MATHEMATICS OF OPERATIONS RESEARCH
Vol. 32, No. 1, February 2007, pp. 215–232
ISSN 0364-765X | EISSN 1526-5471 | 07 | 3201 | 0215
DOI 10.1287/moor.1060.0229

Topological Uniqueness of the Nash Equilibrium for Selfish Routing with Atomic Users

Oran Richman, Nahum Shimkin
Department of Electrical Engineering, Technion–Israel Institute of Technology, Haifa 32000, Israel
{roran@tx.technion.ac.il, shimkin@ee.technion.ac.il, http://www.ee.technion.ac.il/people/shimkin/}

We consider the problem of selfish routing in a congested network shared by several users, where each user wishes to minimize the cost of its own flow. Users are atomic, in the sense that each has a nonnegligible amount of flow demand, and flows may be split over different routes. The total cost for each user is the sum of its link costs, which, in turn, may depend on the user's own flow as well as the total flow on that link. Our main interest here is network topologies that ensure uniqueness of the Nash equilibrium for *any* set of users and link cost functions that satisfy some mild convexity conditions. We characterize the class of two-terminal network topologies for which this uniqueness property holds, and show that it coincides with the class of *nearly parallel networks* that was recently shown by Milchtaich [Milchtaich, I. 2005. Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.* **30** 225–244] to ensure uniqueness in nonatomic (or Wardrop) routing games. We further show that uniqueness of the link flows holds under somewhat weaker convexity conditions, which apply to the mixed Nash-Wardrop equilibrium problem. We finally propose a generalized continuum-game formulation of the routing problem that allows for a unified treatment of atomic and nonatomic users.

Key words: selfish routing; congested networks; Nash equilibrium; Wardrop equilibrium; continuum games
MSC2000 subject classification: Primary: 90B10; secondary: 91A10
OR/MS subject classification: Primary: multicommodity networks; secondary: noncooperative games
History: Received October 21, 2004; revised September 21, 2005.

1. Introduction. Congestion-prone networks have been an object of interest in engineering and operations research for more than five decades, motivated to a large extent by their applications in transportation science and, more recently, in computer communication networks. In these 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 multiuser 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 Altman et al. [3].

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 [32]. Essentially, it requires all traffic to occupy paths with minimal latency. While this basic concept has been addressed by different names, including the Nash equilibrium for infinitesimal users, small user equilibrium or traffic equilibrium, we shall mostly use the term *Wardrop equilibrium* to distinguish it from the finite-user Nash equilibrium that is in the focus 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 routing protocols focus on shortest path routing. For recent overviews of the extensive literature that concerns the Wardrop equilibrium and its variants, see, for example, Patriksson [26], Nagurney [24], Altman et al. [3], and Roughgarden [29].

When cost-optimizing users have control over nonnegligible amounts of flow, we are led to consider the standard Nash equilibrium for finitely many users. We refer to such users as *large* or *atomic* users. A flow profile is a Nash equilibrium point (NEP) if no user can reduce its own cost through a unilateral change of its own flow profile. One of the first papers to study this problem is Haurie and Marcotte [14], which shows convergence of the Nash equilibrium to the Wardrop equilibrium as the number of users increases to infinity. Existence, uniqueness, and some basic properties of the Nash equilibrium are studied in Orda et al. [25], Altman et al. [2], and Altman and Kameda [1]. The notion of a mixed Nash-Wardrop equilibrium, which combines infinitesimal users with positively sized ones, is considered in Harker [13] and Boulogne et al. [6]. Efficient network design and management are considered in Korilis et al. [17, 18], Korilis et al. [20], Korilis et al. [19], and El Azouzi et al. [12], while Roughgarden and Tardos [30] bound 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 Jiménez et al. [15], while La and Anantharam [21] consider a repeated game version of the routing problem, and Azouzi and Altman [4] consider the addition of side constraints on link flows.

We focus here on the issue of uniqueness of the Nash equilibrium in noncooperative routing with atomic 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 (Orda et al. [25]). This result does not hold for networks of general topology, as demonstrated there through a specific counterexample. 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.

For networks of general topology, uniqueness of the Nash equilibrium requires additional conditions on the cost functions. A general set of conditions is related to the notion of diagonal strict convexity, which is a well-known sufficient condition for uniqueness of the Nash equilibrium in convex (or concave) games (Rosen [28]). These conditions were applied to the Nash routing problem in Haurie and Marcotte [14] and Orda et al. [25]. 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. More specific uniqueness conditions are presented in Altman et al. [2] and Altman and Kameda [1]: the first considers link latencies that are polynomial with 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 multiclass problem (Dafermos and Sparrow [9], Dafermos [8]). Additional conditions on the costs that ensure uniqueness are considered in Dafermos [7, 8], Altman and Kameda [1], and Marcotte and Wynter [22]. In a recent paper, Milchtaich [23] provides a complete characterization of all two-terminal network topologies (called *nearly parallel networks*) for which uniqueness is guaranteed under the basic monotonicity requirement. These results are most relevant for the present paper and will be further discussed in the sequel.

The goal of this paper 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. We will also show how these uniqueness results may be extended to the joint Nash-Wardrop problem. This will first be done by observing that the Wardrop equilibrium (with a finite number of user classes) can be represented as a Nash equilibrium in our basic finite-user model, under slightly relaxed convexity assumptions. We will also introduce a continuum game model, that allows a unified and general treatment of atomic and nonatomic users.

The paper is organized as follows: In §2, we present the basic model, and repeat the definition of nearly parallel networks from Milchtaich [23]. Section 3 establishes our main results concerning uniqueness of the Nash equilibrium. In §4, we extend our uniqueness analysis under somewhat relaxed conditions, and show how the Wardrop (or mixed Nash-Wardrop) equilibrium problem can be cast in terms of this model. The continuum game model is considered in §5. Section 6 offers some concluding remarks.

2. Model and preliminaries.

2.1. The network model. Let the network topology be specified by an undirected graph $\mathcal{G} = \mathcal{G}(V, E)$, where V is a finite set of vertices (or *nodes*) and 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 terminal vertices, O (for origin) and D (for destination). We further assume that each edge belongs to some simple path (i.e., a path with no repeat vertices) from O to D. We refer to such a graph together with its O-D pair as an *undirected two-terminal network*.

The actual network contains *directed* links, and is obtained from the undirected graph by replacing each edge with two links. More precisely, an edge e between vertices u and v is split into two directed links, one from u to v and the other from v to u. Thus the resulting network is *bidirectional*, in the sense that each link is paired with another link of opposite direction. (We shall comment on the case of general directed networks at the end of §3. For now, note that by imposing large enough costs on some of the links, one may effectively obtain any subnetwork of the original bidirectional one.) The set of links that connect u to v is denoted by L_{uv}, while L is the set of all links in the directed network.

We are given a set $I=\{1,2,\ldots,n_I\}$ of users who share the given 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 $f_l=\sum_{i\in I} f_l^i$ denote the total flow on link l. User i's flow profile is the vector $\mathbf{f}^i=(f_l^i,\ l\in L)$. The system flow profile $\mathbf{f}$ is the vector of all user flow profiles.

Let $\mathrm{IN}(v)$ and $\mathrm{OUT}(v)$ denote the set of input and output links to node v, and let

$$d_v^i=\begin{cases} d^i\colon & v=O\\ -d^i\colon & v=D\\ 0\colon & v\neq O,D \end{cases} \tag{1}$$

denote the external flow of user i to node v. A feasible flow profile for user i must obey the following flow conservation and positivity constraints:

$$\sum_{l\in\mathrm{OUT}(v)} f_l^i=\sum_{l\in\mathrm{IN}(v)} f_l^i+d_v^i,\quad v\in V \tag{2}$$

$$f_l^i\geq 0,\quad l\in L. \tag{3}$$

We denote the set of feasible flow profiles $\mathbf{f}^i$ for user i by F^i. This is clearly a closed convex polyhedron. A system flow profile $\mathbf{f}$ is feasible if $\mathbf{f}^i\in F^i$ for all $i\in I$.

We note that our definition of feasible flows does not exclude the possibility of cyclic flows (hence F^i is unbounded when the network topology includes cyclic paths). However, our assumptions on the user cost functions will exclude the existence of cyclic flows in equilibrium.

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

$$J^i(\mathbf{f})=\sum_{l\in L}\tilde{J}_l^i(\mathbf{f}),\quad \text{where } \tilde{J}_l^i(\mathbf{f})=J_l^i(f_l^i,f_l). \tag{4}$$

Thus the cost incurred by a user on each link depends only on its own flow f_l^i on that link, 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 $J_l^i(f_l^i,f_l)=f_l^iT_l^i(f_l)$, where $T_l^i(f_l)$ represents the cost per unit flow (or *latency*). Note that the link costs J_l^i may depend on the user i, and similarly for the latency T_l^i. The domain of J_l^i is the set $\{(f_l^i,f_l)\colon 0\leq f_l^i\leq d_l^i,\ 0\leq f_l\leq d_l,\ f_l^i\leq f_l\}$, where d_l is the maximal possible flow on link l, and d_l^i is the maximal flow of user i on that link. We usually expect the costs J_l^i to be positive, although this is not necessary. For the time being, we assume that all costs are finite; infinite costs are discussed in the next subsection.

Definition 2.1. A flow profile $\hat{\mathbf{f}}$ is a Nash equilibrium point (NEP) if, for each $i\in I$,

$$J^i(\hat{\mathbf{f}})=\min_{\mathbf{f}^i\in F^i} J^i(\hat{\mathbf{f}}^1,\ldots,\hat{\mathbf{f}}^{i-1},\mathbf{f}^i,\hat{\mathbf{f}}^{i+1},\ldots,\hat{\mathbf{f}}^I). \tag{5}$$

Note that we consider only the pure-strategy Nash equilibrium. Mixed-strategy equilibria do not exist in our model because of the assumed strict convexity of the cost function of each user in its own strategy (or flow vector).

Let

$$K_l^i(f_l^i,f_l)\triangleq\frac{\partial}{\partial f_l^i}\tilde{J}_l^i(\mathbf{f})=\frac{\partial}{\partial f_l^i}J_l^i(f_l^i,f_l)+\frac{\partial}{\partial f_l}J_l^i(f_l^i,f_l) \tag{6}$$

denote the *marginal cost function* of user i on link l. The last term arises since $f_l=\sum_i f_l^i$. We shall impose the following assumptions on the link cost functions:

Assumption A1. *$J_l^i(f_l^i,f_l)$ is a continuous and continuously differentiable function.*

Assumption A2. *$J_l^i(f_l^i,f_l^i+f_l^{-i})$ is strictly increasing in f_l^i (for any $f_l^{-i}\geq 0$).*

Assumption A3. *$K_l^i(f_l^i,f_l)$ is strictly increasing in both arguments.*

The following remarks concern these assumptions:

(1) Assumptions A1–A3 essentially comply with the definition of *type-A* cost functions in Orda et al. [25].

(2) Assumption A2 means that a the link cost for each user is strictly increasing in its own flow on that link. This clearly excludes the existence of routing loops (namely, cyclic flows) in equilibrium, as cancelling such cycles will strictly decrease the cost incurred on each of the links involved.

(3) In our uniqueness proof for nearly parallel networks, Assumption A2 is used only to exclude cyclic flows (see Lemma 3.6 and its proof). Thus, if the network has no cyclic paths, or if cyclic flows are explicitly forbidden, then Assumption A2 may be dispensed with.

(4) Assumption A3 implies that $\tilde{J}_l^i(\mathbf{f})$ is strictly convex in f_l^i. The latter property is essential for uniqueness of the best response flow of each user, without which uniqueness of the equilibrium can hardly be expected.

(5) Assumption A3 can be somewhat relaxed by requiring K_l^i to be only *weakly* increasing in f_l. Uniqueness for nearly parallel networks still holds, as the proof of Proposition 3.2 uses only this weaker monotonicity property. This relaxed requirement allows, for example, to include in the model isolated users, or "high priority" ones who are not affected at all by the flows of the others (but may still affect them). We retain the stronger form of this assumption to not weaken the *necessity* part of Theorem 3.1.

(6) In §4, we consider in detail the consequences of relaxing the monotonicity requirement in Assumption A3 for the *first* argument of K_l^i. We mention already that uniqueness of the per user flows is no longer guaranteed in that case (even though the strict convexity of $\tilde{J}_l^i(\mathbf{f})$ in f_l^i is maintained).

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

Definition 2.2. A *convex network game* over a two-terminal network $\mathcal{G}$ is a triplet $(I, d, \mathcal{J})$ over $\mathcal{G}$, with cost functions that satisfy Assumptions A1–A3 for each link l and user i.

Existence of an NEP in any convex network game essentially follows from standard results on convex games (see Debreu [10] or Rosen [28]), which establish the existence of a pure-strategy NEP for any n-person game with convex compact action sets and continuous cost functions that are convex in the player's own action. See also Orda et al. [25] for details that are specific to the present model. The only fine point to note here is the requirement for compact action sets, while the feasible flow sets F^i are not a priori bounded when the network contains cyclic paths. Still, since the best response flows of all users are devoid of cyclic flows (because of Assumption A2, as explained before), they can be contained within a compact convex subset of F^i. Clearly, then, the set of NEPs with the users' actions restricted to these compact sets coincides the set of NEPs in the unrestricted case, and the conditions for existence of the NEP are satisfied by the former.

Under the above assumptions, the best response flow of each user (against any given flow profile of the others) is obtained through a convex optimization problem. Necessary and sufficient conditions for a flow profile to be a Nash equilibrium are therefore provided by the Karush-Kuhn-Tucker (KKT) conditions, applied to each user in turn (Orda et al. [25]): A feasible flow profile $\mathbf{f}$ is a Nash equilibrium if and only if there exists a set of constants $\{\lambda_u^i\colon i\in I,\ u\in V\}$, so that for every link $l\in L_{uv}$, and for every user i,

$$\lambda_u^i = K_l^i(f_l^i, f_l) + \lambda_v^i \quad \text{if } f_l^i > 0, \tag{7}$$

$$\lambda_u^i \le K_l^i(f_l^i, f_l) + \lambda_v^i \quad \text{if } f_l^i = 0. \tag{8}$$

We refer to $\{\lambda_u^i\}$ as *the marginal cost parameters* that correspond to $\mathbf{f}$. Conditions (7) and (8) 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

$$\lambda_{uv} \triangleq \lambda_u - \lambda_v = \sum_{l\in p} K_l^i(f_l^i, f_l) \le \sum_{l\in p'} K_l^i(f_l^i, f_l), \tag{9}$$

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

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

Definition 2.3. A network $\mathcal{G}$ has the *topological uniqueness property* if the NEP is unique for any convex network game over $\mathcal{G}$.

2.3. Infinite costs. The basic model presented above presumes that the cost functions J_l^i take only finite values. Infinite costs do, however, appear in common cost functions, and are useful for modeling finite capacity limitations. For example, the often-used M/M/1 delay cost is given by $J_l^i(f_l^i, f_l) = f_l^i/(c_l - f_l)$ for $f_l < c_l$, and $J_l^i = \infty$ for $f_l \ge c_l$, where c_l is the link capacity.

Consider then the extended model (as formalized in Orda et al. [25]), where J_l^i takes values in $\mathbb{R}\cup\{\infty\}$. The definition of a NEP $\hat{\mathbf{f}}$ remains the same. However, we shall make here the distinction between a *finite-cost* NEP, for which the equilibrium costs $J^i(\hat{\mathbf{f}})$ are finite for all users, and an *infinite-cost* NEP, where the cost of at least one user is infinite. Evidently, such a user does not have a finite-cost response against the specified flow profiles of the others.

Assumptions A1–A3 remain the same, except that the required differentiability and monotonicity properties are naturally restricted to the *effective* domain of J_l^i; namely, that part of the domain on which J_l^i is finite. We stress that J_l^i is still required to be continuous over its entire domain, so that no discontinuous jumps to infinity are allowed.

With these assumptions in place, the properties of any *finite cost* NEP are exactly the same as in the finite-cost model. Consequently, all claims made in this paper regarding *uniqueness* of the NEP fully apply to the uniqueness of a finite-cost NEP in the infinite-cost model.

Regarding existence, the standard existence results for convex games still apply here and imply the existence of a NEP. However, to ensure existence of a finite-cost NEP, some additional assumptions are needed. Consider, for example, the following requirement (see Orda et al. [25]):

Assumption A4. *For any flow configuration* **f** *at which some user incurs infinite cost, at least one such user can modify its flow configuration so that its cost becomes finite.*

This assumption clearly excludes the existence of infinite-cost NEPs, so that any NEP is a finite-cost one.

For a two-terminal network, Assumption A4 is natural and easy to verify. Suppose that each link has a well-defined capacity (namely, a link flow at which the cost simultaneously becomes infinite for all users who share that link). Then Assumption A4 is satisfied if and only if the network has enough capacity to support the user demand; namely, the total capacity of any cut of the network between the source and destination nodes exceeds the total demand $d = \sum_i d^i$.

2.4. Nearly parallel networks. We briefly repeat here some definitions and results from Milchtaich [23], according to which network topologies can be classified into one of two classes. The class of *nearly parallel networks* essentially contains the networks shown in Figure 1, as well as serial connections of those networks. The complementary class contains all networks in which one of the basic networks shown in Figure 2 is embedded, in the following sense.

Definition 2.4. A network $\mathcal{G}'$ is said to be *embedded in the wide sense* in network $\mathcal{G}''$ if the latter can be obtained from $\mathcal{G}'$ by some sequence of the following three operations:

(i) *Edge subdivision*: An edge is replaced by two edges with a single common end vertex.

(ii) *Edge addition*: The addition of a new edge joining two existing vertices.

(iii) *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.

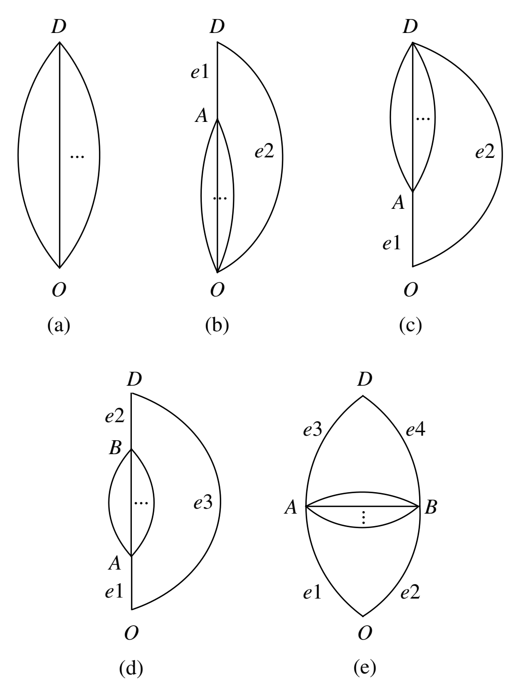

Figure 1. Basic networks that define the class of nearly parallel networks.

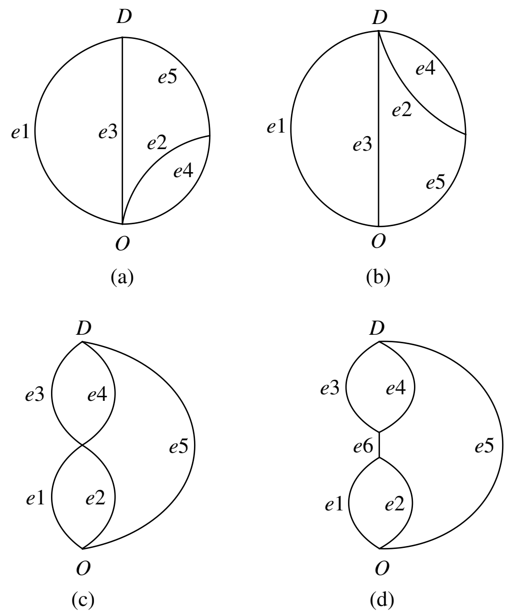

FIGURE 2. Basic networks that are not nearly-parallel

DEFINITION 2.5. 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. Still, the formal definition of nearly parallel networks does require all these basic networks. Note also that only network (e) supports meaningful bidirectional traffic between the same pair of nodes (namely, on the parallel-link network between nodes A and B) given the indicated origin and destination nodes.

PROPOSITION 2.1 (MILCHTAICH [23]). *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 2.5 *and* to serial connections thereof.

3. Uniqueness of the Nash equilibrium. As shown in Orda et al. [25], parallel-link networks possess the topological uniqueness property; namely, uniqueness of the NEP is guaranteed under Assumptions A1–A3. Our main result states that topological uniqueness (in the sense of Definition 2.3) extends to the larger class of nearly parallel networks, and *only* to that class.

THEOREM 3.1. *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. The next subsection establishes some preliminary monotonicity properties that hold in parallel-link subnetworks. Subsection 3.2 establishes uniqueness of the NEP in nearly parallel networks. Subsection 3.3 then presents a basic set of counterexamples to uniqueness, and shows that they can be extended to any network that is not nearly parallel.

3.1. Monotonicity properties for parallel links. A (directed) parallel-link network consists of a set of directed links that share the same start and end node. In the sequel, we shall require certain properties of the equilibrium flows on such networks, when considered as subnetworks of a larger network. For concreteness, we refer to the parallel-link subnetwork that connects node A to B in Figure 1(e), and accordingly denote the start and end nodes by A and B, respectively. Obviously, similar results hold also for the subnetwork of links that lead from B to A.

Consider then a parallel-link network G' (possibly a part of a larger network G), which connects node A to node B. Denote by L_{AB} the set of links in G'. Let

$$f^i_{AB} = \sum_{l \in L_{AB}} f^i_l \tag{10}$$

denote the total flow of user i on that network. Let $\mathbf{f}$ be an equilibrium flow in G. The equilibrium conditions (7) and (8) imply that for every $l \in L_{AB}$,

$$K^i_l(f^i_l, f_l) = \lambda^i_{AB} \quad \text{if } f^i_l > 0 \tag{11}$$

$$K^i_l(f^i_l, f_l) \geq \lambda^i_{AB} \quad \text{if } f^i_l = 0 \tag{12}$$

where $\lambda^i_{AB} = \lambda^i_B - \lambda^i_A$.

Consider two Nash equilibria over G, denoted $\mathbf{f}$ and $\hat{\mathbf{f}}$. Throughout this proof, a hat designates quantities related to $\hat{\mathbf{f}}$, while plain symbols refer to $\mathbf{f}$.

Lemma 3.1. *If $\hat{f}_l \geq f_l$ and $\hat{\lambda}^i_{AB} \leq \lambda^i_{AB}$, then $\hat{f}^i_l \leq f^i_l$.*

Proof. Assume $\hat{f}_l \geq f_l$ and $\hat{\lambda}^i_{AB} \leq \lambda^i_{AB}$. The KKT conditions state that

$$\text{either} \quad \hat{f}^i_l = 0 \leq f^i_l, \quad \text{or} \quad K^i_l(\hat{f}^i_l, \hat{f}_l) = \hat{\lambda}^i_{AB} \leq \lambda^i_{AB} \leq K^i_l(f^i_l, f_l). \tag{13}$$

The assertion is clearly satisfied in the first case, while Assumption A3 implies that it holds in the second case. Note that weak (rather than strict) monotonicity of K^i_l in f_l suffices for the latter implication. □

To proceed, divide the users into two distinct sets

$$I^+_{AB} = \{i \in I \mid \hat{\lambda}^i_{AB} > \lambda^i_{AB}\}, \tag{14}$$

$$I^-_{AB} = \{i \in I \mid \hat{\lambda}^i_{AB} \leq \lambda^i_{AB}\}. \tag{15}$$

Also divide the set of links L_{AB} into two distinct sets

$$L^+_{AB} = \{l \in L_{AB} \mid \hat{f}_l > f_l\}, \tag{16}$$

$$L^-_{AB} = \{l \in L_{AB} \mid \hat{f}_l \leq f_l\}. \tag{17}$$

Lemma 3.2. *Let $\mathbf{f}$ and $\hat{\mathbf{f}}$ be two NEPs, and let (λ^i_{AB}) and $(\hat{\lambda}^i_{AB})$ be the corresponding marginal cost parameters. Then $\sum_{i \in I^+_{AB}} (\hat{f}^i_{AB} - f^i_{AB}) \geq 0$, and the latter inequality is strict if L^+_{AB} is not empty.*

Proof. From Lemma 3.1, if $l \in L^+_{AB}$, then for every $i \in I^-_{AB}$, we have $\hat{f}^i_l \leq f^i_l$. Therefore

$$\sum_{i \in I^-_{AB}} \hat{f}^i_l \leq \sum_{i \in I^-_{AB}} f^i_l. \tag{18}$$

But since $\hat{f}_l > f_l$ for $l \in L^+_{AB}$, this implies

$$\sum_{i \in I^+_{AB}} \hat{f}^i_l > \sum_{i \in I^+_{AB}} f^i_l. \tag{19}$$

Summing up over L^+_{AB} gives

$$\sum_{l \in L^+_{AB}} \sum_{i \in I^+_{AB}} \hat{f}^i_l \geq \sum_{l \in L^+_{AB}} \sum_{i \in I^+_{AB}} f^i_l, \tag{20}$$

where the inequality is strict (and I^+_{AB} not empty) if L^+_{AB} is not empty.

Let l be a link in L^-_{AB}. Then for every $i \in I^+_{AB}$,

$$K^i_l(\hat{f}^i_l, \hat{f}_l) \geq \hat{\lambda}^i_{AB} > \lambda^i_{AB} = K^i_l(f^i_l, f_l), \qquad \text{or} \qquad \hat{f}^i_l \geq f^i_l = 0. \tag{21}$$

Since $\hat{f}_l \leq f_l$ and K^i_l is increasing in both arguments, in either case, it follows that $\hat{f}^i_l \geq f^i_l$. Therefore

$$\sum_{i \in I^+_{AB}} \hat{f}^i_l \geq \sum_{i \in I^+_{AB}} f^i_l. \tag{22}$$

Summing up over L^-_{AB} now gives

$$\sum_{l\in L^-_{AB}}\sum_{i\in I^+_{AB}}\hat f^i_l \geq \sum_{l\in L^-_{AB}}\sum_{i\in I^+_{AB}} f^i_l. \tag{23}$$

Combining (23) and (20) yields

$$\sum_{i\in I^+_{AB}}\hat f^i_{AB} = \sum_{l\in L_{AB}}\sum_{i\in I^+_{AB}}\hat f^i_l \geq \sum_{l\in L_{AB}}\sum_{i\in I^+_{AB}} f^i_l = \sum_{i\in I^+_{AB}} f^i_{AB}. \tag{24}$$

Finally, observe that if L^+_{AB} is not empty, then (20) holds with a strict inequality, so that the inequality in the last equation is strict as well. □

Lemma 3.3. *Let $f_{AB}=\sum_{i\in I} f^i_{AB}$. If $\hat f_{AB}\geq f_{AB}$ and $\hat f^i_{AB} > f^i_{AB}$ for at least one user i, then there is at least one user i for which $\hat f^i_{AB} > f^i_{AB}$ and $\hat\lambda^i_{AB} > \lambda^i_{AB}$.*

Proof. If L^+_{AB} is not empty, then from Lemma 3.2, $\sum_{i\in I^+_{AB}}(\hat f^i_{AB} - f^i_{AB}) > 0$ and for at least one user $i\in I^+_{AB}$ it holds that $\hat f^i_{AB} > f^i_{AB}$. Consider next the case where L^+_{AB} is empty; namely, $\hat f_l \leq f_l$ for all $l\in L_{AB}$. Since $\hat f_{AB}\geq f_{AB}$, this immediately implies that $\hat f_l = f_l$ for every $l\in L_{AB}$ and, consequently, $\hat f_{AB} = f_{AB}$. Now, $\hat f^i_{AB} > f^i_{AB}$ by assumption for at least one user i, so that $\hat f^i_l > f^i_l$ on some link. Since $\hat f_l = f_l$, the strict monotonicity of K^i_l in f^i_l (Assumption A3) together with the equilibrium conditions imply that $\hat\lambda^i_{AB} = K^i_l(\hat f^i_l, \hat f_l) > K^i_l(f^i_l, f_l)\geq \lambda^i_{AB}$. □

3.2. Uniqueness for nearly parallel networks. We next establish the uniqueness of the Nash equilibrium for nearly parallel networks, as asserted in the following proposition:

Proposition 3.1. *If $\mathcal{G}$ is nearly parallel, then the NEP is unique for every convex network game $(I, d, \mathcal{J})$ over $\mathcal{G}$.*

The remainder of this subsection holds the proof of this claim. To start with, if $\mathcal{G}$ is a serial connection of the basic networks in Figure 1, then uniqueness in $\mathcal{G}$ is clearly equivalent to uniqueness in each one of the basic networks. Also, subdividing an edge should not change the property of topological uniqueness because each divided edge is equivalent to a single edge with the sum of the costs of its parts (obviously Assumptions A1–A3 carry over). Hence we need only show that the Nash equilibrium for each of the networks in Figure 1 is unique. Furthermore, it may be seen that networks (a)–(c) are subnetworks of (e), so that multiple equilibria in (a)–(c) may be easily induced in (e) by imposing large enough costs on its additional links. Similarly, network (d) may be reduced to (b) or (c) by unifying links $e1$ and $e2$. It is therefore sufficient to establish uniqueness for the network in Figure 1(e), on which we focus hereafter. In this network, there are two parallel-link subnetworks, one from A to B and the other from B to A. Denote the flows on these subnetworks as follows:

$$f^i_{AB} = \sum_{l\in L_{AB}} f^i_l, \qquad f_{AB} = \sum_{i\in I} f^i_{AB}$$

$$f^i_{BA} = \sum_{l\in L_{BA}} f^i_l, \qquad f_{BA} = \sum_{i\in I} f^i_{BA}$$

Note that $f^i_{AB} > 0$ and $f^i_{BA} > 0$ cannot both hold in equilibrium (for the same user i), since the Nash equilibrium cannot contain cyclic flows (as implied by Assumption A2).

Let $\mathbf{f}$ and $\hat{\mathbf{f}}$ denote two NEPs in network (e). We aim to show that $\mathbf{f} = \hat{\mathbf{f}}$. Denote $\Delta f^i_l = \hat f^i_l - f^i_l$, and similarly $\Delta f^i_{AB} = \hat f^i_{AB} - f^i_{AB}$, etc. Flow conservation yields the following relations (which obviously hold also for the total link flows):

$$\Delta f^i_{e1} = -\Delta f^i_{e2} \tag{25}$$

$$\Delta f^i_{e3} = -\Delta f^i_{e4} \tag{26}$$

$$\Delta f^i_{e1} + \Delta f^i_{BA} = \Delta f^i_{AB} + \Delta f^i_{e3}. \tag{27}$$

The next lemma contains several elementary claims that are needed later on, and follows directly from the monotonicity properties of the marginal costs.

Lemma 3.4. *For the network in Figure 1(e), let $\mathbf{f}$ and $\hat{\mathbf{f}}$ be two NEPs and λ and $\hat\lambda$ be the corresponding marginal cost parameters.*

(a) *If $\Delta f_{e1}\geq 0$, $\Delta f_{e3}\geq 0$ and $\Delta f^i_{e3} > 0$, then $\Delta f^i_{e1}\leq 0$.*

(b) *If $\Delta f_{e1}\geq 0$, $\hat\lambda^i_{AB} > \lambda^i_{AB}$ and $\hat f^i_{AB} > 0$, then $\Delta f^i_{e1}\leq 0$.*

(c) *If* $\Delta f_{e1} \geq 0$, $\hat{\lambda}^i_{BA} < \lambda^i_{BA}$ *and* $f^i_{BA} > 0$, *then* $\Delta f^i_{e1} \leq 0$.
(d) *If* $\Delta f_{e3} \leq 0$, $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$ *and* $\hat{f}^i_{AB} > 0$, *then* $\Delta f^i_{e3} \geq 0$.
(e) *If* $\Delta f_{e3} > 0$, $\Delta f^i_{e3} > 0$ *and* $f^i_{AB} > 0$, *then* $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$.

Proof.
(a) Assume $\Delta f_{e1} \geq 0$, $\Delta f_{e3} \geq 0$ and $\Delta f^i_{e3} > 0$. Suppose, in contradiction, that $\Delta f^i_{e1} > 0$. Then

$$\hat{\lambda}^i_{OD} = K^i_{e1}(\hat{f}^i_{e1}, \hat{f}_{e1}) + K^i_{e3}(\hat{f}^i_{e3}, \hat{f}_{e3}) > K^i_{e1}(f^i_{e1}, f_{e1}) + K^i_{e3}(f^i_{e3}, f_{e3}) \geq \lambda^i_{OD}, \tag{28}$$

where the equality sign follows from the observation that $\hat{f}^i_{e1} > f^i_{e1} \geq 0$, $\hat{f}^i_{e3} > f^i_{e3} \geq 0$ together with (7), the first inequality follows from the monotonicity of K, and the second inequality from (7) and (8).

However, using (25) and (26) yields $f^i_{e2} > \hat{f}^i_{e2} \geq 0$, $f^i_{e4} > \hat{f}^i_{e4} \geq 0$, so that

$$\hat{\lambda}^i_{OD} \leq K^i_{e2}(\hat{f}^i_{e2}, \hat{f}_{e2}) + K^i_{e4}(\hat{f}^i_{e4}, \hat{f}_{e4}) < K^i_{e2}(f^i_{e2}, f_{e2}) + K^i_{e4}(f^i_{e4}, f_{e4}) = \lambda^i_{OD}, \tag{29}$$

which contradicts (28).

(b) Assume $\Delta f_{e1} \geq 0$, $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$, and $\hat{f}^i_{AB} > 0$. Suppose, in contradiction, that $\Delta f^i_{e1} > 0$ (so that $\hat{f}^i_{e1} > 0$). From (25), we conclude that $\Delta f^i_{e2} < 0$ (so that $f^i_{e2} > 0$). Using (9), yields

$$K^i_{e1}(\hat{f}^i_{e1}, \hat{f}_{e1}) + \hat{\lambda}^i_{AB} \leq K^i_{e2}(\hat{f}^i_{e2}, \hat{f}_{e2}) \leq K^i_{e2}(f^i_{e2}, f_{e2}) \leq K^i_{e1}(f^i_{e1}, f_{e1}) + \lambda^i_{AB}. \tag{30}$$

Since $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$, it follows that

$$K^i_{e1}(\hat{f}^i_{e1}, \hat{f}_{e1}) < K^i_{e1}(f^i_{e1}, f_{e1}). \tag{31}$$

On the other hand, since K^i_{e1} is monotone and we assume $\Delta f_{e1} \geq 0$ and $\Delta f^i_{e1} > 0$,

$$K^i_{e1}(\hat{f}^i_{e1}, \hat{f}_{e1}) \geq K^i_{e1}(f^i_{e1}, f_{e1}), \tag{32}$$

hence a contradiction.

(c) The proof is very similar to (b): using (25) and repeating the former argument for $e2$ (with A and B interchanged and $\mathbf{f}$ and $\hat{\mathbf{f}}$ interchanged) implies that $\Delta f^i_{e2} > 0$. Using (25) again yields $\Delta f^i_{e1} < 0$.

(d) The claim and its proof are symmetric to (b).

(e) Assume $\Delta f_{e3} > 0$ $\Delta f^i_{e3} > 0$, and $f^i_{AB} > 0$. From (26), we conclude that $\Delta f^i_{e4} < 0$ and $\Delta f_{e4} < 0$. So $\hat{f}^i_{e3} > 0$ and $f^i_{e4} > 0$ and using (9) gives

$$\lambda^i_{AB} + K^i_{e4}(f^i_{e4}, f_{e4}) \leq K^i_{e3}(f^i_{e3}, f_{e3}) < K^i_{e3}(\hat{f}^i_{e3}, \hat{f}_{e3}) \leq \hat{\lambda}^i_{AB} + K^i_{e4}(\hat{f}^i_{e4}, \hat{f}_{e4}). \tag{33}$$

But since K^i_l is monotone, then $K^i_{e4}(\hat{f}^i_{e4}, \hat{f}_{e4}) < K^i_{e1}(f^i_{e4}, f_{e4})$, so that $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$. □

The next lemma establishes identity of any two NEPs that have equal flows on the two parallel-link subnetworks (from A to B and from B to A).

Lemma 3.5. *If* $\mathbf{f}$ *and* $\hat{\mathbf{f}}$ *are two different NEPs, there is at least one user* i *for which* $f^i_{AB} \neq \hat{f}^i_{AB}$ *or* $f^i_{BA} \neq \hat{f}^i_{BA}$.

Proof. Assume to the contrary that $f^i_{AB} = \hat{f}^i_{AB}$ and $f^i_{BA} = \hat{f}^i_{BA}$ for all i. By (27), it follows that $\Delta f^i_{e1} = \Delta f^i_{e3}$ and $\Delta f_{e1} = \Delta f_{e3}$. Assume further that $\Delta f_{e1} \geq 0$ (the proof is symmetric for $\Delta f_{e1} \leq 0$). Then $\Delta f_{e1} = \Delta f_{e3} \geq 0$. Hence, if $\mathbf{f}$ and $\hat{\mathbf{f}}$ are different, there must be a user i for which $\Delta f^i_{e1} = \Delta f^i_{e3} > 0$. But Lemma 3.4(a) now implies that $\Delta f^i_{e1} < 0$, a contradiction, so that $\mathbf{f}$ and $\hat{\mathbf{f}}$ must coincide. □

The following two lemmas exclude the possibility of two distinct NEPs with $f_{e1} \geq \hat{f}_{e1}$:

Lemma 3.6. *Let* $\mathbf{f}$ *and* $\hat{\mathbf{f}}$ *be two NEPs. If* $\Delta f_{e1} \geq 0$ *and* $\Delta f_{e3} \leq 0$, *then* $\mathbf{f} = \hat{\mathbf{f}}$.

Proof. Assume that $\mathbf{f}$ and $\hat{\mathbf{f}}$ are distinct. Observing Lemma 3.5, we may assume that the following property holds:

$$f^i_{AB} \neq \hat{f}^i_{AB} \quad \text{for some user } i. \tag{34}$$

This entails no loss of generality since, if needed, A and B can be relabelled to satisfy (34), and then $\mathbf{f}$ and $\hat{\mathbf{f}}$ may be swapped to maintain $\Delta f_{e1} \geq 0$ and $\Delta f_{e3} \leq 0$. Now there are two possible cases:

(a) $\Delta f_{AB} \geq 0$. Applying (34) and Lemma 3.3 to the parallel links network from A to B, we conclude that there is at least one user i for which $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$ and $\hat{f}^i_{AB} > f^i_{AB} \geq 0$. Furthermore, from Lemma 3.4(b) $\Delta f^i_{e1} \leq 0$, and from Lemma 3.4(d) $\Delta f^i_{e3} \geq 0$. Using this and (27), we conclude that $\Delta f^i_{BA} > 0$, so that $\hat{f}^i_{BA} > 0$. But since the NEP cannot contain routing loops (as explained before), we obtain a contradiction.

(b) $\Delta f_{AB} \le 0$. In this case, (27) implies that $\Delta f_{BA} \le 0$. Furthermore, by (25) and (26), it holds that $\Delta f_{e2} \le 0$ and $\Delta f_{e4} \ge 0$. This is completely symmetrical to the previous case, and we can repeat the argument above to obtain a contradiction.

It follows that $\mathbf{f}$ and $\hat{\mathbf{f}}$ cannot be different; namely, $\hat{\mathbf{f}} = \mathbf{f}$. $\square$

Lemma 3.7. *Let $\mathbf{f}$ and $\hat{\mathbf{f}}$ be two equilibrium points. If $\Delta f_{e1} \ge 0$, then $\Delta f_{e3} \le 0$.*

Proof. If $\mathbf{f} = \hat{\mathbf{f}}$, then we are done. Otherwise, it may be assumed as above that (34) holds. Assume further, by way of contradiction, that $\Delta f_{e1} \ge 0$ and $\Delta f_{e3} > 0$. Consider the following group of users:

$$I_1 = \{i \in I \mid \Delta f^i_{e3} > 0\} \tag{35}$$

and the following subsets thereof:

$$I_2 = \{i \in I \mid \Delta f^i_{e3} > 0,\ f^i_{AB} > 0\} \tag{36}$$

$$I_2' = \{i \in I \mid \Delta f^i_{e3} > 0,\ f^i_{AB} = 0\} \tag{37}$$

$$I_3 = \{i \in I \mid \Delta f^i_{e3} > 0,\ \hat{f}^i_{BA} > 0\} \tag{38}$$

$$I_3' = \{i \in I \mid \Delta f^i_{e3} > 0,\ \hat{f}^i_{BA} = 0\}. \tag{39}$$

Note that $I_1 = I_2 \cup I_2' = I_3 \cup I_3'$. Recall the definition of $I^+_{AB} = \{i \in I \mid \hat{\lambda}^i_{AB} > \lambda^i_{AB}\}$ in (14), and similarly define

$$I^-_{BA} = \{i \in I \mid \hat{\lambda}^i_{BA} < \lambda^i_{BA}\}. \tag{40}$$

Furthermore, define

$$I_4 = \{i \in I \mid \hat{\lambda}^i_{AB} > \lambda^i_{AB},\ \hat{f}^i_{AB} > 0\} \subseteq I^+_{AB} \tag{41}$$

$$I_5 = \{i \in I \mid \hat{\lambda}^i_{BA} < \lambda^i_{BA},\ f^i_{BA} > 0\} \subseteq I^-_{BA}. \tag{42}$$

From Lemma 3.4(a), if $i \in I_1$, then $\Delta f^i_{e1} \le 0$. According to Lemma 3.4(b) if $i \in I_4$, then $\Delta f^i_{e1} \le 0$, and according to Lemma 3.4(c) if $i \in I_5$, then $\Delta f^i_{e1} \le 0$. We can now conclude that

$$\sum_{i \in I_1 \cup I_4 \cup I_5} \Delta f^i_{e1} \le 0. \tag{43}$$

From Lemma 3.4(e) if $i \in I_2$, then $i \in I^+_{AB}$, which together with (41) yields

$$I_2 \cup I_4 \subseteq I^+_{AB}. \tag{44}$$

From Lemma 3.2, it is known that $\sum_{i \in I^+_{AB}} \Delta f^i_{AB} \ge 0$. Moreover, by (41), for $i \in I^+_{AB} \setminus (I_2 \cup I_4)$, it holds that $\hat{f}^i_{AB} = 0$, and therefore $\Delta f^i_{AB} \le 0$, resulting in

$$\sum_{i \in I_2 \cup I_4} \Delta f^i_{AB} \ge \sum_{i \in I_2 \cup I_4} \Delta f^i_{AB} + \sum_{i \in I^+_{AB} \setminus (I_2 \cup I_4)} \Delta f^i_{AB} = \sum_{i \in I^+_{AB}} \Delta f^i_{AB} \ge 0. \tag{45}$$

Notice that from (37) and (42), for $i \in I_2' \cup I_5$, it holds that $f^i_{AB} = 0$, and therefore $\Delta f^i_{AB} \ge 0$. Combining this with (45) and noting that $I_1 = I_2 \cup I_2'$ gives

$$\sum_{i \in I_1 \cup I_4 \cup I_5} \Delta f^i_{AB} = \sum_{i \in I_2 \cup I_2' \cup I_4 \cup I_5} \Delta f^i_{AB} \ge 0. \tag{46}$$

Repeating the argument after interchanging A with B and $\mathbf{f}$ with $\hat{\mathbf{f}}$ yields

$$\sum_{i \in I_1 \cup I_4 \cup I_5} \Delta f^i_{BA} \le 0. \tag{47}$$

Substituting (46), (47), and (43) in (27) yields

$$\sum_{i \in I_1 \cup I_4 \cup I_5} \Delta f^i_{e3} \le 0. \tag{48}$$

But by definition of I_1

$$\Delta f_{e3} = \sum_{i \in I} \Delta f^i_{e3} \le \sum_{i \in I_1 \cup I_4 \cup I_5} \Delta f^i_{e3} \le 0, \tag{49}$$

which is a contradiction to the assumption that $\Delta f_{e3} > 0$. $\square$

Lemmas 3.6 and 3.7 together prove uniqueness of the equilibrium point in the network in Figure 1(e) when $\Delta f_{e1} \ge 0$. The opposite case of $\Delta f_{e1} \le 0$ is easily reduced to the previous one by interchanging $\mathbf{f}$ and $\hat{\mathbf{f}}$. Thus the proof of Proposition 3.1 is complete.

3.3. Counterexamples to uniqueness. For general network topologies, the NEP need not be unique. In this section, we show that a counterexample (namely, a convex network game with multiple Nash equilibria) may be found for any network that is not nearly parallel. Recall from Proposition 2.1 that in any network that 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). The cost functions are 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 user demands d^i specified 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

Here

$$f_1(x) = \begin{cases} x & \text{if } x < 6 \\ \dfrac{1}{3}(e^{3(x-6)} + 17) & \text{if } x \geq 6 \end{cases} \tag{50}$$

and

$$f_2(x) = \begin{cases} x & \text{if } x < 4 \\ \dfrac{1}{3}(e^{3(x-4)} + 11) & \text{if } x \geq 4 \end{cases}. \tag{51}$$

Note that these functions are continuously differentiable. The infinity symbols in the table stand for large enough functions, so that the relevant user will never choose to use these links.

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 obtained if each user diverts all of its flow to its first option, and another Nash equilibrium is obtained if each user diverts all of its flow to its second option. For example, in the first NEP, (9) may be verified by observing that for User 1,

$$K_{e3}^1(6,6) < K_{e4}^1(0,4) + K_{e5}^1(0,8) \tag{52}$$

and in the second NEP

$$K_{e4}^1(6,6) + K_{e5}^1(6,6) < K_{e3}^1(0,4). \tag{53}$$

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 demands and latency functions $T_l^i(x)$ for each user are now

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	$f_2(x)$	"∞"	$6x$

$f_1(x)$ 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 transfers 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 chooses the path $e1$–$e3$.

These two examples show that multiple equilibria exist in the networks of Figure 2(a) and 2(c). We now need 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 addition of serial links with small enough cost does not alter the equilibrium. We use the following definition:

DEFINITION 3.1. 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

$$\sum_{l\in p} K_l^i(f_l^i, f_l) < \sum_{l\in p'} K_l^i(f_l^i, f_l) \tag{54}$$

for any other path p' that connects O and D.

Note that in a strong NEP, each user utilizes a single path from origin to destination. It may be verified that the NEPs in Examples 1 and 2 are strong.

LEMMA 3.8. *Let $\mathcal{G}$ be a network over which there exists a convex network game with two different 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 Milchtaich [23]. Let $\mathbf{f}$ and $\hat{\mathbf{f}}$ be the two strong NEPs in $\mathcal{G}$ and denote by p^i and $\hat{p}^i$ the unique paths of user i in $\mathbf{f}$ and $\hat{\mathbf{f}}$, respectively. Because 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 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 nonzero cost function for that link. Since the two equilibria in $\mathcal{G}$ are strong, we can set the marginal costs $K_l^i(f_l^i, f_l)$ on that link sufficiently small so that $\mathbf{f}$ and $\hat{\mathbf{f}}$ are still NEPs. Specifically, choose J_l^i, so that $K_l^i(f_l^i, f_l)$ is smaller (for all $f_l^i \le d^i$) than the cost difference between any pair of used and unused routes. □

PROPOSITION 3.2. *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 3.8). Lemma 3.8 now proves the proposition. □

Propositions 3.1 and 3.2 together with Proposition 2.1 provide the proof of Theorem 3.1.

REMARK. Consider a directed network that is not necessarily bidirectional (namely, links do not necessarily appear in pairs of opposite direction as assumed so far). A sufficient condition for uniqueness is easily obtained: Obviously, if the directed network is nearly parallel (in the sense that it is a subnetwork of a bidirectional nearly parallel network), then topological uniqueness holds by the sufficiency part of Theorem 3.1. Indeed, any link that is required to complete a pair may be added with large enough costs, so that the NEPs in the two networks coincide. Moreover, it is clear that this sufficient condition can be applied after removing from our directed network any link that is not part of some (simple) directed path from O to D, as such links do not carry any flow in equilibrium, and thus do not affect its uniqueness. To illustrate this reduction, consider a directed version of network (a) in Figure 2, and suppose that link $e2$ points into O. Then this link can be removed, and the network reduces to a parallel link network. As for the necessity part of Theorem 3.1, some additional care is required, as there exist directed networks that are not nearly parallel but still possess topological uniqueness; such an example can be found in Milchtaich [23, Figure 7]. While reasonable extensions to the definition of nearly parallel networks that cover this example may be given, a detailed study of the necessity part for directed networks will not be pursued here.

4. 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 Wardrop and Nash equilibria in our the network model.

4.1. Multiclass Wardrop equilibrium as an atomic Nash equilibrium. 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

$$\sum_{l \in p^i} T_l^i(f_l) = \min_{p} \sum_{l \in p} T_l^i(f_l) \quad \text{for every } i \in I, \tag{55}$$

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

It was shown in Milchtaich [23] that the Wardrop equilibrium is unique for any choice of (nonnegative, 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 NEP when the number of users is increased to infinity while their individual flow demands decrease accordingly (Haurie and Marcotte [14]). 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 Beckmann et al. [5] for the Wardrop equilibrium with a single user class, which is well known to be equivalent to a convex optimization problem (see also Dafermos and Sparrow [9]). Later, Devarajan [11] 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 class-user i, and let the link costs for that user be given by

$$J_l^i(f_l^i, f_l) \triangleq \int_0^{f_l^i} T_l^i(f_l - f_l^i + x)\, dx. \tag{56}$$

Recalling (6), 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 of 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 (9) with the definition of the Wardrop equilibrium above.

It is readily seen that the cost functions in (56) 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.

4.2. 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 required to be only* weakly *increasing in f_l^i.*

We define a *weakly convex network game* similarly to a convex network game, with Assumption A3 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 hold true:

Theorem 4.1. *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 over a set of NEPs*:
(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.*

Remarks.
(1) Part (iii) of this theorem can be seen to imply Proposition 3.1; namely, the uniqueness of the equilibrium point under Assumption A3. The following proof thus provides an alternative (although somewhat more involved) proof for that proposition.

(2) Assumption A3′ is not sufficient to ensure uniqueness of the per user flows. This is easily seen by considering a weakly convex game that corresponds to a Wardrop equilibrium problem, where users with similar latency functions can simply switch links. We note, however, that uniqueness of the Wardrop equilibrium was shown in Milchtaich [23] to be a generic property in the space of latency functions. A similar property is likely to hold for the present model.

(3) Similarly, while the *marginal* costs $K_l^i(f_l^i, f_l)$ are unique, Assumption A3′ is not sufficient to guarantee uniqueness of the link costs $J_l^i(f_l^i, f_l)$. This is easily verified by a simple example (say, with two parallel links, two users, and link costs of the form (56), so that $K_l^i(f_l^i, f_l) = T_l^i(f_l)$).

(4) As noted in the previous subsection, a multiclass Wardrop equilibrium problem with nonnegative, strictly increasing latency functions $T_l^i(f_l)$ may be represented as a weakly convex network game. Theorem 4.1 thus recovers the uniqueness of the link flows for the Wardrop equilibrium for nearly parallel networks. In fact, it generalizes this result by allowing link latencies of the form $T_l^i = K_l^i(f_l^i, f_l)$, subject to Assumptions A2 and A3′ (where Assumption A2 is equivalent to $K_l^i > 0$ for $f_l^i > 0$).

(5) Similarly, the last result can also be applied to the mixed Nash-Wardrop problem (Boulogne et al. [6]). Thus, Theorem 4.1 holds for any combination of large (atomic) users and small user classes, provided that the cost function of all large users satisfy Assumptions A1, A2, and A3′, while the latency functions $T_l^i = K_l^i(f_l^i, f_l)$ of the small user classes are subject to Assumptions A2 and A3′.

(6) In case the latency functions of some small user class have the standard form $T_l^i(f_l)$, the established uniqueness of the total link flows $\{f_l\}$ together with the strict monotonicity of T_l^i imply that the equilibrium link latencies for this user class are unique. Hence the total latency experienced by that user class over any route it uses is uniquely determined (see Equation (55)).

We proceed with the proof of Theorem 4.1, starting with part (i). Assumptions A1, A2, and A3′ are henceforth assumed to hold.

LEMMA 4.1. *Lemmas 3.1, 3.2, 3.4, and 3.7 continue to hold for weakly convex games, with the following modifications:*
(i) *In Lemma 3.1, we require that one of the inequalities $\hat{f}_l \geq f_l$ or $\hat{\lambda}_{AB}^i \leq \lambda_{AB}^i$ should be strict.*
(ii) *In Lemma 3.4(a), we require $\Delta f_{e3} > 0$ instead of $\Delta f_{e3} \geq 0$.*

PROOF. Same as the original proof of these claims, except that the requirement for strict inequalities in (i) and (ii) replaces the need for the strict monotonicity requirement in Assumption 3. □

Lemma 3.3 needs to be refined as follows:

LEMMA 4.2. *Consider a weakly convex game $(\mathcal{I}, d, \mathcal{J})$ over a parallel-link network from A to B. If $\hat{f}_{AB} \geq f_{AB}$ and L_{AB}^+ is not empty, then there is at least one user i for which $\hat{f}_{AB}^i > f_{AB}^i$ and $\hat{\lambda}_{AB}^i > \lambda_{AB}^i$.*

PROOF. Immediate from Lemma 3.2 and the definitions of I_{AB}^+ and L_{AB}^+ in (14) and (16), respectively. □

Lemma 3.5 needs to be refined as well.

LEMMA 4.3. *If f and $\hat{f}$ are two NEPs so that there exists a link $l \in L$ for which $\hat{f}_l \neq f_l$, then there is at least one link $l \in L_{AB} \cup L_{BA}$ for which $\hat{f}_l \neq f_l$.*

PROOF. Contrary to the claim, suppose that $\hat{f}_l = f_l$ for every $l \in L_{AB} \cup L_{BA}$. Since there is an $l \in L$ for which $\hat{f}_l \neq f_l$, we may assume without loss of generality that $\Delta f_{e1} > 0$ (relabelling *O-D* and *A-B* if necessary). From (27), $\Delta f_{e3} > 0$. Lemma 3.7 now leads to the desired contradiction. □

We can now assert a weaker version of Lemma 3.6.

LEMMA 4.4. *If $\Delta f_{e1} \geq 0$ and $\Delta f_{e3} \leq 0$, then $f_l = \hat{f}_l$ for every $l \in L$.*

PROOF. The proof is similar to that of Lemma 3.6, using Lemmas 4.2 and 4.3 in place of Lemmas 3.3 and 3.5, respectively. □

Lemma 4.4 together with Lemma 3.7 establish the uniqueness of the link flows when $\Delta f_{e1} \geq 0$. The same clearly holds when $\Delta f_{e1} \leq 0$ (as can be seen by interchanging f and $\hat{f}$). Thus Theorem 4.1(i) is established.

Part (ii) of Theorem 4.1 is verified in the following lemma:

LEMMA 4.5. *Let f and $\hat{f}$ be two NEPs with $f_l = \hat{f}_l$, $l \in L$. Then $K_l^i(f_l^i, f_l) = K_l^i(\hat{f}_l^i, \hat{f}_l)$ for all l and i.*

PROOF. Introduce the shorthand notation $K_l^i = K_l^i(f_l^i, f_l)$ and $\widehat{K}_l^i = \widehat{K}_l^i(\hat{f}_l^i, \hat{f}_l)$. Since $f_l = \hat{f}_l$, Assumption A3′ implies that

$$\widehat{K}_l^i > K_l^i \Longrightarrow \Delta f_l^i > 0 \quad (\text{namely, } \hat{f}_l^i > f_l^i),$$
$$\Delta f_l^i \geq 0 \Longrightarrow \widehat{K}_l^i \geq K_l^i. \tag{57}$$

Our argument follows in outline, but not in detail, the proofs of Lemmas 3.5 and 3.6. We fix a user i, and consider separately the following two cases:

(a) $K^i_l = \widehat{K}^i_l$, $l \in L_{AB} \cup L_{BA}$. Suppose there exists some $l \in \{e1, e2, e3, e4\}$, so that $\Delta K^i_l \neq 0$. For concreteness, suppose that $\widehat{K}^i_{e1} > K^i_{e1}$ (the other possibilities can be handled similarly). From (57) $\Delta f^i_{e1} > 0$, hence $\Delta f^i_{e2} < 0$ (by flow conservation) and $\widehat{K}^i_{e2} \leq K^i_{e2}$. We now argue that $\Delta f^i_{e3} > 0$, by first showing that $\hat{f}^i_{AB} = 0$ and $f^i_{BA} = 0$. Indeed, by the equilibrium condition (9) applied to all possible paths from O to B, we obtain

$$K^i_{e2} = \lambda^i_{OB} \leq K^i_{e1} + K^i_l, \quad l \in L_{AB},$$

where the first equality follows since $\Delta f^i_{e2} < 0$ implies that $f^i_{e2} > 0$. Using the relations established between K and $\widehat{K}$, this gives

$$\widehat{K}^i_{e2} < \widehat{K}^i_{e1} + \widehat{K}^i_l, \quad l \in L_{AB}.$$

But, again by (9), this implies that $\hat{f}^i_{AB} \triangleq \sum_{l \in L_{AB}} \hat{f}^i_l = 0$. Using a symmetric argument, starting with $\Delta f^i_{e1} > 0$, we obtain $f^i_{BA} = 0$. If follows that $\Delta f^i_{AB} \leq 0$ and $\Delta f^i_{BA} \geq 0$, and from (27) and $\Delta K^i_{e1} > 0$, we obtain that $\Delta f^i_{e3} > 0$.

We now claim that the relations $\widehat{K}^i_{e1} > K^i_{e1}$, $\Delta f^i_{e1} > 0$ and $\Delta f^i_{e3} > 0$ together lead to contradiction. The argument is the same as in the proof of Lemma 3.4(a), except that $\widehat{K}^i_{e1} > K^i_{e1}$ is now used in (28) to establish the strict inequality, while a weak inequality suffices in (29). We conclude that $\widehat{K}^i_{e1} > K^i_{e1}$ cannot hold in this case. Using the same argument for the opposite inequality and the other three links $e2$, $e3$, and $e4$, we obtain the required conclusion of the lemma in this case.

(b) Contrary to case (a), assume that $K^i_l \neq \widehat{K}^i_l$ for some $l \in L_{AB} \cup L_{BA}$. For concreteness, assume that $\widehat{K}^i_l > K^i_l$ for some $l \in L_{AB}$ (the other cases are handled similarly). By (57), this implies that $\hat{f}^i_l > f^i_l$, hence $\hat{f}^i_l > 0$. Observing (9), we conclude that $\hat{\lambda}^i_{AB} > \lambda^i_{AB}$, and consequently, that $\hat{f}^i_{AB} > f^i_{AB}$. We can now proceed to obtain a contradiction exactly as in case (a) in the proof of Lemma 3.4.

We conclude that only case (a) is possible, and in that case, the assertion of the lemma has indeed been verified. □

We finally observe that part (iii) of Theorem 4.1 follows immediately from the first two and the strict monotonicity of the marginal costs $K^i_l(f^i_l, f_l)$. Thus the proof of Theorem 4.1 is complete.

5. A unified continuum-game model. We next introduce an extension of our selfish routing model, that treats in a unified manner both *large* (atomic) users, each having a positive flow demand, and a continuum of *small* (nonatomic) users, each of which commands an infinitesimal flow demand. As in Milchtaich [23], we use the framework of nonatomic games (Schmeidler [31]) to model small users. Thus, each user (large or small) is explicitly modeled as a rational decision maker with an individual cost function. This is in contrast to the definition of the Wardrop equilibrium in (55), which specifies the behavior of small users via an aggregate flow condition. Besides the unified treatment of small and large users, the model allows for a *continuum* of small user classes (namely, a continuum of different cost functions) alongside a discrete population of large users.

As our focus in this paper is on two-terminal networks, we present the model in this context. The extension to a general network is straightforward, the only difference being that the feasible action set for each user would then depend on its origin-destination pair. After the model is introduced, we describe the relevant uniqueness results.

Consider a two-terminal network as before. The user population is represented by a measure space $(I, \mathcal{I}, \rho)$, where I is a set of users, endowed with a σ-algebra $\mathcal{I}$ and a finite measure ρ. Thus, each element $i \in I$ corresponds to an individual user. The user population may be finite countably infinite or uncountable, according to the cardinality of I. The mass $r^i \triangleq \mu(\{i\})$ of each user specifies its flow requirement; namely, the amount of flow that this user wishes to ship over the network. We refer to users with $r^i = 0$ as small users, while those with $r^i > 0$ are large (or atomic) users. By way of interpretation, zero mass does not imply that a user has no flow to ship, but rather that the effect of its flow on the total link flows is negligible (see Equation (58)).

Two specific examples of possible user populations are: (a) $I = [0, 1]$ endowed with the Lesbegue measure, which is often used as the canonical choice for nonatomic games and (b) $I = \{i_1, i_2, \ldots, i_n\}$, a finite set, with $r_i > 0$ for $i \in I$. The latter reduces to the finite-user game of the previous sections. Taking I as the union of these two sets leads to a mixed user population.

Each user should choose how to ship its flow over available network paths. Routing decisions are specified by *normalized* routing variables ϕ^i_l, which represent the *fraction* of user i's flow to be shipped over link l. Thus the vector $\boldsymbol{\phi}^i = (\phi^i_l,\ l \in L)$ is *feasible* if it satisfies the flow constraints (1)–(2) with $r^i = 1$. Note that the use of normalized routing variables is essential for small users. For large users, the normalized routing variable ϕ^i_l is, of course, equivalent to the flow assignment $f^i_l = r^i \phi^i_l$.

Assuming that the map from i to ϕ_l^i is measurable, the total flow on link l can be defined by integration over the population measure; namely,

$$f_l = \int \phi_l^i \rho(di). \tag{58}$$

A flow configuration $\boldsymbol{\phi} = \{\boldsymbol{\phi}^i,\ i \in \mathcal{I}\}$ is thus said to be feasible if each $\boldsymbol{\phi}^i$ is feasible, and in addition, the map from $i \in \mathcal{I}$ to ϕ_l^i is measurable for every link l.

The cost function for user i is specified by

$$\bar{J}^i(\boldsymbol{\phi}) = \sum_{l \in L} \bar{J}_l^i(\phi_l^i, f_l), \tag{59}$$

where the *normalized link cost* $\bar{J}_l^i(\phi_l^i, f_l)$ is the cost incurred by user i for shipping a fraction ϕ_l^i of its flow on link l. We assume throughout that the map $(i, \phi_l^i, f_l) \mapsto \bar{J}_l^i(\phi_l^i, f_l)$ is measurable for every link l. A feasible flow configuration $\widehat{\boldsymbol{\phi}}$ is an NEP if, for ρ-almost every i,

$$\bar{J}^i(\widehat{\boldsymbol{\phi}}) = \min_{\phi^i} \bar{J}^i(\boldsymbol{\phi}^i, \widehat{\boldsymbol{\phi}}^{-i}), \tag{60}$$

where the minimization is over the feasible set of flow configurations of user i.

To specify our assumptions on the cost functions, define the normalized marginal cost

$$\bar{K}_l^i(\phi_l^i, f_l) \triangleq \frac{d}{d\phi_l^i} \bar{J}(\phi_l^i, f_l) = \frac{\partial}{\partial \phi_l^i} \bar{J}_l^i(\phi_l^i, f_l) + r^i \frac{\partial}{\partial f_l} \bar{J}_l^i(\phi_l^i, f_l) \tag{61}$$

(compare with (6), and note that $r^i = \rho(\{i\})$ is the relative contribution of ϕ_l^i to f_l). We can now state Assumptions A1–A3 and A3′ of the previous sections in the form appropriate for the present model: this simply means replacing f_l^i, J_l^i, and K_l^i by their normalized versions ϕ_l^i, $\bar{J}_l^i$, and $\bar{K}_l^i$, respectively. We shall still refer to these assumptions by their original names, with the understanding that they now pertain to their normalized version.

Let us briefly elaborate on the imposed cost structure. For a large user, the relative link costs may be related to the standard link costs via

$$J_l^i(f_l^i, f_l) = r^i \bar{J}_l^i(\phi_l^i, f_l) \quad \text{with } f_l^i = r^i \phi_l^i, \tag{62}$$

so that $\bar{J}^i = J^i / r^i$. The normalization by the constant factor r^i does not affect the resulting equilibrium, of course. It is now easy to verify that $\bar{K}_l^i(\phi_l^i, f_l) = K_l^i(f_l^i, f_l)$, with K_l^i as defined in (6). Thus Assumptions A1–A3 and A3′, as they apply to a large user in the present model, are equivalent to the corresponding assumptions of the basic model.

For a small user with $r^i = 0$, the relation in (62) is meaningless, and for comparison with the standard nonatomic model, we need to look at the marginal costs. First, note that the second term in (61) is null for such users, and we obtain $\bar{K}_l^i = \partial \bar{J}_l^i / \partial \phi_l^i$. In particular, for $\bar{J}_l^i(\phi_l^i, f_l) = \phi_l^i T(f_l)$, we get $\bar{K}_l^i(\phi_l^i, f_l) = T_l(f_l)$, which is the standard latency function used in the Wardrop model. Assumptions A2 and A3′ then reduce to the standard requirement on T_l to be positive and increasing. Observe though that the present model accommodates a more general form for $\bar{J}_l^i$, even for small users.

We can now turn to the basic questions of existence and uniqueness of the equilibrium. Existence of an NEP under Assumptions A1, A2, and A3′ follows from standard results on continuum games with compact action spaces (possible unboundedness of the feasible action set may be dealt with exactly as in §2.2). In particular, the required result is already outlined in Schmeidler's [31] seminal paper as an extension to his Theorem 1. See, e.g., Theorem 4.1 in Khan [16] for a precise statement and proof of this result. We note that a simpler proof along the lines of Rath [27] may also be applicable to the present model. We further note that a basic purification argument, as in Theorem 2 of Schmeidler [31], can be used to establish the existence of an equilibrium point, where small users with cost functions of the standard form $\bar{J}_l^i(\phi_l^i, f_l) = \phi_l^i T(f_l)$ need not split their flow over different routes.

We next state our final and most general uniqueness result for nearly parallel networks, which is just Theorem 4.1 adapted to the present model.

THEOREM 5.1. *Consider the unified selfish routing model of this section, under Assumptions* A1, A2, *and* A3′. *If the network $\mathcal{G}$ is nearly parallel, then in any pair of Nash equilibria,*

(i) *The link flows $(f_l,\ l \in L)$ are unique.*

(ii) *The marginal link costs $K_l^i(\phi_l^i, f_l)$ are unique for all l and ρ-almost all i.*

(iii) *For ρ-almost every user i whose costs satisfy the stronger Assumption* A3, *this user's flows $(\phi_l^i,\ l \in L)$ are unique. Consequently, the link costs for such a user are unique as well.*

The proof of these claims follows precisely that of Theorem 4.1, with the following semantic changes: J_l^i, K_l^i, f_l^i should be replaced by $\bar{J}_l^i$, $\bar{K}_l^i$, ϕ_l^i; $\sum_i$ replaced by $\int \rho(di)$; "for some i" replaced by "a subset of I with positive measure"; and "for every i" replaced by "for ρ-almost every i."

6. 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 tool for this purpose, that can be 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 network 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, it is often much easier to induce desirable operating conditions (through pricing, for example) when the resulting equilibrium point is unique.

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 somewhat 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.

We have not dealt in this paper with multiterminal 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.

When uniqueness cannot be inferred from the network topology alone, further properties of the cost functions and user characteristics clearly come into play. The well-known diagonal strict convexity conditions (Haurie and Marcotte [14], Orda et al. [25]) on the cost functions are often restrictive, and one may hope to find a middle ground that combines cost function properties with other network and user characteristics. This remains an interesting direction for further research.

Finally, the proposed continuum-game model offers a convenient framework for studying various issues related to mixed Nash-Wardrop equilibria. A particular problem of interest is a general result on the convergence of the many-user Nash equilibrium to the Wardrop equilibrium.

Acknowledgments. The authors thank Igal Milchtaich for helpful discussions and advice, as well as for suggesting the unified model of §5. The authors also thank two anonymous referees and the associate editor for many useful comments that led to considerable improvement of this paper. The reported research was supported by Grant 2002112 from the United States–Israel Binational Science Foundation (BSF).

References

[1] Altman, E., H. Kameda. 2001. Equilibria for multiclass routing in multi-agent networks. *Proc. 40th IEEE Conf. Decision and Control*. Orlando, FL, 604–609.

[2] Altman, E., T. Başar, T. Jiménez, N. Shimkin. 2002. Competitive routing in networks with polynomial cost. *IEEE Trans. Automatic Control* **47** 92–96.

[3] Altman, E., T. Boulogne, R. El Azouzi, T. Jiménez, L. Wynter. 2005. A survey on networking games in telecommunications. *Comput. Oper. Res.* **33**(2) 286–311.

[4] Azouzi, R., E. Altman. 2003. Constrained traffic equilibrium in routing. *IEEE Trans. Automatic Control* **48**(9) 1656–1660.

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

[6] Boulogne, T., E. Altman, H. Kameda, O. Pourtallier. 2002. Mixed equilibrium (ME) for multiclass routing games. *IEEE Trans. Automatic Control* **47**(6) 903–916.

[7] Dafermos, S. C. 1971. An extended traffic assignment model with application to two-way traffic transportation science. *Transportation Sci.* **5** 366–389.

[8] Dafermos, S. C. 1972. The traffic assignment problem for multiclass-user transportation networks. *Transportation Sci.* **6** 73–87.

[9] Dafermos, S. C., F. T. Sparrow. 1969. The traffic assignment problem for a general network. *J. Res. National Bureau Standards* **37B** 91–118.

[10] Debreu, G. 1952. A social equilibrium existence theorem. *Proc. National Acad. Sci.* **38** 886–893.

[11] Devarajan, S. 1981. A note on network equilibrium and noncooperative games. *Transportation Res. Part B* **15B** 421–426.

[12] El Azouzi, R., E. Altman, O. Pourtallier. 2003. Avoiding paradoxes in multi-agent competitive routing. *Comput. Networks* **43** 133–146.

[13] Harker, P. 1988. Multiple equilibrium behaviors on networks. *Transportation Sci.* **22** 39–46.

[14] Haurie, A., P. Marcotte. 1985. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* **15** 295–308.

[15] Jiménez, T., E. Altman, T. Başar, N. Shimkin. 2001. Routing into two parallel links: Game-theoretic distributed algorithms. *J. Parallel Distributed Comput.* (Special Issue on Routing in Computer and Communication Systems) **61** 1367–1381.

[16] Khan, M. A. 1985. On extensions to the Cournot-Nash theorem. *Advances in Equilibrium Theory. Lecture Notes Economic Mathematical Systems* Vol. 244. Springer-Verlag, 79–106.

[17] Korilis, Y., A. Lazar, A. Orda. 1995. Architecting noncooperative networks. *IEEE J. Selected Areas Comm.* **13**(7) 1241–1251.

[18] Korilis, Y., A. Lazar, A. Orda. 1997. Capacity allocation under noncooperative routing. *IEEE Trans. Automatic Control* **42**(3) 309–325.

[19] Korilis, Y., A. Lazar, A. Orda. 1999. Avoiding the Braess paradox in non-cooperative networks. *J. Appl. Probab.* **36** 211–222.

[20] Korilis, Y. A., T. A. Varvarigou, S. R. Ahuja. 1998. Incentive-compatible pricing strategies in noncooperative networks. *Proc. IEEE INFOCOM'98, San Francisco, CA*, 439–446.
[21] La, R. J., V. Anantharam. 2002. Optimal routing control: Repeated game approach. *IEEE Trans. Automatic Control* **47** 437–450.
[22] Marcotte, P., L. Wynter. 2004. A new look at the multiclass network equilibrium problem. *Transportation Sci.* **38**(3) 282–292.
[23] Milchtaich, I. 2005. Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.* **30** 225–244.
[24] Nagurney, A. 1999. *Network Economics: A Variational Inequality Approach*, 2nd ed. Kluwer Academic, Dordrecht, The Netherlands.
[25] Orda, A., R. Rom, N. Shimkin. 1993. Competitive routing in multi-user communication networks. *IEEE/ACM Trans. Networking* **1**(5) 510–521.
[26] Patriksson, M. 1994. *The Traffic Assignment Problem: Models and Methods*. VSP, Utrecht, The Netherlands.
[27] Rath, K. P. 1992. A direct proof of the existence of pure strategy equilibria in games with a continuum of players. *Econom. Theory* **2** 427–433.
[28] Rosen, J. B. 1965. Existence and uniqueness of equilibrium points for concave n-person games. *Econometrica* **33**(3) 520–534.
[29] Roughgarden, T. 2005. *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, MA.
[30] Roughgarden, T., E. Tardos. 2002. How bad is selfish routing? *J. ACM* **49** 236–259.
[31] Schmeidler, D. 1973. Equilibrium points of nonatomic games. *J. Statist. Phys.* **7**(4) 295–300.
[32] Wardrop, J. G. 1952. Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engrg.* **2** 325–378.