A survey on networking games in telecommunications

E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez
INRIA, B.P. 93, 2004 Route des Lucioles, 06902 Sophia-Antipolis Cedex, France

Nov. 2000

Abstract

In this survey, we summarize different modelling and solution concepts of networking games, as well as the different applications in telecommunications. We identify some of the mathematical challenges and methodologies that are involved in these problems. We include here work that have relevance to networking games in telecommunications from other areas, and in particular from road traffic.

1 Introduction

With the deregulation of the telecommunication companies and with the rapid growth of the Internet, the research area of networking games has experienced a remarkable development. It has been recognized that methodology used for optimization problems with a single decision maker are not appropriate to model and analyze decentralized networks in which decisions are taken by different agents (service providers or users). It has been recognized that many modelling aspects from networking games in road traffic (which has been an active research area since the fifties), could be successfully used in telecommunications as well. Yet other concepts needed to be developed and new challenges are being posed by the developing telecommunications technology. Research in networking games has been flourishing in different communities (mathematics, economy, information sciences and road-traffic engineering), sometime working with no communication between them. As a result, some concepts and results have appeared independently in different communities.

In preparing this survey on networking games in telecommunications, we attempted to summarize the different modelling and solution concepts, to highlight the different types of applications in which networking games are useful in telecommunications and to identify some of the mathematical challenges and methodologies that are involved in these problems. We include in this survey work that have relevance to networking games in telecommunications from other areas, and in particular from road traffic.

In preparing this survey, we have encountered a rich literature in flow control [3, 4, 13, 14, 16, 17, 18, 15, 49, 65, 66], network routing [6, 7, 8, 11, 22, 23, 20, 25, 32, 33, 38], [55, 60, 61, 62, 63, 68, 1, 91, 96, 97, 105], file allocation [67], load balancing [51, 52, 53, 54, 43], multi-commodity flow [106, 107], resource allocation [27, 46, 71, 111] and quality of service provisioning [111], see also [5, 12, 10]. Some papers have considered the combination of flow and routing in non-cooperative setting, see [9, 45, 79, 94] as well as [91] and references therein. As shown in the latter reference, when the objective functions of the players are the sum of link costs plus a reward which is a function of the throughput, then the underlying game can be transformed into one involving only routing decisions.

In this survey we describe some of the above issues focusing especially on non-cooperative games. We mention different solution (equilibria) concepts, discuss qualitative as well as quantitative issues. In particular, we describe in more details the problem of uniqueness of equilibrium, the Braess paradox, managing the equilibria through design parameters or pricing as well as the Stackelberg framework for hierarchical equilibrium. We finally survey some work on equilibria in cooperative games that are related to resource allocation, to pricing and to the Stackelberg framework.

The structure of the survey is as follows. We begin in Section 2 by presenting basic notions of game theory related to this survey. We present there in more details notions of multi-criteria and hierarchical equilibria as well as potential games. We then describe in Section 3 the state of the art in non-cooperative service provisioning and routing in networks. In Section 4 we discuss the work on non-cooperative flow-control. In Section 5 we discuss the uniqueness of equilibrium and in Section 6 we describe issues related to convergence to equilibrium from initial non-equilibria strategies. Then we survey in section 7 issues related to some properties of equilibria and the way they can be influenced by network architecting and administration.

2 Basic game concepts

In this section we introduce some definitions of equilibria that are typical in optimization problems arising in communication networks. We put a special emphasis on multi-criteria and on hierarchical equilibria.

An important observation in networks is that even when the architecture of the network is fixed, and it operates under a given fixed and known rules of the network administrator, optimization issues cannot be handled in the framework of a single decision maker. In practice, each user has his own criteria to evaluate the network and his utilities. By his decision, a user (or a class of users) influences not only his own utilities but also those of other users. In practice, users can often take decisions concerning routing, as well as the type and amount of resources they wish to obtain. For example, in an ATM environment [35] (one of the mostly used architecture of high speed networks), the users can decide what type of service to require: CBR (Constant Bit Rate), VBR (Variable Bit Rate), ABR (Available Bit Rate, which in contrast to the two previous ones, is an elastic service, i.e. the user adapts his transmission rate to the state of the network) etc. The user may furthermore negotiate the parameters of quality of service, PCR (Peak Cell Rate), CLR (Cell Loss Ratio), maximum delay etc. Different sets of parameters may suit the service requirements of a user. However, the performance measures (such as throughput, CLR, delay) depend not only on ones choices and decisions in establishing the communication, but also on the decisions of other connected users (this dependence is often described as a function of some network "state"). For example, the available resources and the delay of a best-effort type connections (such as ABR)will depend not only on his own choices, but also on the choices made by other users.

In most of the survey, we shall consider the non-cooperative framework, which models the situation in which we cannot expect different users to act in a cooperative way, but rather to optimize in an individual way. A key optimization concept to describe this situation is the Nash equilibrium.

2.1 Nash equilibrium and its multi-criteria variants

Assume that there are N users in the network, each of which aiming to maximise his own unique utility function. Let u_n denote the decision of user n, and $J^n(\mathbf{u}, x)$, his utility, where $\mathbf{u} = (u_1, ..., u_N)$ is the vector of strategies of all N users, and where x is some parameter representing actions of the network manager (or designer), such as pricing parameters, topological parameters and resource allocation decisions. For x fixed, we say that $\mathbf{u}^*(x) = (u_1^*, ..., u_N^*)$ is a Nash equilibrium if no one of the users can improve his utility function by unilateral deviation. More precisely, for each $n \in \{1, 2, \ldots, N\}$, it should satisfy

$$J^n(\mathbf{u}^*(x), x) = \max_{u_n} J^n(u_1^*, ..., u_{n-1}^*, u_n, u_{n+1}^*, ..., u_N^*, x). \tag{1}$$

In practice, a user may have several criteria or utility functions which he may wish to optimise. The utility function of user n is then vector valued: $J^n = (J_1^n, ..., J_1^k)$ (k may be a function of n, which is then written as k_n). We can then extend the above definition in two ways.

(i) **Multi-criteria equilibrium, or a Nash-Pareto equilibrium:**

The maximum in equation 1) is understood in the Pareto sense. More precisely, we say that a vector x of dimension k dominates a vector y of the same dimension if for any $i = 1, ..., k$ we have: $x_i \geq y_i$ with strict inequality for at least one i. In this case we write $x >_{\mathcal{D}} y$. $\mathbf{u}^*$ is called a *multi-criteria equilibrium* if no user can gain by unilaterally deviating (in the sense of the order $>_{\mathcal{D}}$) from his policy: for each n, there does not exist u_n that gives better performances to user n, i.e. such that

$$J^n(u_1^*, ..., u_{n-1}^*, u_n, u_{n+1}^*, ..., u_N^*, x) >_{\mathcal{D}} J^n(\mathbf{u}^*, x).$$

Although the above formulation seems quite natural in telecommunications, we are not aware of any work that actually used it, except for [90] which will be discussed in more details in the next section. In that reference, however, the precise definition of the above equilibrium is not given.

(ii) **Constrained Nash equilibrium:**

Here we assume that user n maximizes his the first component J_1^n of his utility vector, keeping all other components $J_2^n \ldots J_k^n$ bounded by some given constants. One seeks therefore to obtain a Nash equilibrium under constraints. This notion is quite natural in the context of ATM where users express their requirements for quality of service by bounds they wish to have on delays, CLR, etc. For example, for interactive audio applications, the quality of the communication is insensitive to delay as long as it is below approximately 100msec. An audio application could try to minimize losses subjects to a delay on the constraint.

Denote by $\Pi_n(x) = (\mathbf{u} : J_i^n(\mathbf{u}, x) \leq V_i^n, i = 2, ..., k_n)$ the set of N-tuple actions of the N users that satisfy the $k_n - 1$ for user n, where V_i^n are the bounds defining these constraints. $\mathbf{u}^*$ is then said to be a constrained Nash equilibrium if for all n, $\mathbf{u}^* \in \Pi_n(x)$, and in addition,

$$J^n(\mathbf{u}^*, x) = \max_{u_n} \Big(J^n(u_1^*, ..., u_{n-1}^*, u_n, u_{n+1}^*, ..., u_N^*, x) \text{ such that } (u_1^*, ..., u_{n-1}^*, u_n, u_{n+1}^*, ..., u_k^*) \in \Pi_n(x) \Big). \tag{2}$$

We should mention that constrained Nash has been studied in that form in telecommunications applications (e.g. [65]); however it is a special case of a more general way of defining games under constraints, see [95]. The above setting corresponds to the so called "orthogonal constraints" in [95].

2.2 Hierarchical multi-criteria optimization

We now extend the framework to the case in which the network manager (who may represent the network administrator, or the network designer, or a service provider) has a utility, possibly a vector utility, which he wishes to optimize. Among the components of this utility there may be elements that coincide with the users' utilities, as the manager may be interested that the users are satisfied. However, the manager is typically concerned also with efficiency of the use of resources as well as purely economic considerations. Denote by $R(\mathbf{u}, x)$ the utility of the manager. It depends on one hand on the architecture as well as the administration policy of the network, which we denote by x, and on the other hand, on the users' policy $\mathbf{u}$.

In case where the equilibrium $\mathbf{u}^*(x)$ defined in Subsection 2.1 exists and is unique, the objective of the network manager is to determine x that maximizes the his utilities, assuming that the users choose the equilibrium actions $\mathbf{u}^*$. In other words, the objective of the manager is to find x^* that satisfies

$$R(\mathbf{u}^*(x^*), x^*) = \max_x R(\mathbf{u}^*(x), x). \tag{3}$$

When the users and the network manager have scalar utility functions J^n and R (each one has a single utility) then the set $(\mathbf{u}^*(x^*), x^*)$ is the Stackelberg equilibrium [103]. More generally, it may be appropriate in telecommunications applications to allow J^n and R to be vectors. In that case (3) means that there does not exist x such that $R(\mathbf{u}^*(x), x) >_{\mathcal{D}} R(\mathbf{u}^*(x^*), x^*)$.

In the case that there exists a set $U^*(x)$ that contains several equilibria for the users (which is typically occurs in the case of a Pareto-Nash equilibrium), the objective for the network is to guarantee the best performance for any possible equilibrium, i.e. to find x^* that satisfies

$$R(\mathbf{u}^*(x^*), x^*) = \max_x \min_{\mathbf{u}^* \in U^*} R(\mathbf{u}^*(x), x). \tag{4}$$

One may also consider the case of competition between several networks. This can give rise to a more complex hierarchical game; taking into account the reactions $\mathbf{u}^*(\mathbf{x})$, of the N users to the decisions $\mathbf{x} = (x_1, x_2, \dots x_M)$ of M network managers, the solution concept becomes an extension of equation (3) of the form

$$R^i(\mathbf{u}^*(\mathbf{x}^*), \mathbf{x}^*) = \max_{x^i} R^i(\mathbf{u}^*(\mathbf{x}^*_{-i}, (x_i), \mathbf{x}^*_{-i}, x_i)), \tag{5}$$

where $\mathbf{x}^*_{-i} = (x^*_1, \dots, x^*_{i-1}, x^*_{i+1}, \dots, x^*_M)$, R^i represent the utility (scalar or vector) of manager i, and x_i his decisions.

2.3 Potential games

Recently, Monderer and Shapely [84] have identified a class of games which is called "Potential games". This class includes in particular several types of routing games in networks, such as the congestion games introduced by Rosenthal in [96] as well as the routing games introduced by Wardrop [105] (see in particular [98, 99]). A potential game with N players is characterized by a Potential function $\Phi(\mathbf{u})$ such that for any player i, we have

$$J^i(u_i, \mathbf{u}^{-i}) - J^i(v_i, \mathbf{u}^{-i}) = \Phi(u_i, \mathbf{u}^{-i}) - \Phi(v_i, \mathbf{u}^{-i}).$$

The definition has been extended in [98, 99] to a finite number of classes each of which has an infinite population of individuals. This covers as special case the setting of road traffic for which the Wardrop equilibrium is the solution concept. We delay the definition of Wardrop equilibrium and the discussion on that to subsection 3.3. Potential games have nice properties, such as uniqueness of equilibrium and convergence of greedy algorithms to the equilibrium. This is discussed later in the context of networks in more details.

3 Non-cooperative service provisioning and routing in networks

Below we start by presenting a specific basic structure that many games in networks have in common. We then mention several examples. This framework will allow us, in particular, to define yet another equilibrium concept, the Wardrop equilibrium, which is probably the mostly used equilibrium notion in road traffic games, and has recently gained interest also in telecommunications.

3.1 Framework of service provisioning game

Many games arising in networks may be modelled as follows. There are N applications or classes, and m service classes. Application i has a traffic of rate Λ^i, and has to determine how to split it between a subset S^i of service classes available to that user. A strategy of class i is given by an allocation vector $u_i = (u^i_1, u^i_2, \dots, u^i_m)$ where u^i_j is the amount of traffic that user i assigns to service class j. The set of policies for application i is given by the simplex $\sum_{j=1}^m u^i_j = \Lambda^i$, $u^i_j \geq 0, j = 1, ..., m$.

This framework has been used in many different contexts, such as service provisioning [90] and routing games [45, 1]. We start by briefly mentioning the framework in [90] where utilities are defined as follows.

- There is a QoS (Quality of Service) q_k defined per *service class* k, which is a monotone function of the summation $\sum_{i=1}^{N} u_i^k$.
- The utility for class i of assigning u_j^i to class j is given by $J_i(u_i^k, q_k)$ which is assumed to be monotone in both arguments.

In [90], we have $S^i = \{1, ..., m\}$ for all users, and the utility function for using any service class is binary valued. This corresponds to a rough distinction between acceptable or unacceptable QoS. The goal of a user is to maximize the fraction of the traffic that receives acceptable QoS. This gives rise to non-concave utilities and hence to cases of nonexistence of an equilibrium. Conditions are given in [90] for existence of equilibria, and there is some extension to multidimensional QoS per user (Section 3.4). Moreover, there are some results on the convergence of greedy update policies to the equilibrium.

3.2 Routing games with finitely many users

A problem somewhat related to [90] discussed above, yet with much more complex utility functions, is that of routing in networks. Each user has a given amount of flow to ship and has several paths through which he may split that flow. Such a routing game may be handled by models similar to [90] in the special case of a topology of parallel links. This type of topology is studied in details in the first part of [1] as well as in [11]. However, the model [90] does not cover other topologies. Indeed, in more general topologies, the delay over a *path* depends on how much traffic is sent by other users to any other path that shares common links. Such routing games with general topologies have been studied in [45], in the second part of [1], as well as in [11]. A related model is studied by Rosenthal in [96, 97], yet in a discrete setting. Rosenthal proposes a discrete approach to the network equilibrium model. There are n players and each of one has a unit to ship from a source to a destination and wants to minimize its transport cost (which is the sum of the link costs used by the player). It is shown that in such a model there always exists a pure strategy Nash equilibrium. He introduces a kind of discrete potential function for computing the equilibrium. Nevertheless if a player has more than 1 unit to ship such an equilibrium doesn't always exist.

Whereas the existence of equilibrium has been established under general topologies and fairly general costs, using standard tools, the question of uniqueness has been much more delicate; we address this issue in a later section.

3.3 Wardrop and mixed Equilibrium

Routing games have also been studied in the context of road traffic, already since the fifties [105], where a continuous population of *atomless* users is assumed, each user having a negligible impact on the performance of other users. (When there are, in contrast, finitely many players as in the previous section, they are called *atomic* players.) The equilibrium concept for atomless players in routing games is known as Wardrop equilibrium. This setting with an infinite population of atomless players turns out to be much simpler to handle than the finite case: it has uniqueness properties, and equilibria can be computed through a simple mathematical program (see [78] as well as [91] and references therein). These type of results have recently been obtained in more general settings. The models for which these properties hold have been extended in [70] to games with constraints and in [11] to games with different traffic classes (where delays in nodes or in links may depend on the traffic class). This notion of equilibrium could be a good model for telecommunication networks as well, and even in cases of finite (but large) number of users. Indeed, it has been shown in [45] that limits of equilibria obtained as the number of users goes to infinity converge to the Wardrop equilibrium (see also [98, 110]).

To introduce the notion of Wardrop equilibrium, we use again the framework of the first paragraph of Subsection 3.1. In the context of networks (both road traffic as well as telecommunications), each user is identified with the population of individuals (packets or sessions or drivers, etc) that have a given origin $s(i)$ and a given destination $d(i)$. The set S^i is then identified with all the paths available between $s(i)$ and $d(i)$. The choice of a path is made by each one of the individuals. The strategy u_i now describes the behavior of all individuals in class i, so that u_i^j is the flow of individuals of class i that choose a path $j \in S^i$. In the context of Wardrop equilibrium we speak about costs (delay) rather than utilities. Denote by $D_k(\mathbf{u})$, $k \in \{1, ..., m\}$ the delay (or cost) of path k. Wardrop introduced two principles that characterize an equilibrium: (i) all paths actually used between a source and a destination have the same delay, and (ii) this delay is minimal (i.e. it is smaller than or equal to delays over paths not used between the same source and destination). More precisely, let $S_*^i \subset S^i$ be the subset of paths actually used by user i, i.e. the indices j such that $u_i^j > 0$. Then $\mathbf{u}^*$ is a Wardrop equilibrium if and only if it satisfies

$$\min_{k \in S^i} D_k(\mathbf{u}^*) = D_j(\mathbf{u}^*), \qquad \forall j \in S_*^i, \forall i.$$

Wardrop equilibrium is known to exist under fairly general conditions.

On the importance of the concept of Wardrop equilibrium we can learn from the numerous times that it has been reinvented. The results on Wardrop equilibrium were in fact obtained independently almost fifty years later in a context of mobile telecommunications in [42] and in the context of Potential games in [98]. Wardrop principles have also been obtained independently around thirty years before Wardrop in an economical context [92].

Nash equilibrium and Wardrop equilibrium are two extreme cases that can be modelled in networks. But also the combination of these may occur: some agents may have a large quantity of flow to ship (service providers that may control the routing decisions of all their users) while others agents (individual users who determine directly their routing) may have infinitesimal amount of flow to ship. This scenario along with the corresponding equilibrium notion has been formalized and studied in [22, 44, 53, 108].

4 Non cooperative flow control and flow-routing games

Flow control problems have been considered in different settings, both in dynamic as well as static contexts. By "dynamic" we mean that the decisions of users may depend also on some observed state of the network, which might vary dynamically.

4.1 Static flow control

The static flow control problem is related to the question of what should be the average transmission rate of a user. It is known that this type of problem can often be handled as a particular case of routing problems: the question then become how much traffic to route to the network. Moreover, actual routing decisions and flow control decisions can be done simultaneously in the same framework as discussed before, i.e. of routing games. For references, see [45, 91] and references therein. Another more recent example for that approach is [79].

An important feature in all above references is that costs are given in terms of the sum of link costs. There are models in which this is not the case, and in which the basic building blocks of the overall utilities are not link delays (or link costs), but instead ratio between overall throughput (or some power of it) to overall delay. This is the well known *power criterion*, and it has been frequently used in flow control problems. In the non-cooperative context, some variant of the power criterion has been used in [94] in which the utility is related to the sum of powers over the links. The part of the utility in [94] that corresponds to the delay is given by the sum of all link capacities minus all link flows, multiplied by some entropy function. Thus, this utility does not directly reflect the actual expected delay, but has the advantage that it leads to computable Nash equilibrium in the case of parallel links for the combined flow-routing game.

In [9] the actual power criterion is considered, i.e. the ratio between (some increasing function of) the total throughput of a user and the average delay experienced by traffic of that user. The equilibrium for the flow-routing game is obtained for the limiting case as the number of users become very large. The limit is obtained explicitly; there are cases, however, where two equilibria are obtained.

4.2 Dynamic models

Several non-cooperative flow control models have been proposed and analyzed in a dynamic context. Important references are [49, 65] who consider a network with a general topology in which each source has a window end-to-end flow control. The available information for a user is thus the number of packets within the network not yet acknowledged. Each user wishes to maximize the throughput for his own flow, and at the same, wishes that the delay be bounded by some given value. Thus each user has a multi-criteria optimization problem which is formalized as a constrained optimization problem. The equilibrium obtained is decentralized since each user has only local information on his own unacknowledged packets. Hsiao and Lazar [49] obtained threshold equilibrium policies for this problem using the product form of the network as well as Northon's equivalent. The threshold policy is then obtained through coupled Linear Programs. The existence of an equilibrium is established in [65]. A more general theoretical framework for equilibria with constraints in stochastic games is proposed in [19].

In [14], a rate-based flow control is considered in which each user has to dynamically vary his transmission rate. The available information is assumed to be the queue length (or equivalently, the delay) at the bottleneck queue. The total available bandwidth to all controlled sources at this node is assumed to be the node capacity minus the bandwidth used by higher priority traffic. Typical performance measures can be maximizing throughput and minimizing overflow. Note that we may loose in throughput if the queue is empty, and loose packets if it is full. A good tradeoff between these can be obtained by setting an appropriate target queue length and trying to track it. Another possible performance measure may be related to how well the input rate of a connection tracks its share of the available bandwidth. By considering an immediate cost per user which is a weighted sum of the square of the difference between the queue size and its target value, and the square of the difference between the input rate of a connection and its available bandwidth, the problem is cast into the framework of linear quadratic dynamic games. An equilibrium policy is shown to exist and to be unique; moreover it is explicitly computed along with the resulting performance measures.

Another type of dynamic flow control (combined with routing control) is considered in [108]. The players have to ship a given amount of flow within a certain period and can decide dynamically at what rate to ship an each instant. A dynamic mixed equilibrium (that combines both atomless and atomic users) is computed.

5 Uniqueness of the equilibrium

The uniqueness of an equilibrium is quite a desirable property, if we wish to predict what will be the network behavior. This is in particular important in the context of the network administration and management, in which we wish to predict what impact would design parameters have on the performance in equilibrium.

For routing games in networks in the context of an infinite population of atomless users, the uniqueness of the equilibrium (Wardrop equilibrium [105]) has long been known, see e.g. [91] and references therein). The uniqueness holds in some restrictive sense: it is the *link utilisation* that is unique (and only in special cases it is also the flow in each path that is uniquely determined). The uniqueness result has been extended to multi-class networks (i.e. networks in which there are several classes of users, and the delay in a node or a link may depend on the class) in [11, 55].

As shown in [98, 99], the setting of Wardrop turns out to be a potential game. As such it is not surprising that we have the uniqueness of the Wardrop equilibrium since the uniqueness of equilibrium in potential games has been established in [88]; the equilibrium is shown in fact to be unique not only among the Nash equilibria but also among the larger class of correlated equilibria. Note however, that the uniqueness in that reference was established for finitely many users, and thus it does not cover the framework of Wardrop.

Uniqueness of Wardrop type equilibrium have been obtained in some other related problems. Cominetti and Correa considered in [34] a network with an origin, a destination and n bus lines between them (common-line problem), they analyze a general Wardrop equilibrium for this model. Although this is a problem in road-traffic, we mention it here as we believe that related problems could model situations in telecommunication networks as well. In their model, a bus line is characterized by its in-vehicle travel time and its frequency. Passengers choose a set of lines (the common lines), boarding the first bus (which is not congested) in such a set to reach the destination. Due to the congestion the decision of each passenger depends on the decision of the other passengers and the problem is an equilibrium problem. The players' strategy is called a "strategy flow" (since it is a set of decisions for a noncountable infinite population of users); this determines the "lines flows" which is the flows over each bus line. Under general assumptions (the frequency of a bus is a continuous strictly decreasing function with respect to the flow, a condition to obtain relation between strategy flows and line flows), the authors obtain existence and uniqueness of the equilibrium (two sets of common-lines are used) and moreover they show that an increment in passenger flow does not induce an increase in travel time (except at the bound).

The model for which we have uniqueness of the equilibrium in the case of a finite number of users is quite restricted. In fact, a simple counter-example of a network with four nodes is given in [1] in which two different equilibria exist. It is thus not surprising that much effort has been given to understand the conditions under which there is uniqueness of the equilibrium.

A quite powerful tool for establishing uniqueness is the framework of Rosen [95] who introduced the concept of DSC (Diagonal Strict Concavity) which extends concavity to a multi-user setting. However, this condition typically does not hold in routing games, and there are only few cases in which it has been shown to hold: (i) The problem of two users routing into two parallel queues; the DSC conditions are shown in [1] to hold in the case of light traffic. (ii) A network with general topology with, however, some polynomial costs. The DSC has been established in [7].

In the absence of other general tools for establishing uniqueness, and in view of counterexamples that show that there are cases in which it fails, the study of uniqueness has become a complex case by case study. For some topologies, uniqueness has been obtained (for quite general cost functions): for the case of parallel links [1] and for topologies arising from distributed computing with communication lines, see [51] and references therein. Uniqueness has also been established for symmetrical users [1].

Another interesting result related to uniqueness is the following. Assume that there are two equilibria each having the following property: a user sends positive flow over some link if and only if all other users also send positive flow over that link. Then the two equilibria coincide. This has been established in [1] and further extended in [11].

Although the study of equilibria is more involved in the case of a finite number of users than in the infinite case, the uniqueness results obtained (in all the above references) for the finite case are stronger than for the infinite case: the uniqueness is in the sense of the amount of flow that is sent by each user through each path.

Finally, some recent uniqueness results have been established in [22] (for a general topology) and in [53] (in some particular topology) for the mixed equilibrium case, i.e. the setup that combines atomic and unatomic players.

6 Convergence to the equilibrium

The equilibrium has a meaning in practice only if one can assume that it is actually reached from non-equilibria states, since there is no reason to expect a system to be initially at equilibrium. Several approaches have been proposed in the literature to obtain convergence. Some rely on update policies that have some centralised characteristics (in terms of synchronization between the order of update; An example is the round robin update order), whereas other establish convergence under asynchronous best response mechanism which seem more appropriate for describing a real decentralized non-cooperative system.

Rosen [95] who considered the case of a finite number of players, has established the convergence of a dynamic scheme in which the policies are updated continuously (in time) by all users so as to move in the direction of the gradient of the performance measure. In case of a unique Nash equilibrium this scheme is shown to converge to that equilibrium. In case of multiple equilibria, this procedure converges to one of the equilibria and it is possible to predict to which equilibrium it will converge. As already mentioned, the conditions under which Rosen's setting holds in networks are quite restrictive.

An alternative approach for the dynamic convergence of greedy policies to an equilibrium (even in the absence of a unique equilibrium) is in the class of submodular games and supermodular games [104, 112]. In [23], the authors study a load balancing problem. It is shown that depending on the parameters, the costs are either submodular or supermodular; in both cases this then implies the convergence of the greedy algorithm to the unique equilibrium. Examples of convergence in both submodular setting as well as supermodular games (and their combination) in simple queueing problems are presented in [112].

In [100], Shenker considers a model with a single server (exponential) and several sources (the users- they are Poisson), the utility of a user is function of the amount of service and queue length (i.e. congestion). The users are selfish. It is shown that in such a model no service discipline can guarantee optimal efficiency ant that a service discipline called *Fair Share* guarantees fairness, uniqueness of Nash equilibrium and robust convergence.

For routing games with an infinite population of players, it has been shown that greedy updates converge for quite general costs and for general topology; this was shown in fact for the larger class of potential games [98, 99].

In [8] a very simple case of convergence is considered: that of n users routing to two parallel links. The link costs considered there are linear. Both random (asynchronous) greedy as well as round-robin policies are shown to converge to the equilibrium. However, it is also shown that if more than three players update simultaneously their routing strategy then this results in diverging oscillations. To avoid such oscillations in the case of simultaneous updates one has to use relaxation, i.e. each user should apply at each update some linear combination between the previous strategy and the best response one.

Recently, greedy update have been shown in [23] to converge in a simple setting of distributed computing, a network represented by three nodes and three links (two nodes represent two sources of arrival of tasks, or of jobs, and the third node represent the destination; the links between the sources and the destination represent computers, whereas the link between the sources represent a communication line).

We also mention here the paper [67], that considers the problem of how to split a file between several computers; the decisions are taken in a distributed way by the computers themselves (this involves processing and communication delays). Although there is one global objective that is optimized, this problem has some interesting features of a game (or a team) problem since decisions are distributed. The paper compares three distributed gradient based algorithms based on models that have been developed in the domain of mathematical economy. The algorithms compared belong to the class of resource directed approaches, where at each iteration the marginal value of the resource is computed using the current allocation, by each computer in parallel, then an exchange of this computed value is done between all the computers (using either a broadcast or a central computer that receives and propagates the information). In a third stage the allocation of the resource is changed in a way that the computers with a marginal utility below the average will receive less of the resource and those who have above the average marginal utility will receive more. The authors assume a model of the distributed system where the arrival rate is a Poisson process with parameter λ, and the service is exponentially distributed with parameter μ. They goal is to compute the optimal allocation $(x_1, ..., x_n)$, where the x_i is the fraction of the file resource stored at node i. The authors compare a first derivative algorithm, a second derivative algorithm and a pairwise interaction algorithm. These algorithms have four nice properties: optimality, feasibility, monotonicity and convergence. The first one means that when the algorithm converges, then the solution is optimal. The feasibility implies that even before convergence the solutions are feasible, so the allocation proposed at a stage before convergence can be used and the system will be operational. The Monotonicity property implies that at each iteration the utility function monotonically increases. And the final property is that the algorithm converges for small values of a parameter α. The authors show that the second derivative algorithm converges faster than the first derivative one, and that to avoid problems of communication cost between computers the pairwise interaction algorithm can be used. This performs better in time than the others two, even if many more iterations are needed for convergence.

7 Braess paradox, pricing, and Stackelberg Equilibrium

7.1 The Braess paradox

The service providers or the network administrator may often be faced with decisions related to upgrading of the network. For example, where should one add capacity? or where should one add new links? Decisions related to the network capacity and topology have direct influence on the equilibrium that would be attained.

A frequently used heuristic approach for upgrading a network is through *Bottleneck Analysis*. A system bottleneck is defined as “a resource or service facility whose capacity seriously limits the performance of the entire system” [59, p. 13]. Bottleneck analysis consists of adding capacity to identified bottlenecks until they cease to be bottlenecks. In a non-cooperative framework, however, this heuristic approach may have devastating effects; adding capacity to a link (and in particular, to a bottleneck link) may cause delays of all users to increase; in an economic context in which users pay the service provider, this may further cause a decrease in the revenues of the provider. The first problem has already been identified in the road-traffic context by Braess [24] (see also [36, 101]), and have further been studied in the context of queueing networks [20, 25, 32, 33]. In the latter references both queueing delay as well as rejection probabilities have been considered as performance measure. The focus of Braess paradox on the bottleneck link in a queueing context, as well as the paradoxical impact on the service provider have been studied in [79]. In all the above references, the paradoxical behavior occurs in models in which the number of users is infinitely large, and the equilibrium concept is that of Wardrop equilibrium [105]. Yet the problem may occur also in models involving finite number of players (e.g. service providers) for which the Nash equilibrium is the optimality concept. This has been illustrated in [60, 63]. The Braess paradox has further been identified and studied in the context of distributed computing [51, 52] where arrivals of jobs may be routed and performed on different processors. (An updated list of references on Braess paradox is kept in Braess' home page at http://homepage.ruhr-uni-bochum.de/Dietrich.Braess/# paradox).

7.2 Architecting equilibria and network upgrade

The Braess paradox illustrates that the network designer or service providers, or more generally, whoever is responsible to the network topology and link capacities, have to take into consideration the reaction of non-cooperative users to his decisions. This is in particular important when upgrading the network. Some upgrading guidelines have been proposed in [60, 62, 63] so as to avoid the Braess paradox or so as to obtain a better performance. Another way to answer the question of which link in a network should be upgraded is given in [79] who computes the gradient of the performance with respect to link capacities.

A more ambitious aim is to drive the equilibrium to a social optimal solution. In [60] this is carried out under the assumption that a central manager of the network has some small amount of proper flow to be shipped in the network. It is then shown that the manager's routing decision concerning his own flow can be taken in a way that the equilibrium corresponding to the remaining flows attain a socially optimal solution.

7.3 Pricing

An alternative approach to obtaining efficient operating solutions are through pricing. The goal of the pricing design is that each user faces an individual optimization problem whose solution gives rise to a socially optimal solution.

A naive approach for pricing could be to compute a social optimal policy for each user and simply impose a high fine on any user that deviates from it. This approach would require, however, too much centralized computations and signalling and would be difficult to implement. Therefore, the research on pricing schemes in recent years focused on methods to charge locally (at each link or node) for the resources used, which are easy to measure. It is not clear a-priori that local decentralized schemes can indeed yield the same goal as the naive pricing approach, which explains the large amount of research in this area.

It has been shown in [64] that if price at each link is chosen to be proportional to the congestion level at the link then efficient equilibria are obtained. This was done in the context of a finite number of users. A similar approach had been used in [56, 80] in the context of Wardrop equilibrium.

Pricing has already been used as a tool for obtaining efficient equilibria in the case that demand are controlled (rather than routing) in [37, 82, 86] in queueing context. Recently, with the rapid growth of the Internet and its evolution from a heavily subsidized network to a commercial enterprise, much attention has been given to pricing the demand, see [39, 57, 73, 74, 77, 69].

Pricing schemes for attaining efficient equilibria, where both demand and routing are controlled, have been considered in [79, 83].

Pricing is used in another context in [73], where the authors model an ATM network using a microeconomic paradigm. The network offers bandwidth and buffers for rent. The users have to ask, and pay, for the amount of these resources that can provide them the QoS they required. The authors assume that each user knows a bound of the burstiness of his connection and he also knows the minimum bandwidth μ required for the connection. The authors propose an algorithm that converges to an optimal allocation and service provisioning procedure that prevent cell loss. They prove uniqueness of the obtained equilibrium point. The problem with this algorithm is that once the resources are allocated to a connection, these resources cannot be shared by others connections. A useful feature of this approach is that it gives the users the responsibility of purchasing the resources that will best meet their quality requirements and the network is able to direct, using pricing strategies, the users decisions for obtaining an optimal equilibrium point.

Orda and Shimkin study in [89] the a situation of many users who are selfish, each one has to ship its traffic and to this aim has to choose some service classes (according to its traffic types). It is then assumed that the intent of the service provider is to have a unique allocation of each traffic type to one of the service classes; such an allocation is called the *nominal flow allocation*. The reason for this is that the service provider may wish to optimize the use of resources and thus allocate rare resources only to applications that really need them. Pricing is thus required so as to incite users to choose the service class which is adapted to their needs (QoS) and moreover which corresponds to the intent of the network service provider. The pricing mechanisms done by the network (or equivalently the service provider) cannot be complex in order to avoid a large overhead of signalling. The intuitive simplistic pricing that would depend on both the application and the resources is thus excluded, as it is assumed that within the network it is not possible to distinguish between traffic originating from different applications. Given a price an allocation is *user-optimal* if for every traffic type the cost of the class(es) associated to this traffic is minimal, that is to say that each user has no interest to choose (unilaterally) another class for (a part of) its traffic. Prices which induce that the user-optimal flow allocation is unique and coincides with the nominal flow allocation are called *optimal prices*. Orda and Shimkin establish in [89] a necessary and sufficient condition for the existence of optimal prices in a model which doesn't take into account the congestion. They present computation of optimal prices. Then they consider a congestion model. For the uniqueness of the equilibrium corresponding to individually optimality, they have a condition of the kind of the DSC of Rosen, and they establish for some specific costs the existence of optimal prices.

A related problem is studied in [30], where Chen & Park focus on QoS (Quality of Service) and propose an architecture for noncooperative multiclass QoS provision. They assume that routing is given and is not chosen by the users. There are N users which have some QoS requirements and the network provides service classes at each switch; with each service class is associated a price. Then each user chooses a service class (at each switch) such that its requirements will be satisfied at the lowest price (this choice could be done by the network in a *network centric realization*). They transform this problem (NP-complete) into an unscontrained one (polynomial) with Lagrange multipliers, where the Lagrange multipliers have to be chosen such that the QoS requirements will be satisfied. The authors propose an algorithm to adjust these multipliers in an end-to-end feedback loop. Finally they present a performance evaluation of their model.

Low and Lapsley consider in [76] a model where S sources share a network. Each source (i.e. user) s has a path (set of links that source s uses) and a utility function; the source s chooses its transmission rate in order to maximize its own utility. The goal is to propose a pricing that induce the maximization of the global utility in a non-cooperative distributed environment where the sources calculate their rates so as to maximize their own utility. To this aim they consider the dual problem, where each source maximizes its own utility and the service provider minimizes the dual-objective function, the argument of this function may be seen as a price vector (the l^th coordinate represents the price per unit of bandwidth at link l). Once the optimal price is found, each individual source can maximize (separately) its utility and the union of all decisions will maximize the global utility. Then they present two algorithms which converge to the primal-dual optimality, where the provider adapts prices and sources adapt their rates of transmission. In the first one the informations are synchronous, that is to say the rate from all sources that go through link l arrive at the same time t to the provider and price of all links arrive at the same time at the source. The second one is asynchronous.

In [75], Low considers in this reference a network with a single node with an allocation scheme that provides each user with a fixed minimum and a random extra amount of bandwidth and buffer capacity. The network sets prices on the resources. Low considers two models. In the first one each user has an initial allocation and wants to obtain the allocation which maximize its own utility under the constraint that the price of the new allocation is the same as the initial allocation. In the second one the above constraint is absent. It is shown, for the first model, that at equilibrium all users have positive variable allocation in bandwidth and buffer capacity. For the second model, some properties of the equilibrium are exhibited.

Some other references on pricing in networks are [2, 21, 28, 31, 47, 50, 72, 102, 109].

7.4 Stackelberg equilibrium

One further step in the interaction between the manager (which represents the network designer or operator) and the users, is to assume that the former is interested not just in attaining an efficient equilibria for the latter, but may have also some proper utilities (such as maximizing revenues). The situation is then modelled as a a hierarchical optimization one, in which for each action of the manager correspond one or more possible equilibria of the users (obtained by optimizing their responses). One can then quantify the optimization problem faced by the manager, whose solution concept is known as Stackelberg equilibrium [103]. This framework has been studied in [61, 68, 79, 111] in the telecommunication context (see also [29] in the context of road traffic).

8 Cooperative games and resource sharing

Questions of how to share common resources, or how to share the price for constructing network are typical questions that fall into cooperative games (see e.g. [40, 41, 48, 81]. There is a very rich literature in these areas, and instead of presenting a comprehensive survey on these, we refer to [93, 113], both of which contain surveys on cooperative games. We restrict below to some very recent papers in this area, not surveyed in the above references.

In [111] the problem of bandwidth sharing between different users is considered. A general topology of a network is given, and the question is how much bandwidth is to be allocated by the network to each user at each link. More precisely, it is assumed that each user has a minimum lower bound for his demand, and the question is how to allocate extra link capacity. The paper proposes the Nash Bargaining concept [85, 87] for assigning this capacity. This concept is characterized by the following properties: it is Pareto-optimal, it is scale invariant, i.e. the bargaining solution is unchanged if the performance objectives are affinely (i.e. of the form $au + b$) scaled, the solution is not affected by enlarging the domain if agreement can be found on a restricted domain, and the bargaining point is symmetric i.e. does not depend on the specific labels: users with the same lower bounds and objectives receive the same share. It is shown that this sharing of the bandwidth has the proportional fairness property introduced in [58], and is unique. The authors of [111] then address the issue of pricing and its relation to bandwidth allocation. It is shown that based on a user's budget or willingness to pay as well as his bandwidth demands, a bargaining framework can be developed to allocate the network bandwidths to the users in a way which is optimal in the Pareto sense and is fair. Furthermore, based on this, they develop a pricing scheme based on the congestion in the network for which network revenue is maximized when the network operates at the allocations corresponding to the bargaining solution. This pricing scheme has the property that a user is never charged more than his declared budget but could be charged less if the amount of congestion in the network links used by the connection is low. The computation of the Nash bargaining solution can be implemented using a decentralized algorithm that solves a simple mathematical programming.

The Nash bargaining concept has been recently used in [27] for pricing purposes; the solution concept is used to identify a pricing in which the two players are the service provider and the users, respectively. Unlike [111], only simple topologies for the network are considered. However, the analysis in [27] considers also the case of several user priorities which models the possibility for the service provider to offer several quality of services at different prices.

An alternative concept in cooperative games for sharing resources is the Aumann-Shapely pricing, which also has desirable properties (and is Pareto optimal). Haviv proposes in [46] this approach to allocating congestion costs in a single node under various queueing disciplines. Other concepts for arbitration (which are closely related to the Nash bargaining concept) are the Raiffa-Kalai-Smorodinsky and the modified Thomson solutions. Details on these are discussed in [26] which also provides a general framework that unifies the concept of fairness with different arbitration concepts.

We finally mention work that deals with resource sharing, although it does not cooperative game methodology. The paper [71] considers a multiuser network that is shared by noncooperative users. Each user reserves some of the resource in order to establish a virtual path for its incoming calls. A virtual path is a set of preestablished virtual circuit, dedicated to a specific source-destination. Users are noncooperative: each user seek to optimize its own selfish utility, which accounts for the guaranteed call level quantity of service, as well as for the cost incurred for reserving the resource. The authors investigate the Nash equilibrium of such systems. For the case of a shared resource (the resource available to users modelled by a single link), existence and uniqueness of the Nash equilibrium is proven under reasonable conditions. The paper establishes the convergence to this unique equilibrium under Gauss-Seidel and Jacobi schemes. For general network, users may be sharing more than one resource and one may not be able to model the resource available to users by a single link. Each user would have preferences among several links. The paper extends the results of one resource to various general networks topologies. Finally, the formal results are tested by emulating the schemes on an experimental network.

References

[1] N. Rom A. Orda and N. Shimkin. Competitive routing in multi-user communication networks. *IEEE/ACM Transaction on Networking*, 1:614–627, 1993.

[2] C. E. Agneiv. The theory of congestion tolls. *Journal of Regional Science*, 17:381–393, 1977.

[3] E. Altman. Flow control using the theory of zero-sum Markov games. *IEEE Transactions on Automatic Control*, 39:814–818, 1994.

[4] E. Altman. Monotonicity of optimal policies in a zero sum game: A flow control model. *Advances of Dynamic Games and Applications*, 1:269–286, 1994.

[5] E. Altman. Non zero-sum stochastic games in admission, service and routing control in queueing systems. *Queueing Systems*, 23:259–279, 1996.

[6] E. Altman. A Markov game approach for optimal routing into a queueing network. *Anals of Dynamic Games Vol 5: Stochastic and Differential Games, Theory and Numerical Methods, M. Bardi, T.E.S. Raghavan and T. Parthasarathy (Editors), Birkhauser Boston, Basel, Berlin*, pages 359–376, 1999.

[7] E. Altman, T. Başar, T. Jiménez, and N. Shimkin. Competitive routing in networks with polynomial cost. In *Proceedings of IEEE INFOCOM*, Tel-Aviv, Israel, March 2000.

[8] E. Altman, T. Basar, T. Jimenez, and N. Shimkin. Routing into two parallel links: Game-theoretic distributed algorithms. *to appear in J. Distributed and Parallel Comupting*, 2000.

[9] E. Altman, T. Basar, and R. Srikant. Nash equilibria for combined flow control and routing in networks: Asymptotic behavior for a large number of users. In *Proceedings of IEEE CDC'99*, Phoenix, Arizona, USA, Dec. 1999.

[10] E. Altman and A. Hordijk. Zero-sum Markov games and worst-case optimal control of queueing systems. *Queueing Systems*, 21:415–447, 1995.

[11] E. Altman and H. Kameda. Equilibria for multiclass routing problems in multi-agent networks. *submitted*, 2000.

[12] E. Altman and G. M. Koole. Stochastic scheduling games and Markov decision arrival processes. *Computers and Mathematics with Applications*, 26(6):141–148, 1993.

[13] E. Altman and T. Başar. Optimal rate control for high speed telecommunication networks. In *Proc. of the 34th IEEE Conference on Decision and Control*, New Orleans, Louisiana, USA, Dec. 1995.

[14] E. Altman and T. Başar. Multi-user rate-based flow control. *IEEE Trans. on Communications*, pages 940–949, 1998.

[15] E. Altman, T. Başar, and N. Hovakimian. Worst-case rate-based flow control with an arma model of the available bandwidth. *Anals of Dynamic Games*, 6, 1999.

[16] E. Altman, T. Başar, and R. Srikant. Multi-user rate-based flow control with action delays: a team-theoretic approach. In *Proc. of the 36th IEEE Conference on Decision and Control*, San Diego, California, Dec. 1997.

[17] E. Altman, T. Başar, and R. Srikant. Robust rate control for abr sources. In *IEEE INFOCOM, San-Fransisco, California, USA*, 1998.

[18] E. Altman, T. Başar, and R. Srikant. Congestion control as a stochastic control problem with action delays. *Automatica*, 1999.

[19] E. Altman and A. Shwartz. Constrained markov games: Nash equilibria. *to appear in Annals of the International Society of Dynamic Games*, 2000.

[20] N. G. Beans, F. P. Kelly, and P. G. Taylor. Braess's paradox in a loss network. *J. Appl. Prob.*, 34:155–159, 1997.

[21] P. Bergendorff, D. W. Hearn, and M. V. Ramana. Congestion toll pricing of traffic networks. In *Network Optimization, Lecture Notes in Economics and Mathematical Systems*, Berlin, 1997. Springer-Verlag.

[22] T. Boulogne, E. Altman, H. Kameda, and O. Pourtallier. Mixed equilibrium for multiclass routing games. In *Proc. of the 9th International Symposium of Dynamic Games*, Adelaide, Australia, Dec. 2000.

[23] T. Boulogne, E. Altman, and O. Pourtallier. On the convergence to equilibrium in a problem of parallel computing. In *Proc. of the 9th International Symposium of Dynamic Games*, Adelaide, Australia, Dec. 2000.

[24] D. Braess. Uber ein paradoxen der werkehrsplannung. *Unternehmenforschung*, 12:256–268, 1968.

[25] B. Calvert, W. Solomon, and I. Ziedins. Braess's paradox in a queueing network with state-dependent routing. *J. Appl. Prob.*, 34:134–154, 1997.

[26] X. R. Cao. Preference functions and bargaining solutions. In *Proceedings of IEEE Conference on Decision and Control*, pages 164–171, December 1982.

[27] X. R. Cao and U. X. Shen. Internet pricing: comparison and examples. In *Proceedings of IEEE Conference on Decision and Control*, Sydney, Australia, December 2000.

[28] M. Carey and A. Srinivasan. Externalities, average and marginal costs and tolls on congested networks with time-varying flows. *Operations Research*, 41:217–231, 1993.

[29] O. J. Chen and M. Ben-Akiva. Game-theoretic formulations of the interaction between dynamic traffic control and dynamic traffic assignment. *Transportation Research Record*, (1617), 1998.

[30] S. Chen and K. Park. An architecture for noncooperative QoS provision in many-switch systems. In *IEEE INFOCOM*, 1999.

[31] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang. Pricing in computer networks: Motivation, formulation and example. *IEEE/ACM Transactions on Networking*, 1:617–627, October 1993.

[32] J. E. Cohen and C. Jeffries. Congestion resulting from increased capacity in single-server queueing networks. *IEEE/ACM Trans. on Networking*, 5(2):1220–1225, april 1997.

[33] J. E. Cohen and F. P. Kelly. A paradox of congestion in a queuing network. *J. Appl. Prob.*, 27:730–734, 1990.

[34] R. Cominetti and J. Correa. The common-line problem in congested transit networks. *submitted to Transportation Science*, 2000.

[35] The ATM Forum Technical Committee. *Traffic Management Specification*, Version 4.0, af-tm-0056, April 1996.

[36] S. Dafermos and A. Nagurney. On some traffic equilibrium theory paradoxes. *Transportation Research B*, 18:101–110, 1984.

[37] S. Dewan and H. Mendelson. User delay costs and internal pricing for a service facility. *Management Science*, 36:1502–1517, 1990.

[38] A. A. Economides and J. A. Silvester. Multi-objective routing in integrated services networks: a game theory approach. In *Proc. of INFOCOMt'91*, pages 1220–1225, 1991.

[39] R. J. Edel, N. McKeown, and P. P. Varaiya. Billing users and pricing for TCP. *IEEE Journal on Selected Areas in Communications*, 13:1162–1175, Sept. 1995.

[40] D. Granot and G. Huberman. Minimum cost spanning tree games. *Mathematical Programming*, 21, 81.

[41] D. Granot and M. Maschler. Spanning network games. Discussion paper, Center for Rationality and Ineractive Decision Theory, march 1997.

[42] P. Gupta and P. R. Kumar. A system and traffic dependent adaptive routing algorithm for Ad Hoc networks. In *Proceedings of the 36th IEEE Conference on Decision and Control*, pages 2375–2380, San Diego, USA, Dec. 1997.

[43] C. Kim H. Kameda, J. Li and Y. Zhang. *Optimal Load Balancing in Distributed Computer Systems.* Springer, 1997.

[44] P. T. Harker. Multiple equilibrium behaviors on networks. *Transportation Science*, 22:39–46, 1988.

[45] A. Haurie and P. Marcotte. On the relationship between nash-cournot and wardrop equilibria. *Networks*, 15:295–308, 1985.

[46] M. Haviv. The Aumann-Shapely pricing mechanism for allocating congestion costs. *In press*, 2001.

[47] J. V. Henderson. Road congestion: a reonsideration of pricing theory. *Journal of Urban Economics*, 1:346–365, 1974.

[48] S. Herzog, S. Shenker, and D. Estrin. Sharing th cost of multicast trees: an axiomatic analysis. *IEEE/ACM Transactions on Networking*, 5(6):847–860, december 1997.

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

[50] H. J. Huang and H. Yang. Optimal variable road-use pricing on a congested network of parallel routes with elastic demand. In *Proceedings of the 13th International Symposium on the Theory of Traffic Flow and Transportation*, pages 479–500, 1996.

[51] H. Kameda, E. Altman, and T. Kozawa. A case where a paradox like braess's occurs in the nash equilibrium but does not occur in the wardrop equilibrium - a situation of load balancing in distributed computer systems. In *Proceedings of IEEE CDC'99*, Phoenix, Arizona, USA, Dec. 1999.

[52] H. Kameda, E. Altman, T. Kozawa, and Y. Hosokawa. Braess-like paradoxes in distributed computer systems. *to appear in IEEE Transaction on Automatic control*, 2000.

[53] H. Kameda, E. Altman, and O. Pourtallier. Analytic study of mixed optima in symmetric distributed computer systems. In *Proceedings of the 9th International Symposium on Dynamic Games and Applications*, Adelaide, Australia, Dec. 18-21 2000.

[54] H. Kameda, E. Altman, O. Pourtallier, J. Li, and Y. Hosokawa. Paradoxes in performance optimization of distributed systems. In *Proceedings of SSGRR 2000 Computer and ebusiness conference*, l'Aquila, Italy, Jul 31 - Aug 6 2000.

[55] H. Kameda and Y. Zhang. Uniqueness of the solution for optimal static routing in open bcmp queueing networks. *Mathematical and Computer Modelling*, 22:119–130, 1995.

[56] F. P. Kelly. Network routing. *Philosophical Transactions of the Royal Society, Series A*, 337:343–367, 1991.

[57] F. P. Kelly. tariffs and effective bandwidth in multiservice networks. In *Proceedings of ITCť94*, pages 401–410, 1994.

[58] F. P. Kelly. charging and rate control for elastic traffic. *European Transactions on Telecommunications, special issue on Elastic Services over ATM networks*, 8(1), 1997.

[59] H. Kobayashi. *Modeling and Analysis: AN Introduction to System Performance Evaluation Methodology.* Addison Wesley, 1978.

[60] Y. A. Korilis, A. A. Lazar, and A. Orda. Architecting noncooperative networks. *J. on Selected Areas in Communications*, 13(7):1241–1251, Sept. 1995.

[61] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5:161–173, Feb. 1997.

[62] Y. A. Korilis, A. A. Lazar, and A. Orda. Capacity allocation under non-cooperative routing. *IEEE Transactions on Automatic Control*, 42(3):309–325, March 1997.

[63] Y. A. Korilis, A. A. Lazar, and A. Orda. Avoiding the braess paradox in non-cooperative networks. *Journal of Applied Probability*, 36:211–222, 1999.

[64] Y. A. Korilis, T. A. Varvarigou, and S. R. Ahuja. Pricing noncooperative networks. *Bell Laboratories Technical Memorandum, Submitted to the IEEE/ACM Transactions on Networking*, May 1997.

[65] Y.A. Korilis and A. Lazar. On the existence of equilibria in noncooperative optimal flow control. *Journal of the ACM*, 42 No. 3:584–613, 1995.

[66] Y.A. Korilis and A. lazar. Why is flow control hard: optimality, fairness, partial and delayed information. *preprint*, 1995.

[67] J. Kurose and R. Simha. A microeconomic approach to optimal resource allocation in distributed computer systems. *IEEE Trans. on Computers*, 38(5):705–717, 1989.

[68] R. J. La and V. Anantharam. Optimal routing control: game theoretic approach. In *Proc. of the 36th IEEE Conference on Decision and Control*, San Diego, CA, USA, Dec. 1997.

[69] R. J. La and V. Anantharam. Window-based congestion control with heterogeneous users. In *IEEE INFOCOMt2001*, Ankorage, Alaska, April 2001.

[70] T. Larsson and M. Patriksson. Side constrained traffic equilibrium models – analysis, computation and applications. *Transportation Research*, 33B:233–264, 1999.

[71] A. A. Lazar, A. Orda, and D. E. Pendarakis. Virtual path bandwidth allocation in multi-user networks. *IEEE/ACM Transactions on Networking*, 5(6):861–871, Dec 1997.

[72] P. J. Lederer. A competitive network design problem with pricing. *Transportation Science*, 27(1):25–38, february 1993.

[73] S. Low and P. Varaiya. A new approach to service provisioning in ATM networks. *IEEE/ACM Trans. Networking*, 1(5):547–553, 1993.

[74] S. Low and P. Varaiya. An algorithm for optimal service provisioning using resource pricing. In *Proc. IEEE INFOCOM*, pages 363–373, 1994.

[75] S. H. Low. Equilibrium bandwidth and buffer allocations for elastic traffics. *IEEE/ACM Transactions on Networking*, 8(3):373–383, 2000.

[76] S. H. Low and D. E. Lapsley. Optimization flow control–I: basic algorithm and convergence. *IEEE/ACM Transactions on Networking*, 7(6):861–874, 1999.

[77] J. K. MacKie-Mason and H. R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 13:1141–1149, Sept. 1995.

[78] T. L. Magnati. Models and algorithms for predicting urban traffic equilibria. In *Transportation Planning Models*. North Holland, 1984.

[79] Y. Masuda. Braess's paradox and capacity management in decentralized networks. *manuscript*, 1999.

[80] Y. Masuda and S. Whang. Dynamic pricing for network service: equilibrium and stability. *Management Science*, 45:867–869, 1999.

[81] N. Megido. Cost allocation for Steiner trees. *Networks*, 8, august 1978.

[82] H. Mendelson. Pricing computer services: queueing effects. *Communications of the ACM*, 28:312–321, 1985.

[83] H. Mendelson and S. Whang. *Pricing for communication network services.* PhD thesis, Graduate School of Business, Stanford University, 1994.

[84] D. Monderer and L. S. Shapley. Potential games. *Games and Econ. Behavior*, 14:124–143, 1996.

[85] A. Muthoo. *Bargaining theory with applications.* Cambridge University Press, 1999.

[86] P. Naor. On the regulation of queueing size by levying tolls. *Econometrica*, 37:15–24, 1969.

[87] J. Nash. The bargaining problem. *Econometrica*, 18:155–162, 1950.

[88] A. Neyman. Correlated equilibrium and potential games. *Int. J. of Game Theory*, 26:223–227, 1997.

[89] A. Orda and N. Shimkin. Incentive pricing in multi-class communication networks. In *Proceeding of the IEEE INFOCOM'97*, Kobe, Japan, April 1997.

[90] K. Park, M. Sitharam, and S. Chen. Quality of service provision in noncooperative networks with diverse user requirements. *submitted*, 1999.

[91] M. Patriksson. *The Traffic Assignment Problem: Models and Methods.* VSP BV, P.O. Box 346, 3700 AH Zeist, The Netherlands, 1994.

[92] A. C. Pigou. *The Economics of Welfare.* McMillan & Co., London, 1920.

[93] Nicolas Preux. *Pricing, Cooperation and Competition in Telecommunication Networks (in french).* PhD thesis, University of Blaise Pascal, Ecole Doctorale Science pour l'Ingénieur Clermont-Ferrand, France, december 1998.

[94] S. H. Rhee and T. Konstantopoulos. Optimal flow control and capacity allocation in multi-service networks. In *IEEE Conference on Decision and Control*, Tampa, Florida, USA, Dec. 1998.

[95] J. B. Rosen. Existence and uniqueness of equilibrium points for concave N-person games. *Econometrica*, 33:153–163, 1965.

[96] R. W. Rosenthal. A class of games possessing pure strategy nash equilibria. *Int. J. Game Theory*, 2:65–67, 1973.

[97] R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3:53–59, 1973.

[98] W. H. Sandholm. *Evolutionary justification of Nash Equilibrium.* PhD thesis, Northwestern University, 1998.

[99] W. H. Sandholm. Potential games with continuous player sets. *submitted*, 2000.

[100] S. Shenker. Making greed work in networks: a game-theoretic analysis of switch service disciplines. In *Proc. ACM SIGCOMM*, pages 47–57, 1994.

[101] M. J. Smith. In a road network, increasing delay locally can reduce delay globally. *Transportation Res.*, 12:419–422, 1978.

[102] M. J. Smith. The marginal cost taxation of a transportation network. *Transportation Research-B*, 13:237–242, 1979.

[103] H. V. Stackelberg. *Marktform und Gleichgewicht.* Julius Springer, Vienne, 1934.

[104] D. Topkis. Equilibrium points in nonzero-sum n-person submodular games. *SIAM J. Control and Optimization*, 17:773–787, Nov. 1979.

[105] J. G. Wardrop. Some theoretical aspects of road traffic research communication networks. *Proc. Inst. Civ. Eng.*, Part 2, 1:325–378, 1952.

[106] M. P. Wellman. A market-oriented programming environment and its application to distributed multicommodity flow problems. *Journal of Artificial Intelligence Research*, 1:1–23, 1993.

[107] M. P. Wellman. Market-oriented programming: some early lessons. In S. Clearwater, editor, *Market-Based Control: A Paradigm for Distributed Resource Allocation*, Honk-Kong, 1996. World Scientific.

[108] B. W. Wie. A differential game approach to the dynamic mixed behavior traffic network equilibrium problem. *European Journal of Operational Research*, 83:117–136, 1995.

[109] B. W. Wie and R. L. Tobin. Dynamic congestion pricing models for general traffic models. *Transportation Res. -B*, 32(5):313–327, 1998.

[110] B. W. Wie and R. L. Tobin. On the relationship between dynamic nash and instantaneous user equilibria. *Networks*, 28:141–163, 1998.

[111] H. Yaiche, R. Mazumdar, and C. Rosenberg. A game theoretic framework for bandwidth allocation and pricing of elastic connections in broadband networks: Theory and algorithms. *IEEE/ACM Transactions on Networking*, 8(5):667–678, Oct. 2000.

[112] D. D. Yao. S-modular games with queueing applications. *Queueing Systems*, 21:449–475, 1995.

[113] V. Zakharov. Game theory approach in communication networks. In W. Resenstiel A. Patel, L. Petrosian, editor, *OASIS: DIstributed Searcg System in the Internet*. St. Petersburg State University Press, 1999.