

Congestion pricing and user adaptation

Ayalvadi Ganesh, Koenraad Laevens

Microsoft Research

1 Guildhall Street, Cambridge CB2 3NH, UK

Richard Steinberg

Judge Institute of Management Studies

University of Cambridge, Cambridge CB2 1AG, UK

Abstract—The problem of sharing bandwidth in a communication network has been the focus of much recent research aimed at guaranteeing an appropriate quality of service to users. This is particularly challenging in an environment with a great diversity of users and applications, which makes it difficult, if not impossible, to tightly constrain user attributes and requirements. This motivates shifting the burden of rate allocation from the network to the end-systems. We propose a decentralized scheme for user adaptation and study its dynamics. The proposed scheme uses congestion prices as a mechanism for providing both feedback and incentives to end-systems.

Keywords—Pricing, quality of service, congestion control, Internet

I. INTRODUCTION

Quality of service (QoS) provisioning in the Internet has attracted considerable interest in recent years. One of the earliest attempts to address this problem was Intserv [2], which proposed the reservation of bandwidth and buffers along the route for each connection which required a QoS guarantee. The approach of flow based reservations and service provisioning has spawned a considerable body of research both in the ATM and the Internet communities (see, e.g., [5], [14]). However, there are doubts about the scalability of Intserv because it requires network routers to maintain per-flow state, at least for connections requiring a QoS guarantee.

This has prompted interest in an alternative proposal called Diffserv [1], which is based on distinguishing a small number of service classes and using mechanisms such as priority queueing at the core routers to provide a differentiated service. The core network thus needs to allocate resources only on a per class rather than a per flow basis. It is possible to guarantee QoS if each connection conforms to pre-agreed traffic specifications. For example, Parekh and Gallager [11] have shown that, if the input traffic is shaped by leaky buckets, then the network can guarantee delay bounds by using weighted fair queueing at the switches. The task of ensuring conformance can be moved from the core network to the edges, which typically operate at lower speeds and aggregate fewer connections, thus making it feasible for them to police traffic on a per-flow basis. A number of open questions remain, such as how many traffic classes to have, how to define them and how to set appropriate charges for each class.

The approach studied in this paper is different and does not require the explicit definition of service classes. It is based on a simple and innovative idea explicated in Kelly, Maulloo and Tan [7] and Gibbens and Kelly [4], which is to mark each packet entering a switch during a ‘congestion epoch’. Marks reflect the fact that these packets imposed a cost in terms of delay or loss on some other packet. End users are informed of whether their

packet was marked when acknowledgements flow back to them and are then free to decide how to adapt their transmission rates. One way to implement marking would be to use the ECN bit to carry the mark. In practice, it may be necessary to associate monetary charges with marks in order to create the incentive for users to respond to marks. A variant of this scheme has been studied by Low and Lapsley [8].

In this paper, we consider a modification of the above proposal wherein the switch assigns each packet a price rather than a mark. The price is a real number as opposed to a mark which is one bit. This assumption is to simplify the analysis; we believe a small number of bits of price feedback will suffice but this needs to be validated by simulation. In a network, the price for a route is set as the sum of the prices at each of its links. In other words, when a packet traverses several switches on its route, each switch adds its price to the price currently carried by the packet. The price reflects the degree of congestion encountered by the packet and end users are informed of how much they were charged when their packets are acknowledged. They are free to use this information to adapt their sending rates in any manner they see fit. The fact that prices reflect the ‘social cost’ imposed by a packet on other packets ensures that this mechanism provides users with the right incentives to adapt to congestion.

The price mechanism studied in this paper differs in an important respect from a number of earlier proposals, which were based on pricing individual flows (see, e.g., [6], [12]). It does not require the core network to maintain per flow state. The task of associating prices with end users is pushed to the network edges where it is easier to deal with. We propose a mechanism for users to respond to price information which is motivated by the assumption that they are trying to maximise their individual utilities. A related model has been studied by Massoulié and Key [10] in the limit where the number of users increases to infinity. In this asymptotic regime, the price turns out to be constant and so the model does not capture dynamics. In this paper, we deal with a finite user population and study the dynamics of the network corresponding to our model of user behaviour.

It is worth emphasizing that our pricing model does not preclude Diffserv, for example. On the contrary, it can provide guidance on how to price different service classes. We have proposed a mechanism for the core network to assign each packet the marginal cost it imposes on other users. It is not essential that this charge be passed on in its entirety to end users. A plausible service model is one of intermediaries absorbing these charges while charging end users on the basis of a specified ser-

vice class as in Diffserv. The per packet charges considered here might then provide a basis for setting prices for different service classes.

An important sense in which pricing schemes differ from mandated user adaptation is that users are free to choose how to respond and the price mechanism should be designed to encourage appropriate user behaviour. In [7], [4], the authors analyze a mandated scheme for users to respond to prices and show the system converges to the social optimum. They suggest that this should be true even when user responses aren't mandated. In this paper, we establish convergence to the social optimum when user responses are modelled as being motivated by self-interest rather than by mandate. Our user model is admittedly very simple and it remains an open problem to establish similar results for more complex user behaviours.

The rest of the paper is organized as follows. In Section 2 we model users as attempting to maximize a utility function that encapsulates their valuation of bandwidth. While it is arguable that users can be expected to know their value of bandwidth, it is our contention that most models of user choice implicitly maximize some utility function and there is no loss in making this explicit. The question of how to characterize user utilities is not treated here but the results are derived in the context of a fairly general class of utility functions. The interaction between utility maximizing users can be modeled as a game, and motivates the adaptive expectations framework that we study in Section 3. This is a decentralized framework for adaptation by individual users. The users receive price feedback from the network which they use to predict future prices and to modify their transmission rates accordingly. In Section 3, we describe conditions under which such a decentralized scheme causes the system to converge to a socially optimal allocation of bandwidth; we also determine the speed of convergence. In Section 4 we present results of simulations in which the model assumptions are relaxed, and conclude in Section 5.

II. INDIVIDUAL OPTIMIZATION

We consider a discrete time model of a single link shared by N users. In each time slot n , user i transmits a quantity $x_i(n)$ of data packets on to the link. The unit price of bandwidth in a time slot is determined as a function of the aggregate data arriving on the link in that slot; thus

$$p(n) = \phi(x(n)), \quad \text{where } x(n) = \sum_{i=1}^N x_i(n),$$

and ϕ is a given non-decreasing function. User i derives a utility $u_i(x_i(n))$ in time slot n , which is a non-decreasing function of the bandwidth it uses in that slot (the number of packets it transmits in that time slot). The users' total utility is assumed to be the sum of their utilities in each time slot. User i seeks to maximize

$$V_i(\mathbf{x}) = u_i(x_i) - x_i\phi(x),$$

where $\mathbf{x} = (x_1, \dots, x_N)$ is the vector of bandwidth demands, $x = x_1 + \dots + x_N$ denotes the aggregate demand and $\phi(x)$ the corresponding unit price for bandwidth. We thus have a model

of a game among the users, where $V_i(\mathbf{x})$ denotes the single-stage payoff to user i as a function of the actions of all the users. It is usual in game theory to assume common knowledge of the utility functions of all the players, but such an assumption appears unrealistic in our context where users typically do not even know the number of other players competing with them for resources. Note that all that a player needs to know in order to choose a quantity to send is the price in that time slot. A detailed model of how that price is arrived at through the actions of other players is not required.

The only information available to a user when choosing its transmission rate in time slot n is the history of prices $p(n-1)$, $p(n-2), \dots$, where $p(n) = \phi(x(n))$, and the history of its own actions. This suggests the following natural framework for user adaptation. Each user i forms its own estimate, $\hat{p}_i(n)$, of the price in the n^{th} time slot, $p(n)$, based on the information available to it. It then optimizes its transmission rate, $x_i(n)$, based on this estimate. In other words, user i chooses $x_i(n)$ to maximize $u_i(x_i(n)) - x_i(n)\hat{p}_i(n)$. We propose a simple model of how users form expectations of the price and use this to analyse the system.

Ideally, users' beliefs about the price, $p(n)$, should be modeled as probability distributions. In the particular optimization problem considered above, the probability distribution enters only through its expectation (certainty equivalence is valid). This would not be true for risk averse users. Secondly, the optimization problem considered above ignores the impact of the user's own contribution $x_i(n)$ to their price estimate $\hat{p}_i(n)$. If the function ϕ is known, or can be estimated by users, then it is possible for them to incorporate the effect of their own actions. It is possible but cumbersome to carry through the analysis with this modification. Moreover, its practical relevance is questionable. We are interested in networks with large capacity, serving a large number of users simultaneously. The impact of the traffic offered by a single user on the unit price of bandwidth is likely to be small, and therefore it would not be unrealistic for users to ignore this effect.

III. SYSTEM STABILITY

A. User adaptation

The following assumptions will be made in the rest of the paper.

1. The utility functions are of the form

$$u_i(x) = w_i \frac{x^\beta - 1}{\beta}, \quad \beta < 1, \quad w_i > 0, \quad i = 1, \dots, N. \quad (1)$$

If $\beta = 0$, we interpret the formula to mean $u_i(x) = w_i \log x$.

2. The price function is iso-elastic:

$$\phi(x) = \left(\frac{x}{C}\right)^k, \quad (2)$$

where C is a scale parameter which is associated with the physical capacity of the link, and $k \geq 1$ defines the steepness of the penalty for demand in excess of capacity.

3. Players use an exponentially weighted moving average estimator of the price. Specifically, user i 's price estimate in time slot n is given by

$$\hat{p}_i(n) = \hat{p}_i(n - m_i) + \alpha_i [p(n - m_i) - \hat{p}_i(n - m_i)] = \sum_{k=0}^{n_i} (1 - \alpha_i)^k \alpha_i p(n - km_i) + (1 - \alpha_i)^{n_i+1} \hat{p}_i(-j), \quad (3)$$

where m_i denotes the feedback delay for user i , $n_i = \lfloor n/m_i \rfloor$ and $j = m_i + m_i n_i - n$ is in $\{1, \dots, m_i\}$. Player i chooses $x_i(n)$ to maximise $u_i(x) - x \hat{p}_i(n)$, i.e., $x_i(n)$ is the unique solution of

$$u_i'(x_i(n)) = \hat{p}_i(n). \quad (4)$$

The assumed price estimation scheme (3) may appear odd in that price estimates in one time slot do not directly influence those in another time slot unless they share the same modulus relative to m_i . (If users have different feedback delays, there will of course be an indirect influence through the impact that each user's actions have on the price and thus on the choices of other users.) This motivates us to consider the following alternative:

$$\begin{aligned} \hat{p}_i(n) &= (1 - \alpha_i) \hat{p}_i(n - 1) + \alpha_i p(n - m_i) \\ &= \sum_{k=0}^{n - m_i - 1} (1 - \alpha_i)^k \alpha_i p(n - m_i - k) \\ &\quad + (1 - \alpha_i)^{n - m_i} \hat{p}_i(0), \end{aligned} \quad (5)$$

for $n \geq m_i$; if $n < m_i$, we take $\hat{p}_i(n) = \hat{p}_i(0)$.

The application of logarithmic utility functions to communication networks has been considered by Kelly *et al.* [7], for example; they interpret the parameter w_i as a measure of user i 's willingness to pay for bandwidth. Utility functions of the form (1) have been considered by Massoulié and Roberts [9], who provide an engineering interpretation of the welfare maximization problem for certain values of β . The goal of the price mechanism (2) is to keep link utilization close to C , but at the same time to prevent demand exceeding link capacity. The form we consider, which we call the *CPE price mechanism*, corresponds to constant price elasticity, which has empirical support in the context of consumer demand [13].

B. Global stability

In an earlier paper [3], we showed under milder conditions on the utility and price functions that there is a price q^* which is self-consistent in the sense that, if all users had q^* as their price estimate and chose their transmission rates accordingly, then the price of bandwidth would indeed be q^* . If the utility and price functions have the specific form assumed above, there is only one such price, q^* . We also showed that, under the assumed model of user adaptation, the price estimates of individual users converge to q^* provided each α_i is sufficiently small and the users receive instantaneous price feedback ($m = 1$). Moreover, the users can choose α_i suitably if they know k , the steepness parameter in the price function. In this paper, we extend this investigation in two directions. We allow delays in the price

information fed back to users; in practice, a delay of about one round-trip time is to be expected. Second, we use simulations to study the effect of users starting out with a value of α_i which is too large to ensure stability and decreasing it gradually. This mechanism is proposed as an analogue to slow start in TCP for users to quickly reach a neighbourhood of their steady state rate allocation.

We shall assume for ease of analysis that all players have the same feedback delay ($m_i = m$ for all i), start with the same initial price estimates, $\hat{p}_i(-n)$, $1 \leq n \leq m$, and adapt at a common rate ($\alpha_i = \alpha$ for all i), so that they share the same price estimate in every time slot. Our aim is to show that, as n tends to infinity, $p(n)$ and $\hat{p}_i(n)$, for each i , converge to the unique self-consistent price, q^* . The assumption of common price estimates is very strong, but is needed to make the problem analytically tractable. Our intuition is that common expectations exacerbate oscillations and make convergence more difficult; thus, if convergence holds in this setting, it should be expected to hold when expectations are heterogeneous. We validate this intuition through simulations.

Suppose first that users employ the mechanism in (3). Let q_n denote the common value of the price expectations, $\hat{p}_i(n)$, of all the users in time slot n . We have from (4) and (1) that user i 's transmission rate in time slot n is given by $x_i(n) = (w_i/q_n)^{1/(1-\beta)}$, and so the aggregate demand is

$$x(n) = \left(\frac{W_\beta}{q_n} \right)^{\frac{1}{1-\beta}}, \quad \text{where } W_\beta := \left(\sum_i w_i^{1/(1-\beta)} \right)^{1-\beta}.$$

By (2), the actual price in time slot n is seen to be

$$p(n) = \left(\frac{W_\beta}{q_n C_\beta} \right)^{\frac{k}{1-\beta}}, \quad \text{where } C_\beta := C^{1-\beta}.$$

The users' estimate of the price at time $n + m$ is now given by (3) to be

$$q_{n+m} = (1 - \alpha) q_n + \alpha \left(\frac{W_\beta}{q_n C_\beta} \right)^{\frac{k}{1-\beta}}. \quad (6)$$

For the price q^* to be self-consistent, it must be a fixed point of the recursion, (6), i.e., q^* solves

$$q^* = \left(\frac{W_\beta}{q^* C_\beta} \right)^{\frac{k}{1-\beta}}. \quad (7)$$

This can also be shown using (1), (2) and (4). The solution of (7) is unique since the right hand side is a decreasing function of q^* . Let $\delta_n = (q_n - q^*)/q^*$ denote the relative error in the price estimate at time n . Using (6) and (7), we obtain the recursion

$$\delta_{n+m} = f(\delta_n), \quad (8)$$

where

$$f(x) = (1 - \alpha)(1 + x) + \frac{\alpha}{(1 + x)^K} - 1, \quad K := \frac{k}{1 - \beta}. \quad (9)$$

Observe that $f(x) - x$ is convex, $f(0) = 0$,

$$\begin{aligned} f(x) - x &\rightarrow +\infty \quad \text{as } x \rightarrow -1 \\ f(x) - x &\rightarrow -\infty \quad \text{as } x \rightarrow \infty \end{aligned}$$

Therefore f has a unique fixed point at zero. It is also clear from (9) that $f(x) > -1$ for all $x > -1$. Thus, f maps the interval $(-1, \infty)$ into itself and so its iterates f^n are well defined on this interval. We want to find a condition on α such that, starting from any initial condition δ_0 , the iterates, $\delta_{n+1} = f(\delta_n)$, converge to the fixed point, zero.

Lemma 1: If $\alpha \leq \frac{1}{K+1}$, then $f^n(x) \rightarrow 0$ geometrically fast for any initial condition x .

Lemma 2: Let $\alpha \in (\frac{1}{K+1}, 1]$ be such that there exists $\lambda \in [0, 1)$ satisfying $f(f(x)) \leq \lambda^2 x$ for all $x \geq 0$. Then $f^n(x) \rightarrow 0$ geometrically fast, for any initial condition $x \in (-1, \infty)$. Conversely, if $f(f(x)) \geq x$ for any $x > 0$, then there is an initial condition $x \neq 0$ for which $f^n(x)$ does not go to zero as n goes to infinity.

The proofs of these lemmas can be found in [3]. If $f^n(x) \rightarrow 0$ geometrically at rate λ , then, by (8), the relative error of the price estimate, δ_n , goes to zero geometrically at rate $\lambda^{1/m}$. Thus, feedback delays slow convergence to equilibrium substantially if users adapt using (3).

C. Local stability

Is it possible to speed up convergence to equilibrium? This question motivates us to consider (5) as an alternative model of adaptation. We are unable to analytically establish conditions for global stability in this setting. Instead, we derive conditions for local stability and study global convergence via simulation in Section 4B. Proceeding as before, we let q_n denote the value of the price estimate in time slot n , which is the same for all users under our assumptions. The aggregate transmission rate is obtained from (1) and (4) to be

$$x(n) = (W_\beta/q_n)^{1-\beta}.$$

Thus, by (2), the actual price in time slot n is

$$p(n) = \left(\frac{W_\beta}{q_n C_\beta} \right)^K, \quad K = \frac{k}{1-\beta}.$$

Substituting this in (5) yields the recursion,

$$q_n = (1-\alpha)q_{n-1} + \alpha \left(\frac{W_\beta}{q_{n-m} C_\beta} \right)^K. \quad (10)$$

Define $\delta_n = (q_n/q^*) - 1$ to be the relative error in the price estimate at time n . We obtain from (7) and (10) that

$$\delta_n = (1-\alpha)\delta_{n-1} + \alpha \left[(1 + \delta_{n-m})^K - 1 \right]. \quad (11)$$

Clearly, δ_n has a fixed point at zero. Linearising the recursion in a neighbourhood of this fixed point (i.e., for q_n in a neighbourhood of q^*), we get

$$\delta_n \approx (1-\alpha)\delta_{n-1} - \alpha K \delta_{n-m}. \quad (12)$$

The last recursion is stable if the corresponding characteristic polynomial,

$$z^m - (1-\alpha)z^{m-1} + \alpha K, \quad (13)$$

has all its roots strictly within the unit circle in the complex plane. If this condition holds, then the original nonlinear recursions (10), (11) are stable in a neighbourhood of q^* and 0 respectively. In other words, if the price estimates enter a sufficiently small neighbourhood of q^* , then they are guaranteed to converge to q^* . Convergence occurs at a geometric rate which is given by the largest modulus of the roots of the above polynomial.

Lemma 3: Let $m \geq 2$. If $0 < \alpha < 1/((m-1)K)$ then all the zeros of the polynomial (13) lie strictly within the unit circle in the complex plane.

Proof: Define $F(z, \alpha) = z^m - (1-\alpha)z^{m-1} + \alpha K$. If $\alpha = 0$, then the equation $F(z, \alpha) = 0$ has $m-1$ solutions at $z = 0$ and one at $z = 1$. Moreover, the solutions are continuous functions of α . We first show that for small enough $\alpha > 0$, all roots lie strictly within the unit circle. By continuity this is not a problem for the roots that start off at $z = 0$. Using the implicit function theorem, we get

$$\left. \frac{\partial z}{\partial \alpha} \right|_{(z,\alpha)=(1,0)} = - \left. \frac{\partial F/\partial \alpha}{\partial F/\partial z} \right|_{(z,\alpha)=(1,0)} = -(K+1),$$

which is negative. Thus as α increases from zero, the root at $z = 1$ moves inside the unit circle.

Let α_0 be the smallest value of $\alpha > 0$ such that one of the roots of (13) lies on the unit circle. By continuity, for $\alpha \in (0, \alpha_0)$, all roots of (13) lie strictly within the unit circle. Let $e^{i\theta}$ be a root of (13) corresponding to $\alpha = \alpha_0$. We have from (13) that

$$\cos(m\theta) - (1-\alpha_0)\cos((m-1)\theta) + \alpha_0 K = 0 \quad (14)$$

$$\sin(m\theta) - (1-\alpha_0)\sin((m-1)\theta) = 0 \quad (15)$$

Using (15), we can eliminate $(1-\alpha_0)$ from (14). Further simplification using the identity $\cos a \sin b - \sin a \cos b = \sin(b-a)$ yields

$$\frac{1}{\alpha_0 K} = \frac{\sin((m-1)\theta)}{\sin \theta}.$$

It can easily be verified that the maximum value of the right-hand side is $m-1$. Thus the above equation has no solution unless $\alpha_0 K \geq 1/(m-1)$. It follows that $\alpha_0 \geq 1/(K(m-1))$. Recall that for $0 < \alpha < \alpha_0$, the polynomial in (13) has all its roots within the unit circle. This establishes the claim of the lemma. \blacksquare

IV. SIMULATION RESULTS

In this section, we present simulation results to validate the theoretical analysis of the previous sections. We consider a link of capacity $C = 1$ shared by ten users, evolving in discrete time as described in the previous section. The parameter k determining the steepness of the price curve is taken to be 8. Note that all users of the same type, i.e., with the same functional form for utility (same β in (1)) and the same feedback delays, adaptation parameter and initial price estimate, can be aggregated. Thus, our simulation results correspond to a system with ten distinct user types, with the actual number of users being arbitrarily large.

At the beginning of each time slot, each user i chooses the quantity of data, x_i , it is going to transmit on the link in that time slot. The price in the time slot is determined as a function of the aggregate transmission rate, $x = \sum_i x_i$, in that time slot. User i is informed of the price m_i time slots later and employs this information to determine its price estimate and, consequently, its transmission rate $m_i + 1$ time slots later.

If we think of the parameter C not as the physical capacity of the link but as the capacity times some targeted utilization rate, then the network operator can adjust this parameter to achieve a desired performance target described in terms of packet loss or queueing delays, for example.

A. Basic simulation

The users have logarithmic utility functions, i.e., $\beta_i = 0$ for $i = 0, \dots, 9$, with willingness to pay parameters w_i ranging from 0.10 to 0.19 in steps of 0.01. All users share the same feedback delay $m_i = 10$, the same initial price estimates $\hat{p}_i(-9) = \dots = \hat{p}_i(0) = 10$, and update their price using the same parameter $\alpha_i = 0.1$.

The resulting transmission rates are shown in Figure 1, while Figure 2 depicts the actual price as well as the price estimate, which is common to all users. Observe that the prices and transmission rates change only once every ten time slots and so adaptation is slow.

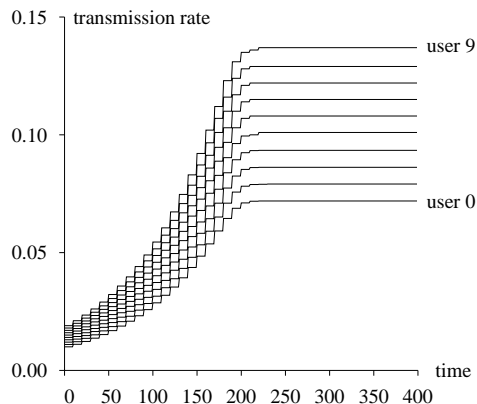


Fig. 1. Simulation A: Transmission rates

The initial price estimate of 10 is pessimistic compared to the equilibrium price, which is 1.39. In practice, we can expect even moderately risk-averse users to start with a pessimistic initial estimate as optimistic initial estimates lead to overcommitting in the first step, resulting in a very high price. Since we have geometric convergence, even if we were to start with an extremely pessimistic guess, say $\hat{p}_i(0) = 100$, the transient would only double.

B. A different model of adaptation

The users have the same utility functions as above and the feedback delay is again 10 time slots for each user. However, the users now employ the estimation mechanism in (5). This results in the value $\alpha = 0.1$ no longer being stable. To ensure

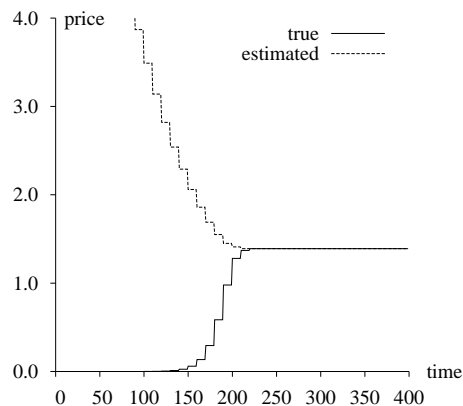


Fig. 2. Simulation A: Price evolution

stability, we take $\alpha_i = 0.02$ for all users. The simulation results, shown in Figures 3 and 4, demonstrate that using (5) instead of (3) enables somewhat faster convergence to equilibrium.

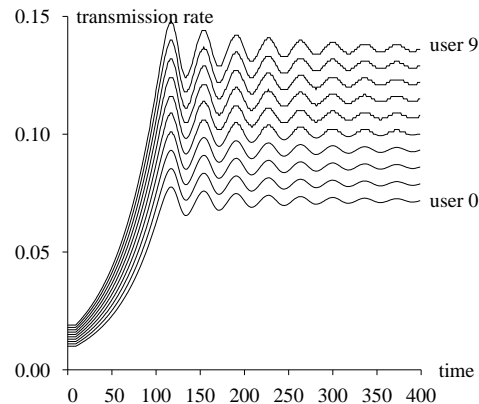


Fig. 3. Simulation B: Transmission rates

C. Heterogeneous price estimates and utility functions

The users now have different initial price estimates, different utility functions and different rates of adaptation. The purpose of this simulation is to show that the stability of the system is not compromised by this heterogeneity. We take $\hat{p}_i(0) = 10/(1+i)$ for $i = 0, \dots, 9$ and β_i ranging from 0 to 0.45 in steps of 0.05 for users 0 through 9 respectively. Once more, the feedback delays m_i are 10 slots. The model of adaptation is given by (5). Since the region of stability for α_i depends on β_i , we take $\alpha_i = 0.02(1 - \beta_i)$, $i = 0, \dots, 9$, which is adequate to guarantee stability. The simulation results, in Figures 5 and 6, show that heterogeneity does not compromise stable convergence to equilibrium.

D. Heterogeneous delays

As illustrated in Figures 7 and 8, heterogeneity in feedback delay does not compromise stability either. The figures are

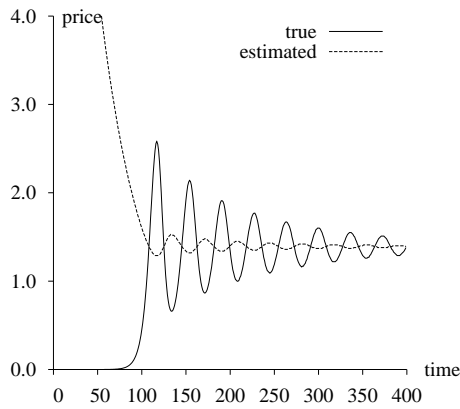


Fig. 4. Simulation B: Price evolution

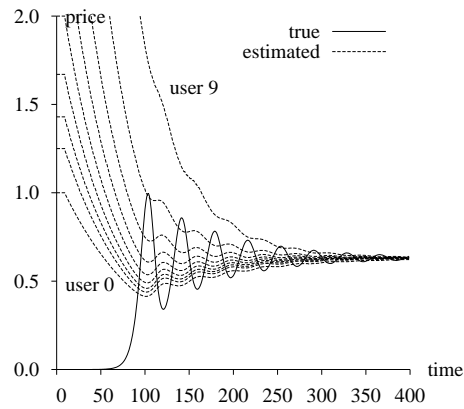


Fig. 6. Simulation C: Price evolution

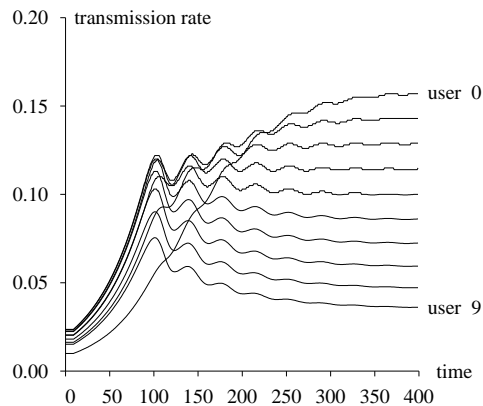


Fig. 5. Simulation C: Transmission rates

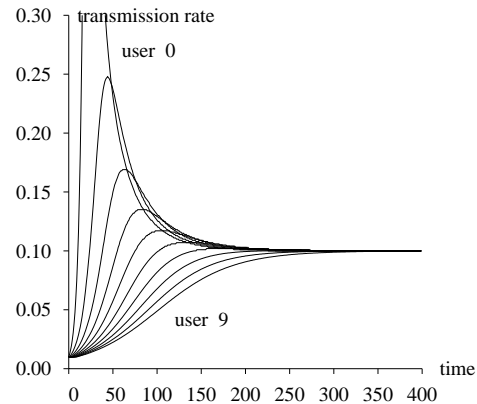


Fig. 7. Simulation D: Transmission rates

based on 10 users all of whose utility functions are logarithmic, with $w_i = 0.1$ and who employ the adaptation mechanism in (5). Their initial price estimate $\hat{p}_i(0)$ is 10. The feedback delays are $m_i = 1 \dots 10$. The adaptation rates α_i are taken to be $0.2/m_i$. Note that, unlike with TCP, the equilibrium bandwidth shares are not influenced by the feedback delays.

E. Time-varying adaptation rate

The last set of simulation results we present illustrates the benefits of using a slow-start-like mechanism. The users start off with adaptation rates $\alpha_i = 0.05$ which are too large to ensure stability. They geometrically decrease their adaptation rates to $\alpha_i = 0.01$, which ensures stability. The rate of decrease is chosen so that it takes about 35 time slots to halve the difference between the initial and final values of α_i . The other parameter values are chosen as in simulation B, which will be the comparator for demonstrating the effect of the slow-start mechanism. Comparing Figures 9 and 10 with Figures 3 and 4 shows that the proposed mechanism does enable users to reach equilibrium more quickly.

V. CONCLUSIONS

This paper was motivated by the problem of providing differentiated quality of service (QoS) to end users in the Internet. We considered the use of congestion prices as a means for achieving a fair and efficient sharing of bandwidth. Users attempt to maximize their individual utility in an environment where prices are determined by the collective actions of all the users. By reflecting the social costs of congestion, such prices enable users to select a grade of service that balances their individual benefit against the costs they impose on other users. Users attempt to predict prices based on their knowledge of the history of the price process, and choose their actions to maximise their utility conditional on their predictions.

We proposed a mechanism for rate adaptation by users in response to price information. The mechanism incorporates delays in the feedback of price information from the network to the users. We showed for this model of adaptation that, under reasonable assumptions, the system converges to an optimal allocation of bandwidth: the users' price predictions converge to the actual price and their bandwidth shares converge to levels which equalize their marginal utility of bandwidth to the price of bandwidth. Finally, we proposed a mechanism akin to slow-

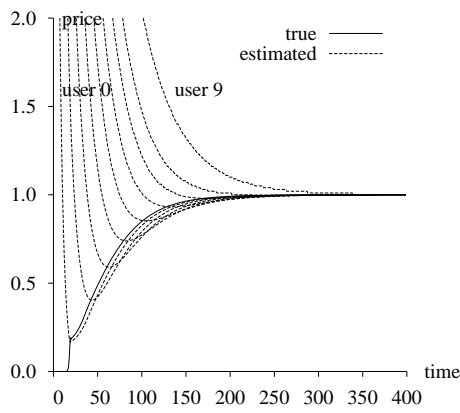


Fig. 8. Simulation D: Price evolution

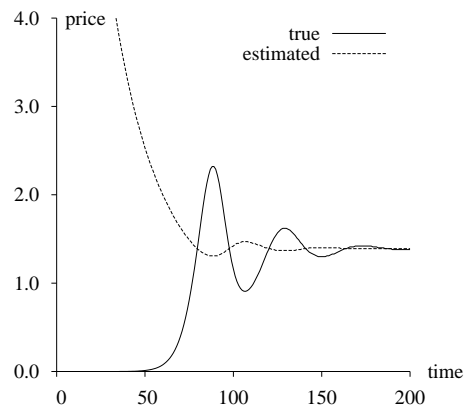


Fig. 10. Simulation E: Price evolution

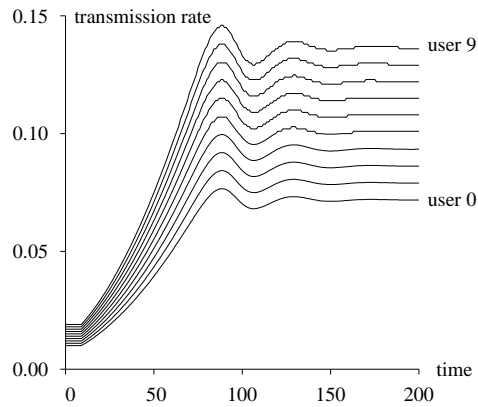


Fig. 9. Simulation E: Transmission rates

start in TCP and showed that it enables users to achieve faster convergence to equilibrium.

REFERENCES

- [1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang and W. Weiss, *An Architecture for Differentiated Services* IETF RFC 2475, 1998.
- [2] R. Braden, D. Clark and S. Shenker, *Integrated Services in the Internet Architecture: an Overview*, IETF RFC 1633, 1994.
- [3] A. Ganesh, K. Laevens and R. Steinberg, "Dynamics of congestion pricing", Microsoft Research Technical Report MSR-TR-2000-70, 2000.
- [4] R. J. Gibbens and F. P. Kelly, "Resource pricing and the evolution of congestion control", *Automatica*, vol. 35, pp. 1969–1985, 1999.
- [5] M. Grossglauser, S. Keshav and D. Tse, "RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic", *IEEE/ACM Transactions on Networking*, vol. 5, pp. 741–755, 1997.
- [6] F. P. Kelly, "Charging and accounting for bursty connections", in *Internet Economics*, Lee W. McKnight and Joseph P. Bailey eds., MIT Press, 1997.
- [7] F. P. Kelly, A. Maulloo and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability", *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [8] S. H. Low and D.E. Lapsley, "Optimization flow control – I: Basic algorithm and convergence", *IEEE/ACM Transactions on Networking*, vol. 7, pp. 861–875, 1999.
- [9] L. Massoulié and J. Roberts, "Bandwidth sharing: objectives and algorithms", *Proceedings of IEEE Infocom*, 1999.
- [10] L. Massoulié and P. Key, "User policies in a network implementing congestion pricing", Preprint, 1999.
- [11] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks: the multiple node case", *IEEE Transactions on Networking*, vol. 2, pp. 137–150, 1994.
- [12] I. Ch. Paschalidis and J. N. Tsitsiklis, "Congestion-Dependent Pricing of Network Services", *IEEE/ACM Transactions on Networking*, vol. 8, pp. 171–184, 2000.
- [13] H. Simon, *Price Management*, New York: North Holland, 1989.
- [14] L. Zhang, S. Deering, D. Estrin, S. Shenker and D. Zappala, "RSVP: A New Resource ReSerVation Protocol", *IEEE Network*, September 1993.