Priority Service and Max-Min Fairness

Peter Marbach

***Abstract*—We study a priority service where users are free to choose the priority of their traffic, but are charged accordingly by the network. We assume that each user chooses priorities to maximize its own net benefit, and model the resulting interaction among users as a noncooperative game. We show that there exists an unique equilibrium for this game and that in equilibrium the bandwidth allocation is weighted max-min fair.**

***Index Terms*—Game theory, max-min fairness, pricing, priority service, quality of service, rate control.**

I. Introduction

PRIORITY services can be used to provide differentiated quality of service (QoS) in packet-based networks (see, for example, Diffserv proposal from the IETF [1]). Clearly, networks offering this service model will charge users based on the priority of their traffic; otherwise, all users would declare their traffic as high priority and the above priority model would degenerate to a best-effort service. In this paper, we consider a model where the prices associated with the different priority classes are static. Users can freely choose the priority of their traffic, but are charged accordingly by the network. The aim of this paper is to study the bandwidth-sharing properties of this pricing scheme.

Users accessing a network might run different applications and, therefore, have different service requirements regarding throughput, packet loss, and delay. For our analysis, we characterize users' service requirements through utility functions. Roughly, submitting more high-priority traffic will increase the utility of an individual user. However, doing so will also increase the cost that the user has to pay to the network. We model this tradeoff by an optimization problem and consider the noncooperative game where users choose a priority allocation to maximize their own net benefit (for an introduction to the basic concepts of game theory, refer to [2]). To formulate an equilibrium for the resulting game, we assume that a change in the traffic allocation of a single user has a negligible effect on the overall performance of the network. This assumption has commonly been adopted for the analysis of priced-based resource allocation schemes when many "small" users share the network (see, for example, [3] and [4]). We would like to point out that a similar assumption is used in economic theory (where it is referred to as the "competitive price-taking assumption") to define the Walras equilibrium for large economies [5] and in transportation networks to define the Wardrop equilibrium ([6], [7]).

Manuscript received April 14, 2002; revised October 25, 2002; approved by IEEE/ACM Transactions on Networking Editor A. Orda.

The author is with the Department of Computer Science, University of Toronto, Toronto, ON M5S 1A4, Canada (e-mail: marbach@cs.toronto.edu).

Digital Object Identifier 10.1109/TNET.2003.818196

Using the above framework, we analyzed in [8] the situation where users access a single link which supports a finite number of priorities. For this case, we showed that that there always exists an equilibrium allocation. Furthermore, we showed that when all users have the same utility function, the bandwidth allocation in equilibrium is max-min fair. (We will provide a definition of max-min fairness in Section III.) In this paper, we extend the results of [8] to the network case. As one would expect, the analysis of the network case is more involved than the single-link case and we employ a simpler service model. In particular, we assume an idealized priority service where the network supports a continuum of priorities $u \in \mathcal{I} = [0, \infty)$. In our main result, we show that for a homogeneous user population (all users have the same utility function), this idealized priority service leads to a max-min fair allocation of the network resources. For the general case of a heterogeneous user population, we show that the resulting allocation is weighted max-min fair, where the weights are derived from the users' utility functions.

Bandwidth-sharing properties of price-based resource allocation schemes have received much attention in recent years ([4], [9]–[11]). Kelly *et al.* propose in [4] a rate-based pricing scheme where prices are set as a function of the traffic load on individual links in the network. Kelly *et al.* show that this pricing scheme achieves (in equilibrium) a socially optimal rate allocation. La and Anatharam extend in [9] the work of Mo and Walrand [11], and propose a family of window-based pricing schemes which can be used to achieve a weighted proportional fair allocation, or approximate arbitrarily close a max-min fair allocation. Based on the dual of [4], Low and Lapsley propose in [10] a rate-based scheme and show that it can achieve various bandwidth sharing goals such as proportional fairness and max-min fairness. Yaiche *et al.* [12] use the concept of a Nash Bargaining Solution to derive a price-based resource allocation scheme that can be applied to the available bit rate (ABR) service in ATM networks. As a special case, this approach can be used to derive the pricing mechanism in [4]. A max-min fair bandwidth allocation can also be obtained by means of packet scheduling [13]–[15] and (explicit) rate allocation [16].

Price-based mechanisms for providing differentiated QoS have been studied in [17]–[21]. Odlyzko proposes Paris Metro Pricing (PMP) as a means for providing differentiated QoS [19], [20]. PMP partitions the network into several logically separated channels where channels differ in the prices paid for using them. This scheme can be implemented using the above price-based priority service. Cocchi *et al.* provide in [18] an experimental study which demonstrates that price-based priority services achieve a higher social welfare than best-effort/flat-rate schemes. Park *et al.* consider in [17] and [21] a class-based service model where traffic is served according to users' QoS requirements. However, there is no price associated with the different traffic classes, i.e., the costs incurred to users are purely performance related. Within this framework, Park *et al.* study the situation where users are free to choose their traffic class [21], as well as where users indicate their QoS requirements and a network controller assigns network resources [17].

For the single-link case, the idealized priority service that we consider here is similar to the smart market proposal by MacKie-Mason and Varian [22]. However, in our framework we assume that users are charged for submitted traffic, where in [22] users are charged only for traffic that is actually delivered. In addition, the price users pay is equal to the priority of the traffic in our idealized pricing service, while in the smart market the price is equal to the priority of the first packet that gets dropped. Despite these differences, we show in Section IV that these two schemes lead to the same bandwidth allocation in the single-link case. To our knowledge, there does not exist an extension of the smart market to the network case.

The rest of this paper is organized as follows. In Section II, we define the problem that we consider. In Section III, we provide the definition of a max-min fair rate allocation and introduce a fairness criteria for a price-based bandwidth allocation schemes. In Section IV, we analyze the idealized priority service for the single-link case, and in Section V, we extend the analysis to the network case.

II. Problem Formulation

Consider a network consisting of a set $\mathcal{L} = \{1,\ldots,L\}$ of links, and let C_l be the capacity of link $l \in \mathcal{L}$. Let $\{1,\ldots,R\}$ be the set of users accessing the network. We associate with each user r a route $(l_{r,1},\ldots,l_{r,K_r})$, where K_r is the length of the route and $l_{r,k} \in \mathcal{L}$ is the kth link traversed by user r. We define $H_l \subset \{1,\ldots,R\}$ as the set of users which pass through link $l \in \mathcal{L}$, user r belongs to the set H_l when the link l belongs to the route of user r.

Suppose that the network uses a priority service to provide differentiated QoS. In particular, we assume that the network supports a continuum of priorities given by the set $\mathcal{I} = [0,\infty)$. At each node in the network, traffic is served according to a strict priority rule in time (service discipline) and space (buffer management): priority u traffic is transmitted only if all traffic with priority $v > u$ has been served (priority in time), and priority u traffic is dropped from the buffer only if there is no traffic with priority $v < u$ left in the buffer (priority in space). Each user can then decide on:

1) its transmission rate, i.e., the amount of traffic submitted per unit time;
2) its traffic allocation, i.e., the priority of the submitted traffic.

We represent the allocation of user r by a function $F_r : \mathcal{I} \mapsto \Re_+$, where the value of the function F_r at u represents the total amount of traffic that user r submits with priority u or higher. Note that the allocation function F_r is nonincreasing and we define $F_r^+(u)$ by

$$F_r^+(u) = \lim_{x \downarrow u} F_r(x), \quad u \in \mathcal{I}.$$

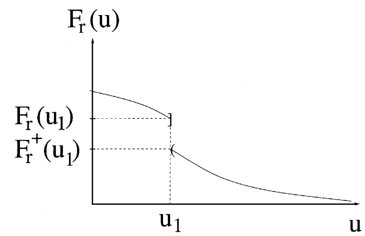

Fig. 1. Allocation function $F_r(u)$ of user r.

Given two priorities u_1 and u_2 such that $0 \leq u_1 < u_2$, the amount of traffic that user submits in the priority range $[u_1, u_2]$ is then given by $F_r(u_1) - F_r(u_2)$. When $F_r(u) > F_r^+(u)$, we can interpret the quantity $(F(u) - F^+(u))$ as the amount of traffic that user r submits in priority u (see Fig. 1). As the function F_r is nonincreasing, we have that

$$F_r^- = \lim_{x \uparrow u} F_r(x) = F_r(u), \quad u \in \mathcal{I}$$

and F_r is continuous from the left [23].

Let $\underline{F} = (F_1,\ldots,F_r)$ be the vector defining the allocation of the individual users $r = 1,\ldots,R$, and suppose that $P_{tr,r}(u,\underline{F})$ is the probability that priority u traffic of user r is delivered to the destination node (and not dropped) under the allocation $\underline{F}$. We will provide a mathematical model for determining the transmission probabilities $P_{tr,r}$ in Section IV; for now, it suffices to note that the above service model implies that high priority traffic is less likely to be dropped at a node, and we have that

$$P_{tr,r}(u,\underline{F}) \geq P_{tr,r}(v,\underline{F}), \quad \text{for } u \geq v.$$

We assume that there exists a feedback mechanism which allows users to determine the transmission probabilities $P_{tr,r}(u,\underline{F})$ under the current allocation $\underline{F}$.

Throughout this paper, we assume that all users perceive the QoS provided by the network solely as a function of their average throughput, i.e., users are insensitive to other QoS parameters such as packet loss and delay. Examples of applications that fall within this framework are email, file transfer, and web browsing. Similar to [8], the analysis presented here can be extended to users which measure QoS as a function of both the average throughput and packet loss. For the idealized priority service that we consider here, it can be shown that both models lead to the same (equilibrium) traffic allocation; we omit a detailed derivation.

For our analysis, we model the service requirements of user r by a utility function U_r which depends on the average throughput x_r of user r, which is given by the Riemann–Stieltjes integral [23]

$$x_r = \int_{\mathcal{I}} P_{tr,r}(u,\underline{F})\,(-dF_r(u)).$$

Assumption I: For every user $r = 1,\ldots,R$, the utility function $U_r : \Re_+ \mapsto \Re_+$ is increasing, bounded, strictly concave, and continuously differentiable. Furthermore, we have that $U_r(0) = 0$.

We note that utility functions with these characteristics are commonly used in the pricing literature (see, for example, [4]). Assumption 1 does not require that all users have the same utility function. Throughout this paper, we assume that the utility function U_r is only known to user r and unknown to all other users and the network.

We assume that users are free to choose the priority for their traffic, but are charged accordingly by the network. In particular, we assume that the network charges a price u per unit traffic submitted to the network in priority u; i.e., the priority u represents the price the network charges per unit traffic. Note that this implies that high-priority traffic is more expensive. Given an allocation F_r, the total cost that user r pays the network is then given by

$$\int_{\mathcal{I}} u \cdot (-dF_r(u)).$$

Assumption 1 implies that a user can increase its utility by submitting more high-priority traffic. However, doing so will also increase the cost that the users has to pay the network. We model this tradeoff by a maximization problem where each user chooses an allocation to optimize its own net benefit, given by utility minus cost. More precisely, given an allocation vector $\underline{F}$, we assume that each user chooses an allocation which solves the following maximization problem:

$$\max_{F_r} \left\{ U_r(x_r) - \int_{\mathcal{I}} u \cdot (-dF_r(u)) \right\}. \quad (1)$$

Note that the above optimization problem assumes that users do not anticipate how their changes will affect the transmission probabilities when they optimize their net benefit. This captures the situation when many "small" users access the network and a change in the traffic allocation of a single user has a negligible effect on the transmission probabilities.

We then consider the situation where all users simultaneously optimize their net benefit. We model this situation as a noncooperative game, and call an allocation vector $(F_1^*, \ldots, F_R^*)$ an equilibrium allocation if for every user r the allocation F_r^* solves the maximization problem given by (1). Under the allocation $(F_1^*, \ldots, F_R^*)$, no user has then an incentive to change its allocation F_r^* as this would decrease the net benefit; hence, an equilibrium is reached. We are interested in the following questions:

1) Does an equilibrium allocation $(F_1^*, \ldots, F_R^*)$ exist?
2) How is the network bandwidth shared among users under an equilibrium allocation $(F_1^*, \ldots, F_R^*)$?

To characterize the bandwidth-sharing properties of an equilibrium allocation, we use the following notation. For every user $r = 1, \ldots, R$, let the function $D_r : \Re_+ \to \Re_+$ be such that

$$D_r(u) = \arg\max_{x \geq 0} \{U_r(x) - xu\}, \quad u \in \Re_+. \quad (2)$$

Under Assumption 1, the maximization problem given by (2) has a unique solution and $D_r(u)$ is given as follows:

$$D_r(u) = \begin{cases} U_n'^{-1}(u) & u \leq u_{r,\max} \\ 0 & u > u_{r,\max} \end{cases}$$

where $u_{r,\max} = U_r'(0)$ and $U_r'^{-1} : [0, u_{r,\max}] \mapsto \Re_+$ is the inverse function of U_r'. Note that Assumption 1 implies that the function D_r is continuous. We refer to D_r as the demand function of user r: The value $D_r(u)$ is the potential demand of user r associated with priority u, i.e., $D_r(u)$ is the amount of traffic users r would submit when the network charges a price u per unit traffic and all traffic is transmitted with probability 1. We define the inverse demand function $G_r(x)$ of user r by

$$G_r(x) = U_r'(x), \quad x \geq 0.$$

In the rest of this paper, we will extensively refer to the demand function D_r and the inverse demand function G_r.

By Assumption 1, we have that

$$\lim_{u \to 0} D_r(u) = \infty.$$

This models that assumption that the unregulated traffic will lead to network congestion, as for every link $l \in \mathcal{L}$ we have $\lim_{u \to 0} \sum_{r \in H_l} D_r(u) = \infty > C_l$. Furthermore, by Assumption 1 we have that

$$\lim_{u \to \infty} D_r(u) = 0.$$

Using this equation, it can be shown that without any loss of generality it suffices to consider allocation functions F_r such that

$$\lim_{u \to \infty} F_r(u) = 0.$$

Let $\mathcal{F}$ be the set of possible allocation functions, i.e., $\mathcal{F}$ is the set of all nonincreasing functions $F : \mathcal{I} \mapsto \Re_+$ that are continuous from the left and for which we have that $\lim_{u \to \infty} F(u) = 0$.

Before we analyze the idealized priority service in Section IV, we briefly review max-min fairness.

III. Bandwidth-Sharing Objectives

An important problem in computer networks is rate allocation, i.e., ensuring that the available network bandwidth is shared among users in a fair manner. A well-established bandwidth-sharing criterion is the max-min fairness criterion as discussed, for instance, by Bertsekas and Gallager [24]. Among others, max-min fairness has been applied to window-based flow control [14], ABR service in ATM networks [16], [25], wireless networks [26], [27], and multicast protocols [28].

A. Max-Min Fairness

One possible way to define fairness for a rate allocation scheme is to require that each user obtains the same transmission rate. For example, when R users access a single link with capacity C, a fair allocation would give each user a transmission rate equal to C/R. Applying this notion to a network consisting of several links could lead to an inefficient use of link resources (see [24]). Instead, one could first allocate the same transmission rate to all users, and then share the remaining network bandwidth to fully utilize the network. More formally, a max-min fair allocation is defined as follows.

Definition I: Let $\{1, \ldots, R\}$ be the set of users accessing the network, let $\mathcal{L}$ be the set of all links in a network, and let C_l be the capacity of link $l \in \mathcal{L}$. Let H_l be the set of all users who pass through link l. We call a rate allocation $(x_1,\ldots,x_R)$ feasible, when for every link $l \in \mathcal{L}$ we have that

$$\sum_{r\in H_l} x_r \le C_l.$$

We call a feasible allocation $(x_1,\ldots,x_R)$ max-min fair, when it is impossible to increase the rate of a user r without losing feasibility or reducing the rate of another user r' with a rate $x_{r'} \le x_r$.

Roughly, this definition states that a max-min fair allocation gives the most poorly treated user (i.e., the user who receives the lowest rate) the largest possible share, while not wasting any network resources.

B. Weighted Max-Min Fairness

When users (applications) have different service requirements, then the network may not want to share bandwidth equally among users. Instead, the network could assign weights (priorities) to users and allocate bandwidth accordingly [24]. Let $B_r(x_r)$, $x \ge 0$, be an increasing function representing the weight assigned to user r at rate x_r. An allocation $(x_1,\ldots,x_R)$ is then weighted max-min fair when for each user r, any increase in x_r would cause a decrease in the transmission rate $x_{r'}$ for some user r', satisfying $B_r(x_{r'}) \le B_r(x_r)$. Definition 1 of a max-min fair allocation is recovered with

$$B_r(x_r) = x_r, \quad r = 1,\ldots,R.$$

C. Max-Min Fairness and Pricing

In [8], we analyzed the above pricing scheme for the situation where users access a single link which supports a finite number of priorities. We showed that this scheme leads to a weighted max-min fair bandwidth allocation with weights

$$B_r(x_r) = -G_r(x_r), \quad r = 1,\ldots,R$$

where $G_r(\cdot)$ is the inverse demand function of user r. This result suggests that there exists a close relationship between a price-based priority services and max-min fairness. To reinforce this relation, we propose in this section a fairness criterion for general pricing schemes which leads to a weighted max-min fair bandwidth allocation, and show in Section IV and V that the idealized price-based priority service considered here satisfies this fairness criterion.

One possible way to define fairness for a pricing scheme is to require that each user is charged the same price u^* per unit transmission rate. For example, when R users access a single link with capacity C, a fair pricing scheme would charge each user the same price u^*, where u^* is such

$$\sum_{r=1}^{R} D_r(u^*) = C.$$

For the network case, insisting that each user pays the same price per unit transmission rate could result in the situation that the network capacity is wasted (see the example given in Fig. 2, which is based on the example in [24, p. 525]).

To avoid this, we could allow that the network charges users different prices as long as the poorest treated users (i.e., the users who are charged that highest price per unit transmission rate) are charged as low a price as possible. More formally, we define a fair pricing scheme as follows.

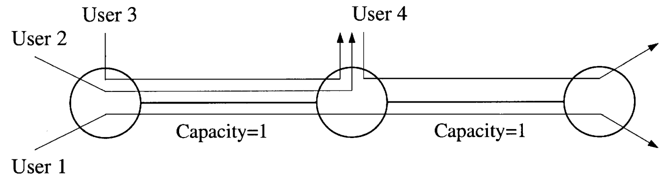

Fig. 2. The above network consists of two links each with a capacity equal to 1. Assume that all users (sessions) have the same utility function and, therefore, the same demand function, which we denote by $D_0(u)$, $u \ge 0$. Requiring that all users pay the same price u^* would lead to an allocation where the transmissions rates are given by $x_1 = x_2 = x_3 = x_r = D_0(u^*) = 1/3$. However, under this allocation only $x_1 = x_2 = 2/3$ units bandwidth are used on the second link, and this allocation wastes network capacity.

Definition 2: Let $\{1,\ldots,R\}$ be the set of users accessing the network, let $\mathcal{L}$ be the set of all links in a network, and let C_l be the capacity of link $l \in \mathcal{L}$. Let H_l, $l \in \mathcal{L}$ be the set of all users who pass through link l. Furthermore, let $D_r(u)$ be the demand function associated with user r. We identify a pricing scheme by the vector $(u_1,\ldots,u_R)$, where u_r is the price that user r is charged per submitted packet. We say that an allocation under the price vector $(u_1,\ldots,u_R)$ is feasible, when for every link $l \in \mathcal{L}$, we have that

$$\sum_{r\in H_l} D_r(u_r) \le C_l.$$

We call a price vector *fair* when the corresponding allocation is feasible, and when it is not possible to reduce the price of a user r without losing feasibility or increasing the price of another user r' with a price $u_{r'} \ge u_r$.

Definition 2 implies that when a price vector $(u_1,\ldots,u_R)$ is fair, then we have

$$u_r \le u_{r,\max} = U_r'(0), \quad r = 1,\ldots,R.$$

We then define $\mathcal{U}$ as the set of all price vectors that satisfy the above constraint, i.e., we have

$$\mathcal{U} = \{(u_1,\ldots,u_R) : 0 \le u_r \le u_{r,\max} \text{ for } r = 1,\ldots,R\}.$$

In the next proposition, we show that pricing scheme is fair as defined above if and only if the resulting allocation is weighted max-min fair with weighting functions $B_r(x_r) = -G_r(x_r)$.

Lemma 1: Let Assumption 1 hold. Consider a price vector $(u_1,\ldots,u_R) \in \mathcal{U}$ and let $(x_1,\ldots,x_R)$ be the corresponding rate vector, i.e., for $r = 1,\ldots,R$, we have that $x_r = D_r(u_r)$ and $u_r = G_r(x_r)$. Then the price vector $(u_1,\ldots,u_R)$ is fair if and only if the rate vector $(x_1,\ldots,x_R)$ is weighted max-min fair for the weight functions $B_r(x_r) = -G_r(x_r), r = 1,\ldots,R$. The above proposition implies that when all users have the same utility function, then a fair pricing scheme results in a max-min fair allocation. Below, we provide a proof for Lemma l.

In the following, we show that the idealized priority service that we introduced in Section II is fair according to Definition 2 and, therefore, leads to a weighted max-min fair rate allocation with weights $B_r(x_r) = -G_r(x_r)$.

Proof: We now prove Lemma l. Let the vectors $(u_1,\ldots,u_R)$ and $(x_1,\ldots,x_R)$ be as in Lemma 1, and consider a new price vector $(\hat{u}_1, \ldots, \hat{u}_R) \in \mathcal{U}$ and rate vector $(\hat{x}_1, \ldots, \hat{x}_R)$ such that for $r = 1, \ldots, R$, we have $\hat{x}_r = D_r(\hat{u}_r)$ and $\hat{u}_r = G_r(\hat{x}_r)$. Assumption 1 implies that the inverse demand function G_r is strictly decreasing and it follows that

$$\hat{u}_r < u_r \quad \text{if and only if} \quad \hat{x}_r > x_r, \quad r = 1, \ldots, R.$$

Using this observation, we prove that the price vector $(u_1, \ldots, u_R)$ is fair if and only if the rate vector $(x_1, \ldots, x_R)$ is weighted max-min fair.

First, we show that when the vector $(u_1, \ldots, u_R)$ is not fair, then the vector $(x_1, \ldots, x_R)$ is not weighted max-min fair with weight functions $B_r(x_r) = -G_r(x_r)$.

Assume that the price vector $(u_1, \ldots, u_R)$ is not fair. Then there exists a user r and a feasible price vector $(\hat{u}_1, \ldots, \hat{u}_R)$ such that $\hat{u}_r < u_r$, and

$$\hat{u}_{r'} \leq u_{r'}, \quad \text{for all users } r' \text{ such that } u_{r'} \geq u_r.$$

Let the rate vector $(\hat{x}_1, \ldots, \hat{x}_R)$ be such that for $r = 1, \ldots, R$, we have $\hat{x}_r = D_r(\hat{u}_r)$ and $\hat{u}_r = G_r(\hat{x}_r)$. As the price vector $(\hat{u}_1, \ldots, \hat{u}_R)$ is feasible, the rate vector $(\hat{x}_1, \ldots, \hat{x}_R)$ is feasible. Furthermore, using the above observation, it follows that $\hat{x}_r > x_r$ and $\hat{x}_{r'} \geq x_{r'}$ for all users r' such that

$$-u_{r'} = B_{r'}(x_r) \leq B_r(x_r) = -u_r.$$

This implies that the rate vector $(x_1, \ldots, x_R)$ is not weighted max-min fair with weight functions $B_r(x_r) = -G_r(x_r)$.

Next, we show that when the vector $(x_1, \ldots, x_R)$ is not weighted max-min, then the vector $(u_1, \ldots, u_R)$ is not fair.

Assume that the price vector $(x_1, \ldots, x_R)$ is not weighted max-min fair. Then there exists a user r and a feasible rate vector $(\hat{x}_1, \ldots, \hat{x}_R)$ such that $\hat{x}_r > x_r$ and $\hat{x}_{r'} \geq x_{r'}$ for all users r' such that

$$-u_{r'} = B_{r'}(x_{r'}) \leq B_r(x_r) = -u_r.$$

Let the rate vector $(\hat{u}_1, \ldots, \hat{u}_R)$ be such that for $r = 1, \ldots, R$, we have $\hat{x}_r = D_r(\hat{u}_r)$ and $\hat{u}_r = G_r(\hat{x}_r)$. As the vector $(\hat{x}_1, \ldots, \hat{x}_R)$ is feasible, the vector $(\hat{u}_1, \ldots, \hat{u}_R)$ is feasible. Furthermore, using the above observation, it follows that $\hat{u}_r < u_r$, and

$$\hat{u}_{r'} \leq u_{r'}, \quad \text{for all users } r' \text{ such that } u_{r'} \geq u_r.$$

This implies that the vector price vector $(u_1, \ldots, u_R)$ is not fair. ■

IV. Single-Link Case

In this section, we analyze the idealized priority service for the single-link case.

A. Link Model

For our analysis, we use a discrete-time model where time is divided into slots of equal length. Let $C > 0$ be a given constant, and consider a single link with a capacity of C units of traffic per time slot. There is no buffering and traffic which does not get transmitted in a given time slot is dropped. Consider a given time-slot, and let

$$F(u) = \sum_{r=1}^{R} F_r(u), \quad u \geq 0$$

be the aggregate amount of traffic submitted during this time-slot in priority u or higher. The link then serves traffic as follows. All traffic with priority u such that $F(u) \leq C$ is transmitted. For priority u^* with

$$F^+(u^*) < C \leq F(u^*)$$

only $C - F^+(u^*)$ units of traffic—chosen at random—are transmitted. Traffic with priority u such that $F^+(u) \geq C$ is dropped. The transmission probability $P_{tr}(u, F)$ for traffic in priority u is then given by

$$P_{tr}(u, F) = \begin{cases} 0, & \text{if } F^+(u) \geq C \\ \frac{C - F^+(u)}{F(u) - F^+(u)}, & \text{if } F(u) > C > F^+(u) \\ 1, & \text{otherwise.} \end{cases} \tag{3}$$

B. Equilibrium Allocation

For the single-link case, we define an equilibrium allocation as follows.

Definition 3: We call an allocation $(F_1^*, \ldots, F_R^*) \in \mathcal{F}^R$ an equilibrium allocation when for every user $r = 1, \ldots, R$, we have that

$$F_r^* = \arg \max_{F_r \in \mathcal{F}} \left\{ U_r(x_r) - \int_{\mathcal{I}} u \cdot (-dF_r(u)) \right\}$$

where $x_r = \int_{\mathcal{I}} P_{tr}(u, F^*)(-dF_r(u))$ and $F^* = \sum_{r=1}^{R} F_r^*$. We have the following result.

Proposition 1: Under Assumption 1, there exists an unique equilibrium allocation $(F_1^*, \ldots, F_R^*)$ which is given as follows. For $r = 1, \ldots, R$, we have

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \leq u^* \\ 0, & \text{otherwise} \end{cases}$$

where u^* is the unique price such that $\sum_{r=1}^{R} D_r(u^*) = C$.

We provide a proof in Appendix A. Proposition 1 states that there exists a unique equilibrium allocation and all users choose the same priority u^* in equilibrium. Furthermore, the aggregated amount of traffic submitted in priority u^* is equal to the link capacity C, and we have

$$\sum_{r=1}^{R} x_r^* = \sum_{r=1}^{R} D_r(u^*) = C.$$

We have the following corollary.

Corollary 1: Under Assumption 1, the bandwidth allocation in equilibrium is weighted max-min fair with weights

$$B_r(x_r) = -G_r(x_r), \quad r = 1, \ldots, R.$$

Proof: Proposition 1 states that in equilibrium only priority u^* is used and the throughput of user r is given by

$$x_r^* = D_r(u^*).$$

Consider the price vector $(u_1, \ldots, u_R) \in \mathcal{U}$ where

$$u_r = \min\{u_{r,\max}, u^*\}, \quad r = 1, \ldots, R.$$

Note that

$$D_r(u_r) = D_r(u^*) = x_r^*, \quad r = 1, \ldots, R$$

and by Proposition 1, we have that

$$\sum_{r=1}^R x_r^* = \sum_{r=1}^R D_r(u^*) = \sum_{r=1}^R D_r(u_r) = C.$$

Using Definition 2 and Lemma 1, this implies that the price vector $(u_1, \ldots, u_R)$ is fair (Definition 2), and the rate allocation $(x_1^*, \ldots, x_R^*)$ is weighted max-min fair with weights $B_r(x_r) = -G_r(x_r)$ (Lemma 1). ■

V. Network Case

In this section, we extend the above analysis to the network case. This case is more involved, as traffic of user r may get dropped by a congested link along its route. As a result, we are cannot explicitly express the link transmission probabilities by an equation similar to (3); instead we have to use a fixed-point equation to determine the link transmission probabilities.

A. Transmission Probabilities $P_{tr,r}$

Consider the network model described in Section II, and let $F_{r,l_k}(\underline{F})$ be the allocation of user r at the kth link along its route under the allocation vector $\underline{F}$; i.e., $F_{r,l_k}(u, \underline{F})$ is the amount of traffic that user submits at link $l_{r,k}$ with priority equal or higher to u. As traffic of user r may get dropped by a congested link along its route, we have that

$$F_r(u) \geq F_{r,l_k}(u, \underline{F}), \quad u \in \mathcal{I}, k = 2, \ldots, K_r.$$

Consider a link $l \in \mathcal{L}$ and let H_l be the sets of users passing through link l. The total amount of traffic submitted on link l in priority u or higher under the allocation $\underline{F}$ is then given by

$$\sum_{r \in H_l} F_{r,l}(u, \underline{F}).$$

We assume that each link $l \in \mathcal{L}$ serves traffic according to the link model of Section II. The transmission probabilities of priority u traffic at link $l \in \mathcal{L}$ is then given by

$$P_l(u, \underline{F}) = P_{tr}\left(u, \sum_{r \in H_l} F_{r,l}(\underline{F})\right), u \in \mathcal{I}$$

where the function $P_{tr}(u, \sum_{r \in H_l} F_{r,l}(\underline{F}))$ is defined by (3).

Note that there is a coupling between the link transmission probabilities $P_{tr,l}$ and the allocation $F_{r,l_{r,k}}$ of user r on link $l_{r,k}$. The link transmission probabilities determine whether traffic of user r gets dropped on the links along the route and, therefore, determine the allocation $F_{r,l_{r,k}}$. However, the link transmission probabilities depend themselves on the allocation of the users who pass through link l. We model this coupling in the next subsection through a fixed-point equation.

We assume that traffic is dropped independently at each link, and the transmission probabilities $P_{tr,r}(u, \underline{F})$, $u \in \mathcal{I}$, for the traffic of user r is given by

$$P_{tr,r}(u, \underline{F}) = \Pi_{k=1}^{K_r} P_{l_{r,k}}(u, \underline{F}), \quad u \in \mathcal{I}. \tag{4}$$

B. Fixed-Point Equation

Consider a given user r with route $(l_{r,1}, \ldots, l_{r,K_r})$ and let $F_{r,l_{r,k}}(u)$ be the amount of traffic that user r submits to link $l_{r,n}$, $n = 1, \ldots, K_r$. For link $l_{r,1}$, the allocation function $F_{r,l_{r,1}}(u)$ is given by

$$F_{r,l_{r,1}}(u) = F_r(u), \quad u \in \mathcal{I}.$$

Let $P_{l_{r,1}}(u, \underline{F})$ be the transmission probability function for the first link $l_{r,1}$ on the route of user r. The (average) amount of traffic that user r submits to link $l_{r,2}$ is then given by

$$\begin{aligned} F_{r,l_{r,2}}(u, \underline{F}) &= \int_u^\infty P_{l_{r,k-1}}(v, \underline{F}) \left(-dF_{r,l_{r,1}}(v)\right) \\ &= \int_u^\infty P_{l_{r,k-1}}(v, \underline{F}) \left(-dF_r(v)\right). \end{aligned}$$

More generally, when $F_{r,l_{r,k}}(u, \underline{F})$ is the amount of traffic that user r submits on link $l_{r,k}$ and $P_{l_{r,k}}(u, \underline{F})$ is the transmission probability function of link $l_{r,k}$, then the amount of traffic that user r submits to the next link $l_{r,k+1}$ along its route is given by

$$F_{r,l_{r,k+1}}(u, \underline{F}) = \int_u^\infty P_{l_{r,k}}(v, \underline{F}) \left(-dF_{r,l_{r,k}}(v)\right).$$

Combining the above results, we obtain the following system of equations between the link transmission probabilities $P_l(u, \underline{F})$ and the link allocations $F_{r,l_k}(u, \underline{F})$. For each user $r = 1, \ldots, R$ we have that

$$F_{r,l_{r,1}}(u, \underline{F}) = F_r(u) \tag{5}$$

$$\begin{aligned} F_{r,l_{r,k}}(u, \underline{F}) &= \ldots \\ &= \int_u^\infty P_{l_{r,k-1}}(v, \underline{F}) \left(-dF_{r,l_{r,k-1}}(v)\right), \\ &\qquad k = 2, \ldots, K_r. \end{aligned} \tag{6}$$

and for each link $l \in \mathcal{L}$ we have that

$$P_l(u, \underline{F}) = P_{tr}\left(u, \sum_{r \in H_l} F_{r,l}(\underline{F})\right), \quad u \in \mathcal{I}. \tag{7}$$

The above system of equations is consistent when there exists a set of link transmission probabilities $P_l(\cdot, \underline{F})$ which satisfies (5)–(7). The following result establishes that there exists such a set of link transmission probabilities.

Proposition 2: Under Assumption 1, for every allocation $\underline{F} = (F_1, \ldots, F_R) \in \mathcal{F}^R$, there exists a set of transmission probability functions $P_l(u, \underline{F}) \in \mathcal{P}$, $l \in \mathcal{L}$ which satisfies the system of equations given by (5)–(7).

We provide a proof in Appendix B.

C. Equilibrium Allocation

For the network case, we define an equilibrium allocation as follows.

Definition 4: We call an allocation $\underline{F}^* = (F_1^*, \ldots, F_R^*) \in \mathcal{F}^R$ an equilibrium allocation when for every user $r = 1, \ldots, R$, we have that

$$F_r^* = \arg\max_{F_r \in \mathcal{F}} \left\{ U_r(x_r) - \int_{\mathcal{I}} u \cdot (-dF_r(u)) \right\}$$

where

$$x_r = \int_{\mathcal{I}} P_{tr,r}(u, \underline{F}^*)\,(-dF_r(u))$$

$$P_{tr,r}(u, \underline{F}) = \Pi_{k=1}^{K_r} P_{l_{r,k}}(u, \underline{F}), \quad u \in \mathcal{I}$$

and the functions $P_l(u, \underline{F}^*)$, $l \in \mathcal{L}$, satisfy the system of equations given by (5)–(7).

We have the following result.

Proposition 3: Under Assumption 1, an allocation $\underline{F}^* = (F_1^*, \ldots, F_R^*) \in \mathcal{F}^R$ is an equilibrium allocation if and only if there exist transmission probabilities $P_l(u, \underline{F}^*)$ such that the following properties hold.

1) The transmission probabilities $P_l(u, \underline{F}^*)$ satisfy the system of equations given by (5)–(7).
2) For every user $r = 1, \ldots, R$, there exists a price $\hat{u}_r$ such that

$$P_{tr,r}(u, \underline{F}^*) = \begin{cases} 1, & u \geq \hat{u}_r \\ 0, & \text{otherwise} \end{cases}$$

where $P_{tr,r}(\hat{u}_r, \underline{F}^*)$ is given by (4).
3) Setting $u_r^* = \min\{\hat{u}_r, u_{r,\max}\}$, we have for every user $r = 1, \ldots, R$ that

$$F_r^*(u) = \begin{cases} D_r(u_r^*), & u \leq u_r^* \\ 0, & \text{otherwise.} \end{cases}$$

This proposition can be proved by the same argument as the one given for Proposition 1, and we omit a detailed proof.

Proposition 3 states that each user r submits traffic exclusively in priority u_r^* and the average throughput x_r^* of user r is given by $F_r(u_r^*) = D_r(u_r^*)$. In addition, the transmission probabilities $P_{tr,r}(u, \underline{F}^*)$, $u \in \mathcal{I}$ for traffic by user r have the following properties. All traffic with priority lower than u_r^* is dropped, i.e., we have

$$P_{tr,r}(u, \underline{F}^*) = 0, \quad \text{for } u < u_r^*.$$

Furthermore, when $D_r(u_r^*)$ is positive then all traffic with priority u_r^* is transmitted, i.e.,

$$P_{tr,r}(u_r^*, \underline{F}^*) = 1.$$

Intuitively, this means that every user passes through a bottleneck link l which drops all traffic with priority less than u_r^*, i.e., at the bottleneck link l we have

$$\sum_{r' \in H_l} D_{r'}(u_r^*) \geq C_l$$

where the above equation holds with equality when $D_r(u_r^*)$ is positive.

By Proposition 3, we have that

$$u_r^* \leq u_{r,\max}, \qquad r = 1, \ldots, R$$

and we can uniquely identify an equilibrium allocation $(F_1^*, \ldots, F_R^*)$ by the price vector $(u_1^*, \ldots, u_R^*) \in \mathcal{U}$. In the following, we show that there exists a unique equilibrium allocation and, moreover, the price vector $(u_1^*, \ldots, u_R^*) \in \mathcal{U}$ associated with this equilibrium allocation is fair (as given by Definition 2).

Let $\mathcal{B} \subset \mathcal{U}$ be the set of all price vectors $(u_1, \ldots, u_R)$ which lead to a feasible allocation, i.e., $\mathcal{B} \subset \mathcal{U}$ is the set of all price vectors $(u_1, \ldots, u_R) \in \mathcal{U}$ such that

$$\sum_{r \in H_l} D_r(u_r) \leq C_l, \quad \text{for all } l \in \mathcal{L}.$$

Given a vector $(u_1, \ldots, u_R) \in \mathcal{B}$, we say that link l is a bottleneck with respect to $(u_1, \ldots, u_R)$ for a user $r \in H_l$, if

$$\sum_{r' \in H_l} D_{r'}(u_{r'}) = C_l$$

and $u_r \leq u'_r$, for all users $r' \in H_l$.

Furthermore, we associate with a price $(u_1, \ldots, u_R) \in \mathcal{B}$ the allocation $(F_1, \ldots, F_R) \in \mathcal{F}$ given by

$$F_r(u) = \begin{cases} D_r(u_r), & u \leq u_r \\ 0, & \text{otherwise.} \end{cases}$$

We have the following results.

Proposition 4: Under Assumption 1, the allocation $\underline{F}^*$ associated with a given price vector $(u_1^*, \ldots, u_R^*) \in \mathcal{B}$ is an equilibrium allocation if and only if every user r has a bottleneck link with respect to $(u_1^*, \ldots, u_R^*)$.

Proof: To prove Proposition 4, we adapt the argument used in [24] to prove the equivalent result for a max-min fairness bandwidth allocation.

Assume that for a given price vector $(u_1, \ldots, u_R) \in \mathcal{B}$, there exists a user r who does not have a bottleneck link with respect to $(u_1, \ldots, u_R)$. This means that for every link $l \in \{l_{r,1}, \ldots, l_{r,K_r}\}$ along the route of user r we have

$$\sum_{r' \in H_l} D_{r'}(u_r) < C_l.$$

Proposition 3 then implies that $(u_1, \ldots, u_R)$ does not correspond to an equilibrium allocation as

$$u_r > \inf\left\{u \in \mathcal{I} : P_{tr,r}(u, \underline{F}^*) = 1\right\}.$$

Moreover, when every user r has a bottleneck link with respect to $(u_1, \ldots, u_R)$, then $(u_1, \ldots, u_R)$ leads to an equilibrium allocation. Indeed, in this case, one can show that there exists for every link $l \in \mathcal{L}$ a price $\hat{u}_l$ such that

$$P_l(u) = \begin{cases} 1, & u \geq \hat{u}_l \\ 0, & \text{otherwise.} \end{cases}$$

Setting

$$\hat{u}_r = \min_{k=1,\ldots,K_r} \hat{u}_{l_{r,k}}, \qquad r = 1, \ldots, R$$

it then follows that for every user $r = 1, \ldots, R$ we have

$$P_{tr,r}(u) = \begin{cases} 1, & u \geq \hat{u}_r \\ 0, & \text{otherwise.} \end{cases}$$

Furthermore, when $D_r(u_r)$ is positive, then one can show that u_r is equal to $\hat{u}_r$. When $D_r(u_r)$ is equal to 0, then one can show that u_r is equal to $u_{r,\max}$, as well as less than or equal to $\hat{u}_r$. Proposition 3 then implies that $(u_1, \ldots, u_R)$ corresponds to an equilibrium allocation. ■

Proposition 5: Under Assumption 1, there exists a unique price vector $(u_1^*, \ldots, u_R^*) \in \mathcal{B}$ such that the corresponding allocation $\underline{F}^*$ is an equilibrium allocation. Furthermore, when the price vector $(u_1^*, \ldots, u_R^*) \in \mathcal{B}$ leads to an equilibrium allocation, then it is fair, i.e., it is not possible to reduce the price u_r^* of a user r without losing feasibility or increasing the price of another user r' with a price $u_{r'}^* \geq u_r^*$.

We provide a proof for Proposition 5 in Appendix C. Proposition 5 states that the price-based priority service is fair as given by Definition 2. Using Lemma 1, this implies that the equilibrium allocation is weighted max-min fair with weights

$$B_r(x_r) = -G_r(x_r), \quad x_r \geq 0.$$

This establishes our main result. The following corollary follows immediately from Proposition 3 and 5.

Corollary 2: Under Assumption 1, there exists an unique equilibrium allocation.

VI. Conclusion

We have analyzed a price-based priority service for which users are charged according to the priority of their traffic. For the case where the network supports a continuum of priorities, we showed that this scheme leads to a weighted max-min fair allocation. Regarding future work, it would be interesting to analyze the situation where the network supports a finite number, rather than a continuum. This case seems to be more involved than the idealized priority service that we consider in this paper; we are currently exploring this extension. Furthermore, we use a simplistic link model for our analysis; it would be interesting extend our analysis which uses discrete-time link model (based on time slots) without buffering, to a continuous-time model with traffic queueing. In addition, it would be interesting to explore whether our assumption that traffic is dropped independently is (asymptotically) true for continuous-time link models with buffering.

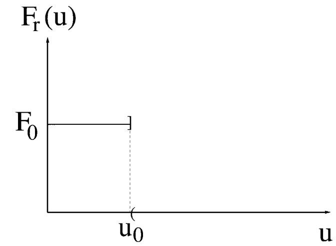

Fig. 3. Step function F_R is an element in $\mathcal{S}$.

Appendix A
Proof of Proposition 1

We use the following notation. Let $\mathcal{S}$ be the set of all step functions $F \in \mathcal{F}$ (see Fig. 3), i.e., $\mathcal{S}$ is the set of all functions $F \in \mathcal{F}$ which have the property that there exist scalars $F_0 \in \Re_+$ and $u_0 \in \mathcal{I}$ such that

$$F(u) = \begin{cases} F_0, & u \leq u_0 \\ 0, & \text{otherwise.} \end{cases}$$

For every user $r = 1, \ldots, R$ let $\mathcal{S}_r$ be the set of all functions $F \in \mathcal{S}$ which have the property that there exists a price $u_0 \in \mathcal{I}$ such that

$$F(u) = \begin{cases} D_r(u_0), & u \leq u_0 \\ 0, & \text{otherwise.} \end{cases}$$

To prove Proposition 1, we first derive a few preliminary lemmas.

Fix a user r and consider a given aggregate allocation $F \in \mathcal{F}$. We then have the following results.

Lemma 2: For every function $F_r \in \mathcal{F}$ there exists a function $F_r^* \in \mathcal{S}$ such that

$$U_r(x_r^*) - \int_0^\infty u\,(-dF_r^*(u)) \geq U_r(x_r) - \int_{\mathcal{I}} u\,(-dF_r(u))$$

where

$$x_r^* = \int_{\mathcal{I}} P_{tr}(u, F)\,(-dF_r^*(u))$$

and

$$x_r = \int_{\mathcal{I}} P_{tr}(u, F)\,(-dF_r(u)).$$

The inequality is strict when $F_r \notin \mathcal{S}$.

Proof: Let the price u^* be given by

$$u^* = \liminf\,\{u \in \mathcal{I} : P_{tr}(u, F) = 1\}$$

and consider the following two cases.

First, assume that $P_{tr}(u^*, F) = 1$ and consider the allocation given by

$$F_r^*(u) = \begin{cases} F_r(u^*), & u \leq u^* \\ 0, & \text{otherwise.} \end{cases}$$

Then we have

$$\begin{aligned} U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) &= \ldots \\ &= U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) + \ldots \\ &\quad + \int_{u^*}^\infty (u-u^*)\left(-dF_r(u)\right) + \int_0^{u^*} u\left(-dF_r(u)\right) \\ &\geq U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) \end{aligned}$$

where the last inequality is strict when $F \notin \mathcal{S}$.

Next, suppose that $P_{tr}(u^*, F) < 1$. By definition, we have that $F_r(u^*) < \infty$ and we can choose a price $u_0 > u^*$ such that

$$U_r\left(F_r(u^*)\right) - U_r(x_r) > \left[F_r(u^*) - F_r(u_0)\right]\left[u_0 - u^*\right].$$

Using u_0, we consider the allocation given by

$$F_r^*(u) = \begin{cases} F_r(u^*), & u \leq u_0 \\ 0, & \text{otherwise.} \end{cases}$$

In this case, we have

$$\begin{aligned} U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) &\geq U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) + \ldots \\ &\quad + \int_0^{u^*} u\left(-dF_r(u)\right) + \ldots \\ &\quad + U_r\left(F_r(u^*)\right) - U_r(x_r) - \ldots \\ &\quad - \left[F_r(u^*) - F_r(u_0)\right]\left[u_0 - u^*\right] \\ &> U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right). \end{aligned}$$

■

Lemma 3: For every function $F_r \in \mathcal{S}$ there exists a function $F_r^* \in \mathcal{S}_r$ such that

$$U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) \geq U_r(x_r) - \int_{\mathcal{I}} u\left(-dF_r(u)\right)$$

where

$$x_r^* = \int_{\mathcal{I}} P_{tr}(u, F)\left(-dF_r^*(u)\right)$$

and

$$x_r = \int_{\mathcal{I}} P_{tr}(u, F)\left(-dF_r(u)\right).$$

The inequality is strict if $F_r \notin \mathcal{S}_r$.

Proof: Let u_r be such that

$$F_r(u) = \begin{cases} x_r, & u \leq u_r \\ 0, & \text{otherwise.} \end{cases}$$

First, we assume that $P_{tr}(u_r, F) = 0$ and let u^* be such that $P_{tr}(u^*, F) = 1$. Using u^*, we consider the allocation

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \leq u^* \\ 0, & \text{otherwise} \end{cases}$$

and let $x_r^* = D_r(u^*)$. Note that x_r^* is the optimal solution to the maximization problem

$$\max_{x \in \mathcal{I}} \left\{U_r(x) - xu^*\right\}$$

and we have that

$$U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) \leq 0$$

and

$$U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) \geq 0.$$

If $D_r(u^*) > 0$, then we have that

$$U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) > 0$$

and the inequality is strict. When $F_r(u_r) > 0$, then we have that

$$U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) < 0$$

and the inequality is strict. Finally, when $D_r(u^*) = F_r(u_r) = 0$, then we have that $F_r = F_r^*$.

Next, we assume that $P_{tr}(u_r, F) = 1$ and let u^* be given by

$$u^* = \liminf\left\{u \in \mathcal{I} : P_{tr}(u, F) = 1\right\}.$$

Consider the allocation

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \leq u^* \\ 0, & \text{otherwise.} \end{cases}$$

Then we have

$$\begin{aligned} U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) &= \ldots \\ &= U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) + \ldots \\ &\quad + U_r\left(x_r^*\right) - U_r(x_r) - u^*\left(x_r^* - x_r\right) + \ldots \\ &\quad + (u_r - u^*)x_r \\ &\geq U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) \end{aligned}$$

where the last inequality is strict when $F_r \notin \mathcal{S}_r$.

Finally, when we have that $0 < P_{tr}(u_r, F) < 1$, then we can choose a price $u^* > u_r$ such that

$$U_r\left(F_r(u^*)\right) - U_r(x_r) > \left[F_r(u_r) - F_r(u^*)\right]\left[u^* - u_r\right]$$

where

$$x_r = P_{tr}(u_r, F)F_r(u_r).$$

In this case, consider the following allocation:

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \le u^* \\ 0, & \text{otherwise.} \end{cases}$$

Then we have

$$\begin{aligned} U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) \\ &\ge \dots \\ &\ge U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) + \dots \\ &\quad + U_r\left(F_r(u^*)\right) - U_r(x_r) - \dots \\ &\quad - \left[F_r(u_0) - F_r(u^*)\right]\left[u^* - u_r\right] + \dots \\ &\quad + U_r\left(x_r^*\right) - U_r\left(F_r(u^*)\right) - \left(x_r^* - F_r(u^*)\right)u^* \\ &> U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right). \end{aligned}$$

■

Lemma 4: Let $u^* = \liminf\{u \in \mathcal{I} : P_{tr}(u, F) = 1\}$. Then there exists optimal allocation $F_r^* \in \mathcal{F}$ for the maximization problem

$$\max_{F_r \in \mathcal{F}} \left\{ U_r(x_r) - \int_{\mathcal{I}} u\left(-dF_r(u)\right) \right\}$$

where

$$x_r = \int_{\mathcal{I}} P_{tr}(u, F)\left(-dF_r(u)\right)$$

if and only if $P_{tr}(u^*, F) = 1$. In this case, the unique optimal allocation is given as

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \le u^* \\ 0, & \text{otherwise.} \end{cases}$$

Proof: By Lemma 2 and 3, it suffices to consider allocations $F_r \in \mathcal{S}_r$, such that

$$F_r(u) = \begin{cases} D_r(u_0), & u \le u_0 \\ 0, & \text{otherwise.} \end{cases}$$

for some $u_0 \in \mathcal{I}$. We can then identify an allocation $F_r \in \mathcal{S}_r$ with the threshold price u_0.

When $u_0 < u^*$, then we have

$$U_r(x_r) - \int_0^\infty u\left(-dF_r(u)\right) \le 0.$$

For every u' such that $P_{tr}(u', F) = 1$, we have for the allocation

$$F_r'(u) = \begin{cases} D_r(u'), & u \le u' \\ 0, & \text{otherwise,} \end{cases}$$

that

$$U_r\left(D_r(u')\right) - \int_0^\infty u\left(-dF_r'(u)\right) \ge 0.$$

Therefore, without loss of generality, it suffices to consider functions $F_r \in \mathcal{S}_r$ such that $u_0 \ge u^*$.

First, suppose that $P_{tr}(u^*, F) = 1$. Let $F_r \in \mathcal{S}_r$ be such that $u_0 > u^*$ and let $F_r^*(u)$ be as given in the statement of the lemma. Then, we have that

$$U_r\left(D_r(u^*)\right) - \int_0^\infty u\left(-dF_r^*(u)\right) > \dots$$
$$> U_r\left(D_r(u_0)\right) - \int_0^\infty u\left(-dF_r(u)\right).$$

This implies that when $P_{tr}(u^*, F) = 1$, then F_r^* is the unique optimal allocation.

Next, suppose that $P_{tr}(u^*, F) < 1$. Let $F_r^*(u)$ be as given in the statement of the lemma, and let $F_r \in \mathcal{S}_r$ be such that $u_0 > u^*$. In particular, choose u_0 such that

$$U_r\left(D_r(u_0)\right) - U_r\left(x_r^*\right) > \left[D_r(u^*) - D_r(u_0)\right]\left[u_0 - u^*\right]$$

where

$$x_r^* = P_{tr}(u^*, F)D_r(u^*).$$

This implies that

$$\begin{aligned} U_r\left(D_r(u_0)\right) - \int_0^\infty u\left(-dF_r(u)\right) \\ &= \dots \\ &= U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right) + \dots \\ &\quad + U_r\left(D_r(u_0)\right) - U_r\left(x_r^*\right) - \dots \\ &\quad - \left[D_r(u^*) - D_r(u_0)\right]\left[u_0 - u^*\right] \\ &> U_r\left(x_r^*\right) - \int_0^\infty u\left(-dF_r^*(u)\right). \end{aligned}$$

Furthermore, note that for every u' such that $u^* < u' < u_0$, we have that

$$U_r\left(D_r(u\prime)\right) - \int_0^\infty u\left(-dF_r'(u)\right) > \dots$$
$$> U_r\left(D_r(u_0)\right) - \int_0^\infty u\left(-dF_r(u)\right)$$

where

$$F_r'(u) = \begin{cases} D_r(u'), & u \le u' \\ 0, & \text{otherwise.} \end{cases}$$

Combining these results, it follows that the maximization problem in the statement of the lemma does not have a solution when $P_{tr}(u^*, F) < 1$. ■

We are now ready to prove Proposition 1.

Proof: By Lemma 4, there exists an equilibrium allocation if and only if there exists a price $u^* \in \mathcal{I}$ such that

$$\sum_{r=1}^{R} D_r(u^*) = C$$

and for every user $r = 1, \ldots, R$, we have

$$F_r^*(u) = \begin{cases} D_r(u^*), & u \le u^* \\ 0, & \text{otherwise.} \end{cases}$$

As by Assumption 1, we have that

$$\sum_{r=1}^{R} D_r(0) > C$$

and it follows that there exists a price $u^* \in \mathcal{I}$ such that $\sum_{r=1}^{R} D_r(u^*) = C$. Furthermore, by Assumption 1 the function $D_r(u)$ is strictly decreasing, and it follows that there exists a unique price $u^* \in \mathcal{I}$ such that $\sum_{r=1}^{R} D_r(u^*) = C$.

Combining these results, we obtain that there exists a unique equilibrium allocation. ■

Appendix B
Proof of Proposition 2

We first derive a preliminary lemma where we show that Proposition 2 holds for the case where the network offers a single priority class.

Assume that the network supports only a single priority class. Let C_l be the link capacities. Furthermore, let d_r be the amount of traffic submitted by user r in this priority class and let α_l be the transmission probability associated with the link $l \in \mathcal{L}$, i.e., let the transmission probabilities α_l, $l \in \mathcal{L}$, be such that

$$F_{r,l_1} = d_r \tag{8}$$

$$F_{r,l_k} = \alpha_{l_{k-1}} F_{r,l_{k-1}}, \qquad k = 2, \ldots, K_r \tag{9}$$

$$\alpha_l = \begin{cases} 0, & \text{if } C_l = 0 \\ \min\left\{ \frac{C_l}{\sum_{r \in A_l} F_{r,l}}, 1 \right\}, & \text{otherwise.} \end{cases} \tag{10}$$

We have the following lemma.

Lemma 5: There exists a solution $(\alpha_1, \ldots, \alpha_L)$ to the system of equations given by (8)–(10).

Proof: We define the mapping $g : [0,1]^L \rightarrow [0,1]^L$ as follows:

$$g_l(\alpha) = \begin{cases} 0, & \text{if } C_l = 0 \\ \min\left\{ \frac{C_l}{\sum_{r \in A_l} F_{r,l}}, 1 \right\}, & \text{otherwise.} \end{cases}$$

where

$$F_{r,l_1} = d_r$$
$$F_{r,l_k} = \alpha_{l_{k-1}} F_{r,l_{k-1}}, \qquad k = 2, \ldots, K_r,$$

Note that the mapping $g : [0,1]^L \rightarrow [0,1]^L$ is continuous, and the set $[0,1]^L$ is convex and compact. By Brouwer's fixed-point theorem [29], it follows that there exists a vector $\alpha^* = (\alpha_1^*, \ldots, \alpha_L^*) \in [0,1]^L$ such that

$$\alpha_l^* = g(\alpha^*), \qquad l = 1, \ldots, L.$$

This implies that there exists a solution $(\alpha_1, \ldots, \alpha_L)$ to the system of equations given by (8)–(10). ■

We can use Lemma 5 to iteratively construct a set of transmission probabilities that solve the system of equations given by (5)–(7). Intuitively, the procedure can be interpreted as "water filling." We start with a high price for accessing the network such that the total demand is small and no link is saturated. We then lower the price (increase the demand) until the first link is saturated. Assume that u^1 is the price that saturates the first link. We then use Lemma 5 to compute the transmission probabilities $P_l(u, \underline{F})$, $l \in \mathcal{L}$, which satisfy (5)–(7) for $u \ge u^1$. We then proceed to lower the price until we saturate the next link at the price u^2 and use Lemma 5 to compute the transmission probabilities $P_l(u, \underline{F})$ for u such that $u^2 \le u < u^1$, and so on.

Note that the transmission probability of priority u_0 traffic depends only on the traffic that users submit in priorities $u \ge u_0$, and not on the traffic in priorities lower than u. Similarly, the demand $F_{r,l}(u_0, \underline{F})$ of user $r \in H_l$ on link $l \in \mathcal{L}$ depends only on the transmission probabilities for priorities $u \ge u_0$ (on previous links along the route of user r), but not on the transmission probabilities for priorities $u < u_0$. We use this property in the following proof for Proposition 2.

Proof: Let A^k be the set of links not saturated at the beginning of iteration k, let B^k be the set of links that get saturated for the first time at iteration k, and let u^k be the (highest) price at which these links get saturated. We then iteratively construct transmission probabilities $P_l(u, \underline{F})$, $l \in \mathcal{L}$, that satisfy (5)–(7) as follows.

Initial conditions: $k = 1$; $A^1 = \mathcal{L}$; $u^0 = \infty$; $u_r^0 = \infty$ for $r = 1, \ldots, R$; and $P_l^0(u, \underline{F}) = 1$ for $l = 1, \ldots, L$ and $u \in \mathcal{I}$.

Then do the following steps.

1) For $r = 1, \ldots, R$, let the function $F_{r,l_j} : \mathcal{I} \mapsto \Re_+$ be given by

$$F_{r,l_j}(u) = \int_u^\infty P_{l_{j-1}}^{k-1}(u)\left(-dF_{r,l_{j-1}}(u)\right), \quad j = 2, \ldots, K_r.$$

2) Set $u^k = \max_{l \in A^k} u_l^k$, where

$$u_l^k = \inf\left\{ u \in \mathcal{I} : \sum_{r \in H_l} F_{r,l}(u) < C_l \right\}.$$

3) Set

$$\tilde{C}_l^k = \begin{cases} C_l - \sum_{r \in D_l} F_{r,l}^+(u^k), & \text{if } l \in A^k \\ 0, & \text{otherwise} \end{cases}$$

$$d_r^k = F_r(u^k) - F_r^+(u^k)$$

and choose transmission probabilities α_l^k, $l \in \mathcal{L}$, which solve

$$d_{r,l_1}^k = d_r^k$$
$$d_{r,l_j}^k = \alpha_{l_{j-1}}^k d_{r,l_{j-1}}^k, \quad j = 2, \ldots, K_r$$
$$\alpha_l^k = \begin{cases} 0, & \text{if } \tilde{C}_l^k = 0 \\ \min\left\{ \frac{\tilde{C}_l^k}{\sum_{r \in H_l} d_{r,l}^k}, 1 \right\}, & \text{otherwise.} \end{cases}$$

4) Set $B^k = \{l \in A^k : \sum_{r \in H_l} d_{r,l}^k \ge \tilde{C}_l^k\}$ where

$$d_{r,l_1}^k = d_r^k$$
$$d_{r,l_j}^k = \alpha_{l_{j-1}} d_{r,l_{j-1}}^k, \quad j = 2, \ldots, K_r.$$

5) When $l \in B^k$, then set

$$P_l^k(u,\underline{F}) = \begin{cases} 0 & u < u^k \\ \alpha_l^k & u = u^k \\ 1 & u > u^k \end{cases}$$

otherwise set $P_{tr,l}^k(u,\underline{F}) = P_{tr,l}^{k-1}(u,\underline{F}),\ u \in \mathcal{I}$.

6) Set $A^{k+1} = A^k \backslash B^k$.

7) If A^{k+1} is empty, or $u^k = 0$, then stop; otherwise set $k = k+1$ and go to 1.

To prove that this procedure indeed provides a solution for the system of equations given by (5)–(7), we proceed in two steps as follows.

a) Let u^k be the price that we pick at Step 2. We show that $u^k < u^{k-1}$.

b) Using this result, we then show that the transmission probabilities that we obtain when the procedure terminates is indeed a solution for the systems of equation given by (5)–(7).

Note that by Lemma 5 we can always carry out Step 3.

We use the following notation. Let $P_l(u,\underline{F})$, $u \in \mathcal{I}$, be the transmission probabilities on link l that we obtain when the procedure terminates. Furthermore, let $F_{r,l}^k(u,\underline{F})$ be the cumulative demand of user r on link l_j on the jth link along its route at iteration k, i.e., we have (see Step 1)

$$F_{r,l_j}^k(u) = \int_u^\infty P_{l_{j-1}}^{k-1}(u)\left(-dF_{r,l_{j-1}}(u)\right), \qquad j = 2,\ldots,K_r$$

and let

$$F_{r,l_j}(u) = \int_u^\infty P_{l_{j-1}}(u)\left(-dF_{r,l_{j-1}}(u)\right), \qquad j = 2,\ldots,K_r$$

be this demand under the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, when the procedure terminates. We make the following observations. By the definition, we have for all users $r = 1,\ldots,R$ that

$$\lim_{u\to\infty} F_{r,l_j}^k(u) = 0, \qquad j = 1,\ldots,K_r.$$

This implies that the set

$$\inf\left\{u \in \mathcal{I} : \sum_{r\in H_l} F_{r,l}^k(u) < C_l\}\right\}$$

is not empty for all links $l \in A^k$, and we can choose a price u^k at Step 2 at every iteration $k = 1,2,\ldots$. By construction, we have that

$$F_{r,l}^k(u) \le F_{r,l}^{k-1}(u), \qquad u \in \mathcal{I}$$

for all users r and for all links l. This implies that for all links $l \in A^k$ we have

$$\sum_{r\in H_l} F_{r,l}^k(u_{k-1}) < C_l.$$

As $u^{k-1} > 0$ (otherwise, the procedure would have already terminated) and $F_{r,l}^k(u)$ is continuous for the left, it follows that $u^k < u^{k-1}$. This implies that the procedure will terminate in a finite number of steps (bounded by L) as at each iteration we saturate at least one new link.

To show the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, that provide a solution to the system of equations given by (5)–(7), we use an induction argument.

By Lemma 5, the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the system of equations given by (5)–(7) for $u \ge u^1$. Now suppose that the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, obtained by the above construction provide a solution for $u \ge u^{k-1}$, where u^{k-1} is the price picked in Step 2 at iteration $k-1$. We wish to show that these probabilities then also provide a solution for $u \ge u^k$. By construction, the following properties hold. For all links $l \in A^k$, we have that

$$P_l(u,\underline{F}) = 1, \qquad u \in (u^{k-1}, u^k)$$

and for all links $l \notin A^k$ we have

$$P_l(u,\underline{F}) = 0, \qquad u \in (u^{k-1}, u^k).$$

Furthermore, for all links $l \in A^k$ we have

$$\sum_{r\in H_l} F_{r,l}(u) = \sum F_{r,l}^k(u) < C_l, \qquad u \in (u^{k-1}, u^k)$$

and for all links $l \notin A^k$ we have

$$\sum_{r\in H_l} F_{r,l}(u) = \sum F_{r,l}^k(u) \ge C_l, \qquad u \in (u^{k-1}, u^k).$$

Combining these results, we obtain that the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, provide a solution for $u > u^k$. By the procedures given in Step 3–5, we have

$$P_l(u^k,\underline{F}) = P_{tr,l}^k(u^k,\underline{F}) = \alpha_l^k$$

and it follows that the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the system of equations given by (5)–(7) for $u \ge u^k$. Let k_0 be total number of iterations needed in the above procedure. Then the above induction argument implies that $P_l(u,\underline{F})$, $l \in \mathcal{L}$, provide a solution to the system of equations given by (5)–(7) for $u \ge u^{k_0}$. When $u^{k_0} = 0$, then we are done. When $u^{k_0} > 0$, then all links are saturated under the price u^{k_0}. Furthermore, in this case we have by construction that the link transmission probabilities for $u \in [0, u^{k_0})$ are given by

$$P_l(u,\underline{F}) = 0, \quad l \in \mathcal{L}.$$

This implies that transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the systems of equations given by (5)–(7) for $[0, u^{k_0})$, which completes the proof. ■

APPENDIX C
PROOF OF PROPOSITION 5

To prove Proposition 5, we again adapt an argument used in [24] to prove the equivalent result for max-min fairness.

Let A^k be the set of links not saturated at the beginning of iteration k, let B^k be the set of links that get saturated for the first time at iteration k, and let u^k be the price for which the links in B^k get saturated. Furthermore, let P^k denote the set of users not passing through any saturated link at the beginning of iteration k.

Initial conditions: $k=1$; $P^1=\{1,\dots,R\}$; $A^1=\mathcal{L}$; $u^0=\infty$; and $u_r^0=\infty$.

1) Set $u^k=\max_{l\in A^k} p_l^k$, where

$$p_l^k=\min\left\{u\in\mathcal{I}:\sum_{r\in H_l:r\notin P^k} D_r(u)\le C_l-\sum_{r\in H_l} D_r\left(u_r^k\right)\right\}.$$

2) Set

$$B^k=\left\{l\in A^k:\sum_{r\in H_l:r\notin P^k} D_r(u^k)=C_l-\sum_{r\in H_l} D_r\left(u_r^k\right)\right\}.$$

3) When B^k is empty, then set $Q^k=P^k$; otherwise, set

$$Q^k=\{r\in P^k: r \text{ passes through a link } l\in B^k\}.$$

4) Set $u_r^{k+1}=\begin{cases}\min\{u_{r,\max},u^k\}, & r\in Q^k\\ u_r^k, & \text{otherwise.}\end{cases}$

5) Set $A^{k+1}=A^k/B^k$.

6) Set $P^{k+1}=P^k/Q^k$.

7) Set $k=k+1$.

8) If P^k is empty, then stop; otherwise, go to 1.

Let k_0 be the iteration when the procedure terminates and let $u_r=u_r^{k_0}$. Note that at each iteration, we decrease the price (and increase the demand) until a new link gets saturated (at price u^k). When the procedure terminates, each user r either passes through a bottleneck link or we have $u_r^{k_0}=0$, and Proposition 3 implies that $(u_1^{k_0},\dots,u_R^{k_0})$ corresponds to an equilibrium allocation.

In fact, $(u_1^{k_0},\dots,u_R^{k_0})$ unique price vector in $\mathcal{B}$ such that every user has a bottleneck link. This can be seen as follows. Assume that $(u_1,\dots,u_R)\in\mathcal{B}$ is an allocation such that each user has a bottleneck link with respect to $(u_1,\dots,u_R)$. Then for every user $r\in Q^1$, we have that

$$u_r=\min\{u_{r,\max},u^1\}.$$

Indeed, if this is not the case, then there exist users r, $r'\in Q^1$, such that $u_r<u_{r'}$ and $u^1<u_{r'}$; otherwise, the vector $(d_1,\dots,d_R)$ would not be a feasible price vector. As for every link we have that

$$\sum_{r\in D_l} D_r(u)<C_l,\quad \text{for } u>u^1$$

it follows that user f' does not have a bottleneck link with respect to $(u_1,\dots,u_R)$. Let $\bar{k}<k_0$ be such that for all $k=1,\dots,\bar{k}$, we have that

$$u_r=\min\{u_{r,\max},u^k\},\quad \text{for all } r\in Q^k.$$

It then follows that for all $r\in Q^{\bar{k}+1}$, we have that

$$u_r=\min\{u_{r.,ax},u^{\bar{k}+1}\}.$$

Indeed, if this is not the case, by the same argument given above, there exists a user $r'\in Q^{\bar{k}+1}$, which does not have a bottleneck link with respect to $(u_1,\dots,u_R)$. Therefore, by induction, we have that $(u_1,\dots,u_R)=(u_1^{k_0},\dots,u_R^{k_0})$. This implies that there exists a unique price vector $(u_1,\dots,u_R)$ such that each user has a bottleneck link with respect to $(u_1,\dots,u_R)$.

Next, we show that a price vector $(u_1,\dots,u_R)\in\mathcal{B}$ has the property that each user r either has a bottleneck link with respect to $(u_1,\dots,u_R)$ or $u_r=0$, if and only if $(u_1,\dots,u_R)$ is fair (as given by Definition 2).

Suppose that $(u_1,\dots,u_R)\in\mathcal{B}$ is fair and, to arrive at a contradiction, assume that there exists a user r with $u_r>0$ that does not have a bottleneck link. Then, for each link l crossed by user r with $\sum D_r(u_r)=C_l$, there must exist a user $r'\neq r$ such that $u_{r'}<u_r$. This implies for each link l on the route of user r, the quantity

$$\delta_l=\begin{cases}C_l-\sum D_r(u_r), & \text{if } \sum D_r(u_r)<C_l\\ u_{r'}-u_r, & \text{if } \sum D_r(u_r)=C_l\end{cases}$$

is positive. Therefore, there exist scalars δ_r, $\delta_{r'}>0$, such that we can decrease the price of user r by setting $u_r'=u_r-\delta_r$ and increase the prices by all users r' with $u_r'<u_r$ by setting $u_{r'}'=u_{r'}'+\delta_{r'}$, without losing feasibility. This contradicts the fairness property of $(u_1,\dots,u_R)$.

Conversely, assume that each user r either has a bottleneck link with respect to a price vector $(u_1,\dots,u_R)\in\mathcal{B}$. Then, to decrease the price of a user r, we must increase the price of some user r' crossing the bottleneck link of user r in order to maintain feasibility. Since $u_r\le u_{r'}$ for all users r' crossing the bottleneck link of user r. This implies that the price vector $(u_1,\dots,u_R)$ is fair.

References

[1] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured Forwarding PHB Group," IETF, RFC 2597, 1999.

[2] R. Gibbons, *Game Theory for Applied Economists*. Princeton, NJ: Princeton Univ. Press, 1992.

[3] A. Ganesh, K. Laevans, and R. Steinberg, "Congestion pricing and user adaptation," in *Proc. IEEE INFOCOM*, 2001, pp. 959–965.

[4] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control for communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237–252, 1998.

[5] H. Varian, *Microeconom. Anal.*. New York: Norton, 1992.

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

[7] J. Wardrop, "Some theoretical aspects of road traffic research," *Proc. Inst. Civil Eng. II*, vol. 1, pp. 325–378, 1952.

[8] P. Marbach, "Pricing priority classes in differentiated services networks," in *37th Annu. Allerton Conf. Communication, Control, and Computing*, Monticello, IL, Sept. 1999.

[9] R. J. La and V. Anantharam, "Utility-based rate control in the internet for elastic traffic," *IEEE/ACM Trans. Networking*, vol. 10, pp. 272–286, Apr. 2002.

[10] S. H. Low and D. E. Lapsley, "Optimization flow control, part I: Basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, pp. 861–875, Dec. 1999.

[11] J. Mo and J. Walrand, "Fair end-to-end window-based congestion control," *IEEE/ACM Trans. Networking*, vol. 8, pp. 556–567, Oct. 2000.

[12] H. Yaiche, R. Mazumdar, and C. Rosenberg, "A game-theoretic framework for rate allocation and charging of elastic connections in broadband networks," *IEEE/ACM Trans. Networking*, vol. 8, pp. 667–678, Oct. 2000.

[13] A. Demers, S. Keshav, and S. Shenker, "Analysis and simulations of a fair queueing algorithm," in *Proc. ACM SIGCOMM*, 1989, pp. 1–12.

[14] E. L. Hahne, "Round-robin scheduling for max-min fairness in data networks," *IEEE J. Select. Areas Commun.*, vol. 9, pp. 1024–1039, July 1991.

[15] M. Katevenis, "Fast switching and fair control of congested flow in broad-band networks," *IEEE J. Select. Areas Commun.*, vol. 5, pp. 1315–1326, Aug. 1987.

[16] A. Charny, “An algorithm for rate allocation in a packet-switching network with feedback,” Master’s thesis, Massachusetts Inst. Technol., Cambridge, 1994.
[17] S. Chen and K. Park, “An architecture for noncooperative QoS provision in many-switch systems,” in *Proc. IEEE INFOCOM*, 1999, pp. 864–872.
[18] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, “Pricing in computer networks: Motivation, formulation, and example,” *IEEE/ACM Trans. Networking*, vol. 1, pp. 614–627, Dec. 1993.
[19] A. M. Odlyzko, “Paris metro pricing: The minimalist differentiated services solution,” in *Proc. 7th Int. Workshop Quality of Service*, 1999, pp. 159–161.
[20] ——, “Paris metro pricing for the Internet,” in *Proc. ACM Conf. Electronic Commerce*, 1999, pp. 140–147.
[21] K. Park, M. Sitharam, and S. Chen, “Quality of service provision in noncooperative networks: Heterogeneous preferences, multi-dimensional QoS vectors, and burstiness,” in *Proc. Int. Conf. Information and Computation Economies*, 1998, p. .
[22] J. K. MacKie-Mason and H. R. Varian, “Pricing the internet,” in *Public Access to the Internet*, B. Kahin and J. Keller, Eds. Englewood Cliffs, NJ: Prentice-Hall, 1994.
[23] W. Rudin, *Principles of Mathematical Analysis*. New York: McGraw-Hill, 1976.
[24] D. P. Bertsekas and R. Gallager, *Data Networks*. Englewood Cliffs, NJ: Prentice-Hall, 1992.
[25] “Traffic Management Specification, Version 4.0,” ATM Forum Traffic Management Working Group, Tech. Rep. AF-TM-0056.000, Apr. 1996.
[26] X. Huang and B. Bensaou, “On max-min fairness and scheduling in wireless ad-hoc networks: Analytical framework and implementation,” in *Proc. 2nd ACM Int. Symp. Mobile Ad-Hoc Networking and Computing*, 2001, pp. 221–231.
[27] L. Tassiulas and S. Sarkar, “Maxmin scheduling in wireless networks,” in *Proc. IEEE INFOCOM*, 2002, pp. 763–772.
[28] D. Rubenstein, J. Kurose, and D. Towsley, “The impact of multicast layering on network fairness,” *IEEE/ACM Trans. Networking*, vol. 10, pp. 169–182, Apr. 2002.
[29] D. Smart, *Fixed Point Theorems*. London, U.K.: Cambridge Press, 1974.

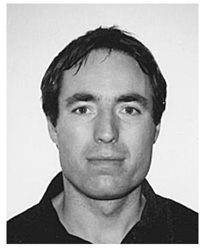

Peter Marbach was born in Lucerne, Switzerland. He received the Eidg. Dipl. El.-Ing. degree from the ETH Zurich, Switzerland, in 1993, the M.S. degree in electrical engineering from Columbia University, New York, NY, in 1994, and the Ph.D. degree in electrical engineering from the Massachusetts Institute of Technology, Cambridge, in 1998.

He joined the Department of Computer Science, University of Toronto, Toronto, ON, Canada, as an Assistant Professor in 2000. He was previously a Post-Doctoral Research Fellow at the Center for Communication Systems Research, University of Cambridge, Cambridge, U.K., and a Visiting Scientist at the Siemens Corporate Research Center, Munich, Germany. His research interests are in the fields of communication networks, stochastic systems, optimization, and control.

Dr. Marbach received the IEEE INFOCOM 2002 Best Paper Award for his paper "Priority Service and Max-Min Fairness."