A Survey of Uniqueness Results for Selfish Routing

Nahum Shimkin

Department of Electrical Engineering
Technion, Israel Institute of Technology
Haifa 32000, Israel
shimkin@ee.technion.ac.il

Abstract. We consider the problem of selfish or competitive routing over a network with flow-dependent costs which is shared by a finite number of users, each wishing to minimize the total cost of its own flow. The Nash Equilibrium is well known to exist for this problem under mild convexity assumptions on the cost function of each user. However, uniqueness requires further conditions, either on the user cost functions or on the network topology. We briefly survey here existing results that pertain to the uniqueness issue. We further consider the mixed Nash-Wardrop problem and propose a common framework that allows a unified treatment of this problem.

1 Introduction

The selfish routing problem involves a number of non-cooperative users, or players, each wishing to ship a certain amount of flow over a shared network, where link costs are flow dependent. A user can choose which route (or routes) to use in order to minimize the total cost of its own flow. This gives rise to a non-cooperative game, with the associated Nash equilibrium as the central solution concept.

Selfish routing was first considered by Wardrop [28] in the context of transportation networks. This paper introduced the notion of shortest-path equilibrium, or Wardrop equilibrium, where only minimal-cost pathes are used between each origin-destination pair. This may be view as the Nash equilibrium of a game between a continuum of infinitesimal users. Recent overviews of the extensive literature that concerns the Wardrop equilibrium and its variants may be found in [22, 20, 3, 25].

The finite-user version of the selfish routing model was introduced in the literature more recently, motivated in part by the non-centralized view of communication networks. The paper [11] shows convergence of the Nash equilibrium (for symmetric users) 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 [21, 2, 4, 23]. The notion of a mixed Nash-Wardrop equilibrium, which combines infinitesimal users with positively-sized ones, is considered in [10, 7]. Efficient network design and management are considered in [13, 14, 16, 15, 9], 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 [12], while [17] considers a repeated game version of the routing problem, and [5] considers the addition of side-constraints on link flows.

Our focus here in on the question of uniqueness of the Nash equilibrium in selfish routing. Besides its theoretical interest, uniqueness is of obvious importance for predicting network behavior in equilibrium. From the computational aspect, efficient procedures that find *all* Nash equilibria are virtually non-existent when the equilibrium is non-unique. Uniqueness is also of particular importance for network management, where regulating the user behavior in a single equilibrium (using pricing, for example) is usually much easier than for several equilibria simultaneously.

Uniqueness is well-known to hold for the basic (single-class) Wardrop equilibrium, assuming only that the link costs are strictly increasing in the link flow. In that case the Wardrop equilibrium has been shown in [6] to be equivalent to convex optimization problem, and hence is unique. However, this simple cost monotonicity requirement no longer suffices for the the finite-user Nash equilibrium, as shown through simple counter-examples, nor for the multi-class Wardrop equilibrium problem (where link costs depend on the user class). Therefore, additional conditions are required to guarantee uniqueness in these cases. Existing conditions may be roughly divided into two types: conditions on link cost functions on the one hand, and conditions on the network topology on the other. In this paper we provide a brief survey of these uniqueness results, focusing on the finite-user model. We will also show that the multi-class Wardrop equilibrium may be embedded within the finite-user problem, and outline a general framework that handles the joint Nash-Wardrop problem in a unified manner.

2 The Game Model

Consider a network which is defined by a directed graph $\mathcal{G} = \mathcal{G}(V, L)$, where V is a finite set of vertices (or nodes) and $L \subset V \times V$ is a set of edges or links. This network is shared by a finite set $I = \{1, 2, \ldots, n\}$ of users, where each user i needs to deliver a given positive amount d^i of flow from its source node O_i to its destination node D_i, and may divide its flow between the set of paths π_i that connect these nodes. Denote by f_l^i the flow of user i on link l, and let $f_l = \sum_{i \in I} f_l^i$ denote the total flow on link l. Furthermore, $\mathbf{f}_l = (f_l^i)_{i \in I}$ is the flow vector over link l, $\mathbf{f}^i = (f_l^i)_{l \in L}$ is user i's flow profile, and $\mathbf{f} = (\mathbf{f}^i)_{i \in I}$ denotes the system flow profile.

The flow profile of each user is subject to the standard positivity and conservation requirements. That is, $f_l^i \geq 0$, and the sum of flows at each node (including external incoming or leaving flow) is null. We denote the set of feasible flow profiles $\mathbf{f}^i$ for user i by F^i, which is clearly a closed, convex polyhedron, and by F the set of feasible system profiles.

Let $J^i(\mathbf{f})$ denote the cost function for user i. We consider additive costs of the form

$$J^i(\mathbf{f}) = \sum_{l \in L} J^i_l(\mathbf{f}_l)\,. \tag{1}$$

Thus, the cost for each user is the sum of its link costs, and the cost of any given link depends only on the flow vector on that link. We further impose here the following assumptions:

Assumption A1: $J^i_l(\mathbf{f}_l) = f^i_l T^i_l(f_l)$.

Assumption A2: The function T^i_l takes values in $[0, \infty]$, and is continuously differentiable, strictly increasing (where finite), and convex.

$T^i_l(f_l)$ is the cost per unit flow for user i on link l. Note that the per-unit costs may differ between users; this may arise, for example, due to user-dependent pricing. A simple consequence of these assumptions is that the link cost function $J^i_l(\mathbf{f}^l)$ is strictly convex in f^i_l, hence the user cost $J^i(\mathbf{f})$ is convex in $\mathbf{f}_l$.

We note that more general cost functions of the form $J^i_l(\mathbf{f}_l) = J^i_l(f^i_l, f_l)$ have been considered in [21] and subsequent literature. However, in this review we will focus on the above-mentioned case, which is of most practical interest.

A cost function that is often used in the context of communication networks is the M/M/1 delay functions, namely $T^i_l(f_l) = \frac{1}{C_l - f_l}$ for $f_l < C_l$, and $T^i_l = \infty$ for $f_l \geq C_l$, where C_l is the link capacity. Note that the above assumptions allows the per-unit costs to assume infinite values, as long as the increase to infinity is continuous.

A flow profile $\hat{\mathbf{f}}$ is a Nash equilibrium point (NEP) if each user's flow profile is a best-response against the combined flows of the others. That is, for each $i \in I$,

$$J^i(\hat{\mathbf{f}}) = \min_{\mathbf{f}^i \in F^i} J^i(\hat{\mathbf{f}}^1, \ldots, \hat{\mathbf{f}}^{i-1}, \mathbf{f}^i, \hat{\mathbf{f}}^{i+1}, \ldots, \hat{\mathbf{f}}^I). \tag{2}$$

A simple consequence of Assumptions A1-A2 above is that the link cost function $J^i_l(\mathbf{f}^l)$ is strictly convex in f^i_l, hence the user cost $J^i(\mathbf{f})$ is strictly convex in $\mathbf{f}_l$. If follows that the above model is a convex game, and existence of the NEP essentially follows from classical results [8, 24]. As the best-response minimization problem faced by each user is a convex program, its solution is unique (whenever finite). However, as is well known, uniqueness of the best response does not guarantee uniqueness of the equilibrium point.

When cost functions take infinite values, some care is needed in distinguishing finite-cost equilibria from infinite-cost ones, where at least one user does not have a finite-cost response to the flow of the others. To exclude existence of infinite-cost equilibria some additional assumptions are required. An fairly straightforward one is the following:

Assumption A3: For any flow configuration $\mathbf{f}$ which involves infinite costs, at least one user whose cost is infinite can modify its flow configuration to obtain a finite cost.

Irrespectively of Assumption A3, our discussion will henceforth focuse on finite-cost equilibria and their uniqueness.

Nonuniqueness: A first counterexample to the uniqueness of the NEP under reasonable convexity assumptions was given in [21], using a two-user four-node network. The user cost functions were not however given in the form of Assumption A1. Counterexamples with cost functions that do comply with A1-A2 are given in [23] for the networks shown the Figure 2 (we return to these networks in Section 5). In all these examples non-uniqueness is essential, in the sense that the user costs are different in the two equilibria.

Elastic demand: The model considered in this paper assumes that flow demands are fixed. Elastic demand can be incorporated into this model by eliminating the demand constraint and subtracting a flow utility term $U^i(d^i)$ from the cost function (1). The utility function is usually assumed to be convex increasing in the flow, which maintains the convexity of the overall cost. One approach to treat the elastic-demand case is to reduce it to the fixed demand model by adding a dedicated link for each user that absorbs its excess flow, with cost that represents the flow utility. A direct proof of uniqueness for the parallel link network may found in [1] and [18].

3 Cost Function Conditions

A general tool for establishing uniqueness of the NEP in convex games is the notion of Diagonal Strict Convexity (DSC) introduced in [24]. This condition may be applied to the network routing problem to obtain per-link sufficient conditions. It then remains to determine what classes of link cost functions satisfy this property.

Let $g^i(\mathbf{f}) = \frac{\partial J^i(\mathbf{f})}{\partial \mathbf{f}^i}$ denote gradient of user i's cost with respect to its flow vector, and for a fixed vector $\rho \in \mathbb{R}^n$ let $g(\mathbf{f}, \rho) = (\rho^i g^i(\mathbf{f}^i))_{i=1}^n$ (arranged as a row vector). Then the cost functions $\{J^i\}$ satisfy the DSC property if $g(\mathbf{f}, \rho)$ is strictly increasing in $\mathbf{f}$ for some positive vector ρ. That is $\rho^i > 0$, and

$$(g(\hat{\mathbf{f}}, \rho) - g(\mathbf{f}, \rho)) \cdot (\hat{\mathbf{f}} - \mathbf{f}) > 0 \quad \text{for all nonequal } \mathbf{f}, \hat{\mathbf{f}} \in F\,. \tag{3}$$

As established in [24], the DSC property implies uniqueness of the equilibrium in the routing game.

The DSC property (3) may be written in scalar notation as

$$\sum_{l \in L} \sum_{i \in I} \rho_i (g_l^i(\hat{\mathbf{f}}_l) - g_l^i(\mathbf{f}_l))(\hat{f}_l^i - f_l^i) > 0 \tag{4}$$

It is now clear that a *sufficient* condition for the DSC property to hold for the overall game is that a DSC-like property holds for each link separately, but with a common weight vector ρ. We summarize this as follows.

Theorem 1. *Suppose there exist numbers $\rho_i > 0$ so that, for each link l,*

$$\sum_{i \in I} \rho_i (g_l^i(\hat{\mathbf{f}}_l) - g_l^i(\mathbf{f}_l))(\hat{f}_l^i - f_l^i) > 0 \tag{5}$$

for any pair of feasible link flows $\hat{\mathbf{f}}_l \neq \mathbf{f}_l$. Then the NEP is unique.

A second-order sufficient condition given in [24] for the DSC property (3) is that the Jacobian matrix G of $g(\mathbf{f}, \rho)$ with respect to $\mathbf{f}$ be positive definite ($G + G^T > 0$) for every feasible $\mathbf{f}$. Applying this condition on to the last result leads to the following result.

Corollary 1 ([21]). *Suppose that matrix $G_l(\mathbf{f}_l, \rho)$ is positive definite for each link l, where*

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

Then the condition of the last theorem holds, and the NEP is unique.

A couple of simple examples from [21] will be useful for illustrating the nature of these conditions, and in particular the effect of the system load and cost function steepness.

Example 1: Assume two users, $I = a, b$, and consider a link l with capacity C_l and M/M/1 costs: $T_l^i = 1/(C_l - f_l)$, where $f_l = f_l^a + f_l^b$. Then for $f_l < C_l$,

$$G_l(\mathbf{f}_l, \rho) = \frac{1}{(C_l - f_l)^3} \begin{pmatrix} 2\rho_a(C_l - f_l^b) & \rho_l(C_l + f_l^a - f_l^b) \\ \rho_b(C_l + f_l^b - f_l^a) & 2\rho_b(C_l - f_l^a) \end{pmatrix}$$

Assume that the total flows are in a rectangle which is bounded away from the link capacity, namely $f_l^a \le r^a$, $f_l^b \le r^b$ where $r^a + r^b < C_l$. It may be easily verified that the DSC condition on the matrix G_l holds with $\rho^a = r^b$ and $\rho^b = r^a$. However, this is not the case under the alternative constraint $f_l^a + f_l^b < C_l$. Indeed, for any fixed vector ρ the matrix $G_l(\mathbf{f}_l, \rho)$ is not positive definite if f_l^a or f_l^b is close enough to C_l. Evidently, the condition in the Corollary is satisfied in lightly loaded networks (flow requirements $d^a + d^b < C_l$ for each link), but not when the feasible total flow on some link exceeds the capacity.

Example 2: Let $T_l^i = P(f_l)$, $i = a, b$, where P is a monic polynomial with degree $m \ge 1$. Then it may be verified that, with $\rho = (1, 1)$, $G_l(\mathbf{f}_l)$ is positive definite over the entire positive quadrant if $m \le 7$, but not if $m \ge 8$. Thus, DSC is implied here if the cost function is “not too steep”.

The next result was established for polynomial-like cost functions of the form

$$T_l^i(f) = a_l f^{p_l} + b_l \, . \tag{6}$$

Such costs find application in the context of road traffic. Let $p^* = \frac{3n-1}{n-1}$, where n is the number of users. Note that $p^* > 3$ for any n.

Theorem 2 ([2]). *Assume the per-link costs (6), with $a_l > 0$ and $0 < p_l < p^*$ for all l. Then the NEP is unique.*

The proof proceeds by demonstrating positive definiteness of the (n by n) matrix $G_l(\mathbf{f}_l, \rho)$, with $\rho_i \equiv 1$.

4 Symmetric Users

A general uniqueness result holds for the case of symmetric users, namely when all users have identical origin-destination pairs, flow demands and link cost functions. In that case the NEP is unique, and indeed turns out to be symmetric (namely with identical link flows) [21]. The proof is by direct analysis, which uses the first-order optimality conditions to show that non-symmetric flows lead to a contradiction.

5 Topological Conditions

Given that uniqueness does not always hold under Assumptions A1-A2 in networks of general topology, the question arises as to whether there exist restricted network topologies for which this general uniqueness property holds. This question was answered in the affirmative in [21] for parallel-link networks. In a recent work, Milchtaich [19] characterized all two-terminal network topologies for which this property holds for the multi-class Wardrop equilibrium. This result was extended in [23] to the finite user model, as described below. We start by defining the following basic property.

Definition 1. *A network $\mathcal{G}$ has the* topological uniqueness property *if the NEP is unique for any routing game over $\mathcal{G}$ that satisfies Assuptions A1-A2.*

The discussion in this section will be focused on *two-terminal networks*, where the source and the destination of all users are the same. The simplest network topology of interest is that of a parallel network: In this case the network has only two nodes, with one serving as the origin node for all users and the other as the destination node. As mentioned, it was shown in [21] that a parallel-link network has the topological uniqueness property.

We proceed to define *nearly parallel networks*, following [19]. As shown there, undirected two-terminal 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. *A network $\mathcal{G}'$ is said to be* embedded in the wide sense *in network $\mathcal{G}''$ if $\mathcal{G}''$ can be obtained from $\mathcal{G}'$ 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.*

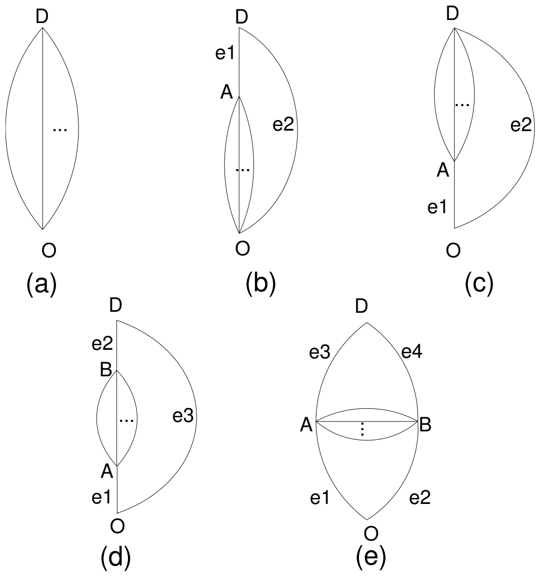

Fig. 1. Basic networks that define the class of nearly-parallel networks

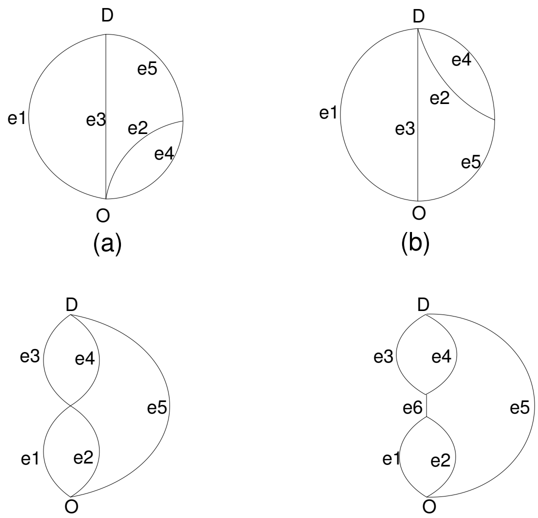

Fig. 2. Basic networks that are not nearly-parallel

Definition 3. *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.*

Theorem 3 ([19]). *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 $\mathcal{G}$.*

The actual (directional) network model is obtained from the non-directional one by replacing each edge with two directional links, one in each direction. Of the five networks in Figure 1, only network (e) supports meaningful bidirectional traffic between some pair of nodes (namely, on the parallel-link network between nodes A and B) given the indicated origin and destination nodes. Indeed, network (e) is the most general of the five, 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.

The following result states that topological uniqueness extends to the class of nearly parallel networks, and *only* to that class.

Theorem 4 ([23]). *A two-terminal network $\mathcal{G}$ has the topological uniqueness property if, and only if, $\mathcal{G}$ is a nearly parallel network or is a series connection of such networks.*

The proof of sufficiency uses specific arguments related to monotonicity properties of the marginal link costs. The proof of necessity proceeds by providing a (three-user) counterexample to uniqueness with cost functions that satisfy A1–A2 for each of the networks shown in Figure 2, and then showing that these basic examples can be extended to any network that is not nearly parallel by using the embedding property in Theorem 3(ii).

6 Mixed Nash-Wardrop Routing

Recall that the Wardrop equilibrium may be considered as the limit of the Nash routing problem, where the user size is infinitesimal. A natural extension to the model is to consider jointly both large (atomic) users and a continuum of infinitesimal users that share the same network, to which we refer as the mixed Nash-Wardrop model [10, 7]. As in the multi-class Wardrop model, we assume that infinitesimal users belong to a (finite) number of user classes, distinguished by their cost functions.

While the equilibrium conditions for atomic and infinitesimal user classes are defined from different perspectives, they actually share common properties and a unified treatment of these two types of users is desirable. In [23] two different approaches for unified treatment are presented, and used in particular to obtain proper extensions of the above topological uniqueness properties to the mixed model. Due to space limitations we do not provide details here. In broad terms, the two proposed approaches are:

1. Reduction to a finite user Nash model: Here each service class is transformed to a single atomic user with an appropriate cost function. This may

be considered a multi-class extension of the well known representation of the single-class Wardrop equilibrium as a (convex) optimization problem.

2. A continuum-game model: Here the framework of non-atomic games [27] is used to model small users. Thus, each user (large or small) is explicitly modelled as a rational decision maker with an individual cost function. This is in contrast to the usual definition of the Wardrop equilibrium, which specifies the behavior of small-user classes via an aggregate flow condition. As opposed to the previous approach, the model allows for a *continuum* of infinitesimal-user classes alongside the discrete population of large users.

In either case, the cost functions obtained for the infinitesimal users or infinitesimal user classes satisfy somewhat weaker conditions than Assumptions A1–A2 and their natural extensions. Still, these conditions do allow to obtain unified uniqueness results for this model, which recover the known topological uniqueness results for both the Nash and Wardrop equilibrium.

7 Conclusion

While new grounds have been gained recently in the analysis of the uniqueness issue in selfish routing, it appears that much remains to be done. Without further conditions on the cost functions, uniqueness results are limited to a fairly restricted class of network topologies. On the other hand, the sufficient conditions that have been explored so far based on diagonal convexity are link-based and do not bring the network topology into play at all. One may hope to find a middle ground that combines cost function properties with other network and user characteristics. This remains a challenging direction for further research.

Acknowledgment. This research was supported by Grant No. 2002112 from the United States–Israel Binational Science Foundation (BSF).

References

1. E. Altman, R. El Azouzi, T. Basar, and R. Srikant, *Combined competitive flow control and routing games*, Workshop on Networking Games and Resource Allocation, Petrozavodsk, July 2002.
2. E. Altman, T. Başar, T. Jiménez, and N. Shimkin, *Competitive routing in networks with polynomial cost*, IEEE Trans. Automat. Contr. **47** (2002), 92–96.
3. E. Altman, T. Boulogne, R. El Azouzi, T. Jiménez, and L. Wynter, *A survey on networking games in telecommunications*, Comput. Oper. Res. **33** (2005), no. 2, 286–311.
4. E. Altman and H. Kameda, *Equilibria for multiclass routing in multi-agent networks*, 2001, In Proc. 40th IEEE Conf. on Decision and Control, Orlando, Dec. 2001, pp. 604–609.
5. R. Azouzi and E. Altman, *Constrained traffic equilibrium in routing*, IEEE Trans. Automat. Contr. **48** (2003), no. 9, 1656–1660.
6. M. Beckmann, C. B. McGuire, and C. B. Winsten, *Studies in the economics of transportation*, Yale University Press, New Haven, 1956.

7. T. Boulogne, E. Altman, H. Kameda, and O. Pourtallier, *Mixed equilibrium (ME) for multiclass routing games*, IEEE Trans. Automat. Contr. **47** (2002), no. 6, 903–916.
8. G. Debreu, *A social equilibrium existence theorem*, Proc. Natl. Acad. Sci. **38** (1952), 886–893.
9. R. El Azouzi E. Altman and O. Pourtallier, *Avoiding paradoxes in multi-agent competitive routing*, Computer Networks **43** (2003), 133–146.
10. P. Harker, *Multiple equilibrium behaviors on networks*, Transportation Sci. **22** (1988), 39–46.
11. A. Haurie and P. Marcotte, *On the relationship between Nash-Cournot and Wardrop equilibria*, Networks **15** (1985), 295–308.
12. T. Jiménez, E. Altman, T. Başar, and N. Shimkin, *Routing into two parallel links: game-theoretic distributed algorithms*, J. Parallel Distrib. Comput. (Special Issue on Routing in Computer and Communication Systems) **61** (2001), 1367–1381.
13. Y. Korilis, A. Lazar, and A. Orda, *Architecting noncooperative networks*, IEEE J. Selected Areas in Communication **13** (1995), no. 7, 1241–1251.
14. ———, *Capacity allocation under noncooperative routing*, IEEE Trans. Automat. Contr. **42** (1997), no. 3, 309–325.
15. ———, *Avoiding the braess paradox in non-cooperative networks*, J. Appl. Probab. **36** (1999), 211–222.
16. Y. A. Korilis, T. A. Varvarigou, and S. R. Ahuja, *Incentive-compatible pricing strategies in noncooperative networks*, Proc. INFOCOM'98 (San-Francisco, CA), April 1998, pp. 439–446.
17. R. J. La and V. Anantharam, *Optimal routing control: Repeated game approach*, IEEE Trans. Automat. Contr. **47** (2002), 437–450.
18. I Menache and N. Shimkin, *Capacity management for and equilibrium for proportional QoS*, Technical Report CCIT No. 575, Department of Electrical Engineering, Technion, Israel. March 2006. Submitted to IEEE Trans. Networking.
19. I. Milchtaich, *Topological conditions for uniqueness of equilibrium in networks*, Math. Oper. Res. **30** (2005), 225–244.
20. A. Nagurney, *Network economics: A variational inequality approach*, 2nd ed., Kluwer Academic, Dordrecht, The Netherlands, 1999.
21. A. Orda, R. Rom, and N. Shimkin, *Competitive routing in multi-user communication networks*, IEEE/ACM Trans. Networking (1993), 510–521.
22. M. Patriksson, *The traffic assignment problem: Models and methods*, VSP, Utrecht, The Netherlands, 1994.
23. O. Richman and N. Shimkin, *Topological uniqueness of the nash equilibrium for selfish routing with atomic users*, Math. Oper. Res. (2007), in press.
24. J. B. Rosen, *Existence and uniqueness of equilibrium points for concave n-person games*, Econometrica **33** (1965), no. 3, 520–534.
25. T. Roughgarden, *Selfish routing and the price of anarchy*, MIT Press, Cambridge, MA, 2005.
26. T. Roughgarden and E. Tardos, *How bad is selfish routing*, J. ACM **49** (2002), 236–259.
27. D. Schmeidler, *Equilibrium points of nonatomic games*, Journal of Statistical Physics **7** (1973), no. 4, 295–300.
28. J. G. Wardrop, *Some theoretical aspects of road traffic research*, Proc. Inst. Civ. Eng. **2** (1952), 325–378.