero the NEPs would obviously not be affected. However, since a null cost violates Assumption A3, we choose a small non-zero cost function for that link. Since the two equilibria in $\mathcal{G}$ are strong we may set the marginal costs $K_l^i(f_l^i, f_l)$ on that link sufficiently small f and $\hat{f}$ are still NEPs. Specifically we select J_l^i so that $K_l^i(f_l^i, f_l)$ is less then the difference between used routes and unused routes for $f_l^i \leq d^i$. □

Proposition 3: For every network $\mathcal{G}$ in which one of the networks in Figure 2 is embedded in the wide sense, one can find a convex network game for which the equilibrium is not unique.

Proof: Example 1 demonstrates the claim for the network in Figure 2(a). A symmetric example can be used for the network in Figure 2(b). Example 2 shows the same for the network in Figure 2(c). We can apply Example 2 to the network in Figure 2(d) by imposing small enough costs on the additional link $e6$ without affecting the two strong equilibria (as argued in Lemma 1). Lemma 1 now proves the proposition. □

Propositions 2 and 3 together with Proposition 1 complete the proof of Theorem 1.

IV. Weakly Convex Network Games

In this section we consider the uniqueness of the Nash equilibrium under slightly weaker conditions on the link cost functions. These weaker conditions enable us to embed the *Wardrop* equilibrium (with a finite number of user classes) within the finite-user game model. We start by delineating the relation between the (multiclass) Wardrop equilibrium and the Nash equilibrium in our model.

A. Wardrop and Nash Equilibria

Consider the same network model as defined above, except that the user index $i \in I$ now designates a *user class*. Each user class may be thought of as a continuum of infinitesimal users, all sharing the same cost characteristics. The latency of link l for class-i users is given by $T_l^i(f_l)$, which we assume to be a positive and strictly increasing function. A flow profile $\mathbf{f}$ is a (multiclass) Wardrop equilibrium if

$$\sum_{l\in p^i} T_l^i(f_l) = \min_{p} \sum_{l\in p} T_l^i(f_l) \quad \text{for every } i \in I\,, \tag{13}$$

where p^i is any route employed by user-class i, and p' is any other feasible route for that class.

It was shown in [20] that the Wardrop equilibrium is unique for any choice of (non-negative, strictly increasing) latency functions $T_l^i(f_l)$ *if and only if* the network is nearly parallel.

The Wardrop equilibrium and the Nash equilibrium with finitely-many (atomic) users may be related from two different viewpoints. First, the Wardrop equilibrium may be obtained as the limit of the Nash equilibrium point when the number of users is increased to infinity while their individual flow demands decrease accordingly [13]. More relevant here, the Wardrop equilibrium is mathematically equivalent to a finite-user Nash equilibrium with properly defined costs, to be specified shortly. This relation was already observed in [5] for the Wardrop equilibrium with a single user class, which is well known to be equivalent to a (single-user) optimization problem; see also [9]. Later, [10] indicated the equivalence of the (still single-class) Wardrop equilibrium to the Nash equilibrium in a routing game where a distinct player is assigned to each origin-destination pair.

Returning to our Wardrop equilibrium problem with link latencies $T_l^i(f_l)$, consider a corresponding routing game where each user i corresponds to user class i, and let the link costs for that user be given by

$$J_l^i(f_l^i, f_l) \triangleq \int_0^{f_l^i} T_l^i(f_l - f_l^i + x)dx\,. \tag{14}$$

Recalling (4), it is easily verified that $K_l^i(f_l^i, f_l) = T_l^i(f_l)$, namely T_l^i is the marginal cost for this cost function. As T_l^i is strictly increasing in f_l by assumption, it follows that the cost of each user is strictly convex in its own decision variables. The equivalence between the Nash equilibrium of this routing game and the Wardrop equilibrium in the original model follows immediately by comparing the optimality conditions for the Nash equilibrium in (7) with the definition of the Wardrop equilibrium above.

It is readily seen that the cost functions in (14) satisfy our basic Assumptions A1-A3, *except* for the fact that $K_l^i(f_l^i, f_l)$ is not strictly increasing in f_l^i (as it is only a function of f_l), which violates Assumption A3. This motivates us to consider a weaker version of this assumption.

B. Uniqueness of the NEP for Weakly Convex Games

Consider the following relaxed version of Assumption A3.

Assumption A3': Same as Assumption A3, except that the cost functions $K_l^i(f_l^i, f_l)$ are only required to be *weakly* increasing in f_l^i.

We define a *weakly convex network game* similarly to a convex network game, except that Assumption A3 is replaced by Assumption A3'. Under this weaker assumption, the Nash equilibrium need no longer be unique even for nearly-parallel networks. Still, the following uniqueness results can be established; for a proof and further remarks see [25].

Theorem 2: Let $(I, d, \mathcal{J})$ be a weakly convex game over a network $\mathcal{G}$. If $\mathcal{G}$ is nearly parallel, then the following uniqueness properties hold for any Nash equilibrium point:

(i) The link flows (f_l, $l \in L$) are unique.
(ii) The marginal link costs (namely $K_l^i(f_l^i, f_l)$ for all l, i) are unique.
(iii) For any user i whose costs satisfy the stronger Assumption A3, this user's flows (f_l^i, $l \in L$) are unique. Consequently, the link costs for this user are unique as well.

We note that the last part of this theorem in fact implies Proposition 2. Further note that Assumption A3' does not suffice to ensure uniqueness of the per-user flows, as can be seen by a simple counter-example.

As outlined in the previous subsection, a multiclass Wardrop equilibrium with non-negative, strictly increasing latency functions $T_l^i(f_l)$ may be represented as a weakly convex network game model. Theorem 2 thus recovers the uniqueness of the link flows for the Wardrop equilibrium for nearly parallel networks. Furthermore, this theorem may be applied in an obvious manner to the mixed Nash-Wardrop problem [6]. Thus, Theorem 2 applies to any combination of large (atomic) users and (a finite number of) infinitesimal-user classes.

V. CONCLUSION

As networks become larger and less centralized, it is usually hard to give theoretical predictions regarding the precise operating conditions of the network. Equilibrium analysis provides a useful proxy for this purpose, which has been used both for the qualitative understanding of basic phenomena, as well as for setting up the quantitative models that are essential for network management. Uniqueness of the equilibrium is important both for analysis and management. When the equilibrium is not unique, the network behavior becomes less predictable. Simulation results, for example, cannot be relied on to give a complete picture of the network operation. From the management point of view, several schemes have propose the use of pricing and related management tools to enforce the efficiency of the equilibrium. It is usually easy to maintain the efficiency of some equilibrium point, but when the equilibrium is not unique it is not clear that the intended equilibrium point will indeed be the one to take effect, and the management task tends to become much harder.

This paper provides a complete characterization of two-terminal network topologies for which the Nash equilibrium is unique, under broad conditions on the cost functions, and for any number and size of network users. Unfortunately, the class of networks for which this broad sense of uniqueness holds is restricted. Thus, alongside the verification of uniqueness for nearly parallel networks, the result also points out those network configuration that might bring about multiple equilibria.

Our analysis in Section IV applies to the mixed Nash-Wardrop equilibrium problem, provided that the number of infinitesimal-user classes is finite. A more general model, which allows a continuum of infinitesimal-user classes alongside the large (atomic) users, is considered in [25].

We have not dealt in this paper with multi-terminal networks, in the sense that different flows (of different users, or even of the same user) may correspond to different source and destination pairs. While either necessary or sufficient conditions may be extracted from our results, it remains open whether a complete characterization of topological uniqueness may be given for this case.

VI. ACKNOWLEDGMENTS

The authors gratefully acknowledge Igal Milchtaich for useful discussions and suggestions regarding this work. This research was supported by Grant No. 2002112 from the United States–Israel Binational Science Foundation (BSF).

REFERENCES

[1] E. Altman, T. Başar, T. Jiménez, and N. Shimkin. Competitive routing in networks with polynomial cost. *IEEE Trans. Automat. Control*, 47:92–96, 2002.
[2] E. Altman, T. Boulogne, R. El Azouzi, T. Jiménez, and L. Wynter. A survey on networking games in telecommunications. *Computers and Operations Research*, 33(2):286–311, 2005.
[3] E. Altman and H. Kameda. Equilibria for multiclass routing in multi-agent networks, 2001. In Proc. 40th Annual IEEE Conference on Decision and Control.
[4] R. Azouzi and E. Altman. Constrained traffic equilibrium in routing. *IEEE Trans. Automat. Control*, 48(9):1656–1660, 2003.
[5] M. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, 1956.
[6] T. Boulogne, E. Altman, H. Kameda, and O. Pourtallier. Mixed equilibrium (ME) for multiclass routing games. *IEEE Trans. Automat. Control*, 47(6):903–916, 2002.
[7] S. C. Dafermos. An extended traffic assignment model with application to two-way traffic transportation science. *Transportation Sci.*, 5:366–389, 1971.
[8] S. C. Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transportation Sci.*, 6:73–87, 1972.
[9] S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *J. Res. Nat. Bur. Standards*, 37B:91–118, 1969.
[10] S. Devarajan. A note on network equilibrium and noncooperative games. *Transportation Res. Part B*, 15B:421–426, 1981.
[11] R. El Azouzi E. Altman and O. Pourtallier. Avoiding paradoxes in multi-agent competitive routing. *Computer Networks*, 43:133–146, 2003.
[12] P. Harker. Multiple equilibrium behaviors on networks. *Transportation Sci.*, 22:39–46, 1988.
[13] A. Haurie and P. Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.
[14] T. Jiménez, E. Altman, T. Başar, and N. Shimkin. Routing into two parallel links: game-theoretic distributed algorithms. *J. Parallel Distrib. Comput.*, 61:1367–1381, 2001. Special Issue on Routing in Computer and Communication Systems.
[15] Y. Korilis, A. Lazar, and A. Orda. Architecting noncooperative networks. *IEEE J. Selected Areas in Communication (Special Issue on Advances in the Fundamentals of Networking)*, 13(7):1241–1251, 1995.
[16] Y. Korilis, A. Lazar, and A. Orda. Capacity allocation under non-cooperative routing. *IEEE Trans. Automat. Control*, 42(3):309–325, 1997.
[17] Y. Korilis, A. Lazar, and A. Orda. Avoiding the braess paradox in non-cooperative networks. *J. Appl. Probab.*, 36:211–222, 1999.
[18] Y. A. Korilis, T. A. Varvarigou, and S. R. Ahuja. Incentive-compatible pricing strategies in noncooperative networks. In *Proc. IEEE INFOCOM'98*, pages 439–446, San-Francisco, CA, April 1998.
[19] R. J. La and V. Anantharam. Optimal routing control: Repeated game approach. *J. Appl. Probab.*, 36:211–222, 1999.
[20] I. Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.*, 30:225–244, 2005.
[21] A. Nagurney. *Network Economics: A Variational Inequality Approach*. Kluwer Academic, Dordrecht, The Netherlands, 2nd edition, 1999.
[22] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multi-user communication networks. *IEEE/ACM Trans. Networking*, pages 510–521, 1993.
[23] M. Patriksson. *The Traffic Assignment Problem: Models and Methods*. VSP, Utrecht, The Netherlands, 1994.
[24] J. B. Rosen. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica*, 33(3):520–534, 1965.
[25] O. Richman and N. Shimkin. Topological Uniqueness of the Nash Equilibrium for Atomic Selfish Routing. Submitted to *Math. Oper. Res.*, revised, Sept. 2005 (available online).
[26] T. Roughgarden and E. Tardos. How bad is selfish routing. *J. ACM*, 49:236–259, 2002.
[27] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proc. Inst. Civ. Eng.*, 2:325–378, 1952.