Pricing Differentiated Services Networks: Bursty Traffic

Peter Marbach

Abstract—**We study the role of pricing in Differentiated Services (DiffServ) networks. We model DiffServ as a priority service, where users are given the freedom to choose the priorities of their traffic, but are charged accordingly. Using a game theoretic framework, we study the case where users choose an allocation of priorities to packets in order to optimize their net benefit. For the case where users with bursty traffic access a single link, we show that there always exists an equilibrium for the corresponding non-cooperative game. Furthermore we show that pricing can be used to provide relative QoS guarantees.**

Keywords—**Differentiated Services, Quality-of-Service, Pricing.**

I. Introduction

In this paper we study the role of pricing in differentiated services networks. We adopt an approach similar to [7], however while the analysis in [7] was carried out for smooth traffic (such as multiplexed voice connections, or long file transfers), the focus in this paper is on bursty traffic (such as interactive Web traffic).

The differentiated services (DiffServ) proposal of the IETF (Internet Engineering Task Force) aims to provide Quality-of-Service (QoS) in packet-based networks (such as the Internet) in a simple and scalable way (see [6]). The basic idea is the following. DiffServ networks provide a (small) number of different service classes. Users accessing the network assign (mark) their traffic to one of these classes, and the network forwards packets at each node using class-specific rules (called per-hop behaviors). This mechanism can be used to provide differentiated QoS; for example, DiffServ could be used to implement a premium and a best-effort service, where premium traffic always gets preferential treatment (and therefore experiences less delay and is less likely to be dropped) over best-effort traffic. In this framework, it is natural to assume that users are charged based on the service classes that they use, i.e. a high-quality (premium) service would be more expensive than a low-quality (best-effort) service. Also, one would expect that users choose service classes according to their service requirements, i.e. users will try to meet their requirements at the lowest possible cost. In this paper we analyze this situation for users with bursty traffic.

For our analysis, we model DiffServ as a priority service, and consider a network which offers a finite set $\mathcal{I} = \{1, ..., N\}$ of priority levels. At each node, packets are served according to a priority rule, so that packets with high priorities will experience less delay and are less likely to be dropped. Users are free to choose the priorities for their traffic, but are charged accordingly by the network. In particular, the network charges a price u_i per unit traffic submitted with priority i, where for $i, j \in \mathcal{I}$ we have

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

The fact that users are charged for every submitted (rather than for transmitted) packets, makes this scheme easy to implement: the network forwards packets according to their priorities, and there is no need to track dropped packets (or link congestion levels) and send pricing information back to the users. In addition, our analysis shows that charging users for all submitted packets will provide the right incentives (back-pressure) for users to reduce their rate during periods of congestion. Another important feature of this model is that the prices u_i, $i \in \mathcal{I}$, are static and do not change (on a fast time-scale) to respond to changes in the network load. However users can change their allocation dynamically (on a fast time-scale) to maximize their net benefit from using the network. This property of the above pricing scheme allows to keep the operation of the network simple("dumb network") by pushing the complexity towards the edge of the network ("intelligent users/end hosts").

To study this priority service, we adopt an economic framework where we characterize the users' service requirements by utility functions. 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. It is then natural to assume that users will choose an allocation of packets to priorities to maximize their net benefit, i.e. their utility minus cost. We model this situation as a non-cooperative game for which we wish to determine whether there exists an equilibrium, and how pricing decision influence the characteristic of an equilibrium. To keep the analysis simple, we consider only the single link case.

This setting is similar to the one in [7], however, in [7] it was assumed that traffic is smooth and long-lived, so that users are able to accurately estimate (and predict) the QoS in the different priority levels, and adapt their allocation accordingly. In this paper, we assume that the traffic is bursty, i.e. we assume that users submit traffic in short-lived, high-volume bursts (as it is typically the case for interactive Web traffic). As a result, users do not know with certainty whether a burst will be transmitted in a particular priority, but can only infer the transmission probability (distribution) of a burst. Compared with [7], we obtain qualitatively different results. In [7], it was showed that (for smooth traffic) pricing can be used to provide absolute QoS guarantees, i.e. pricing allows to upper bound the data loss probability in the different priority levels. This is no longer true for bursty traffic. However, for bursty traffic, we can still use pricing to provide relative QoS guarantees, i.e. when $P_{tr}(i)$ is the probability that a burst gets transmitted (in equilibrium) in priority $i \in \mathcal{I}$, then

P. Marbach is with the Department for Computer Sciences, University of Toronto; email: marbach@cs.toronto.edu.

Part of this research was carried out at the Centre for Communication Systems Research, University of Cambridge, UK. This research was supported by a contract with Alcatel Bell, Belgium.

we have the following relationship,

$$\frac{P_{tr}(i)}{P_{tr}(i+1)} \geq \frac{u_i}{u_{i+1}}.$$

The framework that we consider here is closely related to the Paris Metro Pricing (PMP) proposal [8] for providing QoS in the Internet. PMP partitions the network into several logically separated channels where channels differ in the prices paid for using them. This can be implemented by using the priority scheme, combined with pricing, described above. However, [8] does not analyze the existence of an equilibrium and the role of pricing. [3] provides simulation results to demonstrate that priorities can be an efficient tool for providing QoS in packet-based networks. In particular, [3] illustrates that using priorities, combined with pricing, achieves a higher a social welfare than a best-effort/flat-rate scheme. Gupta et. al. investigate in [5] a priority service similar to the one proposed here, however they consider the situation where the network dynamically changes the prices associated with the different priority classes in order to track a socially optimal allocation. This approach is theoretically appealing, but expensive to implement. The scheme considered in this paper is simpler as prices associated with the different service classes are fixed. Park et. al. also consider in [2], [9] a service where packets get marked according to customer's QoS requirements. However, there is no price associated with the different traffic classes, i.e. the costs incurred to customers are purely performance related. They study the situation where users are free to choose their traffic class [9], as well as where users indicate their QoS requirements and a network controller assigns network resources [2].

The rest of the paper is organized as follows. In Section II we summarize the results of [7]. In Section III we present the traffic and link model that we use for bursty traffic. In Section IV, we analyze the case where users with elastic traffic access the link. In Section V, we outline how the analysis can be extend to the case where users with tight QoS requirements access the link. In Section VI we present a case study to illustrate the results that we obtained.

II. Background: Smooth Traffic

In this section, we give a summary of [7], where we present the traffic model (which is a fluid-flow model) and link model, and state the main results (without proofs).

We use a discrete-time 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 packets per time slot, where packets can be arbitrarily small. There is no buffering and packets which do not get transmitted in a given time slot are dropped. The link supports a finite set $\mathcal{I} = \{1, ..., N\}$ of priorities, where 1 is the lowest priority - 2 the second lowest, and so on. Assume that R users share the link, and assume that each user submits packets at a constant and deterministic rate. In particular, let $d_r = (d_r(1), ..., d_R(N)) \in \Re_+^N$, $r = 1, ..., R$, be the allocation of user r, where $d_r(i)$ is the number of priority i packets user r submits in each time slot, and let $(d_1, ..., d_R)$ be the corresponding allocation vector. Furthermore, let $d = (d(1), ..., d(N))$ be

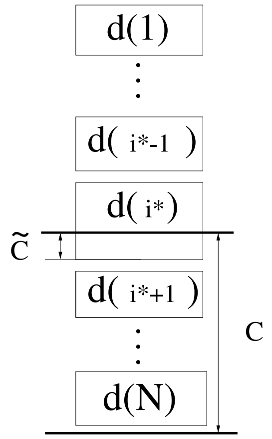

Fig. 1. The link transmits all packets with a priority higher than i^*. From the packets with priority i^*, only $\tilde{C} = \left(C - \sum_{i=i^*+1}^{N} d(i)\right)$ packets are transmitted, which are chosen at random. All other packets are dropped.

the aggregated allocation (over all users), where

$$d(i) = \sum_{r=1}^{R} d_r(i), \qquad i = 1, ..., N,$$

is the total number of packets submitted in priority i in each time slot. The link then serves packets as follows (see Figure 1 for an illustration). In each time slot, all packets with priority i such that

$$\sum_{k=i}^{N} d(k) < C$$

are transmitted. For priority i^* with

$$\sum_{i=i^*+1}^{N} d(i) < C \quad \text{and} \quad \sum_{i=i^*}^{N} d(i) \geq C,$$

only $\left(C - \sum_{i=i^*+1}^{N} d(i)\right)$ packets, which are chosen at random, get transmitted. All packets with priority $i < i^*$ are dropped. For the aggregated allocation $d \in \Re_+^N$, the (packet) transmission probability $P_{tr}(i, d)$ in priority i is then given as follows,

$$P_{tr}(i,d) = \begin{cases} 1 & \text{if } C > \sum_{k=i}^{N} d(k), \\ \frac{C - \sum_{k=i+1}^{N} d(k)}{d(i)} & \text{if } \sum_{k=i}^{N} d(k) \geq C > \ldots \\ & > \sum_{k=i+1}^{N} d(k), \\ 0 & \text{otherwise.} \end{cases}$$

Let $P_{tr}(d) = (P_{tr}(1,d), ..., P_{tr}(N,d))$ be the corresponding probability vector. The average throughput of user r is then given by

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

Users are charged a price u_i per submitted packet with priority i, where for $i, j \in \mathcal{I}$, $i < j$, we have $0 < u_i < u_j$. In addition, there exists a feedback mechanism that allows users to detect whether their packets were transmitted and to estimate the packet transmission probabilities in each priority.

The above model describes a priority service which charges a fixed price u_i is per submitted packet with priority i. However the price for actually transmitting a packet will depend on the total demand: in order to transmit a packet in priority i, users have to submit a packet (on average) $1/P_{tr}(i,d)$ times and the (average) price for transmitting a packet in priority i is then equal to $u_i/P_{tr}(i,d)$. In the following, we study this service.

A. Users with Elastic Traffic

We analyze the above service model for users with elastic traffic [11], i.e. users who perceive quality of service solely as a function of the average throughput. For this case, we can characterize the service preferences of each user through a quality indicator (utility function) which depends on a single variable: the average throughput. We will assume that these utility functions are increasing in the average throughput. Examples of users with elastic traffic are users sending email, transferring files, and browsing the Web.

Consider a fixed user $r \in \{1,...,R\}$. Furthermore, suppose that $d_r = (d_r(1),...,d_r(N)) \in \Re_+^N$ is the allocation of user r and $d = (d(1),...,d(N)) \in \Re_+^N$ is the aggregated allocation. Then, we associate with user r the utility function $U_r(x_r)$, where

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

is the average throughput of user r. We make the following assumption.

Assumption 1: For each user $r = 1,...,R$, the function $U_r : \Re_+ \mapsto \Re_+$, is increasing, bounded, strictly concave, and twice differentiable.
Assumption 1 implies that for every user $r = 1,...,R$, the function $U_r : \Re_+ \mapsto \Re_+$, can be written as

$$U_r(x_r) = \int_0^{x_r} G_r(x)dx,$$

where $G_r(x) : \Re_+ \mapsto \Re_+$ is bounded, strictly decreasing and differentiable with

$$G_r(0) > 0 \quad \text{and} \quad \lim_{x\to\infty} G_r(x) = 0.$$

As we will show below, the function G_r can be interpreted as the inverse demand of user r with utility U_r. Note that Assumption 1 does not require that all users have the same utility function.

Assumption 1 implies that a user can achieve a high utility by submitting high priority packets; however, doing so will also increase the cost. Users therefore face a trade-off between achieving a high utility and keeping their cost low. We capture this situation as follows. Given an aggregated demand $d = (d(1),...,d(N)) \in \Re_+^N$, and an allocation $d_r = (d_r(1),...,d_r(N)) \in \Re_+^N$ of user r, we associate with user r a net benefit function given by

$$U_r(x_r) - \sum_i d_r(i)u_i,$$

where

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

Note that $\sum_i d_r(i)u_i$ is the total cost user r pays the link under the allocation d_r. In the following, we assume that each user chooses an allocation that maximizes its own net benefit. However, the net benefit functions of the individual users are couple through the transmission probabilities $P_{tr}(i,\cdot)$. Therefore, changing d_r will not only affect the net benefit of user r, but the net benefit functions of all users. In the next subsection, we study this interaction among users.

We introduce some more notation. For $r = 1,...,R$, let the function $D_r : \Re_+ \to \Re_+$ be defined such that $D_r(u)$ is the optimal solution to the maximization problem, i.e. we have

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

Under assumption 1, the above maximization problem has a unique solution. $D_r(u)$ 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. Note that the function $G_r(x)$ that we defined above is the inverse of the function $D_r(u)$. Finally, we use $D(u)$, $u \in \Re_+$, given by

$$D(u) = \sum_{r=1}^{R} D_r(u),$$

to denote the aggregated demand over all users under the price u. In the following, we will extensively use refer to the demand function D_r, and the inverse demand function G_r, of user r, as well as the total demand D.

B. Formulation as a Non-Cooperative Game

In this subsection, we study the case where all users simultaneously try to maximize their own net benefit. We model this situation as a non-cooperative game. For an introduction to the basic concepts of game theory we refer to [4]. For the analysis, we assume that users ignore how they influence the transmission probabilities when optimizing their own 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.
(b) it is impossible (or too expensive) for a customer to determine how the choice of an allocation may affect the transmission probabilities.
We then define an equilibrium allocation as follows.

Definition 1: Let $(d_1^*,...,d_R^*)$ be a given allocation vector such that $d_r^* \in \Re_+^N$, $r = 1,...,R$, and let $d^* = (d^*(1),...,d^*(N))$,

$$d^*(i) = \sum_{r=1}^{R} d_r^*(i), \qquad i = 1,...,N,$$

be the corresponding aggregated allocation. We call $(d_1^*,...,d_R^*)$ an equilibrium allocation when for every user $r = 1,...,R$, we have

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

where

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

Note that in the maximization problem given in the Definition 1, the transmission probabilities to compute the average throughput x_r do not depend on the decision variable d_r, but only on 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^*, $r = 1, ..., R$, 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. This implies that once an equilibrium is reached, no user will deviate.

In the following, we show that there always exists an equilibrium allocation (Proposition 2). In addition, we characterize an equilibrium allocation; in particular, we show how pricing decisions influence the properties of an equilibrium (Proposition 1). For the analysis, we assume that the demand in the lowest priority level exceeds the link capacity and the total demand $D(u_i)$ for priority $i \in \mathcal{I}$ is strictly positive. The case where the link capacity exceeds the the demand in the lowest priority level leads to the (trivial) situation where all users choose in equilibrium the rate $x_r = D_r(u_1)$, $r = 1, ..., R$, and all packets are transmitted with probability 1. When the total demand $D(u_i)$ in priority i is equal to zero, then priority i will not be used in equilibrium, and we can ignore this priority.

Assumption 2: We have

$$D(u_1) > C$$

and

$$D(u_i) > 0, \qquad \text{for } i = 1, ..., N.$$

In the next proposition, we will derive the necessary conditions for an allocation vector to be an equilibrium allocation. This will allow us to characterize an equilibrium allocation and determine how pricing decisions influence the properties of an equilibrium.

Proposition 1: Let Assumption 1 and 2 hold. Consider an allocation vector $(d_1^*, ..., d_R^*)$ and let

$$d^* = \sum_{r=1}^{R} d_r^*$$

be the corresponding aggregated allocation. Furthermore, let

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

be the average throughput of user r, $r = 1, ..., R$ under $(d_1^*, ..., d_R^*)$, and define

$$u_{N+1} = \max_{r=1,...,R} G_r(0).$$

Then the following holds. When $(d_1^*, ..., d_R^*)$,$d_r \in \Re_+^N$, $r = 1, ..., R$, is an equilibrium allocation, then there exists a priority $i_0 \in \{1, ..., N\}$ such that, for $r = 1, ..., R$, we have

(a) $d_r^*(i) = 0$, for all $i \notin \{i_0, i_0 + 1\}$;

(b) $P_{tr}(i_0, d^*) \geq \frac{u_{i_0}}{u_{i_0+1}}$;

(c) $x_r^* = D_r(u^*)$ where

$$u^* = \frac{u_{i_0}}{P_{tr}(i_0, d^*)}.$$

Note that u_{N+1} is lowest price such that the total demand vanishes, i.e. we have that $D(u) > 0$, for $u \in [0, u_{N+1})$ and $D(u) = 0$, for $u \geq u_{N+1}$. Using Proposition 1, we can characterize an equilibrium allocation as follows.

(a) states that, under every equilibrium allocation, at most two priorities are used. This result has the following interpretation. Let $(d_1^*, ..., d_R^*)$ be a given allocation vector, let $d^* = \sum_{r=1}^{R} d_r^*$ be the corresponding aggregated demand, and let the priority i^* be such that

$$P_t(i, d^*) = 0, \text{ for all } i < i^*$$

and

$$P_t(i, d^*) > 0, \text{ for all } i \geq i^*.$$

Under the link model given in Section III, packets with a priority lower than priority i^* will never be transmitted. As users are charged for every submitted packet, it is not economical to use these priorities. Furthermore, packets with a priority higher than $i^* + 1$ will be transmitted with the same probability as packets with priority i^*+1, but at a higher cost; and these priorities won't be used either.

(b) states that when priority i is used under a given equilibrium allocation $(d_1^*, ..., d_R^*)$, then the probability that a submitted packet with priority i gets transmitted is lower-bounded by u_i/u_{i+1}. This result has the following interpretation. In equilibrium, the average price for transmitting a packet in priority i is equal to $u_i/P_{tr}(i, d^*)$. When

$$P_{tr}(i, d*) < \frac{u_i}{u_{i+1}},$$

then the average price for transmitting a packet in priority i is higher than for priority $i + 1$, i.e. we have that

$$u_i/P_{tr}(i, d^*) > u_{i+1}.$$

In this case it is more economical for users to switch their packets from priority i to priority $i + 1$, and therefore, Property (b) will hold in equilibrium.

(c) implies that the average throughput of user r is

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

where

$$u^* = \frac{u_{i^*}}{P_{tr}(i^*, d^*)}.$$

Here, u^* can be interpreted as the average price users pay for transmitting a packet; and the average throughput of user r is equal to the demand D_r under u^*. This result then suggests that the pricing scheme is fair in the sense that all users receive a throughput according to the same price u^*.

Based on Proposition 1, one can derive sufficient conditions for an $(d_1^*, ..., d_R^*)$ to be an equilibrium allocation; which can then be used to construct an equilibrium allocation and to obtain the following result.

Proposition 2: Let Assumption 1 and 2 hold. Then there exists an equilibrium allocation. Furthermore, when there exists a priority $i_0 \in \mathcal{I}$ such that

$$D(u_{i_0}) > C > D(u_{i_0+1}),$$

then the equilibrium is unique.
Proposition 2 states that we have a unique equilibrium allocation when there exists a priority $i_0 \in \mathcal{I}$ such that

$$D(u_{i_0}) = \sum_{r=1}^{R} D_r(u_{i_0}) > C > \sum_{r=1}^{R} D_r(u_{i_0+1}) = D(u_{i_0}).$$

When there exists a priority i_0 such that

$$D(u_{i_0}) = C,$$

then it can been shown that there is an infinite number of equilibrium allocations

C. Discussion

The analysis of [7] shows that (for smooth traffic) pricing can be used to provide absolute QoS guarantees, i.e. the probability that a submitted packet in priority i gets transmitted is lower bounded by the price ratio u_i/u_{i+1}.

To derive these results, a fluid-flow model was employed where the demand in each time slot is given by the (deterministic) allocation $(d_1, ..., d_R)$. While this is a reasonable assumption for smooth traffic, it is clearly not realistic for bursty traffic. In the following, we develop a traffic model for bursty traffic, and carry out the above analysis for this new model.

III. Bursty Traffic: Traffic and Link Model

In this section, we present the traffic and link model that we use for bursty traffic. The goal is to formulate a model that is simple, but still captures the features of bursty traffic that are relevant to our analysis. For the notation, we use $\lambda_r = (\lambda_r(1), ..., \lambda_r(N))$ to denote the allocation of a user with bursty traffic.

A. Traffic Model

For our analysis, we use the following traffic model. Consider a fixed time slot and, for a given user r, a fixed allocation $\lambda_r = (\lambda_r(1), ..., \lambda_r(N)) \in \Re_+^N$. User r will then submit in this time slot a burst in priority i with probability $\lambda_r(i)$, and submit no traffic in priority i with probability $1 - \lambda_r(i)$. We call $\lambda_r = (\lambda_r(1), ..., \lambda_r(N))$ the rate vector of user r. Let $\lambda = (\lambda(1), ..., \lambda(N)$ be the aggregate rate vector (over all users), i.e.

$$\lambda(i) = \sum_{r=1}^{R} \lambda_r(i), \qquad i = 1, ..., N,$$

and let

$$\Lambda(i) = \sum_{n=i}^{N} \lambda(n), \qquad i = 1, ..., N,$$

be the rate at which burst get submitted in priority i or higher. In the following, we assume that all bursts are of the same size B, and set B equal to 1.

In addition, we assume that many "small" users share the link (i.e. R is large), so that $\lambda_r(i) << 1$, for $r = 1, ..., R$, and $i = 1, ..., N$. This assumption allows us to use the Poisson approximation for the probability distribution of total number of bursts submitted in priority i and higher in a given time slot. Let M_i $i = 1, ..., N$, be the total number of bursts submitted in priority i and higher in a given time slot; using the Poisson approximation, we then obtain the following probability distribution for M_i,

$$P(M_i = m) = \Lambda(i)^m \frac{\exp(-\Lambda(i))}{m!}, \qquad m = 0, 1, 2, ...$$

B. Link Model

For the link model, we assume that a burst in priority i is lost whenever M_i (which is the total number of bursts submitted in priority i and higher), exceeds the link capacity. Note that this is a simplifying assumption: when M_i exceeds the link capacity C and M_{i+1} is strictly lower than C, then only a faction, but not all, of the priority i bursts are lost. The two main reasons for making this assumption are the following.
1. Using this assumption simplifies the analysis and makes the results more transparent.
2. This assumption leads to a good approximation of the loss rate when the burst size B is large compared with the link capacity C. In fact, for the case where the burst size B is equal to the link capacity C, and in each time slot at most one burst can be transmitted, the above assumption leads to the exact expression for the loss rate. Therefore, this model captures the case where traffic is very bursty, i.e. the burst size is large compared with the link capacity.
The (fluid flow) traffic and link model of Section II, and the above model, capture the situation where the traffic is very smooth, and very bursty, respectively. In reality, the traffic characteristic will often lie between this two extreme cases, and results that are based on either model will not accurately predict the actual behavior of the system. However, as it is often the case, knowing what happens in the two extreme cases will provide valuable insight into the general system behavior. In Section VI, we use a numerical example to illustrate this point.

The above assumptions then lead to the following link model. Consider a fixed time slot and a fixed allocation $(\lambda_1, ..., \lambda_R)$. Let M_i, $i = 1, ..., N$, be the total number of bursts submitted in priority i and higher in this time slot. Furthermore, let $m_0 = \lfloor C \rfloor$ be the maximal number of bursts that the link can transmit per time slot. Using the above traffic model and assuming that a burst in priority i is lost whenever M_i exceeds m_0, the probability that a burst is transmitted in priority i is then given by

$$P_{tr}(i, \lambda) = 1 - P_{loss}(\Lambda(i)),$$

where $P_{loss}(\Lambda(i))$ is

$$P_{loss}(\Lambda(i)) = \sum_{m=m_0+1}^{\infty} \Lambda(i)^m \frac{\exp(-\Lambda(i))}{m!}.$$

IV. Users with Elastic Traffic

Again, we analyze the above priority scheme for users with elastic traffic, and assume that the utility functions of the users depend only on the average throughput. More precisely, when $\lambda_r = (\lambda_r(1), ..., \lambda_r(N))$ is the rate vector of users r, and $\lambda = (\lambda(1), ..., \lambda(N))$ is the total rate vector over all users, then we associate with users r the utility $U_r(\kappa_r)$, where

$$\kappa_r = \sum_{i=1}^{N} P_{tr}(i, \lambda)\lambda_r(i)$$

is the average throughput of users r. As in Section II, let $D_r(u)$ be the demand function associated with the utility U_r, i.e. we have

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

and let

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

be the total demand over all users under price u.

Again, we assume that each users wants to maximize its own net benefit. Given an aggregated rate vector $\lambda = (\lambda(1), ..., \lambda(N))$ and a rate vector $\lambda_r = (\lambda_r(1), ..., \lambda_r(N))$ of a given user r, the net benefit of user r is equal to

$$U_r(\kappa_r) - \sum_{i=1}^{N} \lambda_r(i) u_i \tag{1}$$

where

$$\kappa_r = \sum_{i=1}^{N} P_{tr}(i,\lambda)\lambda_r(i).$$

A. Formulation as a Non-Cooperative Game

Consider now the non-cooperative game where all users simultaneously optimize their net benefit. As in the previous section, we assume that users ignore how they influence the transmission probabilities when optimizing their own net benefit. Similarly, we then define an equilibrium rate vector as follows.

Definition 2: We call an rate vector $(\lambda_1^*, ..., \lambda_R^*)$, $\lambda_r^* \in \Re_+^N$, $r = 1, ..., R$, an equilibrium allocation when, for every user $r = 1, ..., R$, we have

$$\lambda_r^* = \arg\max_{\lambda_r\in\Re_+^N} \left\{ U_r(\kappa_r) - \sum_{i=1}^{N} \lambda_r(i) u_i \right\} \tag{2}$$

where

$$\kappa_r = \sum_{i=1}^{N} P_{tr}(i,\lambda^*)\lambda_r(i)$$

and $\lambda^* = \sum_{r=1}^{R} \lambda_r^*$.
We then have the following results.

Proposition 3: Let Assumption 1 and 2 hold. Consider a rate vector $(\lambda_1^*, ..., \lambda_R^*)$ and let

$$\lambda^* = \sum_{r=1}^{R} \lambda_r^*,$$

be the corresponding aggregated rate vector. Furthermore, let

$$\kappa_r^* = \sum_{i=1}^{N} \lambda_r^*(i) P_{tr}(i,\lambda^*), \qquad r = 1, ..., R,$$

be the average throughput of user r under $(\lambda_1^*, ..., \lambda_R^*)$, and let

$$u_{N+1} = \max_{r=1,\ldots,R} G_r(0).$$

Then the following holds. When $(\lambda_1^*, ..., \lambda_R^*)$ is an equilibrium allocation then there exists a priority $i_0 \in \{1, ..., N\}$ such that for $r = 1, ..., R$ we have
(a) $\lambda^*(i) > 0$ for all $i \leq i_0$, and $\lambda^*(i) = 0$ for all $i > i_0$.
(b) For all $i \leq i_0$ we have

$$P_{tr}(1,\lambda^*) = \frac{u_1}{u_i} P_t r(i,\lambda^*),$$

(c) $P_{tr}(i_0,\lambda^*) \geq \frac{u_{i_0}}{u_{i_0+1}}$,
(d) $\kappa_r^* = D_r(u^*)$ where

$$u^* = u_1/P_{tr}(1,\lambda^*).$$

Let us comment on the Properties (a)-(d) of Proposition 3, and compare Proposition 3 with Proposition 1 of the previous section.
(a) states that a range of priorities classes are used in equilibrium, and priority 1 is always used. This in contrast to the result in the previous section where at most two priorities are used in equilibrium.
(b)+(c) states the probability that a submitted packet in priority $i \leq i_0$ gets transmitted is given by

$$P_{tr}(i,\lambda^*) = \frac{u_i}{u_1} P_{tr}(1,\lambda^*).$$

Furthermore, for priority i_0 we have that

$$P_{tr}(i_0,\lambda^*) \geq u_{i_0}/u_{i_0+1}.$$

In contrast with Proposition 1, here the transmission probabilities (for the the priorities $i < i_0$) do not only depend on pricing, but also on the transmission probability for priority i_0. Therefore, pricing can no longer be used to provide absolute QoS guarantees. However, pricing can still be used to set the ratio between the transmission probabilities in the different priorities, and to provide relative QoS guarantees.
(d) states the average throughput of user r is

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

where

$$u^* = \frac{u_{i_0}}{P_{tr}(i_0,d^*)}.$$

As in the previous section, this result suggests that the pricing scheme is fair in the sense that all users pay the same (congestion) price u^*.

Similar to the previous section, we can use Proposition 3 to derive sufficient conditions for an $(\lambda_1^*, ..., \lambda_R^*)$ to be an equilibrium allocation; which can then be used to construct an equilibrium allocation and to obtain the following result.

Proposition 4: Let Assumption **??** and 2 hold, then there always exists an equilibrium allocation.
For the case of bursty traffic, we were not able to derive a simple condition under a unique equilibrium allocation exists.

Proof: Here we only give the outline of a proof. One can construct an equilibrium allocation by showing (using monotonicity) that there exists a $\Lambda^* \geq 0$, such that

$$\sum_{i=1}^{N} \lambda(i,\Lambda^*)\Big(1 - P_{loss}(\Lambda^*)\Big)\frac{u_i}{u_1} = \sum_{r=1}^{R} D_r\left(\frac{u_1}{1 - P_{loss}(\Lambda^*)}\right)$$

where

$$\lambda(i,\Lambda^*) = P_{tr}^{-1}\left((1-P_{loss}(\Lambda^*))\frac{u_i}{u_1}\right) - \ldots \quad (3)$$
$$-P^{-1}\left((1-P_{loss}(\Lambda^*))\frac{u_{i+1}}{u_1}\right), \quad i=1,...,N-1,$$
$$\lambda(N,\Lambda^*) = P_{tr}^{-1}\left((1-P_{loss}(\Lambda^*))\frac{u_N}{u_1}\right), \quad (4)$$

and $P_{tr}^{-1} : \Re_+ \mapsto \Re_+$ is the inverse of $1 - P_{loss}(\cdot)$. Let

$$\Lambda(i,\Lambda^*) = \sum_{k=i}^{N} \lambda(i,\Lambda^*).$$

When Eq. (3) and (4) hold, then we have for $i = 1, ..., N$,

$$1-P_{loss}(\Lambda(i,\Lambda^*)) = \Big(1-P_{loss}(\Lambda^*)\Big)\frac{u_i}{u_N}, \quad \text{if } \lambda(i,\Lambda^*) > 0,$$

and

$$u_i \geq \frac{u_1}{\Big(1-P_{loss}(\Lambda^*)\Big)}, \qquad \text{otherwise.}$$

It can then be verified that when the above conditions hold, then the following rate vector is an equilibrium allocation: for $i = 1, ..., N$ set

$$\lambda_r^*(i) = \lambda(i,\Lambda^*)\frac{D_r\left(u_1/\Big(1-P_{loss}(\Lambda^*)\Big)\right)}{\sum_{l=1}^{R} D_l\left(u_1/\Big(1-P_{loss}(\Lambda^*)\Big)\right)}.$$

■

V. Users with QoS-Sensitive Traffic

For the traffic and link model of Section II, we showed in [7] how the analysis can be extended to users who measure QoS not only as a function of the average throughput, but also the probability that packets get lost. A similar analysis can be done for the traffic and link model for bursty traffic. We model this situation by adding into the users objective function a penalty term which is function of the transmission probabilities P_{tr}.

Consider a fixed user r with utility function U_r. In addition, assume that $\lambda_r = (\lambda r(1), ..., \lambda r(N))$ is the rate vector of user r and $\lambda = (\lambda(1), ..., \lambda(N))$ is the aggregated rate vector over all users. We then associate with user r the total utility

$$U_r(\kappa_r) - \sum_{i=1}^{N} c_r\Big(P_{tr}(i,\lambda)\Big)\lambda_r(i)$$

where

$$\kappa_r = \sum_{i=1}^{N} P_{tr}(i,\lambda_r)\lambda_r(i),$$

and the function $c_r : [0,1] \rightarrow \Re_+$, $r = 1, ..., R$, is non-increasing, twice continuously differentiable (with bounded first and second derivatives), and convex, with

$$c_r(1) = 0.$$

The function c_r models the QoS requirements of user r with respect to packet loss. When c_r is equal to 0 for all $\rho \geq 0$, then we obtain the utility function of the previous section where users only cared about throughput. A user with with tight QoS requirements for packet loss can be modeled by letting the function $c_r(\rho)$ increase sharply when ρ falls below 1.

Under some additional technical assumptions (which capture the case were many small users share the link, i.e. the demand $D_r(u_1)$ of each user $r = 1, ..., R$ is small compared with the link capacity C), we can analyze this situation and obtain results that are the analog to the ones in [7] for the case with smooth traffic. In particular, we have that there always exists an equilibrium allocation and pricing can be used to provide relative QoS. However for the case where users also measure QoS as a function of the packet loss, it is no longer true that all users see the same (congestion) price u^*: but users with with tight QoS requirements will pay more, to get a better QoS.

VI. Case Study

In this section, we use a case study to illustrate the results of the previous section. The purpose of this study is the following.
1. The analysis of the previous section was based on a simple traffic and link model. Using a $M/M/1$ priority queue with a finite buffer for the case study, we want to study the predictive power of the results of Section IV for more realistic traffic and link models.
2. The way we modeled the transmission probabilities of a burst is very simple and one would expect it to lead to accurate results only for the case where the traffic is very bursty, i.e. the burst size B is large compared with the link capacity. Here, we want to illustrate this point.

A. Traffic and Link Model

We use the following set-up for our case study. Consider a single link with a finite buffer which supports N priorities. There are R users sharing the link. Users submit packets in the different priorities according to independent Poisson processes, where $\lambda_r = (\lambda_r(1), ..., \lambda_r(N)$ is the rate vector of customer $r \in \{1, ..., R\}$. Let $m = 1, 2, ...$ be a scaling factor, and assume that packet service times are independently and identically distributed according to an exponential distribution with rate $m\mu$. The buffer holds up to m packets (including the packet in service). Note that for small m, the packet size is large compared with the link capacity and the buffer size; for large m, the packet size is small.

We model that link as a priority queue with preemption (see [1]). There are two types of priority operations for the above link model: priority for admittance/space (with preemption/pushout) and priority for service (also with preemption); where the space-priority classification and the delay-priority classification are the same and within each priority service order is FIFO. Priority for admittance works as follows. If a packet of priority i arrives to a full buffer, and there are no packets of priority $j < i$ in the buffer, then the arriving packet is lost. If a packet of priority i arrives to a full buffer, and there are packets of priority $j < i$ in the buffer, then the arriving packet joins the buffer by pushing out one of the lowest priority packets from the buffer. The arriving packet then joins the tail of the priority i subqueue. Priority for service is as follows. The server always works on the packet at the head of the highest-priority subqueue.

If a packet of even higher priority arrives, the current service is preempted. The packet whose service is preempted is not lost, and the partial work already done on it is not wasted. There are two ways for a packet to be "lost" or "dropped": if the packet arrives to a full buffer when there are no lower priority packets to push out; or if the packet was admitted to the buffer when it first arrived, but it is subsequently pushed out by a higher priority arrival. If the packet is not lost/dropped, it will eventually be "transmitted". Note that similar to the model of Section IV, in this model packets of lower priority are invisible to packets of higher priority.

To derive the probability $\tilde{P}_{tr}(i,\lambda)$ that a burst in priority i gets transmitted for the above link model, we use the following result for a $M/M/1$ queue (see [1]). Consider a $M/M/1$ queue with a queue size m. When the burst arrival rate is equal to $\Lambda \in \Re_+$ and the burst service time is equal to $m\mu$, then a burst is lost (because the queue is full) with probability

$$P_{loss}(\Lambda) = \left(\frac{\Lambda}{m\mu}\right)^m \frac{1-\frac{\Lambda}{m\mu}}{1-\left(\frac{\Lambda}{m\mu}\right)^{m+1}}, \qquad \Lambda \in \Re_+.$$

Note that the rate at which bursts are lost is equal to $P_{loss}(\Lambda)\Lambda$.

Using this result, we obtain for the above priority queue that the the rate at which packets are dropped in priority i and higher is equal to

$$P_{loss}(\Lambda(i))\Lambda(i) = P_{loss}(\Lambda(i))\Big(\Lambda(i+1)+\lambda(i)\Big).$$

Note that the loss rate $P_{loss}(\Lambda(i))\Lambda(i)$ includes losses of priority i bursts, and losses of bursts in priority $i+1$ and higher. Let $\tilde{P}_{tr}(i,\lambda)$ be the probability that a burst in priority i gets transmitted, then we can rewrite the loss rate $P_{loss}(\Lambda(i))\Lambda(i)$ for $i = 1,...,N-1$ as follows,

$$\begin{aligned} P_{loss}(\Lambda(i))\Lambda(i) &= P_{loss}(\Lambda(i))\Big(\Lambda(i+1)+\lambda(i)\Big) \\ &= P_{loss}(\Lambda(i+1))\Lambda(i+1) + ... \\ &\quad +\Big(1-\tilde{P}_{tr}(i,\lambda)\Big)\lambda(i), \end{aligned}$$

and for $i = N$ we obtain

$$P_{loss}(\Lambda(N))\Lambda(N) = \Big(1-\tilde{P}_{tr}(N,\lambda)\Big)\lambda(N).$$

Using this equation, we obtain the following expressions for probability $\tilde{P}_{tr}(i,\lambda)$ that a burst gets transmitted in priority $i = 1,...,N-1$. For $\lambda(i) > 0$, we have

$$1-\tilde{P}_{tr}(i,\lambda) = \frac{\Big(P_{loss}(\Lambda(i)) - P_{loss}(\Lambda(i+1))\Big)\Lambda(i+1)}{\lambda(i)} + ... + P_{loss}(\Lambda(i+1)),$$

and, for $\lambda(i) = 0$, we have

$$\begin{aligned} 1-\tilde{P}_{tr}(i,\lambda) &= P'_{loss}(\Lambda(i+1))\Lambda(i+1) + ... \\ &\quad + P_{loss}(\Lambda(i+1)), \end{aligned}$$

where $P'_{loss}(\cdot)$ is the derivative of $P_{loss}(\cdot)$.

Finally, for $i = N$, we have that

$$1-\tilde{P}_{tr}(N,\lambda) = P_{loss}(\Lambda(N)).$$

Again, we assume that users get charged a price u_i for each burst that they submit in priority i.

B. Case Study

We use the following set-up for the case study. We set the number of priorities to $N = 8$, and the number of users to $R = 5$. Furthermore, we choose the same utility function U_r for all users $r = 1,...,5$, where the corresponding demand functions is given by

$$D_r(u) = \begin{cases} 35 - 0.001u^2 & 0 \le u \le \sqrt{\frac{35}{0.001}} \\ 0 & \text{u} > \sqrt{\frac{35}{0.001}} \end{cases}$$

We set the prices for the different priorities are as follows,

$$u_1 = 1, \quad \text{and } u_{i+1} = 2u_i, \quad i = 1,...,N-1.$$

We then wish to identify an equilibrium allocation $(\lambda_1^*,...,\lambda_R^*)$ such that for $r = 1,...,R$, we have

$$\lambda_r^* = \arg\max_{\lambda_r \in \Re_+} \left\{ U_r(\kappa_r) - \sum_{i=1}^N \lambda_r(i)u_i \right\} \qquad (5)$$

where

$$\kappa_r = \sum_{i=1}^N \tilde{P}_{tr}(i,\lambda^*)\lambda_r(i)$$

and $\lambda^* = \sum_{r=1}^R \lambda_r^*$. To compute an equilibrium, we use an iterative algorithm where we update the allocation $\lambda_r^{(k)}$ at the kth step according to

$$\lambda_r^{(k+1)}(i) = \left[\lambda_r^{(k)}(i) + \gamma\Big(U_r'(\kappa_r^{(k)})\tilde{P}_{tr}(i,\lambda^{(k)}) - u_i\Big)\right]^+,$$

where

$$\kappa_r^{(k)} = \sum_{i=1}^N \tilde{P}_{tr}(i,\lambda^{(k)})\lambda_r^{(k)}(i),$$

$$[x]^+ = \max\{0,x\}, \quad x \in \Re,$$

and $\gamma > 0$ is a small stepsize parameter. When, for $r = 1,...,R$, we have that

$$\lambda_r^{(k)}(i) = \lambda_r^{(k+1)}(i), \quad i = 1,...,N,$$

then $(\lambda_1^{(k)},...,\lambda_R^{(k)})$ is an equilibrium allocation as defined by Eq. 5 (to see this, note that, for a given $\lambda^* \in \Re_+$, the objective function in Eq. 5 is concave in λ_r). Note, that the analysis of the previous section does not guarantee that such an equilibrium exists, nor that the above iterative algorithm converges.

We implemented the above algorithm for $m = 5, 10, 20, 50, 100$ and $\gamma = 0.01$. In these cases, the algorithm converged to an equilibrium. Furthermore, different values for the initial allocation $(\lambda_1^{(0)},...,\lambda_R^{(0)})$ lead to same equilibrium. Table I gives the equilibrium allocations for the different values of m, and Table II gives the transmission probabilities under an equilibrium.

We make the following observations. For $m = 5$, the equilibrium is as predicted by Proposition 3, i.e. we have that all priorities $i \le i^* = 6$ are used. Furthermore, we have

$$\begin{aligned} P_{tr}(i,\lambda^*) &= \frac{u_i}{u_{i+1}} P_{tr}(i+1,\lambda^*) \\ &= \frac{1}{2} P_{tr}(i+1,\lambda^*), \qquad i = 1,...,5, \end{aligned}$$

$$P_{tr}(i^*,\lambda^*) \ge \frac{u_{i^*}}{u_{i^*+1}} = \frac{1}{2}.$$

TABLE I
EQUILIBRIUM ALLOCATION.

m	5	10	20	50	100
$\lambda^*(1)$	2.13	0.466	2.13	2.10	0
$\lambda^*(2)$	0.88	1.81	0.26	0	0
$\lambda^*(3)$	1.70	3.02	1.85	0	0
$\lambda^*(4)$	1.01	0.27	0.76	0	0
$\lambda^*(5)$	1.77	0	0	0	0
$\lambda^*(6)$	3.94	9.84	21.26	54.06	106.02
$\lambda^*(7)$	0	0	0	0	0
$\lambda^*(8)$	0	0	0	0	0

TABLE II
TRANSMISSION PROBABILITIES.

m	5	10	20	50	100
$\tilde{P}_{tr}(1, \lambda^*)$	0.0288	0.0288	0.0288	0.0288	0.01
$\tilde{P}_{tr}(2, \lambda^*)$	0.0575	0.0575	0.0575	0.0575	0.01
$\tilde{P}_{tr}(3, \lambda^*)$	0.115	0.09	0.115	0.06	0.01
$\tilde{P}_{tr}(4, \lambda^*)$	0.23	0.23	0.23	0.06	0.01
$\tilde{P}_{tr}(5, \lambda^*)$	0.46	0.46	0.28	0.06	0.01
$\tilde{P}_{tr}(6, \lambda^*)$	0.92	0.92	0.92	0.92	0.93
$\tilde{P}_{tr}(7, \lambda^*)$	1	1	1	1	1
$\tilde{P}_{tr}(8, \lambda^*)$	1	1	1	1	1

For $m = 100$, the equilibrium is as predicted by Proposition 1, i.e. for $i^* = 6$, we have that $\lambda^*(i) = 0$, $i \neq i^*$, and $P_{tr}(i^*, \lambda^*) \geq \frac{1}{2}$. For $m = 10, 20$, and 50, neither Proposition 1, nor Proposition 3, accurately predict the equilibrium. However, similar to Proposition 3, we have for $\lambda^*(i) > 0$ and $\lambda^*(i) > 0$, $i < j$, that

$$P_{tr}(i, \lambda^*) = \frac{u_i}{u_j} P_{tr}(j, \lambda^*) = \left(\frac{1}{2}\right)^{j-i} P_{tr}(j, \lambda^*).$$

This case study suggests that the traffic model of Section III captures the case where traffic is very bursty (i.e. the burst size tends to be large compared with the link capacity and buffer size); and that the fluid-flow model of Section II captures the case where traffic is very smooth (i.e. burst size tends to be small compared with the link capacity and the buffer size). Neither model is able to accurately capture the case where the traffic is neither very bursty and nor very smooth. Nevertheless, by characterizing the two limiting cases, the above analysis is able to provide insight about how users will allocate their traffic in equilibrium.

VII. CONCLUSION

We modeled DiffServ as a priority service. Users are free to choose the priorities they attach to their traffic, but are charged accordingly. Assuming that many users share a single link, we showed that there always exists an equilibrium for the corresponding non-cooperative game. Furthermore, we showed that pricing can be used to provide relative QoS guarantees, i.e. the factor by which the transmission probability is decreased by moving from priority $i + 1$ to i is lower bounded by the price ratio u_i/u_{i+1}.

For our analysis, we used a simplistic link model. We illustrated through a case study that the model is accurate enough when the (traffic) bursts tend to be large (compared with the buffer size).

For future work, one could carry out the above analysis for a more realistic link model than the one of Section III. Also, it would be interesting to consider the situation where users dynamically change their allocation, and to determine whether the allocations indeed converge to an equilibrium allocation. In particular, one could investigate the convergence properties of the algorithm in Section VI. Finally, it would be interesting to study the above priority scheme for the network case. However, at the moment, it is not clear to what extend this is possible.

ACKNOWLEDGMENTS

The authors thanks the reviewers of the paper: their suggestions and comments were very helpful, and have been included into the final version of the paper. The author also would like to acknowledge the suggestions and comments of Frank Kelly, Richard Gibbens, Damon Wischik, Stephen Turner (Statistical Laboratory, University of Cambridge, UK); Peter Key, Laurent Massoulie, Koenraad Laevens (Microsoft Research, Cambridge, UK); Stewart Lee, Emanuele Giovannetti (Center for Communications Systems Research, Cambridge, UK); and Bart Van Doorselaer (Alcatel-Bell, Antwerp, Belgium).

REFERENCES

[1] D. P. Bertsekas and R. Gallager, *"Data Networks"*, Prentice-Hall, Englewood-Cliffs, New Jersey, 1992.
[2] S. Chen and K. Park, *"An Architecture for Noncooperative QoS Provision in Many-Switch Systems,"* Proc. IEEE INFOCOM '99, 1999.
[3] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, *"Pricing in Computer Networks: Motivation, Formulation, and Example,"* IEEE/ACM Transactions on Networking, 1, 1993.
[4] R. Gibbons, Game Theory for Applied Economists, Princeton University Press, 1992.
[5] A. Gupta, D. O. Stahl, and A. B. Whinston, *"A Stochastic Equilibrium Model of Internet Pricing,"* Journal of Economics Dynamics and Control, 21:697–722, 1997.
[6] IETF, *"Differentiated Services Working Group,"* http://www.ietf.org/html.charters/diffserv-charter.html.
[7] P. Marbach, *"Pricing Priority Classes in a Differentiated Services Network"*, Proceedings of the 37th Annual Allerton Conference on Communications, Control, and Computing, 1999.
[8] A. Odlyzko, *"Paris Metro Pricing for the Internet"*, http://www.research.att.com/~amo, 1998.
[9] K. Park, M. Sitharam, and S. Chen, *"Quality of Service Provision in Noncooperative Networks: Heterogenous Preferences, Multi-Dimensional QoS Vectors, and Burstiness,"* Proc. International Conference on Information and Computation Economies, 1998.
[10] J.B. Rosen, *"Existence and Uniqueness of Equilibrium Points for Concave N-Player Games"*, Econometrica, 33(3):520-534, July 1965.
[11] S. Shenker, *"Fundamental design issues for the future Internet"*, IEEE J. Selected Areas Comm., 13 (1995), 1176-1188, 1995.