

Tariffs, Mechanisms and Equilibria at a Single Internet Link

Costas Courcoubetis and Antonis Dimakis

Department of Informatics
Athens University of Economics and Business
Patission 76, Athens 10434, Greece
{courcou, dimakis}@aueb.gr

Abstract. We analyze the interplay between the demand for downloads, choice of congestion control mechanism, and tariff structure at a single link, when users have preferences in terms of average download delay and they are charged according to the number of ECN marked packets they receive. Our model involves a timescale separation approach, where in the fast timescales active flows compete for instantaneous bandwidth share by optimally tuning congestion control parameters in a noncooperative fashion. This is modeled by letting flows choose utility functions within the network utility maximization framework laid down by Kelly [1]. On a slower timescale, users selfishly change their otherwise unrestricted demand for downloads based on the average experienced download delay and charges incurred. We study the equilibrium of this loop of interactions from the point of view of social welfare.

For homogeneous users we find that optimal equilibria are induced when they choose among linear utility functions, while this is not the case for logarithmic, i.e., proportionally fair congestion controllers. We next consider two types of users, web-browsing and bittorrent, where the latter are much less sensitive to download delay than the former. If bittorrent users respond to congestion according to proportionally fair utility functions the charge induced by ECN marked packets does not provide the correct signal for service differentiation and the resulting equilibrium is suboptimal. On the other hand, if flows are charged for the volume of bits they transfer as well, then social welfare maximum is attained for some price per bit. This reveals a new role for bit volume pricing: to provide bittorrent users the correct incentives for choosing congestion controllers that effectively give priority to web-browsing users.

Key words: pricing, congestion control, congestion marks, noncooperative equilibrium.

1 Introduction

In analyzing the performance of congestion control mechanisms through either simulation or theory, one assumes a given level of demand. But if we are asked to answer questions such as

- What is the impact of different tariffs on user-perceived performance and Internet service provider (ISP) profits?
- Given a specific tariff, what congestion control algorithm should a user use in order to improve performance and charges his traffic incurs?
- Is it beneficial for an ISP to buy more link capacity from a wholesale provider?
- What is the outcome of competition between ISPs?

then we need to precisely determine the equilibrium level of demand that results from the complex interactions between ISPs, users and congestion control algorithms.

Indeed, a user may improve his download delays by using a more aggressive congestion control algorithm but will be charged more by a tariff scheme that penalizes congestion. Hence this user must optimally balance his preference for short download delays with the level of charges he incurs. If he discovers that despite the optimal choice of congestion control he pays more than he ought to for the download delay he experiences, he might lower his frequency with which he is initiating downloads. On a much slower timescale, the ISP may notice a drop in demand and may need to lower its price in order to maintain its previous level of profits. In turn, the user will observe the price drop and will again reconsider his choice of congestion control aggressiveness and so on.

Hopefully, given enough time this loop of interactions will converge to some noncooperative equilibrium level of demand as a result of the tariff structure employed by the ISP, and the choice of congestion control algorithms users are optimizing their aggressiveness over.

In search for a tractable model for exploring equilibria of this sort, we isolate three timescales and consider the relevant stakeholders and their actions involved in each of these timescales.

- ISP chooses prices (of marks/bits or something else), expands capacity [slow timescale]
- Users decide whether to download or not, what congestion control parameters to use. [medium timescale]
- Congestion control reacts on instantaneous congestion. [fast timescale]

There are decisions taking place on even slower timescales, e.g., users change congestion control algorithms, switch between providers, operating systems vendors publish their congestion control software, but we have decided to omit those in order to make the analysis tractable. In fact, in this paper we consider the case of a single bottleneck link with fixed capacity.

Our main goal is to understand the incentive properties of pricing when this is based on primal congestion control algorithms, i.e., based on volume of ECN marked packets (congestion marks) [3], versus pricing based on data volume or on combinations of both. Our basic modelling assumption that is used to carry out the above analysis is that network users are separated into two different classes: the ‘web-browsing’ users class, where latency in downloading files is important, and the p2p users, the ‘bittorrent users’, for which downloading

delay is orders of magnitude less important. Using this separation assumption we look at the sequence of systems when delay importance on bittorrent users becomes negligible and study the efficiency properties of the equilibria as this limit is reached. A main result of this analysis is that if bittorrent users choose among α -fair utility functions with $\alpha > 0$ [2], such as logarithmic, the charge induced by congestion marks alone does not provide the correct signal for service differentiation and the resulting equilibrium is suboptimal. This results from our assumption that the demand for downloads is potentially unrestricted, and sensitive to the download delay. Under this assumption, as we will see later, it is possible that an increasingly high number of bittorrent users joins as their delay sensitivity decreases. Thus bittorrent flows accumulate a non-negligible number of congestion marks. In contrast, consider what happens when demand for bittorrent downloads is fixed. Bittorrent users will not be willing to incur any congestion marks as they become less sensitive to delay.

Now if flows are charged for the volume of bits they transfer as well, then social welfare maximum is attained for some price per bit. This reveals a new role for bit volume pricing: to provide bittorrent users the correct incentives for choosing congestion controllers that effectively give priority to web-browsing users.

We must finally comment that our results are asymptotic results when Internet users can be abstracted into two classes regarding their delay sensitivity. Probably our results may not hold when we add another user type whose delay sensitivity goes also to zero but at a slower rate than the bittorrent type. However, it is under the two type model where the properties we bring out become most apparent.

The paper is organized as follows. In Section 2 we relate our work to previous work done in the area. In Section 3 we introduce the model and notation to be used throughout the paper, and define the concept of equilibrium. The main results concerning social welfare of equilibria under a single and multiple user types are contained in Sections 4 and 5 respectively. Finally, we summarize in Section 6.

2 Related Work

In an important paper Naor [4] has posed the question of noncooperative equilibrium arrival rate in a FCFS queue, and considered the case where an arriving customer knows the queue length and is allowed to balk. He showed that the equilibrium is inefficient in terms of social welfare if no pricing is imposed, and this is fixed when an appropriate toll is imposed on non balking customers. The same model under the assumption that customers do not observe the queue length is considered in [5] and similar tolls to [4] were proposed. Such tolls that internalize externalities were well known to economists and queueing theorists [6, 7]. In Section 5 the bit volume charge will have a similar internalizing role, in the case of two extreme user types.

Another way to internalize externalities without the imposition of tolls is by changing the service discipline. In an intriguing result, [8] has shown that the LCFS service discipline provides the correct incentives for customers to optimally balk. Another important result by [9] has shown that if customers bid for their priority so that the higher bidder gets full priority, then the resulting equilibrium is efficient. Such a combination of bids and service discipline, as we will see, corresponds in our framework to a particular choice of congestion control parameters by the users. One can view the results of this paper as an extension of the demand models of [4, 5, 9–11], where users are allowed to bid for instantaneous service in a very general way.

Kelly and Gibbens in [3] view the choice of congestion controller by the end-to-end users as a game between them, where users have preferences in terms of instantaneous bandwidth. Additionally, they propose flows to be charged proportionally to the amount of ECN marked packets they generate. In contrast to [4, 5, 9–11], the flow arrival rate is fixed, i.e., balking is not allowed, but flows affect the service discipline by tuning the parameters of congestion control algorithms. The ECN mark charge offers a way to provide service differentiation since at times of high congestion, bandwidth is allocated to the flows that value it most. Interestingly, when balking is allowed a mark charge may not be sufficient to achieve the maximum social welfare (see Section 5). Inefficiency of Nash equilibria resulting from the noncooperative choice of congestion control parameters, is observed in [12, 13] when users have preferences on instantaneous bandwidth. In contrast to the models of [12, 13], here inefficiency can arise even when all users are of the same type.

References [14, 15] consider whether class differentiated pricing at a queueing system offering two classes of priority, gives the correct incentives for customers to select the appropriate class. It is found that this is possible under identical service requirements by all customers. Although we seek for similarly defined efficient equilibria as in [14, 15], here prices are not set by a central planner but they result from user choices. Thus, potentially there is an infinite number of service classes.

3 Model

Consider a single link of capacity C bits/s shared by n flow types. We assume that there is an arbitrarily high number of flows where each decides whether to join the system or balk in a random (e.g., after a coin toss) and independent manner. We further assume that flows possess no information about the actual state of the system, e.g., the number of active downloads. The only information available to them is the number of flow types along with their defining parameters. That is, we consider a nonobservable model according to the terminology in [11]. Thus after each flow has made its decision, the resulting arrival process of type- i flows forms a Poisson process with rate λ_i flows/s. Each nonbalking type- i flow terminates once the download of an exponentially distributed file size of mean

$1/\mu_i$ bits completes. We use θ_i to denote the demand λ_i/μ_i (in bit/s) of type i -flows.

A nonbalking flow, at the instant of its arrival, also decides which congestion control mechanism to use throughout its duration. Since we are interested in the (fast timescale) equilibrium service rates, congestion control mechanisms are modeled by nondecreasing concave functions and the use of the convex optimization framework of Kelly [1] to determine the equilibrium allocations which we assume are reached instantaneously. We allow a randomized choice of congestion control mechanisms, so the choice is over the set $\mathcal{P}(\mathcal{U})$ of distributions on \mathcal{U} , a subset of concave, nondecreasing, continuously differentiable utility functions. We will frequently consider $\mathcal{U} = \mathcal{U}_\alpha$, the set of α -fair utilities $\{u : u(x) = -wx^{1-\alpha}/(1-\alpha), \forall x > 0, \text{ and some } w \geq 0\}$ with $\alpha \geq 0, \alpha \neq 1$ [2]. The case $\alpha = 1$ corresponds to instantaneous bandwidth allocations that are (weighted) proportionally fair [1], and can be defined by $\mathcal{U}_1 = \{u : u(x) = w \log x, \forall x > 0, \text{ and some } w \geq 0\}$; here, the gain parameter w is interpreted as the willingness to pay per unit of time [1]. Since for fixed α , the elements of \mathcal{U}_α are indexed by the gain parameter $w \geq 0$, we can take $\mathcal{P}(\mathcal{U}_\alpha)$ to be the set of probability distributions on \mathbb{R}_+ . We also introduce the relation $>$ on $\mathcal{P}(\mathcal{U})$, defined by

$$\pi_1 > \pi_2 \Leftrightarrow (\pi_1 \times \pi_2)(\{(u_1, u_2) \in \mathcal{U} \times \mathcal{U} | u'_1(x) > u'_2(y) \text{ for all } x, y > 0\}) = 1.$$

This is defined to signify that the flows of some type take full priority over flows of another type. Moreover, $\mathcal{P}(\mathcal{U})$ is equipped with the Prohorov metric, i.e., $\pi_n \rightarrow \pi$ means weak convergence of probability distributions.

If at any instant t , a flows are active with utilities u_1, \dots, u_a , then the instantaneous bandwidth share $x_j(t)$ (in bit/s) of flow $j = 1, \dots, a$ is determined by the optimum solution of

$$\max_{x_1, \dots, x_a \geq 0} \sum_{j=1}^a u_j(x_j) \text{ over } \sum_j x_j \leq C,$$

and the instantaneous charge rate is $\lambda(t)x_j(t)$ per second, where $\lambda(t)$ is the Lagrange multiplier of the capacity constraint. In [3] it is demonstrated how the rate of congestion marks conveys the value of the Lagrange multiplier $\lambda(t)$, so, as in [3], in the sequel we identify congestion mark pricing with the imposition of a charge $\lambda(t)$ on each unit of allocated bandwidth. Thus, during the lifetime of a flow an average charge per bit equal to

$$m_i(\theta, \pi) = E^0 \left(\int_0^T \lambda(t)x_j(t)dt \right) \mu_i,$$

is incurred, if flow j is of type- i and T its download delay. Here, E^0 is expectation under the Palm distribution P^0 on arrivals of type- i flows, i.e., P^0 characterizes statistically what is seen by a typical arriving type- i flow.

Let $d_i(\pi'; \theta, \pi)$ be the average download delay of a type- i flow divided by the average file size, i.e., $1/\mu_i$, when this flow chooses congestion control mechanism

according to $\pi' \in \mathcal{P}(\mathcal{U})$ and all other flows (including those of the same type) choose according to $\pi = (\pi_1, \dots, \pi_n)$. ($m_i(\pi'; \theta, \pi)$ is defined similarly.) Thus $d_i(\pi_i; \theta, \pi)$ is the average delay per bit of a typical type- i flow, which we denote simply by $d_i(\theta, \pi)$. For α -fair utility functions we conveniently write $d_i(w; \theta, \pi)$ for $d_i(\delta_w; \theta, \pi)$ where $\delta_w \in \mathcal{P}(\mathcal{U}_\alpha)$ assigns unit probability to the utility function in \mathcal{U}_α with parameter w . (The value of α will be clear from the context.) Each type- i download brings a reward $r_i - c_i d_i(\theta, \pi)$ per bit, where r_i represents a delay independent reward, and c_i is the cost of each second of delay per bit. Consequently, the net value per bit of a type- i flow is

$$r_i - c_i d_i(\theta, \pi) - m_i(\theta, \pi) \quad (1)$$

Since the maximum potential demand for each flow type is arbitrarily high, if (1) is positive more type- i flows will join and ultimately drive (1) to zero after delay increases sufficiently. This gives rise to the following definition.

Definition 1. For each i let \mathcal{A}_i be a subset of $\mathcal{P}(\mathcal{U})$, where \mathcal{U} is the set of nondecreasing, concave, continuously differentiable functions on \mathbb{R}_+ . The vector (θ, π) is an equilibrium if the following hold for all i :

1. $r_i - c_i d_i(\theta, \pi) - m_i(\theta, \pi) \leq 0$. If the inequality is strict then $\theta_i = 0$.
2. π_i maximizes $r_i - c_i d_i(\pi'; \theta, \pi) - m_i(\pi'; \theta, \pi)$ over $\pi' \in \mathcal{A}_i$.

In our model θ_i is the result of a Wardrop type of equilibrium that is induced by (1), hence it is outside the control of the individual user; the choice of π_i is strategic. We have not developed a theory for existence and uniqueness of equilibria in general. In what follows we deal with specific \mathcal{U} where explicit calculation of equilibria is possible.

To any vector (θ, π) irrespective of it being an equilibrium or not, we associate a level of social welfare equal to

$$\mathcal{W}(\theta, \pi) = \sum_{i=1}^n \theta_i (r_i - c_i d_i(\theta, \pi)). \quad (2)$$

A vector (θ, π) is called *efficient* if it maximizes (2) over $\theta \geq 0$, $\sum_{i=1}^n \theta_i \leq C$, $\pi_i \in \mathcal{A}_i$.

In the following section we consider the simplest case of a system consisting of a single user type. We consider this simpler case first in order to highlight some key efficiency results.

4 Single User Type

Here we consider the case $n = 1$. The number of active flows behaves as an M/M/1 queue, so $d_1(\theta, \pi) = 1/(C - \theta_1)$ for any \mathcal{U} and $\pi \in \mathcal{P}(\mathcal{U})$. Thus the optimal demand θ_1^* maximizes $\theta_1(r_1 - c_1/(C - \theta_1))$, which gives

$$\theta_1^* = \max \left(C - \sqrt{\frac{c_1 C}{r_1}}, 0 \right). \quad (3)$$

Contrary to the case where the demand for downloads is fixed [1, 3, 13, 12], not all congestion controllers possess efficient equilibria when users selfishly optimize demand.

Proposition 1. *1. If $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_0)$ there exists a unique efficient equilibrium.
2. If $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_1)$ the only symmetric pure equilibrium is inefficient.*

Proof. Part 1 is shown in [9]. The intuition of why this should hold is the following. An arriving flow that picks a parameter value p will take preemptive priority over all active flows with lower parameter values. Moreover, if the instantaneous bandwidth at any time during the lifetime of the flow is nonzero, then the shadow price will be exactly p . Thus, the equilibrium choice of parameter is such that it does not pick any particular value with a positive probability, i.e., does not have atoms, since an infinitesimally small parameter increase results into taking priority over a positive number of flows, with positive probability. Hence, any parameter p within some range will result into identical net values per bit (1). Since this is true for the smallest p within this range, i.e., for the smallest priority, it is optimal for an arriving flow to join if and only if it is optimal to join under the LCFS policy. But under such a policy a nonbalking flow does not pose any delay penalties to subsequently arriving flows, and thus will join if it is socially optimal to do so [8].

For the second part, let $f(w)$ be the average delay per bit of a flow with willingness to pay w when all other flows have a unit willingness to pay. By [10],

$$f(w) = \frac{1 + w - w\theta/C}{1 + w - \theta/C} \frac{1}{1 - \theta/C} \frac{1}{C}. \quad (4)$$

If w^* is the equilibrium willingness to pay then the maximum net value per bit $r_1 - c_1 f(w/w^*) - w f(w/w^*)$ must be achieved at $w = w^*$, so

$$\left. \frac{d}{dw} (c_1^* f(w/w^*) + w f(w/w^*)) \right|_{w=w^*} = 0 \implies w^* = -\frac{c_1 f'(1)}{f'(1) + f(1)} = \frac{c_1 \theta_1}{2(C + \theta_1)}.$$

Such payment will induce a demand θ_1 with

$$r_1 - \frac{c_1}{C - \theta_1} = \frac{c_1 \theta_1 \mu_1}{2(C + \theta_1)(C - \theta_1)},$$

which clearly is not the optimal demand (3).

We conjecture that a similar inefficiency outcome as for \mathcal{U}_1 holds for any \mathcal{U}_α with $\alpha > 0$. Congestion control mechanisms that correspond to the set \mathcal{U}_0 of linear utility functions have been proposed in [1, 3, 16, 17].

5 Two Extreme User Types

Here we consider the efficiency of equilibria for two flow types, where one of them is nearly insensitive to download delay. To do this, we consider a sequence

of systems as the coefficient c_2 of type-2 converges to zero and determine the limit (θ^0, π^0) of equilibria $(\theta^{c_2}, \pi^{c_2})$ of the c_2 -systems. Thus (θ^0, π^0) provides an asymptotically exact approximation for $c_2 > 0$, and it is easier to determine than $(\theta^{c_2}, \pi^{c_2})$ directly. We believe that the case $n = 2, c_2 = 0$ provides valuable intuition for more general systems because the key properties become most apparent.

We first determine the limit of the efficient allocation vectors as $c_2 \downarrow 0$. As expected, it is optimal for the delay sensitive type to take full preemptive priority at the limit. (To emphasize the dependence on c_2 we write $\mathcal{W}(\theta, \pi, c_2)$ for the social welfare when the linear coefficient of type-2 users is c_2 .)

Proposition 2. *Let \mathcal{U} be as in Definition 1. Then,*

$$\max_{\theta \geq 0, \theta_1 + \theta_2 \leq C, \pi \in \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{U})} \mathcal{W}(\theta, \pi, c_2) \rightarrow \mathcal{W}^*, \text{ as } c_2 \downarrow 0, \quad (5)$$

where

$$\begin{aligned} \mathcal{W}^* &= \max_{\theta \geq 0, \theta_1 + \theta_2 \leq C, \pi \in \mathcal{P}(\mathcal{U}) \times \mathcal{P}(\mathcal{U})} [\theta_1(r_1 - c_1 d_1(\theta, \pi)) + r_2 \theta_2] \\ &= \max_{0 \leq \theta_1 \leq C} \left[\theta_1 \left(r_1 - \frac{c_1}{C - \theta_1} \right) + r_2(C - \theta_1) \right]. \quad (6) \end{aligned}$$

Moreover, if maximizing vectors $(\theta^{c_2}, \pi^{c_2})$ of $\mathcal{W}(\cdot, \cdot, c_2)$ have a limit (θ^0, π^0) it must satisfy $\theta^0 = \theta^*$. Also, $\theta_1^0 > 0 \implies \pi_1^0 > \pi_2^0$, where θ^* is the optimizing demand in (6).

Proof. For small but fixed $\epsilon > 0$,

$$\max_{\theta_1 + \theta_2 \leq C - \epsilon, \pi} \mathcal{W}(\theta, \pi, c_2) \rightarrow \max_{\theta_1 + \theta_2 \leq C - \epsilon, \pi} \mathcal{W}(\theta, \pi, 0)$$

as $c_2 \downarrow 0$. Moreover, the right hand side converges to \mathcal{W}^* as $\epsilon \downarrow 0$. Now, since

$$\max_{\theta_1 + \theta_2 \leq C - \epsilon, \pi} \mathcal{W}(\theta, \pi, c_2) \leq \max_{\theta_1 + \theta_2 \leq C, \pi} \mathcal{W}(\theta, \pi, c_2) \leq \mathcal{W}^*,$$

we conclude that (5) is true.

Now any limit demand θ^0 equals θ^* because (6) possesses a unique solution. If $\theta_1^0 > 0$ but $\pi_1^0 > \pi_2^0$ does not hold then $d_1(\theta^{c_2}, \pi^{c_2}) \rightarrow +\infty$ since $\theta_1^0 + \theta_2^0 = C$. But this contradicts the fact that $\mathcal{W}(\theta^{c_2}, \pi^{c_2}, c_2) \geq 0$, hence $\pi_1^0 > \pi_2^0$.

Next, we would like to determine whether the tariff based on congestion marks [3] is adequate to enforce socially optimal equilibria $(\theta^{c_2}, \pi^{c_2})$ as $c_2 \downarrow 0$, i.e., whether $(\theta^0, \pi^0) = \lim_{c_2 \downarrow 0} (\theta^{c_2}, \pi^{c_2}) = (\theta^*, \pi^*)$ holds. As the next proposition shows, (θ^0, π^0) is not efficient when $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_\alpha)$ for any $\alpha \geq 0$, and $\mathcal{A}_2 = \mathcal{P}(\mathcal{U}_1)$. We conjecture that the same result holds for more general \mathcal{A}_2 except the case $\mathcal{A}_2 = \mathcal{P}(\mathcal{U}_0)$. Although we do not fully determine (θ^0, π^0) we find that no limit of the equilibria sequence $(\theta^{c_2}, \pi^{c_2})$ can be efficient. This is surprising since if demand was fixed, i.e., the maximum potential demand of both types is strictly

below C , then as c_2 becomes smaller, type-2 flows should be less willing to pay for instantaneous bandwidth because these will be served anyway when no type-1 flows are present, and this extra delay is bounded. Thus, type-2 flows will become second priority at the limit $c_2 \downarrow 0$.

However if demand varies, then in systems with small c_2 , one may observe a higher number of type-2 flows on the average, since more flows will join if their payments draw close to zero. Thus at equilibrium, the net value per bit will be close to zero as well. In other words, a tragedy of the commons situation is plausible among type-2 flows. If on the other hand such a situation does not arise and payments from type-2 flows come to the system at a positive rate, then type-2 flows will take up some of the bandwidth during times when type-1 flows are present. As $c_2 \downarrow 0$ this amount of bandwidth becomes non-negligible and the system does not behave as a priority system as it is supposed to in order to be efficient.

Proposition 3 (Inefficiency of tariff based on congestion marks alone).
 Let $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_\alpha)$, $\alpha \geq 0$ and $\mathcal{A}_2 = \mathcal{P}(\mathcal{U}_1)$. If $(\theta^{c_2}, \pi^{c_2})$ is equilibrium for each $c_2 > 0$ then

$$\limsup_{c_2 \downarrow 0} \mathcal{W}(\theta^{c_2}, \pi^{c_2}, c_2) < \mathcal{W}^*,$$

where \mathcal{W}^* is as in Proposition 2.

Proof. We follow the discussion prior to the statement of this proposition, and first show that type-2 flows achieve a decreased level of welfare if their rate of payments is not r_2 . Since $0 \leq m_2(\theta^{c_2}, \pi^{c_2}) \leq r_2$ we can take a subsequence of $(\theta^{c_2}, \pi^{c_2})$, which by abuse of notation we denote again by $(\theta^{c_2}, \pi^{c_2})$, such that $\lim_{c_2 \downarrow 0} m_2(\theta^{c_2}, \pi^{c_2}) = \delta$, and $(\theta^{c_2}, \pi^{c_2}) \rightarrow (\theta^0, \pi^0)$.

As in the proof of Proposition 2 one can show that $\lim_{c_2 \downarrow 0} \mathcal{W}(\theta^{c_2}, \pi^{c_2}, c_2) = \mathcal{W}^*$ implies $\lim_{c_2 \downarrow 0} \theta^{c_2} = \theta^*$. Moreover by Definition 1, $r_2 - c_2 d_2(\theta^{c_2}, \pi^{c_2}) \leq m_2(\theta^{c_2}, \pi^{c_2})$ for each c_2 , so if $\delta < r_2$ we have that

$$\lim_{c_2 \downarrow 0} \mathcal{W}(\theta^{c_2}, \pi^{c_2}, c_2) \leq \theta_1^*(r_1 - c_1 d_1(\theta^*, \pi^*)) + \delta \theta_2^* < \mathcal{W}^* .$$

i.e., on this subsequence the system is inefficient.

Now consider the case where the payment rate is exactly r_2 , i.e.,

$$\delta = \lim_{c_2 \downarrow 0} m_2(\theta^{c_2}, \pi^{c_2}) = \lim_{c_2 \downarrow 0} \frac{E_{c_2} \left(\sum_{i=1}^{n_2(0)} w_i \right)}{\theta_2^{c_2}} = r_2 ,$$

where $n(t) = (n_1(t), n_2(t))$ is the vector of active type-1 and type-2 flows at time t , and w_i is the congestion control parameter chosen by the i -th active type-2 flow.

We will show that type-1 flows see strictly larger average delays than what they would observe at a priority system. Since a typical type-1 flow sees time averages upon arrival, we have

$$\liminf_{c_2 \downarrow 0} P_{c_2}^0 \left(\sum_{i=1}^{n_2(0-)} w_i > \frac{r_2 \theta_2^*}{2} \right) > 0 ,$$

where $P_{c_2}^0$ is the Palm probability on arrivals of type-1 flows and the dependence on c_2 is shown explicitly. Now, define the following events:

$$F_1 = \left\{ \sum_{i=1}^{n_2(0-)} w_i > \frac{r_2 \theta_2^*}{2} \right\}, \quad F_2 = \{n_1(0-) = 0\},$$

$$F_3 = \{\text{no type-1 flows arrive during } (0,1]\},$$

$$F_4 = \{\text{no flows depart from the system during } (0,1]\}.$$

Now we can assume that $\lim_{c_2} P_{c_2}^0(F_2) > 0$ holds, otherwise the limit is not efficient and we are done. Moreover, $\lim_{c_2} P_{c_2}^0(F_3) = e^{-\theta_1^* \mu_1 - \theta_2^* \mu_2}$ and

$$\liminf_{c_2 \downarrow 0} P_{c_2}^0(F_4 | F_1, F_2, F_3) \geq e^{-(\mu_1 + \mu_2)C}.$$

Thus,

$$\liminf_{c_2 \downarrow 0} P_{c_2}^0(F_1 \cap F_2 \cap F_3 \cap F_4) > 0. \quad (7)$$

Let $x_1(t)$ be the instantaneous bandwidth at time $t \geq 0$, allocated to the type-1 flow arriving at time 0. On $\omega \in F_1 \cap F_2 \cap F_3 \cap F_4$ we have $x_1(t) = x_1(0)$, $n(t) = n(0)$ for all $t \in (0, 1]$ and $x_1(0)$ is given by the optimum x_1 in

$$\max_{x_1, x_2 \geq 0} u(x_1) + \sum_{i=1}^{n_2(0)} w_i \log x_2, \quad \text{over } x_1 + x_2 \leq C,$$

which is less than the optimum x'_1 in

$$\max_{x'_1, x_2 \geq 0} u(x'_1) + \frac{r_2 \theta_2^*}{2} \log x_2, \quad \text{over } x'_1 + x_2 \leq C. \quad (8)$$

Now let's compare the rate allocated to the type-1 flow that arrived at time 0 when type-2 flows give up priority. On the same probability space, consider a second system with the same type-1 flow arrival instances, same file sizes, and the same choice of utility functions but where type-1 takes full preemptive priority. If we denote by $\hat{n}(t) = (\hat{n}_1(t), \hat{n}_2(t))$ the vector of active flows at time t , and by $\hat{x}_1(t)$ the instantaneous bandwidth allocated to the type-1 flow that arrives at time 0, then it is evident that $n_1(t) \geq \hat{n}_1(t)$, $x_1(t) \leq \hat{x}_1(t)$ at any t . Furthermore at any time instant, any type-1 flow that is present at the priority system is also present at the system without full priority. Now, on $\omega \in F_1 \cap F_2 \cap F_3 \cap F_4$ we have $\hat{x}_1(t) = C$ for all $t \in (0, 1]$. Moreover, we need consider only $u \in \mathcal{U}_\alpha$ with (uniformly in c_2) bounded gain parameters, since the charge per bit never exceeds r_1 . So for any utility function u in (8), $\hat{x}_1(t) - x_1(t) > \phi$ for some $\phi > 0$ not dependent on c_2 . Thus for such ω , the delay of arriving type-1 flows is greater than that achieved by full priority, by more than ϕ/C . As $c_2 \downarrow 0$, a non-negligible proportion of type-1 flows encounters a delay difference of ϕ/C , by (7). For the rest of type-1 flows we know that their delay is at least as large as the one under full priority in the coupled system. Therefore, the average delay of a typical type-1 flow as $c_2 \downarrow 0$ is strictly greater than what the optimal allocation (6) would result in.

The inefficiency could be eliminated if prioritization and payments are linked differently than by the congestion mark tariff. Indeed, in Proposition 4 we show that when $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_0)$, $\mathcal{A}_2 = \mathcal{P}(\mathcal{U}_1)$ and a charge is imposed on the volume of downloaded bits as well as on the volume of congestion marks, the system is efficient in the limit as $c_2 \downarrow 0$. The intuition of why this is true is the following. Suppose a charge v^{c_2} is imposed on every bit carried by the link. Then, the net value per bit of type- i becomes,

$$r_i - c_i d_i(\theta, \pi) - v^{c_2} - m_i(\theta, \pi), i = 1, 2,$$

under any demand/parameter vector (θ, π) . Also redefine the definition of equilibrium such that part 1 in Definition 1 becomes “ $r_i - c_i d_i(\theta, \pi) - v^{c_2} - m_i(\theta, \pi) \leq 0$. If the inequality is strict then $\theta_i = 0$.” Now, if $v^{c_2} \approx r_2$ then type-2 flows have a very small margin for extracting value out of their bits. Thus, the charges incurred by congestion marks should be kept small, and so type-2 flows have the incentive of choosing low priority. On the other hand, since they pay a nonzero price v for each bit, a tragedy of the commons situation does not arise.

We also observe that Proposition 4 may not hold for other \mathcal{A}_1 beyond $\mathcal{P}(\mathcal{U}_0)$. This is because, by Proposition 2, efficiency implies that the system asymptotically behaves as a priority system, so the equilibrium demand reached by type-1 flows should be the same as in the single type case (see Section 4) for a delay independent reward equal to $r_1 - v^{c_2} \approx r_1 - r_2$ instead of r_1 . But as Proposition 1 shows, a choice of $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_1)$ would not generate the maximum welfare for type-1 flows; one should have $\mathcal{P}(\mathcal{U}_0)$ instead.

Proposition 4 (Efficiency of two part tariff). *Let $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_0)$ and $\mathcal{A}_2 = \mathcal{P}(\mathcal{U}_1)$. For each $c_2 > 0$ let $(\theta^{c_2}, \pi^{c_2})$ be an equilibrium under bit price v^{c_2} with $v^{c_2} \rightarrow r_2$, and $(\theta^{c_2}, \pi^{c_2}) \rightarrow (\theta^0, \pi^0)$. The following hold:*

1. *If $r_1 > r_2$ then $\pi_1^0 > \pi_2^0$, $\pi_2^0(\{0\}) = 1$.*
2. *$\mathcal{W}(\theta^{c_2}, \pi^{c_2}, c_2) \rightarrow \mathcal{W}^*$ as $c_2 \rightarrow 0$, where \mathcal{W}^* is as in Proposition 2.*

Proof. Since $(\theta^{c_2}, \pi^{c_2})$ is an equilibrium for each c_2 , we have $0 \leq m_2(\theta^{c_2}, \pi^{c_2}) = r_2 - v^{c_2} - c_2 d_2(\theta^{c_2}, \pi^{c_2}) \leq r_2 - v^{c_2}$. But $v^{c_2} \uparrow r_2$, so $E_{c_2}^0(wT)\mu_2 = m_2(\theta^{c_2}, \pi^{c_2}) \rightarrow 0$, where w, T is the willingness to pay and download delay of an arriving type-2 flow at time 0. Since $T \geq Z/C$ where Z is the random file size, we have $E_{c_2}^0(w)/(\mu_2 C) = E_{c_2}^0(wZ)/C \leq E_{c_2}^0(wT) \rightarrow 0$. Thus, $\pi_2^{c_2} \rightarrow \pi_2^0$ where $\pi_2^0(\{0\}) = 1$, i.e., as type-2 users become less sensitive to delay they tend to be less aggressive. We now show that this is not the case for any significant proportion of type-1 users by arriving at a contradiction. If $\pi_1^0(\{0\}) > 0$ then for any $\gamma > 0$ there exists a sequence $p^{c_2} \downarrow 0$ with

$$r_1 - c_1 d_1(p^{c_2}; \pi^{c_2}, \theta^{c_2}) - v^{c_2} - p^{c_2} = r_1 - c_1 d_1(\pi^{c_2}, \theta^{c_2}) - v^{c_2} - m_1(\theta^{c_2}, \pi^{c_2}) = 0,$$

$$\text{and } \pi_1^{c_2}((p^{c_2}, +\infty)) \geq 1 - \gamma.$$

We show that a type-1 user is better off by choosing a constant parameter $\epsilon > 0$ for all c_2 small enough. By Lemma 1 and the fact $d_1(p^{c_2}; \theta^{c_2}, \pi^{c_2}) \geq$

$d_1(0; ((1 - \gamma)\theta_1^{c_2}, \theta_2^{c_2}), \pi^0)$, we have

$$\begin{aligned} \limsup_{c_2 \downarrow 0} [c_1 d_1(\epsilon; \theta^{c_2}, \pi^{c_2}) + \epsilon - c_1 d_1(p^{c_2}, \theta^{c_2}, \pi^{c_2}) - p^{c_2}] \\ \leq \frac{c_1 C}{(C - \theta_1^0 \pi_1^0((\epsilon, +\infty)))^2} + \epsilon - \frac{c_1 C}{(C - \theta_1^0(1 - \gamma))^2}. \quad (9) \end{aligned}$$

But $\pi_1^0((\epsilon, +\infty)) \leq 1 - \pi_1^0(\{0\}) < 1$, so the right hand side of (9) is less than or equal to

$$\frac{c_1 C}{(C - \theta_1^0(1 - \pi_1^0(\{0\})))^2} + \epsilon - \frac{c_1 C}{(C - \theta_1^0(1 - \gamma))^2},$$

which is strictly negative for sufficiently small $\epsilon > 0$ and some $\gamma > 0$. Hence for such ϵ, γ

$$\begin{aligned} r_1 - c_1 d_1(\theta^{c_2}, \pi^{c_2}) - v^{c_2} - m_1(\theta^{c_2}, \pi^{c_2}) \\ < r_1 - c_1 d_1(\epsilon; \theta^{c_2}, \pi^{c_2}) - v^{c_2} - m_1(\epsilon; \theta^{c_2}, \pi^{c_2}) = 0, \end{aligned}$$

for sufficiently small c_2 , which contradicts the fact that $(\theta^{c_2}, \pi^{c_2})$ is an equilibrium for all c_2 .

We now proceed to prove part 2. If $r_1 < r_2$ then for all values of c_2 for which $r_1 < v^{c_2} < r_2$, the equilibrium allocation $\theta^{c_2} = (0, C)$ is efficient. So assume $r_1 > r_2$ in what follows. If $0 < r_1 - r_2 < c_1/C$ then again $\theta^{c_2} = (0, C)$ is efficient since type-1 flows will never find it beneficial to join as it is optimal to balk even at an empty system, so we are left with $r_1 - r_2 > c_1/C$. By part 1 and Lemma 2, (θ_1^0, π_1^0) is characterized as the equilibrium that results when type-2 flows are absent, and the delay independent reward in (1) for type-1 users is $r_1 - r_2$ (instead of r_1 originally). Since $\mathcal{A}_1 = \mathcal{P}(\mathcal{U}_0)$, by Proposition 1 we have that θ_1^0 maximizes

$$\theta_1 \left(r_1 - r_2 - \frac{c_1 \mu_1}{C - \theta_1} \right) = \mathcal{W}(\theta_1, C - \theta_1, 0) - r_2 C,$$

with respect to θ_1 , hence $\theta_1^0 = \theta_1^*$.

The next lemma shows that if the average rate of payments by type-2 flows is negligible and they tend to choose less aggressive parameters, then they do not interfere with type-1 flows in the limit.

Lemma 1. *Under the assumptions of Proposition 4, for every $p > 0$*

$$d_1(p; \theta^{c_2}, \pi^{c_2}) \rightarrow d_1(p; \theta^0, \pi^0) = \frac{C}{(C - \theta_1^0 \pi_1^0((p, +\infty)))^2}, \text{ as } c_2 \downarrow 0.$$

Proof. For fixed c_2 we calculate the instantaneous bandwidth $x_i(t)$ allocated to all type- i flows at time t when $n_i(t)$ type- i flows are present, for $i = 1, 2$. Furthermore, let $p_1, \dots, p_{n_1(t)}$ and $w_1, \dots, w_{n_2(t)}$ be the congestion control parameters

chosen by the active type-1 and type-2 flows respectively. $x_1(t), x_2(t)$ are given by the optimal solution of the optimization problem

$$\max_{x_1, x_2 \geq 0} \left(\max_{i=1, \dots, n_1(t)} p_i x_1 + \sum_{j=1}^{n_2(t)} w_j \log x_2 \right), \text{ where } x_1 + x_2 \leq C.$$

For fixed but arbitrary $\delta > 0$, $x_1(t)$ dominates $\hat{x}_1(t)$ the optimum x_1 in

$$\max_{x_1, x_2 \geq 0} \left(\max_{i=1, \dots, n_1(t)} I(t) p_i x_1 + \delta \log x_2 \right), \text{ where } x_1 + x_2 \leq C.$$

where $I(t) = 1$ if $\sum_{j=1}^{n_2(t)} w_j \leq \delta$ and zero otherwise.

Now,

$$\begin{aligned} P_{c_2}(I(t) = 1) &= 1 - P_{c_2} \left(\sum_{j=1}^{n_2(t)} w_j > \delta \right) \\ &\geq 1 - \frac{E_{c_2} \left(\sum_{j=1}^{n_2(t)} w_j \right)}{\delta} = 1 - \frac{m_2(\theta^{c_2}, \pi^{c_2})}{\delta} \rightarrow 1, \end{aligned}$$

as $c_2 \downarrow 0$. Thus, most of the time the departure rate is no less than $\mu_1 \hat{x}_1(t)$, which for small δ is arbitrarily close to C when $n_1(t) > 0$. Hence, as $c_2 \downarrow 0$ the distribution of $n_1(t)$ at stationarity converges to that of a birth-death process with birth rate $\theta_1^0 \mu_1$ and death rate C , and an arriving type-1 flow with parameter p will see an M/M/1 queue with arrival rate $\mu_1 \theta_1^0 \pi_1^0((p, +\infty))$ and departure rate $\mu_1 C$. That is, it will have the least priority in this queue, and so

$$d_1(p; \theta^0, \pi^0) = \frac{C}{\left(C - \theta_1^0 \int_p^\infty d\pi_1^0 \right)^2},$$

using standard results for the M/M/1 queue.

The next lemma shows that the limit (θ^0, π^0) of equilibria as $c_2 \downarrow 0$, is itself an equilibrium of a system where type-2 flows are completely insensitive to delay. This permits a narrowing down of the possible limits by analyzing a simpler system.

Lemma 2. *Under the assumptions of Proposition 4, (θ^0, π^0) satisfies,*

1. $r_1 - r_2 - c_1 d_1(\theta^0, \pi^0) - m_1(\theta^0, \pi^0) = 0$.
2. π_1^0 maximizes $r_1 - r_2 - c_1 d_1(\pi_1'; \theta^0, \pi^0) - m_1(\pi_1'; \theta^0, \pi^0)$ over $\pi_1' \in \mathcal{A}_1$.

Proof. By Lemma 1 and since $\pi^0(\{0\}) = 0$, we have

$$d_1(\theta^{c_2}, \pi^{c_2}) = \int_0^\infty d_1(p; \theta^{c_2}, \pi^{c_2}) d\pi_1^{c_2}(p) \rightarrow \int_0^\infty d_1(p; \theta^0, \pi^0) d\pi_1^0(p) = d_1(\theta^0, \pi^0),$$

as $c_2 \downarrow 0$. Moreover,

$$0 = r_1 - c_1 d_1(\theta^{c_2}, \pi^{c_2}) - v^{c_2} - m_1(\theta^{c_2}, \pi^{c_2}),$$

for all c_2 , and $m_1(\theta^{c_2}, \pi^{c_2}) = \int p d\pi_1^{c_2}$. Thus,

$$0 = r_1 - c_1 d_1(\theta^0, \pi^0) - m_1(\theta^0, \pi^0).$$

We now show part 2. We have,

$$\begin{aligned} r_1 - c_1 d_1(\pi'_1; \theta^{c_2}, \pi^{c_2}) - v^{c_2} - m_1(\pi'_1; \theta^{c_2}, \pi^{c_2}) \\ \leq r_1 - c_1 d_1(\theta^{c_2}, \pi^{c_2}) - v^{c_2} - m_1(\theta^{c_2}, \pi^{c_2}) \\ \leq r_1 - c_1 d_1(\theta^{c_2}, \pi^0) - v^{c_2} - \int_0^\infty p d\pi_1^{c_2}(p). \end{aligned}$$

Taking $c_2 \downarrow 0$ and using Lemma 1 yields

$$r_1 - c_1 d_1(\pi'_1; \theta^0, \pi^0) - r_2 - m_1(\pi'_1; \theta^0, \pi^0) \leq r_1 - c_1 d_1(\theta^0, \pi^0) - r_2 - m_1(\theta^0, \pi^0).$$

6 Discussion

In the simple model of this paper we find that congestion control mechanisms coming from linear utilities have better properties than those based on logarithmic ($\alpha = 1$). The linear utility functions \mathcal{U}_0 correspond to the FileTransfer algorithm in [3] and are used in [16, 17] along with slow timescale mechanisms that vary the slope of the linear utility according to user preferences over average rates rather than instantaneous one. But in practice, users decide whether their flows join or balk on the basis of some knowledge of the state, e.g., congestion level. Models allowing this type of information are surveyed in [11] and may prove helpful in the context of this paper.

Notice that logarithmic utility functions produce inefficient outcomes when users have linear valuations in average download delay. But this does not immediately imply the inadequacy of similarly behaving algorithms such as TCP. Since users might value higher moments of delay as well, e.g., variance, it will be interesting to account for these as well.

The existence and uniqueness of equilibria when utility functions vary over larger sets is an open problem, even in the single type case. For linear and logarithmic utility functions existence is shown by explicitly computing the effect of different parameters on net value per bit. For general utility functions we expect that more general methods such as fixed point theorems will prove fruitful.

Finally, in this paper we have not determined the welfare loss of inefficient outcomes. It is an open problem whether similar statements as in [12] regarding the price of anarchy, hold in the present model.

Acknowledgements. The authors would like to thank Bob Briscoe and Damon Wischik for useful discussions and suggestions on the subject of this paper, as well as the anonymous reviewers for their helpful comments. This research was supported by Trilogy (<http://www.trilogy-project.org>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme.

References

1. Kelly, F.P.: Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications* 8, 33–37 (1997)
2. Mo, J., Walrand, J.: Fair End-to-End Window-Based Congestion Control. *IEEE/ACM Transactions on Networking* 8, 556–567 (2000)
3. Gibbens, R., Kelly, F.P.: Resource Pricing and the Evolution of Congestion Control. *Automatica* 35, 1969–1985 (1999)
4. Naor, P.: The Regulation of Queue Size by Levying Tolls. *Econometrica* 37, 17–24 (1969)
5. Edelson, M., Hildebrand, K.: Congestion Tolls for Poisson Queueing Processes. *Econometrica* 43, 81–92 (1975)
6. MacKie-Mason, J., Varian, H.: Pricing Congestible Network Resources. *IEEE Journal on Selected Areas in Communications* 13, 1141–1149 (1995)
7. Gallager, R.: A Minimum Delay Routing Algorithm Using Distributed Computation. *IEEE Transactions on Communications* 25, 73–85 (1977)
8. Hassin, R.: On the Optimality of First Come Last Served Queues. *Econometrica* 53, 201–202 (1985)
9. Hassin, R.: Decentralized Regulation of a Queue. *Management Science* 41, 163–173 (1995)
10. Haviv, M., Van der Wal, J.: Equilibrium Strategies for Processor Sharing and Random Queues with Relative Priorities. *Probability in the Engineering and Informational Sciences* 11, 403–312 (1997)
11. Hassin, R., Haviv, M.: *To Queue or not to Queue: Equilibrium Behavior in Queueing Systems*. Kluwer Academic Publishers, Boston/Dordrecht/London (2005)
12. Johari, R., Tsitsiklis, J.: Efficiency Loss in a Network Resource Allocation Game. *Mathematics of Operations Research* 29, 407–435 (2004)
13. Hajek, B., Gopalakrishnan, G.: Do Greedy Autonomous Systems Make for a Sensible Internet? In: *Conference on Stochastic Networks*, Stanford University (2002)
14. Mendelson, H., Whang, S.: Optimal Incentive-Compatible Priority Pricing for the M/M/1 Queue. *Operations Research* 38, 870–883 (1990)
15. Mandjes, M.: Pricing Strategies Under Heterogeneous Service Requirements. In: *IEEE INFOCOM 2003*, pp. 1210–1220. IEEE Press, San Francisco (2003)
16. Key, P., Massoulié, L., Vojnovic, M.: Farsighted Users Harness Network Time-Diversity. In: *IEEE INFOCOM 2005*, pp. 2382–2394. IEEE Press, Miami (2005)
17. Courcoubetis, C., Dimakis, A.: Congestion Control and Pricing of Volume Intensive Applications. To appear at the 21st International Teletraffic Congress. Full version available at <http://nes.aueb.gr/research/volume.pdf> (2008)