



Working Paper Series

14/2006

Congestion pricing and non-cooperative games
in communication networks

Ganesh, A., Laevens, K. and Steinberg, R.



CAMBRIDGE
Judge Business School

These papers are produced by Judge Business School, University of Cambridge. They are circulated for discussion purposes only. Their contents should be considered preliminary and are not to be quoted without the authors' permission.

Author contact details are as follows:

Ayalvadi Ganesh
Networking Group
Microsoft Research, Cambridge
ajg@microsoft.com

Koenraad Laevens
SMACS Research Group
Ghent University
kl@telin.UGent.be

Richard Steinberg
Judge Business School
University of Cambridge
r.steinberg@jbs.cam.ac.uk

This paper was presented at the
Dagstuhl Seminar # 06132 on "Internet Economics" in Wadern, Germany
and is forthcoming in *Operations Research*

Please address enquiries about the series to:

Research Support Manager
Judge Business School
Trumpington Street
Cambridge CB2 1AG, UK
Tel: 01223 760546 Fax: 01223 339701
E-mail: research-support@jbs.cam.ac.uk

Congestion pricing and non-cooperative games in communication networks

Ayalvadi Ganesh* Koenraad Laevens[†] Richard Steinberg[‡]

April 2006

To appear in *Operations Research*

Abstract

We consider congestion pricing as a mechanism for sharing bandwidth in communication networks, and model the interaction among the users as a game. We propose a decentralized algorithm for the users that is based on the history of the price process, where user response to congestion prices is analogous to “fictitious play” in game theory, and show that this results in convergence to the unique Wardrop equilibrium. We further show that the Wardrop equilibrium coincides with the welfare maximizing capacity allocation.

Subject classifications:

Communications: computer networks. Games: non-cooperative.

*Microsoft Research, 7 J J Thompson Avenue, Cambridge CB3 0FB, UK

[†]Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Ghent, Belgium

[‡]University of Cambridge, Judge Business School, Cambridge CB2 1AG, UK

1 Introduction

The problem of sharing bandwidth among users in a communication network has been the focus of much recent research. For early work, see Kelly, Maulloo, and Tan (1998), Gibbens and Kelly (1999), and Low and Lapsley (1999), and for a recent overview see Srikant (2004). Centralized solutions to this problem are impractical in large networks serving very diverse users. This motivates shifting the burden of rate allocation from the network to the end-systems. We propose a new decentralized scheme for user adaptation.

In this paper, the term *user* refers to an instance of an application like e-mail or web transfer running on a computer (the end-system) connected to the Internet. The majority of applications on the Internet employ TCP (Transmission Control Protocol) to adjust their transmission rates. In TCP, the receiver sends an acknowledgement of each received packet to the sender. The network will drop packets when it is congested. Each sender continually increases its transmission rate until it fails to receive an acknowledgement, whereupon it assumes the network is congested and cuts back its sending rate.

Dropped packets are both inefficient and late as indicators of congestion. This forces users to sharply adapt their transmission rates, resulting in rate oscillation and reduced throughput. An indicator of incipient congestion that avoids packet drops could achieve higher network utilization. This is the motivation of ECN (Explicit Congestion Notification) (Ramakrishnan and Jain (1990), Floyd (1994)), wherein the network routers and switches provide early congestion feedback by marking packets. The marks are returned to the sender with the acknowledgement of the receipt of a packet.

A broader problem with TCP, which is not addressed by ECN, is that it makes

no explicit attempt to discriminate between users on the basis of differing application requirements; nevertheless, there is implicit discrimination on the basis of network characteristics (such as round-trip time for the connection) because of properties of the feedback loop employed by TCP. In practice, users have very different requirements and it would be desirable for the network to provide a differentiated service that is responsive to these different requirements. It is hard to see how to do this in a coherent manner without introducing differentiated charges. Indeed, in the absence of such charges, users would have no incentive to honestly reveal their requirements. By reflecting the social costs imposed by a user, charges can also serve to discourage rate adaptation strategies that may be individually beneficial but socially harmful.

Gibbens and Kelly (1999) have proposed a simple and innovative mechanism to implement usage-based charging. Their scheme is scalable in the sense that it does not require core network routers to keep track of individual source-destination pairs but only of aggregate traffic. Prices are set on the basis of aggregate traffic and communicated periodically to users, who can then decide for themselves how to best satisfy their requirements at the given price. One way to communicate price feedback is by modifying ECN to carry prices instead of marks indicating congestion.

Kelly, Maulloo, and Tan (1998) proposed a scheme for users to individually adapt their rates based on price feedback and showed that, under certain conditions, this mechanism converges in the long run to a socially optimal allocation of bandwidth. Users were modelled as having a utility function that is additively separable over time, and increasing and concave in the instantaneous bandwidth they receive (see Shenker (1995)). Such users have been termed *elastic*. A key

idea in the work of Kelly et al. is to view user adaptation as a distributed gradient ascent algorithm for maximizing social welfare. Consequently, the social welfare function serves as a Lyapunov function for the dynamics. This work was extended by Johari and Tan (2001), who studied the stability of the same dynamics but including feedback delays.

An approach based on solving the *dual* to the welfare maximization problem has been studied by Low and Lapsley (1999). Here, the network adapts prices based on observed aggregate demand, and users attempt to maximize their instantaneous utility based on the price feedback they receive. Users are myopic in that they attempt to maximize their utility without taking account of the likely response of other users to the common price information. Low and Lapsley (1999) show that, if the network adjusts prices sufficiently slowly, the system again converges to the welfare maximizing allocation. Kunniyur and Srikant (2001) considered a model wherein the users adapt their transmission rates on a fast timescale while the network adapts its marking function (which could be interpreted as a price) on a slow timescale. They studied convergence subject to an assumption of separation of timescales.

In this paper, we model the *interaction* between users as a game, and show that a Nash equilibrium of this game coincides (in a large system) with the welfare maximizing solution of the frameworks studied by Kelly, Maulloo, Tan (1998) and Low and Lapsley (1999). We also present a model of *user response* to congestion prices that is very similar to “fictitious play” in game theory, and show that this user behavior results in convergence to the aforementioned Nash equilibrium. In our model, users attempt to maximize instantaneous utility based on their *expectations* of the price, where expectations are formed adaptively based

on the history of the price process. In fictitious play, each player selects a best response to the empirical distribution of actions of his opponents (Fudenberg and Levine (1998)). Here, these actions only impact each player through the price; if, in addition, players are risk neutral, then it suffices that they choose their actions as a best response to a (weighted) average of past prices, as we suggest in this paper. The practical implementation of pricing schemes might involve bandwidth brokers who act as intermediaries between network service providers and end users. Anderson, Kelly, and Steinberg (2006) describe one way to set up such an intermediary service.

The game-theoretic formulation in this paper differs from the work cited above, and is intended to reflect strategies that might be adopted by self-interested users. In order to model feedback delays, we employ a discrete-time formulation. The resulting dynamical system is significantly different from that in Kelly et al. and also that in Low and Lapsley; in particular, it does not possess a natural candidate for a Lyapunov function, and it is not the gradient of any scalar function. (Indeed, Lyapunov functions cannot generally be found for systems with delayed feedback; e.g., Massoulié (2002) and Vinnicombe (2002), who studied extensions of Kelly et al.'s model with feedback delays, rely on techniques other than Lyapunov functions to establish stability. Srikant (2004) also discusses the technical difficulties in proving convergence when feedback is not instantaneous.) Hence, we rely on different techniques based on contraction mappings to prove convergence of the dynamics; furthermore, we show that the unique equilibrium for the dynamics generated by the users' adaptation to prices, starting from arbitrary initial conditions, is the welfare-maximizing allocation, which is also the Wardrop equilibrium of the game; it coincides with the Nash equilibrium in a large system

limit.

The rest of the paper is organized as follows. In Section 2, we model as a game the utility maximization problem of individual users. We also discuss some of the informational issues, make some observations about equilibria, and place our work within the context of well-known models of competition. This framework helps to motivate the *adaptive expectations* approach that we develop in Section 3. In Section 3, we derive conditions under which this model of adaptation leads the users to converge to a socially optimal allocation of bandwidth; in addition, we determine the speed of convergence. The results of this paper in more detail are as follows. Lemma 1 uses Brouwer's fixed point theorem to show *existence* and *uniqueness* of a price that is self-consistent in the following sense: if all the players have it as their *predicted* price, and choose their transmission rate accordingly, then the resulting price coincides with the prediction. We also show how the functional form of the pricing function is connected with queueing phenomena at the link. We then establish conditions that ensure convergence from arbitrary initial conditions to the self-consistent price. In particular, Lemma 2 shows that the price expectations of the players will be bounded and will converge to the (unique) fixed point. As we are taking expectations over the history of the process, we then in Theorems 1 and 2 provide conditions on the associated averaging parameter such that geometric convergence to the fixed point is ensured from any initial condition. Finally, Lemma 3 gives a sufficient and another necessary condition such that the hypothesis of the previous Theorem does indeed hold. In Section 4, we present conclusions along with suggestions for future work. All proofs of lemmas and theorems appear in the appendix.

2 Model

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). 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.

If the utility functions u_i of all players and the price function ϕ are common knowledge, then this is the well-known model of *Cournot competition*; see, for example, Tirole (1998), Section 5.4. A vector of demands is called a *Nash equilibrium* if no player has an incentive to unilaterally change its demand. Thus a Nash equilibrium \mathbf{x} is characterized by the property that for each i , x_i maximizes $V_i(x, \mathbf{x}_{-i})$ over $x \in \mathbb{R}_+$; here \mathbf{x}_{-i} denotes the vector $(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_N)$. Under suitable differentiability assumptions, a necessary condition for \mathbf{x} to be a Nash equilibrium is that

$$u'_i(x_i) - \phi(x) - x_i\phi'(x) = 0 \quad \forall i. \tag{1}$$

This is also sufficient if, for each i , $V_i(\mathbf{x})$ is concave in x_i . There can be multiple Nash equilibria and sufficient conditions for uniqueness of the Nash equilibrium do not appear to be simple or intuitive (Fudenberg and Tirole (1991)).

If the function ϕ is known to users, then they can incorporate the effects of their own actions on the price. However, this effect is typically small in large systems, and it is more common to assume that users choose x_i such that

$$u'_i(x_i) - \phi(x) = 0. \quad (2)$$

Such a vector of x_i is called a *Wardrop equilibrium*. The Wardrop equilibrium approximates the Nash equilibrium for large N (Haurie and Marcotte (1985)).

The motivation for the approach we adopt in this paper is as follows. The assumption that all the utility functions are common knowledge is unrealistic. It is impractical in networks for users to be aware of even the *number* of other users sharing resources with them, let alone knowing all of their utility functions. 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 $u_i(x_i(n)) - x_i(n)\hat{p}_i(n)$. Note that users' beliefs about the price, $p(n)$, should be modelled as probability distributions. However, since we are seeking to maximize $[u_i(x_i(n)) - x_i(n)p(n)]$, the probability distribution on the price $p(n)$ enters only through its expectation (i.e., certainty equivalence is valid). Hence, we are justified in adapting this expectation on the basis of the price history. This would not be true for risk averse users.

As noted earlier, the model we study in this paper corresponds to Cournot competition. Another commonly used model of oligopolistic competition is the Bertrand model, in which each producer of an undifferentiated good sets a price and is willing to supply any demand at that price. Consumers choose the quantity demanded based on the market price, which is obviously the lowest of the prices offered by the different producers. In our context, this model corresponds to each user specifying a price it is willing to pay per unit bandwidth, i.e., per packet. If the user can generate arbitrarily high demand at this price, then the network accepts only the user with the highest bid. More realistically, the demand from each user would be bounded, in which case the network would accept packets from a collection of highest bidders. This corresponds precisely to the smart market model that was proposed by Mackie-Mason and Varian (1996), where users attach a bid to each packet that is the maximum price they are willing to pay for its transmission. The network admits packets in decreasing order of bids until capacity is reached. Each admitted packet is charged the amount of the highest rejected bid (or zero if all packets are admitted). While this scheme is efficient in an economic sense, it is impractical for a number of reasons. One is the sheer difficulty of conducting repeated auctions at the speeds dictated by high-capacity networks. Another issue is the fairness of comparing bids on a packet that has freshly entered the network with one that has been waiting its turn for a long time. The model of congestion pricing we study does not suffer from these problems.

3 User Adaptation and Convergence

The following assumptions will be made in the rest of the paper. The utility functions u_i are increasing, strictly concave and differentiable for each i , while the price function ϕ is strictly increasing, convex and differentiable.

Let $\hat{p}_i(n)$ denote player i 's expectation of the price of bandwidth in time slot n , which is formed, in a manner to be specified, on the basis of all information available to him up to time $n-1$. This information consists of the history of prices $p(k)$, and his own actions $x_i(k)$, for all times k up to and including $n-1$. Player i chooses his action $x_i(n)$ to maximize his payoff conditional on his expected price, $\hat{p}_i(n)$, i.e., he seeks an $x_i(n)$ that achieves

$$\max_x \{u_i(x) - x \hat{p}_i(n)\}.$$

Since u_i is strictly concave and differentiable, $x_i(n)$ is the solution of

$$u'_i(x_i(n)) = \hat{p}_i(n). \quad (3)$$

We define $h_i = (u'_i)^{-1}$, noting that the inverse exists since u'_i is strictly decreasing by the assumption that u_i is strictly concave. If $u'_i(0)$ is finite, then h_i is only defined on the interval $[0, u'_i(0)]$. We extend its definition to $\mathbb{R}_+ \cup \infty$ by setting $h_i(p) = 0$ for all $p > u'_i(0)$. Now $x_i(n) = h_i(\hat{p}_i(n))$, and the actual price in time slot n is given by

$$p(n) = \phi \left(\sum_{i=1}^N x_i(n) \right) = \phi \left(\sum_{i=1}^N h_i(\hat{p}_i(n)) \right). \quad (4)$$

We now ask the following: Is there a price that is self-consistent in the sense that, if all users have it as their predicted price and select their transmission rate accordingly, then the resulting price coincides with the prediction? More precisely,

is there a q^* such that

$$q^* = \phi \left(\sum_{i=1}^N h_i(q^*) \right). \quad (5)$$

The answer is yes, by Brouwer's fixed point theorem (see Varian (1992), for example), as we now show.

Lemma 1 *Suppose the utility functions u_i are strictly concave, increasing and differentiable for each i , and that the price function ϕ is continuous and increasing.*

Let h_i be defined as above. Then, the function $q \mapsto \phi \left(\sum_{i=1}^N h_i(q) \right)$ has a fixed point. In other words, there is a self-consistent price q^ such that $q^* = \phi \left(\sum_{i=1}^N h_i(q^*) \right)$. Moreover, the fixed point q^* is unique. \square*

We have thus established the *existence* and *uniqueness* of a self-consistent price in the sense described above. In general, it is not clear whether the process of user adaptation *converges* to the fixed point from arbitrary initial conditions. We shall now investigate this question in a specific model of adaptation.

Suppose players make use of a one-step error correction mechanism for prediction, which corresponds to using an exponentially weighted moving average estimator. Specifically, let the assumed model of user expectation formation be

$$\begin{aligned} \hat{p}_i(n+1) &= \hat{p}_i(n) + \alpha_i [p(n) - \hat{p}_i(n)] \\ &= \sum_{k=0}^n (1 - \alpha_i)^k \alpha_i p(n - k) + (1 - \alpha_i)^{n+1} \hat{p}_i(0), \end{aligned} \quad (6)$$

where the α_i are given constants, and it is assumed that $1 - \alpha_i \in (0, 1)$ for all i . Our aim is to show that, as n tends to infinity, $p(n)$ and $\hat{p}_i(n)$, for each i , converge to the self-consistent price, q^* .

There are two reasons why players should employ some form of averaging in generating predictions. The first is stability. If players choose their optimal response based on the current price, i.e., $\alpha_i = 1$ for all i , then there are initial conditions for which the dynamics will not converge to the fixed point; see Theorem 2 and Lemma 3 below. Indeed, most work on congestion pricing assumes some form of averaging, either rate smoothing at the users (e.g. Kelly et al. (1998)) or price smoothing in the network (e.g. Low and Lapsley (1999)), as this is required to guarantee stability. The second reason for averaging is to smooth out stochastic fluctuations. We point out later that under suitable assumptions, the packet arrival process in a large system can be modelled as a (inhomogeneous) Poisson process with rate $x(t)$. Thus, the number of packets arriving in a discrete time slot will be a Poisson random variable, and the price in any time slot will also be random, even in a static setting where $x(t) \equiv x$. The objective of the players in this context is to choose their rates in order to maximize utility subject to the mean price; they cannot track the noise in the price process by definition. In practice, of course, the system is not static, and if the α_i are chosen too small, then players won't be able to react fast enough to sudden changes in the network. By characterizing how large α_i can be while ensuring stability, we provide guidance on how to choose α_i in practice.

Assumptions: The *utility functions* are of the form

$$u_i(x) = w_i \frac{x^\beta - 1}{\beta}, \quad \beta < 1, \quad w_i > 0, \quad i = 1, \dots, N. \quad (7)$$

Utility functions of this form are well known in economics, where they are variously termed CRRA (constant relative risk aversion) or CES (constant elasticity of substitution) or “isoelastic utility functions”; see, for example, Blanchard and Fisher (1989). If $\beta = 0$, we interpret the formula to mean $u_i(x) = w_i \log x$.

The *price function* is assumed to be isoelastic

$$\phi(x) = \left(\frac{x}{C}\right)^k, \quad (8)$$

where C is a scale parameter that 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. We assume for simplicity that all players adapt their price expectations at the same rate, that is, $\alpha_i = \alpha$ for all i .

Logarithmic utility functions date back to Daniel Bernoulli's explanation of the St. Petersburg Paradox; see Bernoulli (1738). The application of logarithmic utility functions to communication networks has been considered by Kelly et al. (1998) where they interpret the parameter w_i as a measure of user i 's willingness to pay for bandwidth. Utility functions of the form (7) have been considered by Massoulié and Roberts (1999), who provide an engineering interpretation of the welfare maximization problem for certain values of β . It has been observed that the rate adaptation of TCP approximates the behavior of a user who seeks to maximize such a utility function with $\beta = -1$.

In the Cournot competition analogue of our model, the function ϕ plays the role of a demand function. The form we consider corresponds to constant price elasticity. This function also has a close connection with the loss functions of some queues, as we now show.

Consider first an open-loop queueing system with N independent flows, in which each flow has mean rate x . The *aggregate* arrival process converges to a Poisson process in the following sense: if $A^N(t, u)$ is the total number of packets arriving in the interval (t, u) , then the random process $\tilde{A}^N(u) = A^N(t, t + u/N)$ converges weakly to a Poisson process with rate x (Cao and Ramanan, 2002). They also show that this result carries through to queue size: if $Q^N(t)$ is the queue

size at time t , then the distribution of $Q^N(t)$ converges to that for a queue fed by a Poisson process with arrival rate x and served at constant rate C , in an infinite-buffer system, assuming $x < C$. We expect that this result can be extended to a system with a finite buffer B , and to $x \geq C$. The loss probability for a finite-buffer open-loop queue is thus

$$p = L_B(x/C),$$

where L_B can be calculated by finding the equilibrium distribution of a suitable Markov chain.

Suppose next that the source rates x vary over time, but “slowly.” It is known that $Q^N(t)$ makes excursions of size $O(1)$ in timescale $O(1/N)$ (Raina and Wischik (2005)). Therefore, in any short interval $(t, t + \delta)$, the queue size will repeatedly hit empty and full. This suggests that, if the mean arrival rate $x(t)$ does not change by much in the interval $(t, t + \delta)$, the loss probability is

$$p(t) = L_B(\rho(t)), \tag{9}$$

where $\rho(t) = x(t)/C$. So, for example, if the resource were modelled as an $M/M/1$ queue with service rate C packets per unit time, at which a packet is marked with a congestion signal if it arrives at the queue to find at least B packets already present, this would yield

$$p(x) = \left(\frac{x}{C}\right)^B.$$

Thus, our price function ϕ can be motivated as being proportional to the packet loss probability for a given aggregate arrival rate.

Recall that an objective of the pricing scheme is to avoid packet drops by signalling incipient congestion. Therefore, links should set prices in a manner

that ensures that aggregate demand does not exceed the physical capacity of the link for extended periods of time; this may involve setting C in the price function (8) to be some fraction, such as 90%, of the actual link capacity. The parameter C can be set, for instance, using an adaptive scheme; the idea is to adjust C on a slow time scale in order to achieve a desired trade-off between high link utilization and low packet drop probability. A similar idea has been proposed and analyzed by Kunniyur and Srikant (2001) in a related context.

We shall, for the remainder of the paper, focus on the pricing function (8). So in the context of our model, we can rewrite the fixed point equation (5) as

$$q^* = \left(\frac{W_\beta}{q^* C_\beta} \right)^K \text{ where } K = \frac{k}{1-\beta}, W_\beta = \left(\sum_i w_i^{\frac{1}{1-\beta}} \right)^{1-\beta}, C_\beta = C^{1-\beta}. \quad (10)$$

Note that this equation has a unique solution as the right hand side is a decreasing function of q^* . Let $\hat{p}_i(n)$ denote the price expectation of user i at the beginning of time slot n , and define the weighted average

$$\hat{p}(n) = W_\beta^{-1/(1-\beta)} \sum_{i=1}^N w_i^{1/(1-\beta)} \hat{p}_i(n). \quad (11)$$

We wish to show that, as $n \rightarrow \infty$, each of the users' price expectations $\hat{p}_i(n)$ converges to the unique self-consistent price q^* that solves (10).

We first show that the price expectations of all users remain bounded. We then use this to show that the price expectations of the users converge to a common value, and the common value is the fixed point, q^* .

Lemma 2 *The price expectations of all users remain confined at all times to an interval that is bounded away from zero and infinity. To be precise, let $m = \alpha^{1/K} q^* \min_{i=1}^N (w_i/W_\beta)$. Then, for all $i = 1, \dots, N$, and all $n \geq 1$, $\hat{p}_i(n) \geq (1 - \alpha)m$. Moreover, for any $\delta > 0$ and all n sufficiently large, $\hat{p}_i(n) \leq (1 +$*

$$\delta) q^* \left(\frac{q^*}{(1-\alpha)^m} \right)^K \text{ for all } i. \quad \square$$

Note that the actual price $p(n)$ in time slot n is the common value of the feedback signal received by all users at the end of this time slot, and which they use to form their expectations at the beginning of the next time slot. It is now immediate from (6) that, irrespective of the initial prices $\hat{p}_i(0)$, the expectations $\hat{p}_i(n)$ converge towards each other at a geometric rate. Specifically,

$$|\hat{p}_i(n) - \hat{p}(n)| \leq |1 - \alpha|^n \max_{i,j} |\hat{p}_i(0) - \hat{p}_j(0)| \quad \forall i = 1, \dots, N.$$

Henceforth, we fix $\epsilon > 0$ and consider n sufficiently large that $|\hat{p}_i(n) - \hat{p}(n)| < \epsilon$ for all i .

Recall that, by (3) and (7), user i 's transmission rate in time slot n is given by

$$x_i(n) = \left(\frac{w_i}{\hat{p}_i(n)} \right)^{\frac{1}{1-\beta}} = \left(\frac{w_i}{\hat{p}(n)} \right)^{\frac{1}{1-\beta}} \left[1 + \frac{\hat{p}_i(n) - \hat{p}(n)}{\hat{p}(n)} \right]^{\frac{-1}{1-\beta}}.$$

We showed in Lemma 2 that, for all n sufficiently large, $\hat{p}(n)$ lies in a compact interval bounded away from the origin. Hence, by Taylor's theorem, there is a constant γ such that

$$\left| x_i(n) - \left(\frac{w_i}{\hat{p}(n)} \right)^{\frac{1}{1-\beta}} \left[1 - \frac{1}{1-\beta} \frac{\hat{p}_i(n) - \hat{p}(n)}{\hat{p}(n)} \right] \right| \leq \gamma \epsilon^2.$$

In the following, γ will denote a generic constant, not necessarily the same in each instance. Summing the above equation over i , we find that the aggregate demand satisfies $|x(n) - (W_\beta / \hat{p}(n))^{1/(1-\beta)}| \leq \gamma \epsilon^2$, and so, by (8), the actual price in time slot n satisfies $|p(n) - (W_\beta / (\hat{p}(n) C_\beta))^K| \leq \gamma \epsilon^2$. Consequently, by (6), the mean price estimate at time $n + 1$ satisfies

$$\left| \hat{p}(n+1) - (1-\alpha)\hat{p}(n) - \alpha \left(\frac{W_\beta}{\hat{p}(n) C_\beta} \right)^K \right| \leq \gamma \epsilon^2. \quad (12)$$

Let $\delta_n = (\hat{p}(n) - q^*)/q^*$ denote the error in the mean price estimate at time n relative to the equilibrium price. Using (12) and (10), we obtain the perturbed recursion

$$|\delta_{n+1} - f(\delta_n)| \leq \gamma\epsilon^2, \quad \text{where } f(x) = (1-\alpha)(1+x) + \frac{\alpha}{(1+x)^K} - 1. \quad (13)$$

Observe that $f(x) - x$ is convex, $f(0) = 0$, and that $f(x) - x \rightarrow +\infty$ as $x \rightarrow -1$, while $f(x) - x \rightarrow -\infty$ as $x \rightarrow \infty$. Therefore f has a unique fixed point at zero. It is also clear from (13) 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 δ_n converge to a size ϵ neighborhood of the origin, the fixed point of f . Such a condition is provided in the following two theorems.

Theorem 1 *Suppose $\alpha < \frac{1}{K+1}$ and ϵ is sufficiently small. Then, starting from any initial condition, the relative error, δ_n , of the mean price estimate converges geometrically fast to an ϵ -neighborhood of the origin.*

Theorem 2 *Let $\alpha \in (\frac{1}{K+1}, 1]$ be such that, for every $\epsilon > 0$, there is a $\lambda \in [0, 1)$ satisfying $f(f(x)) \leq \lambda^2 x$ for all $x \geq \epsilon$. Then, for any initial condition, δ_n converges geometrically to an ϵ -neighborhood of the origin. Conversely, if $f(f(x)) > x$ for any $x > 0$, then there is an initial condition for which δ_n stays bounded away from the origin.*

We now derive a simple condition on α that is sufficient for the hypotheses of the above theorem to hold, and a different condition which is necessary.

Lemma 3 Suppose $\alpha \in (0, 1)$, $K > 1$, and that f is defined as in (13). If $\alpha K \leq 1$, then, for any $\epsilon > 0$, there is a $\lambda < 1$ such that $f(f(x)) \leq \lambda^2 x$ for all $x > \epsilon$. On the other hand, if $\alpha(K + 1) > 2$, then there is an $x > 0$ such that $f(f(x)) \geq x$.

Remark: Extensive numerical studies suggest that the condition $\alpha(K + 1) \leq 2$ is both necessary and sufficient for $f \circ f$ to be a contraction, but we have not been able to prove sufficiency. This condition is necessary and sufficient for local stability of the liberalization of the above system around its equilibrium.

4 Conclusions

Congestion pricing has been proposed as a mechanism for reflecting the social costs of congestion to users of the Internet. We considered the behavior of users in the presence of congestion charges, and modelled the resulting interaction as a game. We proposed a model of utility maximizing players in an environment where prices are determined, *ex post*, by the collective actions of all. The players attempt to predict prices based on their knowledge of the history of the price process, and choose their actions to maximize their utility conditional on their predictions. This is analogous to fictitious play in game theory. We found that, under reasonable assumptions, this model of adaptation leads the system to an efficient allocation: players' price predictions converge to the actual price and their bandwidth shares converge to levels that equalize their marginal utility of bandwidth to the price of bandwidth.

A natural next step would be to extend the model and analysis to a general network and, in particular, to model heterogeneous delays and user-specific smoothing parameters for price prediction. It would also be interesting to examine the

effect different pricing functions may have on the interaction between users. Such functions, which also have their genesis in queueing phenomena, serve as templates for “resource design,” which in turn help understand and meet Quality of Service (QoS) requirements. Thus, their impact on network performance and user behavior would be of interest to network service providers as well as software vendors.

Acknowledgements

We are grateful to Frank Kelly and Gaurav Raina for helpful comments. Richard Steinberg acknowledges support in part from the Communications Research Network (CRN), a part of the Cambridge-MIT Institute (CMI).

Appendix

PROOF OF LEMMA 1: Define $f: \mathbb{R}_+^N \rightarrow [0, 1]^N$ by

$$f_i(p_1, \dots, p_N) = \frac{p_i}{1 + p_i}, \quad i = 1, \dots, n,$$

and $\Psi: \mathbb{R}_+^N \rightarrow \mathbb{R}_+^N$ by

$$\Psi_i(p_1, \dots, p_N) = \phi \left(\sum_{i=1}^N h_i(p_i(n)) \right), \quad i = 1, \dots, N.$$

If Ψ has a fixed point, then it must necessarily be of the form (q^*, \dots, q^*) with N identical components, and it is clear that q^* solves (5). Now f is invertible and it is not hard to see that $f \circ \Psi \circ f^{-1}$ is a continuous map from the compact,

convex set $[0, 1]^N$ into itself. Thus, it has a fixed point, $\mathbf{x} \in [0, 1]^N$, by Brouwer's fixed point theorem. Note that $x_i < 1$ for all $i = 1, \dots, N$; if $x_i = 1$, then the i^{th} component of $f \circ \Psi \circ f^{-1}(\mathbf{x})$ is zero, since $h_i(\infty) = 0$ by definition, and so \mathbf{x} cannot be a fixed point. Thus, any fixed point \mathbf{x} of $f \circ \Psi \circ f^{-1}$ is in $[0, 1)^N$, and so $f^{-1}(\mathbf{x}) \in \mathbb{R}^N$. Clearly $f^{-1}(\mathbf{x})$ is a fixed point of Ψ .

Next, if u_i is strictly concave, then u'_i is decreasing, and so is $h_i = (u'_i)^{-1}$. Since ϕ is increasing, the function $x \mapsto \phi(h_1(x), \dots, h_N(x))$ is decreasing. Thus, it can have at most one fixed point, establishing uniqueness. \square

PROOF OF LEMMA 2: Let n denote the first time, possibly 0, at which $\hat{p}_i(n) < m$ for some i . Then, by (3) and (7), $x_i(n) \geq (w_i/m)^{1/(1-\beta)}$. Since the aggregate demand $x(n)$ is no smaller than the demand, $x_i(n)$, of user i , we obtain from (8) that

$$p(n) \geq \left(\frac{w_i}{mC_\beta} \right)^K \geq \frac{1}{\alpha} \left(\frac{W_\beta}{q^*C_\beta} \right)^K = \frac{q^*}{\alpha},$$

where we have used the definition of m to obtain the second inequality, and (10) to obtain the last equality. Thus, $\hat{p}_i(n+1) = (1-\alpha)\hat{p}_i(n) + \alpha p(n) \geq q^*$ for all i .

We have thus shown that, if $\hat{p}_i(n) < m$ for some i and n , then $\hat{p}_i(n+1) > m$. On the other hand, if $\hat{p}_i(n) \geq m$, then $\hat{p}_i(n+1) \geq (1-\alpha)m$. This establishes the first claim of the lemma.

Now, since $\hat{p}_i(n) \geq (1-\alpha)m$ for all i , it follows from (3) and (7) that

$$x(n) \leq \left(\frac{W_\beta}{(1-\alpha)m} \right)^{1/(1-\beta)},$$

so that by (8)

$$p(n) \leq \left(\frac{W_\beta}{(1-\alpha)mC_\beta} \right)^K = \left(\frac{q^*}{(1-\alpha)m} \right)^K q^*.$$

Since this holds for all n , the second claim of the lemma is immediate from (6).

□

PROOF OF THEOREM 1: We have $f'(x) = 1 - \alpha - \frac{\alpha K}{(1+x)^{K+1}}$. Now $\alpha K \leq 1 - \alpha$ by assumption, so $0 < f'(x) < 1 - \alpha$ for all $x > 0$. Since $f(0) = 0$, we get $0 < f(x) < (1 - \alpha)x$ for all $x > 0$. Choose ϵ small enough that $\rho := 1 - \alpha - \gamma\epsilon < 1$. Then, by (13),

$$-\gamma\epsilon^2 \leq \delta_{n+1} \leq \rho\delta_n \quad \forall \delta_n > \epsilon.$$

Thus, for initial conditions $\delta_0 > \epsilon$, δ_n converges geometrically at rate ρ to the interval $[-\gamma\epsilon^2, \epsilon]$.

Since f is convex and $f(0) = 0$, the equation $f(x) = 0$ has at most one other solution, which we denote x_0 . Observe that $x_0 \leq 0$, because $f'(0) \geq 0$. Now, for $x \in (x_0, 0)$, we have by the convexity of f that $0 > f(x) > f(0) + xf'(0) = \lambda x$, where $\lambda = 1 - \alpha(K+1) \in [0, 1)$. Choosing ϵ small enough that $\rho := \lambda + \gamma\epsilon < 1$, we get

$$\gamma\epsilon^2 > \delta_{n+1} > \rho\delta_n \quad \forall \delta_n \in (x_0, -\epsilon).$$

So, if the initial condition is $\delta_0 \in (x_0, -\epsilon)$, then δ_n converges geometrically at rate ρ to the interval $[-\epsilon, \gamma\epsilon^2]$.

If $\delta_0 \in (-1, x_0]$, then $f(\delta_0) \geq 0$ by convexity, so $\delta_1 \geq -\gamma\epsilon^2$. Thus, either $\delta_1 > \epsilon$ and the previous argument applies to the initial condition δ_1 , or δ_1 is already in an ϵ -neighborhood of the origin.

Finally, we wish to show that once the iterates δ_n enter an ϵ -neighborhood of the origin, they remain in this neighborhood. Recall that $f'(0) = 1 - \alpha(K+1) \in [0, 1)$. Hence, for small enough ϵ , $\rho := \sup_{x \in [-\epsilon, \epsilon]} |f'(x)| < 1$, and so $|f(x)| < \rho|x|$ on $[-\epsilon, \epsilon]$. It is now immediate from (13) that, if $\delta_n \in [-\epsilon, \epsilon]$, then

so is δ_{n+1} . This completes the proof of the theorem. \square

PROOF OF THEOREM 2: We have by (13) that $|\delta_{n+1} - f(\delta_n)| \leq \gamma\epsilon^2$ for all n . We also showed in Lemma 2 that $\hat{p}(n)$ is eventually confined to a compact set bounded away from zero. Consequently, δ_n is confined to a compact set bounded away from -1 ; on this set, f is differentiable and its derivative is bounded, so f is Lipschitz. Hence, $|f(\delta_{n+1}) - f(f(\delta_n))| \leq \gamma\epsilon^2$, for a possibly different constant, γ . Thus, by the triangle inequality,

$$|\delta_{n+2} - f(f(\delta_n))| \leq |\delta_{n+2} - f(\delta_{n+1})| + |f(\delta_{n+1}) - f(f(\delta_n))| \leq \gamma\epsilon^2. \quad (14)$$

Observe that $f'(0) = 1 - \alpha(K + 1)$ is negative by assumption. Hence, by convexity, $f(x) \geq 0$ for all $x \leq 0$. Moreover, $f(x) \rightarrow \infty$ as $x \rightarrow \infty$, so there is an $x^0 > 0$ such that $f(x^0) = 0$, $f(x) < 0$ on $(0, x^0)$ and $f(x) > 0$ on (x^0, ∞) . Now, if $\delta_n \in (0, x_0)$, then $f(f(\delta_n)) > 0$; moreover, if $\delta_n \in (\epsilon, x_0)$, then $f(f(\delta_n)) < \lambda^2\delta_n < \lambda^2x_0$. Combining this with (14), we see that $\delta_{n+2} < \rho\delta_n$ for some $\rho < 1$, and that $\delta_{n+2} \in (-\gamma\epsilon^2, x_0)$. This implies that, if the initial condition has $\delta_0 \in (0, x_0)$, then δ_{2n} converges at a geometric rate to an ϵ -neighborhood of zero.

On the other hand, suppose $\delta_n > x_0$. Since $f'(x) < 1 - \alpha$ for all x , $f(x) = f(x) - f(x^0) < (1 - \alpha)(x - x^0)$ for all $x > x^0$. Hence, by (13), $\delta_{n+1} \in (-\gamma\epsilon^2, \rho\delta_n)$, for some $\rho < 1$. Thus, if initially $\delta_0 > x_0$, then δ_n converges at a geometric rate to the interval $(-\gamma\epsilon^2, x_0)$; once δ_n is in this interval, the argument above is applicable. Finally, if $x < 0$, then $f(x) > 0$; hence, if $\delta_0 < 0$, then $\delta_1 \in (-\gamma\epsilon^2, \infty)$, and the arguments above are applicable.

Next, we show that, once the iterates δ_n enter an ϵ -neighborhood of the origin, they remain confined to this neighborhood. We show in Lemma 3 below that, if

$f'(0) = 1 - \alpha(K + 1) < -1$, then the hypothesis of the theorem cannot hold; for small enough $x > 0$, we will have $f(f(x)) > x$. Hence, by assumption, $f'(0) \in (-1, 0]$. We can now use the same argument as in the proof of Lemma 1, where we had $f'(0) \in [0, 1)$.

It remains only to establish the converse. If we consider an initial condition where all the price expectations, $\hat{p}_i(0)$, are equal to each other (and hence to $\hat{p}(0)$), then it is easy to verify that the recurrence (13) holds without the perturbation term, $\gamma\epsilon^2$. Moreover, the price expectations of all users continue to be identical, so $\delta_{n+1} = f(\delta_n)$ for all n . Hence, to establish the converse, we need to show that there is an initial condition x for which $f^n(x)$ remains bounded away from zero for infinitely many n . Now, $f(f(x)) \geq x$ for some $x > 0$ by assumption. Since $f(x)/x \rightarrow 1 - \alpha < 1$ as $x \rightarrow \infty$, it follows that $f(f(x)) < x$ for sufficiently large x . Since $f \circ f$ is continuous, and $f(f(0)) = 0$, there is an $x^* > 0$ such that $f(f(x^*)) = x^*$. Then $f^{2n}(x^*) = x^*$ remains bounded away from zero. This completes the proof of the theorem. \square

PROOF OF LEMMA 3: In the proof of Theorem 2 above, we saw that $f(x) < (1 - \alpha)x$ for all $x > x^0$, where x^0 is the unique strictly positive root of the equation $f(x) = 0$. Thus, $f(f(x)) < (1 - \alpha)^2x$ for all $x > x^0$, and it only remains to establish the claims of the lemma on the interval $(0, x_0)$. In fact, it suffices to establish them on the interval $(0, x_{\min})$, where $x_{\min} \in (0, x^0)$ is the minimizer of the convex function, f . This is because, for any $x \in (x_{\min}, x^0)$, there is a corresponding $\hat{x} \in (0, x_{\min})$ such that $f(\hat{x}) = f(x)$. Thus, if $f(\hat{x}) < \lambda^2\hat{x}$, then $f(x) < \lambda^2\hat{x} < \lambda^2x$.

Since the case where $\alpha(K + 1) < 1$ has already been dealt with in Theorem 1,

we shall henceforth assume that $\alpha(K+1) \geq 1$. It will be more convenient to work with the function

$$g(x) = \log [(1 - \alpha)e^x + \alpha e^{-Kx}]. \quad (15)$$

We have $g(0) = 0$ and $g'(0) = 1 - \alpha(K+1) \leq 0$. Also, $f(x) = \exp(g(\log(1+x))) - 1$, so $f(f(x)) = \exp(g(\log(1+f(x)))) - 1 = \exp(g(g(\log(1+x)))) - 1$.

Suppose first that $\alpha K \leq 1$. To establish the first claim of the lemma, it suffices to show that, for any $\epsilon > 0$, there is a $\kappa > 0$ such that $g(g(y)) \leq y - \kappa$ for all $y \in (\epsilon, y_{\min})$, where

$$y_{\min} = \frac{1}{K+1} \log \frac{\alpha K}{1-\alpha}, \quad (16)$$

is the minimizer of the convex function, g . Indeed, if there is such a κ , then

$$\begin{aligned} f(f(x)) &= \exp(g(g(\log(1+x))) - 1) \\ &\leq \exp(\log(1+x) - \kappa) - 1 \\ &= (1+x)e^{-\kappa} - 1. \end{aligned}$$

Taking $\lambda^2 = e^{-\kappa} < 1$, we have $f(f(x)) < \lambda^2 x$ for all $x \in (e^\epsilon - 1, e^{y_{\min}} - 1)$. It is easy to verify that $e^{y_{\min}} - 1 = x_{\min}$, the minimizer of $f(x)$. We shall now show that there is such a κ .

Suppose y is in $(0, y_{\min})$, so that $g(y) < 0$. Define $\beta = -g'(0) = \alpha(K+1) - 1$. Now, $\beta \in [0, \alpha] \subset [0, 1)$ by the assumption that $\alpha K \leq 1$ and $\alpha(K+1) \geq 1$. Hence, by the strict convexity of $x \mapsto e^x$ and Jensen's inequality, $g(y) > (1 - \alpha)y - \alpha K y = -\beta y$ for all $y \neq 0$. Now, since g is strictly decreasing on $(-\infty, 0]$,

we have for $y \in (0, y_{\min})$ that

$$\begin{aligned}
g(g(y)) &< g(-\beta y) = \log [(1 - \alpha)e^{-\beta y} + \alpha e^{K\beta y}] \\
&= y + \log [(1 - \alpha)e^{-(\beta+1)y} + \alpha e^{(K\beta-1)y}] \\
&= y + \log [(1 - \alpha)e^{-\alpha(K+1)y} + \alpha e^{(\alpha K-1)(K+1)y}]. \tag{17}
\end{aligned}$$

But $\alpha K \leq 1$, so it follows that $g(g(y)) < y + \log [(1 - \alpha) + \alpha] = y$ for all $y > 0$. In particular, $g(g(y)) - y$ is strictly smaller than zero on the compact interval $[\epsilon, y_{\min}]$. Since g is continuous, the maximum of $g(g(y)) - y$ on this interval is attained; denote it by $-\kappa$ and note that $\kappa > 0$. Thus, $g(g(y)) \leq y - \kappa$ for all $y \in [\epsilon, y_{\min}]$, which establishes the first claim of the lemma.

Suppose, on the other hand, that $\alpha(K + 1) > 2$. Since f is continuously differentiable in a neighborhood of the origin, with $f(0) = 0$ and $f'(0) = 1 - \alpha(K + 1) = -\beta$, we get

$$f(x) = -\beta x + O(x^2), \quad f(f(x)) = \beta^2 x + O(x^2).$$

Since $\beta < -1$, it follows that $f(x) > x$ for all $x \neq 0$ in a neighborhood of the origin. This establishes the second claim of the lemma. \square

References

ANDERSON, E., F. KELLY AND R. STEINBERG, (2006): “A contract and balancing mechanism for sharing capacity in a communication network”, *Management Science*, 52, 39–53.

BERNOULLI, D., (1738): “Specimen Theoriae Novae de Mensura Sortis” [“Exposition of a new theory on the measurement of risk”] *Commentarii Academiae Scientiarum Imperialis Petropolitanae*, Tomus V [*Papers of the Imperial Academy of Sciences in Petersburg*, Vol. V], 175–192.

BLANCHARD, O. AND S. FISHER, (1989): *Lectures on Macroeconomics*, Cambridge, MA and London: The MIT Press.

CAO, J. AND K. RAMANAN, (2002): “A Poisson limit for buffer overflow probabilities”, *Proc. IEEE Infocom*.

FLOYD, S., (1994): “TCP and explicit congestion notification”, *ACM Computer Communication Review*, 24(5), 10–23.

FUDENBERG, D. AND D. LEVINE, (1998): *The Theory of Learning in Games*, Cambridge, MA: The MIT Press.

FUDENBERG, D. AND J. TIROLE, (1991): *Game Theory*, Cambridge, MA: The MIT Press.

GIBBENS, R.J. AND F. P. KELLY, (1999): “Resource pricing and the evolution of congestion control”, *Automatica*, 35, 1969–1985.

HAURIE, A. AND P. MARCOTTE, (1985): “On the relationship between Nash-Cournot and Wardrop equilibria”, *Networks*, 15 295–308.

JOHARI, R. AND D. TAN, (2001): “End-to-end congestion control for the Internet: delays and stability”, *IEEE/ACM Trans. Networking*, 9(6), 818–832.

KELLY, F. P., A. MAULLOO AND D. TAN, (1998): “Rate control in commu-

nication networks: shadow prices, proportional fairness and stability”, *Journal of the Operational Research Society*, 49, 237–252.

KUNNIYUR, S. AND R. SRIKANT, (2001): “Analysis and design of an adaptive virtual queue algorithm for active queue management”, *Proc. ACM SIGCOMM*.

LOW, S.H. AND D.E. LAPSLEY, (1999): “Optimization flow control – I: Basic algorithm and convergence”, *IEEE/ACM Transactions on Networking*, 7, 861–875.

MACKIE-MASON, J.K. AND H.R. VARIAN, (1996): “Some economics of the Internet”, in W. Sichel and D.L. Alexander (eds.), *Networks, Infrastructure and the New Task for Regulation*, University of Michigan Press, Ann Arbor.

MASSOULIÉ, L., (2002): “Stability of distributed congestion control with heterogeneous feedback delays”, *IEEE Trans. Autom. Control* 47(6), 895–902.

MASSOULIÉ, L. AND J. ROBERTS, (1999): “Bandwidth sharing: objectives and algorithms”, *Proc. INFOCOM*, 1395–1403.

RAMAKRISHNAN, K.K. AND R. JAIN, (1990): “A binary feedback scheme for congestion avoidance in computer networks”, *ACM Transactions on Computer Systems*, 8, 158–181.

RAINA, G. AND D. WISCHIK, (2005): “Buffer sizes for large multiplexers: TCP queueing theory and instability analysis”, *Proc. EuroNGI Conference on Next Generation Internet Networks*.

TIOLE, J., (1988): *The Theory of Industrial Organization*, Cambridge, MA: The MIT Press.

SHENKER, S., (1995): “Fundamental design issues for the future Internet”, *IEEE Journal on Selected Areas of Communications*, 13, 1176–1188.

SRIKANT, R., (2004): *The Mathematics of Internet Congestion Control*, Birkhauser.

VARIAN, H.R., (1992): *Microeconomic Analysis*, 3rd ed., New York: W.W. Norton & Co.

VINNICOMBE, G, (2002): “On the stability of networks operating TCP-like congestion control”, *Proc. 15th IFAC World Congress on Automatic Control*.