Priority Service and Max-Min Fairness

Peter Marbach

Abstract— **We study a pricing scheme for networks which use priorities to provide differentiated quality of service. We consider the situation where users are free to choose the priority of their traffic, but are charged accordingly. We model this situation as a non-cooperative game, where users behave in a selfish manner and choose an allocation of priorities to packets to optimize their own net benefit. We show that there exists an unique equilibrium for this game and the bandwidth allocation in equilibrium is weighted max-min fair.**

I. Introduction

We study the fairness properties of a price-based resource allocation scheme for data networks. More precisely, we consider a packet-based network that uses a priority service to provide Quality-of-Service (QoS). We assume that users are free to choose the priority of their traffic, but are charged accordingly by the network. We study how the network resources are shared among users under this priority service.

It has been proposed to use a priority service to provide QoS in the Internet (see Diffserv proposal of the IETF [3] and LAN protocol of the IEEE [11]): network traffic is classified into different priorities and packets are served at each node according to a priority rule. This mechanism allows to provide differentiated QoS where high priority traffic will experience less delay and is less likely to be dropped. Clearly, a network offering this service model has to deploy a pricing scheme where user are charged based on the priority of their traffic; otherwise all users would submit their traffic in the highest priority class. We consider a pricing scheme where the network announces to users a menu of static prices associated with the different priority classes. Users are free to choose the priority of their traffic, but are billed by the network accordingly. In this context, pricing becomes effectively a resource allocation mechanism: one would expect that users assign traffic to priorities in order to meet their QoS requirements at the lowest possible cost, and therefore use high priorities only for traffic with strict QoS requirements. In this paper, we study this price-based resource allocation scheme and investigate how network resources are shared among users.

We adopt an economic framework for the analysis of this pricing scheme. We characterize users' service requirements by a utility function which depends on the throughput. Roughly, submitting more high priority packets will increase the utility of an individual user. However, doing so will also increase the cost the user has to pay to the network provider. As a result, users face a trade-off between achieving a high utility and keeping their cost low. It is then natural to assume that users want to choose priorities in order to maximize their net benefit, which is defined as their utility minus cost. We model the situation where all users simultaneously optimize their own net benefit as a (non-cooperative) game. We determine (a) whether there exists an equilibrium for this game and (b) how network resources are shared among users in equilibrium (bandwidth sharing property). The classical bandwidth sharing principle for data networks is the max-min fairness criteria as discussed, for instance, by Bertsekas and Gallager [1]. Simply put, a max-min fair allocation gives the most poorly treated customer (i.e. the customer who receives the lowest rate) the largest possible share, while not wasting any network resources. In our main result, we show that there exists a close relationship between the above pricing mechanism for priority services and max-min fairness. For the case where a homogeneous population of users (all users have the same utility function) access the network, we show that above scheme leads to a max-min fair allocation of the network resources. For the general case of an inhomogeneous user population, we show that the resulting allocation is weighted max-min fair, where the weights associated with a user is a function of its utility function.

Price-based resource allocation schemes, and the analysis of their fairness properties, have received much attention in recent years ([4], [5], [6], and [10]). Broadly, this research can be classified into two categories.

(a) Market-Based Approach: In this approach, utility functions are used to characterize the service requirements of an individual user. It is assumed that the utility functions of individual users are unknown to the network, but users know their own utility function. The goal of the network is to allocate resources to maximize an objective function that depends on the users utilities. For example, the goal of the network could be to allocate resources to maximize the total utility over all users (called the social welfare). To achieve this goal, the network uses pricing to obtain (implicit) information about the users utility functions and to allocate resources accordingly. Using this approach, Kelly, Maulloo and Tan propose in [4] a rate-based scheme, where the network sends a price signal to users which reflect the traffic load on the links along their route, and users choose a transmission rate to optimize their net benefit. In [4], it is shows that the resulting scheme maximizes (in equilibrium) the social welfare and achieves a (weighted) proportional fair allocation of the network resources. Using similar ideas as presented in the work in [4], La and Anatharam extend in [5] the work of Mo and Walrand [10] and propose a family of window-based allocation schemes. They show that their scheme can be used to achieve a (weighted) proportional fair allocation or approximate arbitrarily close a max-min fair allocation.

(b) Control-Based Approach: In this approach, the utility functions are assigned to users (applications) by the network to characterize their service requirements. The users' utility func-

Dept. of Computer Science, University of Toronto, 10 King's College Rd., Toronto, Ontario, M5S 1A4, Canada. E-mail: marbach@cs.toronto.edu

tions can then be interpreted as a control variable, which is used by the network to achieve a desired resource allocation goal. Using this approach, Low and Lapsley proposes in [6] a rate-based flow control algorithm that is closely related to [4], where 'prices' are interpreted as control signals to determine the users' transmission rate (in a distributed manner). They show that by suitably choosing the utility functions, the network can achieve various bandwidth sharing goals such as proportional and max-min fairness.

We pursue a market-based approach. In our context, this approach allows a simple operation of the network: all the network has to do is to forward packets according to their priorities.

The rest of the paper is organized as follows. In Section II we define more precisely the problem that we consider and briefly review the results of [9]. In Section III, we provide the definition of a max-min fair rate allocation, and introduce a fairness criteria for a price-based resource allocation schemes. In Section IV, we carry out the analysis of the single link case, which we extend in Section V to the network case.

II. Problem Formulation

Consider a packet-based network which offers a finite set $\mathcal{I} = \{1, ..., N\}$ of priority levels, where 1 is the lowest priority, 2 the second lowest, and so on. Users are free to choose the priority levels of their traffic, but are charged accordingly by the network. The network charges a price u_i per unit traffic submitted with priority i, where for $i, j \in \mathcal{I}$ we have

$$u_i < u_j, \qquad \text{when } i < j.$$

At each node, packets are served according to a strict priority rule for both time (service discipline) and space (buffer management): packets of priority i are transmitted only if all packets with priority $j > i$ have been served (priority in time), and packets of priority i are dropped from a buffer only if there are no packets of priority $j < i$ left in the buffer (priority in space). This means that high priority packets experience less delay and are less likely to be dropped during periods of congestion. A user accessing the network can then decide on

(a) its transmission rate, i.e. the amount of traffic submitted per unit time, and
(b) its traffic allocation, i.e. the priority of the submitted traffic.

Consider a given user denoted by r and assume that user r submits traffic in priority i at a rate $d_r(i)$, $i \in \mathcal{I}$. Furthermore, suppose that $P_{tr,r}(i)$ is the probability that priority i traffic of user r is transmitted through the network (and not lost). As packets are served according to a priority rule, we have that

$$P_{tr,r}(i) \leq P_{tr,r}(j), \qquad \text{when } i < j.$$

Furthermore, suppose that user r perceives QoS only as a function of the average throughput x_r given by

$$x_r = \sum_{i=1}^{N} d_r(i) P_{tr,r}(i).$$

For this case, we can characterize user r by a quality indicator (utility function) which depends only on x_r. This type of user is often called a user with elastic traffic. Examples of users with elastic traffic are users sending email, transferring files, and browsing the Web. Let $U_r(x_r)$ be the utility of user r under the average throughput x_r. Throughout the paper, we make the following assumption.

Assumption 1: The utility function $U_r : \Re_+ \mapsto \Re_+$, is increasing, bounded, strictly concave, and twice differentiable.

We note that utility functions for elastic traffic with these characteristics are commonly used in the pricing literature as in [4]. Assumption 1 implies that a user can achieve a high utility by increasing the amount of high priority traffic; however, doing so will also increase this cost. We model this trade-off by the following maximization problem

$$\max_{d_r \in \Re^N} \left\{ U_r(x_r) - \sum_i d_r(i) u_i \right\},$$

where

$$x_r = \sum_{i=1}^{N} P_{tr,r}(i) d_r(i).$$

This maximization problem has the following interpretation. Given the transmission probabilities $P_{tr,r}(i)$, $i = 1, ..., N$, user r wants to find an allocation d_r which strikes a balance between achieving a high utility (given by $U_r(x_r)$)and keeping the cost low (given by $\sum_i d_r(i) u_i$).

The solution to the above maximization problem will depend on the transmission probabilities $P_{tr,r}(i)$, which themselves depend on the overall traffic load in the network. Therefore, users will have to monitor the current traffic conditions (and estimate the transmission probabilities in the different priority classes), and then update the traffic allocations to optimize their net benefit. In addition, changing d_r will not only affect the net benefit of user r, but also change the transmission probabilities for other users and therefore affect their net benefit functions. This leads to an interaction (coupling) among users: the allocation decision of one user potentially affects the net benefits of all other users.

A. Non-Cooperative Game

We model this interaction among users as a non-cooperative game (for an introduction to the basic concepts of game theory we refer to [2]). We assume that R users access the network and each user behaves in a selfish way, i.e. each user is only interested in optimizing its own net benefit. In addition, we assume that users do not anticipate how their changes will affect the transmission probabilities $P_{tr,r}(i)$ when they optimize their net benefit. This simplifying assumption corresponds to the standard competitive price-taking assumption of economic theory. In this context, it can be justified when

(a) many "small" users access the link so that the influence of an individual user on the transmission probabilities is negligible small.
(b) it is impossible (or too expensive) for a user to determine how the choice of an allocation may affect the transmission probabilities.

We use the following notation. Let $d_r = (d_r(1), ..., d_r(N))$ be the allocation of user r, let $\underline{d} = (d_1, ..., d_R)$ be the vector containing the allocations of user $r = 1, ..., R$, and let $P_{tr,r}(i, \underline{d})$ be the probability that priority i traffic of user r is transmitted under $\underline{d}$. We then define an equilibrium allocation as follows.

Definition 1: Let $\underline{d}^* = (d_1^*, ..., d_R^*)$ be a given allocation vector. We call $(d_1^*, ..., d_R^*)$ an equilibrium allocation when

$$d_r^* = \arg\max_{d_r \in \Re_+^N} \left\{ U_r(x_r) - \sum_{i=1}^{N} d_r(i) u_i \right\}, \quad r = 1, ..., R,$$

where

$$x_r = \sum_{r=1}^{R} d_r(i) P_{tr,r}(i, \underline{d}^*).$$

In the above maximization problem, the transmission probabilities to compute the average throughput x_r do not depend on the decision variable d_r, but only on $\underline{d}^*$; this reflects the assumption that users ignore how they influence the transmission probabilities when optimizing their own net benefit. Under an equilibrium allocation $(d_1^*, ..., d_R^*)$, the allocation d_r^* maximizes the net benefit of user r. Therefore, no user has an incentive to change its allocation d_r^* as this would decrease the net benefit.

We introduce some additional notation. For every user $r = 1, ..., R$, let the function $D_r : \Re_+ \to \Re_+$ be such that

$$D_r(u) = \arg\max_{x \geq 0} \{U_r(x) - xu\}, \quad u \in \Re_+. \qquad (1)$$

Under Assumption 1, the above maximization problem has a unique solution and D_r is well defined. $D_r(u)$ can be interpreted as the demand function of user r, i.e. $D_r(u)$ is the number of packets users r submits when the price per packet is u and all packets are transmitted with probability 1.

For the maximization problem given by Eq. 1, the first order condition (which is a sufficient condition under Assumption 1) is given by

$$U_r'(x) = u,$$

where U_r' denotes derivative of U_r. This implies that the inverse demand function $G_r(x)$ of user r is given by

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

In the following, we will extensively refer to the demand function D_r and the inverse demand function G_r.

B. Previous Results and Extensions

In [9] we analyzed the case where R users access a single link with capacity C and showed that in this case there always exists an equilibrium allocation. In addition, we showed that there exits a unique price u^* such that in equilibrium the average throughput of user r is given by

$$x_r^* = D_r(u^*), \qquad r = 1, ..., R.$$

Furthermore, in equilibrium we have that

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

When all users have the same utility function, then the bandwidth of the link is shared equally among all users (in equilibrium), i.e. we have that

$$x_r^* = \frac{C}{R}, \qquad r = 1, ..., R.$$

This means that the above priority service leads to a max-min fair bandwidth allocation in the single link case (we will provide a definition of max-min fairness in the next section). It is then natural to ask whether this property also holds for the network case and the above priority scheme leads to max-min fair allocation. In the following, we address this question. 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 than the one in [9] for our analysis. Rather than assuming that the network supports a finite set of different priorities, we consider an idealized priority service where the the network supports a continuum of priorities $u \in \mathcal{I} = [u_{min}, \infty)$. Again, we assume that users are free to choose the priority of their traffic and are charged accordingly. In particular, users are charge a price u per unit transmission rate in priority $u \in \mathcal{I}$. We would like to point out, that for the single link case this idealized priority service is similar to the smart market proposal by [8]. However, here we assume that users pay for each submitted packet, where as in [8] users are charged only for transmitted packets, i.e. packets which are actually delivered. In addition, in our idealized pricing scheme, the price users pay for a packet is equal to the bid attached to the packet, while in the smart market the price is equal to the bid of the first packet that gets dropped. Despite these differences, two schemes lead to the same bandwidth allocation, and the same (congestion) price, in the single link case. To our knowledge, there does not exist an extension of the smart market to the network case.

Before we start with our analysis in Section IV, we briefly review max-min fairness.

III. Bandwidth Sharing Objectives

An important (and difficult) problem in data networks is the question how the available bandwidth should be shared among users. The classical sharing principle is the max-min fairness criteria as discussed, for instance, by Bertsekas and Gallager [1].

A. Max-Min Fairness

One possible way to define fairness for a rate allocation scheme is to require that each user should obtain 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 [1]). Instead, one could first allocate the same transmission rate to all users, and then share the remaining network bandwidth to fully utilize the network. Or ore formally, a max-min fair allocation is defined as follows.

Definition 2: Let $\{1, ..., R\}$ 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 $(x_1, ..., x_R)$ feasible, when for every link $l \in \mathcal{L}$ we have that

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

We call a feasible allocation $(x_1, ..., x_R)$ 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 $x_{r'} \leq x_r$.

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 different weights (priorities) to users capturing their service requirements, and allocate resources accordingly. This situation can be captured as follows. Let $B_r(x_r)$, $x \geq 0$, be an increasing function representing the weight assigned to user r at rate x_r. An allocation $(x_1, ..., x_R)$ 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 $B_r(x_{r'}) \leq B_r(x_r)$. The Definition 2 of a max-min fair allocation is recovered with

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

C. Max-Min Fairness and Pricing

For the single link case, the priority service described above leads to max-min fair allocation (in equilibrium) when all users have the same utility function. When users have different utility functions, one can verify that the equilibrium allocation is weighted max-min fair with weights

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

where $G_r(\cdot)$ is the inverse demand function of user r. These results suggest that there exists a close relationship between a price-based priority service and the concept of max-min fairness. In this subsection, we want to reinforce this relation by developing a fairness criteria for pricing schemes.

One possible way to define fairness for a pricing scheme is to require that each user is charged the same price u^* per unit bandwidth. For example, when R users access a single link with capacity C, a fair pricing scheme would charge each user the same price u^* per unit bandwidth, 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. To see this note the following. In the case of homogeneous utility functions, requiring that all users pay the same price implies that all users obtain the same transmission rate. However, requiring that all users obtain the same transmission rate generally leads to an inefficient use of network capacity as discussed above. 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 is 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.

Definition 3: Let $\{1, ..., R\}$ 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 $(u_1, ..., u_R)$, where u_r is the price user r is charged per submitted packet. We say that an allocation under the pricing scheme $(u_1, ..., u_R)$ is feasible, when for every link $l \in \mathcal{L}$, we have that

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

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

One can show that (under suitable assumptions) a pricing scheme is fair as defined above, if and only if the resulting allocation is weighted max-min fair with weighting functions given by

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

where G_r is the inverse demand function of user r. When all user have the same utility function, then a fair pricing scheme results in a max-min fair allocation.

In the following, we show that the pricing scheme for the priority service that we introduced above is fair, i.e. leads to a weighted max-min fair rate allocation with weights

$$B_r(x_r) = -G_r(x_r).$$

IV. Single Link Case

We first analyze the single link case for the priority scheme described in Section II. We use a discrete-time fluid flow model for our analysis, where time is divided into slots. Let $C > 0$ be a given constant, and consider a single link with a fixed 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. The link supports a continuum $\mathcal{I} = [u_{\min}, \infty)$ of different priorities, where $u_{\min} > 0$ represents the lowest priority. Assume that R users share the link, and assume that each user submits traffic at a constant and deterministic rate. More precisely, let $F_r : \mathcal{I} \mapsto \Re_+$, be the (cumulative) allocation of user r, where $F_r(u)$ is the amount of traffic that user r submits with a priority higher or equal to u (see Figure 1 for an illustration). We then have that for each user $r = 1, ..., R$, the

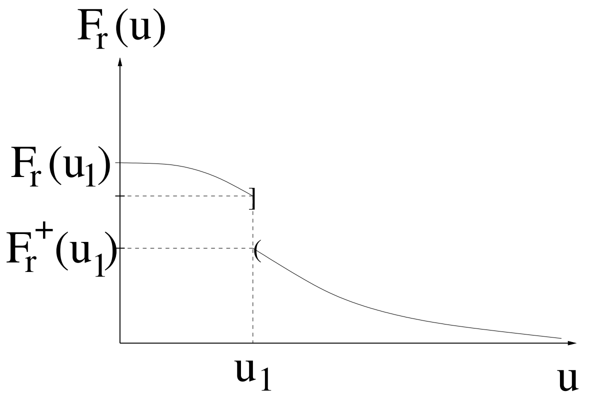

Fig. 1. The (cumulative) allocation function $F_r(u)$ is decreasing and continuous from the left. When $F(u) \neq F^+(u)$, then the quantity $(F(u) - F^+(u))$ as the amount of traffic submitted with priority u

allocation function F_r is non-increasing and continuous from the left, with

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

Let $\mathcal{F}$ be the set of all a possible (cumulative) allocation functions. For a function $F \in \mathcal{F}$, we define $F^+(u)$ by

$$F^+(u) = \lim_{x \downarrow u} F(x).$$

When $F_r(u) \neq F_r^+(u)$, then we can interpret the quantity $(F(u) - F^+(u))$ as the amount of traffic that user r submits with priority u (see Figure 1).

Let

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

be the aggregate allocation function over all users for a given time slot. 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

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

In each time slot, users are charged a price u per unit traffic submitted with priority $u \in \mathcal{I}$. When $F_r \in \mathcal{F}$ is the allocation of user $r \in \{1, ..., R\}$, the total cost of user r has to pay the network in this time slot is given by the Riemann-Stieltjes integral

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

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

A. Formulation as a Non-Cooperative Game

We consider the situation where all users behave in a selfish manner and choose an allocation which optimizes their own net benefit. We define an equilibrium is allocation as follows.

Definition 4: We call an allocation $(F_1^*, ..., F_R^*) \in \mathcal{F}^R$ an equilibrium allocation when for every user $r = 1, ..., R$, 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 $F^* = \sum_{r=1}^{R} F_r^*$ is the cumulative demand over all users.

We make the following assumption to exclude the trivial case where the link capacity exceeds the the demand under the price $u_{\min}$.

Assumption 2: For the interval $\mathcal{I} = [u_{\min}, \infty)$ we have

$$\sum_{r=1}^{R} D_r(u_{\min}) > C.$$

We then have the following result.

Proposition 1: Let Assumption 1 and 2, hold. Then there exists an unique equilibrium allocation $(F_1^*, ..., F_R^*)$ which is given as follows. For $r = 1, ..., R$, 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 in equilibrium only priority u^* is used, and the throughput of user r is given by

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

As $\sum_{r=1}^{R} D_r(u^*) = C$, it follows that the equilibrium allocation is weighted max-min fair with weights

$$B_r(x_r) = -G_r(x_r).$$

In the following section, we extend the analysis to the network case and show that Proposition 1 still holds.

V. Network Case

Let $\mathcal{L}$ be the set of links in the network, and let C_l be the capacity of link $l \in \mathcal{L}$. We associate with each user r a route $(l_{r,1}, ..., l_{r,K_r})$, where K_r is the length of the route and $l_{r,k} \in \mathcal{L}$ is the kth link traversed by user r. Let $H_l \subset \{1, ..., R\}$ be the set of users that pass through link $l \in \mathcal{L}$. We assume that each link serves traffic according to the link model of Section IV, and that traffic is dropped independently at each link. Let $\underline{F} = (F_1, ..., F_R)$ be a given allocation, and let $P_l(u, \underline{F})$ be the resulting transmission probability function for link $l \in \mathcal{L}$. As traffic is dropped independently at each node, the transmission probabilities $P_{tr,r}(u,\underline{F})$, $u \in \mathcal{I}$, for the traffic of user r are given by

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

We have the following relationship between the transmission probability functions $P_l(u,\underline{F})$, $l \in \mathcal{L}$, and the allocation $F_{r,l_{r,k}}(u,\underline{F})$, $k = 1,...,K_r$, of user r on the kth link along its route,

$$F_{r,l_{r,1}}(u,\underline{F}) = F_r(u) \tag{2}$$

$$F_{r,l_{r,k}}(u,\underline{F}) = \int_u^\infty P_{l_{r,k-1}}(v,\underline{F})(-dF_{r,l_{r,k-1}}(v)), \quad k = 2,...,K_r. \tag{3}$$

For this model to be consistent, we need that

$$P_l(u,\underline{F}) = P_{tr}\left(u, \sum_{r\in H_l} F_{r,l}(\underline{F})\right), \text{ for all } l \in \mathcal{L} \text{ and all } u \in \mathcal{I}. \tag{4}$$

Note that the transmission probability of priority u_0 traffic depends only on the traffic that users submit in priorities $u \geq u_0$, but 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 $u \geq u_0$ (on previous links along the route of user r), but not on the transmission probabilities for priorities $u < u_0$. In the appendix, we use this property to construct transmission probabilities that satisfy the system of equation given by Eq. (2)-(4). We have the following result.

Proposition 2: Let Assumption 1 hold. Then, for every allocation $\underline{F} = (F_1,...,F_R) \in \mathcal{F}$, 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 equation given by Eq. (2)-(4).

We provide a proof in Appendix B.

A. Non-Cooperative Game

Again, we study the non-cooperative game where each user tries to optimize its own net-benefit. We call an allocation $\underline{F}^*=(F_1^*,...,F_R^*) \in \mathcal{F}$ an equilibrium allocation when for every user $r = 1,...,R$, 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

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

and the functions $P_l(u,\underline{F}^*)$, $l \in \mathcal{L}$, satisfy the system of equation given by Eq. (2)-(4). We have the following result.

Proposition 3: Let Assumption 1 hold. An allocation $\underline{F}^*=(F_1^*,...,F_R^*) \in \mathcal{F}^R$ is an equilibrium allocation if and only if there exist transmission probabilities $P_l(u,\underline{F}^*)$ that satisfy the system of equation given by Eq. (2)-(4), and the following holds for every user $r = 1,...,R$. For $u_r^* = \liminf\{u \in \mathcal{I} : P_{tr,r}(u,\underline{F}^*) = 1\}$, we have that $P_{tr,r}(u_r^*,\underline{F}^*) = 1$ and

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

This proposition can be proven by the same argument as the one given for Proposition 1, and we omit here a detailed proof. Proposition 3 states that each user r submits traffic exclusively in priority u_r^*. In addition, the transmission probabilities $P_{tr,r}(u,\underline{F}^*)$, $u \in \mathcal{I}$, for user r traffic has the following properties: all traffic of priority u_r^* is transmitted (i.e. we have $P_{tr,r}(u_r^*,\underline{F}^*) = 1$), but all lower priority traffic is dropped (i.e. we have $P_{tr,r}(u,\underline{F}^*) = 0$ for $u < u_r^*$). 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

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

Proposition 3 implies that we can identify an equilibrium allocation $(F_1^*,...,F_R^*)$ by a vector $(u_1^*,...,u_R^*) \in \mathcal{I}^R$: each user r submit all its traffic in priority u_r^* and the average throughput of user r is given by

$$x_r = F_r(u_r^*) = D_r(u_r^*).$$

In the following, we show that there exist a unique equilibrium allocation and, moreover, the price vector $(u_1^*,...,u_R^*)$ associated with this equilibrium allocation is fair (as given by Definition 3).

Let $\mathcal{B} \subset \mathcal{I}^R$ be the set of all price vectors $(u_1,...,u_R)$ which lead to a feasible allocation, i.e. we have

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

Given a vector $(u_1,...,u_R) \in \mathcal{B}$, we say that link l is a bottleneck with respect to $(u_1,...,u_R)$ 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 $r' \in H_l$. We have the following results.

Proposition 4: Let Assumption 1 hold. The allocation $\underline{F}^*$ associated with the vector $(u_1^*,...,u_R^*) \in \mathcal{B}$ is an equilibrium allocation if and only if each user r either has that the user r has a bottleneck link with respect to $(u_1^*,...,u_R^*)$ or $u_r^* = u_{\min}$.

Proposition 5: Let Assumption 1 hold. There exists a unique price vector $(u_1^*,...,u_R^*) \in \mathcal{B}$ such that the corresponding allocation $\underline{F}^*$ is an equilibrium allocation. Furthermore, when the price vector $(u_1^*,...,u_R^*) \in \mathcal{B}$ 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 loosing feasibility, or increasing the price of another user r' with a price $u_{r'}^* \geq u_r^*$.

We provide proofs for Proposition 4 and 5 in Appendix C and D. Proposition 5 states that the price-based priority service is fair as given by Definition 3. This implies that the equilibrium allocation is weighted max-min fair with weights

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

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

Corollary 1: Let Assumption 1 hold. Then there exists an unique equilibrium allocation.

VI. Conclusions

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. The analysis of situation where the network supports finite number, rather than a continuum, of priorities seems to be more involved; we are currently working on this extension.

References

[1] D. P. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, Englewood-Cliffs, New Jersey, 1992.

[2] R. Gibbons, Game Theory for Applied Economists, Princeton University Press, 1992.

[3] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski,*" Assured forwarding phb group,"* IETF: RFC 2597, 1999.

[4] F. P. Kelly, A. Maulloo, and D. Tan, *"Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability,"* Journal of the Operational Research Society 49, pp. 237-252, 1998.

[5] R. J La and V. Anatharam, *"Charge-Sensitive TCP and Rate Control in the Internet,"* prepring, 2001.

[6] S. H. Low and D. E. Lapsley, *"Optimization Flow Control, I: Basic Algorithm and Convergence,"*, *IEEE/ACM* Transactions on Networking, 7(6), pp. 861-75, 1999.

[7] S. Athuraliya and S. H. Low, *"Optimization Flow Control, II: Implementation,"* submitted for publication.

[8] J. K. MacKie-Mason and H. R. Varian. Pricing the Internet. In B. Kahin and J. Keller, editors, *Public Access to the Internet*. Prentice-Hall, Englewood-Cliffs, New Jersey, 1994.

[9] P. Marbach Pricing Priority Classes in Differentiated Services Networks. 37th Annual Allerton Conference on Communication, Control, and Computing, Monticello, IL, September 1999.

[10] J. Mo and J. Walrand, *Fair end-to-end window-based congestion control*. SPIE '98 International Symposium on Voice,Video and Data Communications, October 1998.

[11] Ch. Semeria and F. Fuller, *"3com's Strategy for Delivering Differentiated Service Levels,"* http://www.3com.com/technology/tech_net/white_papers/500652.html, 1999.

Appendix

A. Proof of Proposition 1

We use the following notation. Let $\mathcal{S}$ be the set of all step functions $F \in \mathcal{F}$, 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

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

For every user $r = 1, ..., R$ 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

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

We first state a few preliminary lemmas. The derivation of these lemmas is straight forward and we omit here detailed proofs.

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

Lemma 1: For every $F_r \in \mathcal{F}$ there exists a $F_r^* \in \mathcal{S}$ such that

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

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 $F_r \notin \mathcal{S}$.

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

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

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 if $F_r \notin \mathcal{S}_r$.

Lemma 3: Let

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

Then there exists optimal allocation $F_r^* \in \mathcal{F}$ for the maximization problem

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

where

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

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

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

We are now ready to prove Proposition 1.

Proof: By Lemma 3, 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 $r = 1, ..., R$, we have

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

As by Assumption 2, we have that

$$\sum_{r=1}^R D_r(u_{\min}) > C,$$

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. ∎

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 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

$$\begin{aligned} F_{r,l_1} &= d_r && (5)\\ F_{r,l_k} &= \alpha_{l_{k-1}} F_{r,l_{k-1}}, \quad k = 2, ..., K_r, && (6)\\ \alpha_l &= \begin{cases} 0 & \text{if } C_l = 0, \\ \min\left\{ \frac{C_l}{\sum_{r \in A_l} F_{r,l}}, 1 \right\} & \text{otherwise.} \end{cases} && (7) \end{aligned}$$

We have the following lemma.

Lemma 4: There exists a solution $(\alpha_1, ..., \alpha_L)$ to the system of equations given by Eq. (5)-(7).

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

$$\begin{aligned} F_{r,l_1} &= d_r \\ F_{r,l_k} &= \alpha_{l_{k-1}} F_{r,l_{k-1}}, \quad k = 2, ..., K_r, \end{aligned}$$

Note that this mapping is continuous. As the set $[0,1]^L$ is convex and compact, it follows that there exists a solution to the system of equations above. ∎

We can use Lemma 4 to iteratively construct a set of transmission probabilities that solve the system of Eq. (2)-(4). 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 4 to compute the transmission probabilities $P(u, \underline{F})$ for $u \geq u_1$, i.e. we show that the transmission probabilities $P(u, \underline{F})$ satisfy Eq. (2)-(4) for $u \geq u_1$. We then proceed to lower the price until we saturate the next link at the price u_2 and use Lemma 4 to compute the transmission probabilities $P(u, \underline{F})$ for u such that $u_2 \leq u < u_1$; and so on.

Proof: Let A^k be the set of links not saturated at the beginning of iteration k, let B^k be the set of link that get saturated for the first time at iteration k, and let u^k be the (highest) price at which these links get saturated. Furthermore, let P^k denote the set of users not passing through any saturated link at the beginning of iteration k. We then iteratively construct transmission probabilities $P(u, \underline{F})$, $u \in \mathcal{I}$, that satisfy Eq. (2)-(4) as follows.

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

Then do the following steps.

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

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

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

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

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

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

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

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

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

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

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

To prove that this procedure indeed provides a solution for the systems of equation given by Eq. (2)-(4), 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 Eq. (2)-(4).

Note that by Lemma 4 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 $F^k_{r,l}(u,\underline{F})$ 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)

$$F^k_{r,l_j}(u) = \int_u^\infty P^{k-1}_{l_{j-1}}(u)(-dF_{r,l_{j-1}}(u)), \quad j = 2,...,K_r,$$

and let

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

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 $r = 1,...,R$, that

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

This implies that the set

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

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

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

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

$$\sum F^k_{r,l}(u_{k-1}) < C_l.$$

As $u^{k-1} > u_{\min}$ (otherwise the procedure would have already terminated) and $F^k_{r,l}(u)$ 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 $P_l(u,\underline{F})$, $l \in \mathcal{L}$, provide a solution to the systems of equations given by Eq. (2)-(4) we use an induction argument.

By Lemma 4, the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the system of equations given by Eq. (2)-(4) 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

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

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

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

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

$$P_l(u^k,\underline{F}) = P^k_{tr,l}(u^k,\underline{F}) = \alpha^k_l,$$

and it follows that the transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the system of equations given by Eq. (2)-(4) 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 Eq. (2)-(4) for $u \ge u^{k_0}$. When $u^{k_0} = u_{min}$, then we are done. When $u^{k_0} > u_{min}$, 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 $u \in [u_{min}, u^{k_0})$ are given by

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

This implies that transmission probabilities $P_l(u,\underline{F})$, $l \in \mathcal{L}$, solve the systems of equations given by Eq. (2)-(4) for $[u_{min}, u^{k_0})$, which completes the proof.

■

C. Proof of Proposition 4

To prove Proposition 4, we adapt the argument used in [1] to prove the equivalent result for a max-min fairness bandwidth allocation.

Assume that a given price vector $(u_1,...,u_R)$ leads to a feasible allocation, but there exists a user r with $u_r > u_{\min}$ who does not have a bottleneck link with respect to $(u_1,...,u_R)$. This mean to for every link $l \in \{l_{r,1},...,l_{r,K_r}\}$ along the route of user r we have

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

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

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

Moreover, when for each user r we either have that $u_r = u_{\min}$, or r has a bottleneck link with respect to $(u_1,...,u_R)$, then $(u_1,...,u_R)$ leads to an equilibrium allocation. Indeed, in this case we have that for every user $r = 1,...,R$

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

and Proposition 3 implies that $(u_1,...,u_R)$ corresponds to an equilibrium allocation .

D. Proof of Proposition 5

To prove Proposition 5, we again adapt an argument used in [1] 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 condition: $k = 1$; $P^1 = \{1, ..., R\}$; $A^1 = \mathcal{L}$; $u^0 = \infty$; and $u_r^0 = \infty$.

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

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

2) Set

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

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

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

4) Set $u_r^{k+1} = \begin{cases} u^k & r \in Q^k \\ u_r^k & \text{otherwise} \end{cases}$
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 the procedure terminates each user r either passes through a bottleneck link or we have $u_r^{k_0} = u_{min}$, and Proposition 3 implies that $(u_1^{k_0}, ..., u_R^{k_0})$ corresponds to an equilibrium allocation.

In fact, the allocation that we constructed above is the unique allocation such that every user has a bottleneck link. This can be seen as follows. Assume that $(u_1, ..., u_R)$ is an allocation such that each user has a bottleneck link with respect to $(u_1, ..., u_R)$. Then for every user $r \in Q^1$, we have that

$$u_r = u^1.$$

Indeed if this is not the case, then there exist users $r, r' \in Q^1$, such that $u_r < u_{r'}$ and $u^1 < u_{r'}$. As for every link we have that

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

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

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

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

$$u_r = u^{\bar{k}+1}.$$

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 $(u_1, ..., u_R)$. Therefore, by induction, we have that $(u_1, ..., u_R)$= $(u_1^{k_0}, ..., u_R^{k_0})$. This implies that there exists a unique price vector $(u_1, ..., u_R)$ such that each user has a bottleneck link with respect to $(u_1, ..., u_R)$.

Next, we show that a price vector $(u_1, ..., u_R)$ has the property that each user has a bottleneck link with respect to $(u_1, ..., u_R)$, if and only if it is fair (as given by Definition 3).

Suppose that $(u_1, ..., u_R) \in \mathcal{B}$ is fair and, to arrive at a contradiction, assume that there exists a user r with no bottleneck link. Then, for each link l crossed by user r with $\sum D_r(u_r) = C_l$, there must exists a user $r' \neq r$ such that $u_{r'} < u_r$, thus the quantity

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

is positive. Therefore, there exist scalars $\delta_r, \delta_{r'} > 0$, such that for $u'_r = u_r - \delta_r$ and $u'_{r'} = u'_{r'} + \delta_{r'}$ we have

$$\sum_r D_r(u'_r) \leq C_l$$

this contradicts the fairness property of $(u_1, ..., u_R)$.

Conversely, assume that each session has a bottleneck link with respect to a price vector $(u_1, ..., u_R) \in \mathcal{B}$. Then, to increase the u_r of user r, while maintaining feasibility, we must decrease the rate of some user r' crossing the bottleneck link of user r. Since $u_r \leq u_{r'}$ for all users r' crossing the bottleneck link of user r, the price vector $(u_1, ..., u_R)$ is fair.