CHARGE-SENSITIVE TCP AND
RATE CONTROL IN THE INTERNET

by

Richard J. La and Venkat Anantharam

CHARGE-SENSITIVE TCP AND
RATE CONTROL IN THE INTERNET

by

Richard J. La and Venkat Anantharam

ELECTRONICS RESEARCH LABORATORY

Charge-Sensitive TCP and Rate Control in the Internet

Richard J. La and Venkat Anantharam

Department of Electrical Engineering & Computer Sciences

University of California at Berkeley

{hyongla, ananth}@eecs.berkeley.edu

Abstract

In a communication network a good rate allocation algorithm should reflect the utilities of the users while being fair. We investigate this fundamental problem of achieving the system optimal rates in the sense of maximizing aggregate utility, in a distributed manner, using only the information available at the end hosts of the network. This is done by decomposing the overall system problem into subproblems for the network and for the individual users, by introducing a pricing scheme. The users are to solve the problem of maximizing individual net utility, which is the utility less the amount they pay. We provide algorithms for the network to adjust its prices and the users to adjust their window sizes such that at an equilibrium the system optimum is achieved. Further, the equilibrium prices are such that the system optimum achieves weighted proportional fairness. It is notable that the update algorithms of the users do not require any explicit feedback from the network, rendering them easily deployable over the Internet. Our scheme is incentive compatible in that there is no benefit to the users to lie about their utilities.

i

1 Introduction

As the Internet explodes in size and in the number of users, one of challenging questions network designers face is how to provide a fair and efficient allocation of the available bandwidth. To this end researchers have proposed many different rate allocation mechanisms. In the current Internet, most connections use variants of the Transmission Control Protocol (TCP), which is a window-based congestion control mechanism. It is widely recognized, however, that TCP does not generally lead to a fair or efficient allocation of bandwidth among the connections [3, 17, 5]. The fact that the Internet is now in the public domain, and thus in a potentially noncooperative environment, has stimulated much work on pricing mechanisms to ensure that users do not misbehave and to provide quality of service in accordance with users' willingness to pay. Researchers in this area have proposed different schemes based on time and volume measurements [11] or on per-packet charges [20]. Furthermore, research has shown that flat rate charging leads to inefficiency, where a large number of low-usage users end up subsidizing a small number of high-usage users [4]. This argues that usage-based pricing is desirable. A good discussion on the market structure is given in [15]. They present several different possible market structures for network pricing, such as a competitive market and monopoly, and provide insight as to how pricing schemes should be designed based on the market structure.

Recently Kelly [9] has suggested that the problem of rate allocation should be posed as one of achieving maximum aggregate utility for the users. Since the solution needs to be decentralized and should be incentive compatible in that users should have no incentive to lie about their true utilities, Kelly proposes using pricing to decompose the overall system problem into subproblems for the network and for the individual users, with the goal of designing algorithms such that a system optimum is achieved when users' choices of rates they are willing to pay for and the network's choice of allocated rate are in equilibrium. Kelly *et al.*[10] have proposed two classes of algorithms which can be used to implement proportionally fair pricing and solve a *relaxation* of system optimum problem.

The algorithm in [10] requires explicit feedback from the switches inside network. Due to the size of the Internet, an algorithm that requires an extensive modification inside the network may not be suitable for deployment. Further, many connections in the Internet use TCP to control their transmission rates. In this paper we investigate the fundamental problem of achieving the system optimum that maximizes the aggregate utility of the users, using only the information available at the end hosts. We introduce two algorithms that can be deployed over the Internet without significant modification within the network. These algorithms can be thought of as a natural

extension of an algorithm introduced in [18]. These algorithms effectively decompose the system problem into a network problem and a problem for each user, by introducing a pricing mechanism. In the first algorithm, while the users solve the user optimization problems on a larger time scale, the underlying window-based transmission rate control mechanism solves the network problem on a smaller time scale. The second algorithm does not require users to solve the user optimization problem. Instead, it incorporates the user optimization problem into the window size updating rules. The unique fixed point of the mapping defined from these algorithms is proven to result in the system optimal rates. Further, the price a user pays equals the delay cost it imposes on the other users[1]. We show that our algorithms do not require any explicit feedback from the network, and can be easily deployed in the current Internet. We also demonstrate that these algorithms are incentive compatible: it is always in users' own interest to tell the truth about their utility functions.

The rest of the paper is organized as follows. We first motivate the problem and describe the model used in the analysis and the methodology adopted, followed by some related work. Section 4 explains the pricing scheme, which is followed by the algorithms in sections 5 and 6. We present two numerical examples in section 7. We briefly discuss how the backlog in the network can be controlled using our pricing scheme in section 8 and then finish by drawing some conclusions and indicating some directions for future research in section 9.

2   Motivation, Background & Model

In this section we first motivate our problem by describing some of the currently most popular congestion control mechanism used in the Internet. We discuss some of the problems with the current congestion control mechanisms. We then formulate our problem of maximizing the total utility of the users and describe the model we use to present our congestion control proposal.

2.1   Motivation

From its advent the Internet has been decentralized, relying on disciplined behavior from the users. Without a centralized authority, the network users have a great deal of freedom as to how they share the available bandwidth in the network. The increasing complexity and size of the Internet renders centralized rate allocation control impractical. In view of these constraints, researchers have proposed many rate allocation algorithms to be implemented at the end hosts in a decentralized manner [10, 18, 14, 8]. These algorithms can be roughly categorized into two classes: rate-based

---

[1]This can be also interpreted as an externality cost imposed on the network.

algorithms and window-based algorithms. A rate-based algorithm directly controls the transmission rate of the connection, based on either feedback from the network [10] or on measurements taken at the end host. A window-based algorithm adjusts the congestion window size, which is the maximum number of outstanding packets, in order to control the transmission rate and backlog inside the network associated to the connection.

The most widely deployed flow/congestion control mechanism in the current Internet is the Transmission Control Protocol (TCP), which is a window-based congestion control mechanism. TCP however does not necessarily lead to a fair or efficient rate allocation of the available bandwidth. It is well known that TCP[2] as currently implemented suffers from a high packet loss rate and a delay bias [13, 17, 2]. The high packet loss rate is a consequence of periodic oscillation of the window sizes and aggressive slow start, while the delay bias is the result of a discrepancy in window update rates among different connections. In order to address these issues Brakmo *et al.*[2] have introduced another version of TCP, named TCP Vegas, with a fundamentally different bandwidth estimation scheme. They have demonstrated that TCP Vegas results in a lower packet loss rate, which in turn leads to higher efficiency, and have suggested that TCP Vegas does not suffer from a delay bias, which leads to a more fair rate allocation[3] of the bandwidth. This was later proved by Mo *et al.*[17] in the case that there is a single bottleneck in the network.

Mo and Walrand [18] in another paper, have proposed and studied another fair window-based end-to-end congestion control mechanism, which is similar to TCP Vegas but has a more sophisticated window updating rule. They have shown that proportional fairness[4] can be achieved by their $(p, 1)$-proportionally fair algorithm and that max-min fairness can be achieved as a limit of (p, α)-proportionally fair algorithms as α goes to ∞ [18].

Clearly fairness is a desirable property for a rate allocation mechanism. There is, however, another aspect of a rate allocation problem that has not been addressed by the previous mechanisms. Due to the various requirements of different applications it is likely that the users will have different utility functions. For instance, users running video conferencing tools are likely to receive a higher utility than another user surfing the World-Wide Web (WWW) on its free time. This suggests that although fairness is a desirable property, fairness alone may not be a good objective. We suggest that a good rate allocation mechanism should not only be fair, but should also allocate the available bandwidth in such a way that the overall utility of the users is maximized. In this paper we describe a pricing mechanism that achieves these goals without requiring knowledge of

---

[2]The most popular versions of TCP in the Internet are TCP Tahoe and Reno.

[3]Fairness refers to max-min fairness that is discussed in section 2.2.

[4]Proportional fairness is introduced by Kelly in [9], and is discussed in the next subsection.

the users' utility functions and without requiring any explicit feedback from the network.

2.2 Model and Background

We consider a network with a set $\mathcal{J}$ of resources or links and a set $\mathcal{I}$ of users. Let C_j denote the finite capacity of link $j \in \mathcal{J}$. Each user has a fixed route $\mathcal{J}_i$, which is a non-empty subset of $\mathcal{J}$.[5] We define a zero-one matrix A, where $A_{i,j} = 1$ if link j is in user i's route $\mathcal{J}_i$ and $A_{i,j} = 0$ otherwise. When the throughput of user i is x_i, user i receives utility $U_i(x_i)$. We assume that the utility $U_i(x_i)$ is an increasing, strictly concave and continuously differentiable function of x_i over the range $x_i \geq 0$.[6] Furthermore, we assume that the utilities are additive so that the aggregate utility of rate allocation $x = (x_i, i \in \mathcal{I})$ is $\sum_{i \in \mathcal{I}} U_i(x_i)$. Let $U = (U_i(\cdot), i \in \mathcal{I})$ and $C = (C_j, j \in \mathcal{J})$. In this paper we study the feasibility of achieving the maximum total utility of the users in a distributed environment, using only the information available at the end hosts. Under our model this problem can be formulated as follows.

SYSTEM(U,A,C):

$$
\begin{aligned}
maximize \quad & \sum_{i \in \mathcal{I}} U_i(x_i) \\
subject\ to \quad & A^T x \leq C \\
over \quad & x \geq 0
\end{aligned}
\tag{1}
$$

The first constraint in the problem says that the total rate through a resource cannot be larger than the capacity of the resource. Given that the system knows the utility functions U of the users, this optimization problem may be mathematically tractable. However, in practice not only is the system not likely to know U, but also it is impractical for a centralized system to compute and allocate the users' rates due to the large scale of the network. Hence, Kelly in [9] has proposed to consider the following two simpler problems.

Suppose that given the price per unit flow λ_i, user i selects an amount to pay per unit time, p_i, and receives a flow $x_i = \frac{p_i}{\lambda_i}$.[7] Then, the user's optimization problem becomes the following [9].

USER$_i$(U$_i$; λ_i) :

$$
maximize \quad U_i(\frac{p_i}{\lambda_i}) - p_i
\tag{2}
$$

---

[5]We assume that there is a routing mechanism and do not discuss the issue of routing in this paper.

[6]Such a user is said to have elastic traffic.

[7]This is equivalent to selecting its rate x_i and agreeing to pay $p_i = x_i \cdot \lambda_i$.

$$over \quad p_i \geq 0$$

The network, on the other hand, given $p = (p_i, i \in \mathcal{I})$, attempts to maximize the sum of weighted log functions $\sum_{i \in \mathcal{I}} p_i \log(x_i)$. Then the network's optimization problem can be written as follows [9].

NETWORK(A,C;p) :

$$maximize \quad \sum_{i \in \mathcal{I}} p_i \log(x_i) \tag{3}$$
$$subject\ to \quad A^T x \leq C$$
$$over \quad x \geq 0$$

Note that the network does not require the true utility functions $(U_i(\cdot), i \in \mathcal{I})$, and pretends that user i's utility function is $p_i \cdot \log(x_i)$ to carry out the computation. It is shown in [9] that one can always find vectors $\lambda^* = (\lambda_i^*, i \in \mathcal{I}), p^* = (p_i^*, i \in \mathcal{I})$, and $x^* = (x_i^*, i \in \mathcal{I})$ such that p_i^* solves $USER_i(U_i; \lambda_i^*)$ for all $i \in \mathcal{I}$, x^* solves $NETWORK(A, C; p^*)$, and $p_i^* = x_i^* \cdot \lambda_i^*$ for all $i \in \mathcal{I}$. Further, the rate allocation x^* is also the unique solution to $SYSTEM(U, A, C)$. This suggests that the problem of solving $SYSTEM(U, A, C)$ can be achieved by an algorithm that solves $NETWORK(A, C; p(t))$ for a given $p(t)$ at time t on a smaller time scale, and drives $p(t)$ to p^* on a larger time scale.

Another important aspect of a rate allocation mechanism is fairness. There are many different definitions for fairness, the most commonly accepted one being max-min fairness. A rate allocation is max-min fair if a user's rate cannot be increased without decreasing the rate of another user who is already receiving a smaller rate. In this sense, max-min fairness gives an absolute priority to the users with smaller rates. However, often in order to achieve a max-min fair allocation the optimality of the network needs to be sacrificed as discussed in [16]. In order to handle this tradeoff between fairness and optimality Kelly [9] has proposed another definition of fairness. A vector of rates $x' = (x_i', i \in \mathcal{I})$ is said to be *weighted proportionally fair*[8] with weight vector p if x' is feasible, i.e., $x' \geq 0$ and $A^T x' \leq C$, and for any other feasible vector x, the following holds:

$$\sum_{i \in \mathcal{I}} p_i \frac{x_i - x_i'}{x_i'} \leq 0$$

Note that a rate allocation x solves $NETWORK(A, C; p)$ if and only if it is weighted proportionally fair with weight vector p.

Let us demonstrate the differences between max-min and proportional fairness by an example. Clearly max-min and proportional fairness are the same in a single link case. Thus, we consider a

---

[8]This is called *rates per unit charge are proportionally fair* in [9]

Figure 1: Max-min and proportional fairness.

multiple link case. Consider the example in Figure 1. There are three users that share the network with two links. One can easily see that the unique max-min fair allocation in this example is $(x_1, x_2, x_3) = (\frac{1}{2}, \frac{1}{2}, \frac{1}{2})$, and every user receives the same rate. By solving $NETWORK(A, C; \underline{1})$, where $\underline{1}$ is a vector, whose elements are all 1, one can show that the proportionally fair allocation is $(x_1, x_2, x_3) = (\frac{2}{3}, \frac{2}{3}, \frac{1}{3})$. Note that the sum of rates received from all links is the same for all users under the proportionally fair allocation.

3 Related Work

In this section we briefly describe some of the past work we draw upon and discuss how it is related to the problem we address in this paper.

3.1 Rate-based Algorithms

Based on the idea of proportional fairness Kelly *et al.*[10] have proposed two classes of *rate-based* algorithms that solve a relaxation of $SYSTEM(U, A, C)$ problem. These algorithms are based on the idea of shadow price charged at a resource, which is a function of the total rate going through the resource. We describe one of the algorithms, which they call the *primal algorithm*, in this subsection.

Suppose that every user adopts a rate-based flow control. Let w_i and $x_i(t)$ denote user i's willingness to pay per unit time and its rate at time t, respectively. Now suppose that at time t each resource $j \in \mathcal{J}$ charges a price per unit flow of $\mu_j(t) = p_j(\sum_{i:j \in \mathcal{J}_i} x_i(t))$, where $p_j(\cdot)$ is an increasing function of the total rate going through it. Consider the system of differential equations

$$\frac{d}{dt}x_i(t) = \kappa(w_i - x_i(t) \sum_{j \in \mathcal{J}_i} \mu_j(t)) \tag{4}$$

where

$$\mu_j(t) = p_j\left(\sum_{i:j \in \mathcal{J}_i} x_i(t)\right).$$

6

These equations can be motivated as follows. Each user first sets a price per unit time it is willing to pay. Then, every user adjusts its rate based on the feedback provided by the resources in the network in such a way that at an equilibrium the price it is willing to pay equals its aggregate cost. The feedback from a resource $j \in \mathcal{J}$ can be thought of as a congestion indicator, requiring a reduction in the flow rates going through the resource.

Kelly *et al.* show that, for fixed w_i, under some conditions on $p_j(\cdot), j \in \mathcal{J}$, the above system of differential equations converges to a point that maximizes the following expression

$$\mathcal{U}(x) = \sum_i w_i \cdot \log x_i - \sum_j \int_0^{\sum_{i:j \in \mathcal{J}_i} x_i} p_j(y) dy.$$

Further, if the users update their willingness to pay $w_i(t)$ according to

$$w_i(t) = x_i(t) \cdot U_i'(x_i(t))$$

while $x_i(t)$ evolves according to (4), then $x(t)$ converges to a unique stable point that maximizes the following revised expression

$$\mathcal{U}(x) = \sum_i U_i(x_i) - \sum_j \int_0^{\sum_{i:j \in \mathcal{J}_i} x_i} p_j(y) dy. \tag{5}$$

Note that the first term in (5) is the objective function in our *SYSTEM(U, A, C)* problem. Thus, the algorithm proposed by Kelly *et al.* solves a relaxation of the $SYSTEM(U, A, C)$ problem.

3.2 TCP and a Fair End-to-end Congestion Control Algorithm

Unlike a connection with a rate-based congestion control algorithm, a TCP connection adjusts its rate by updating its window size, based on the estimated congestion state of the network. There are several different versions of TCP, which can be categorized into two classes, based on their bandwidth estimation schemes. TCP Tahoe and Reno use packet losses as an indication of congestion in the network, while TCP Vegas [2] and the algorithm proposed by Mo and Walrand [18] use the estimated queue size to adjust the congestion window size[9].

Over the years researchers have observed that the most widely used versions of TCP, which are TCP Tahoe and Reno, exhibit several undesirable characteristics such as a delay bias and a high packet loss rate [5, 6, 17]. In order to deal with these issues Mo and Walrand [18] have investigated the existence of a fair window-based end-to-end congestion control algorithm that updates the congestion window size more intelligently.

---

[9]We call these loss-based and queue-based TCP, respectively.

7

We first present a fluid model of the network that describes the relationship between the window sizes, rates, and queue sizes. Throughout the paper we assume that the switches exercise the first-in-first-out (FIFO) service discipline. This model can be represented by the following equations:

$$A^T x - C \leq 0 \tag{6}$$

$$Q(A^T x - C) = 0 \tag{7}$$

$$X(AC'q + d) = w \tag{8}$$

$$x \geq 0, q \geq 0, \tag{9}$$

where

$$C = (C_1, ..., C_J)^T, q = (q_1, ..., q_J)^T,$$

$$d = (\bar{d}_1, ..., \bar{d}_I)^T, w = (w_1, ..., w_I)^T,$$

$$X = diag(x), C' = diag(C_1^{-1}, ..., C_J^{-1}), Q = diag(q),$$

w_i is the congestion window size of user i, $\bar{d}_i$ is the propagation delay of route $\mathcal{J}_i$ not including the queueing delay, and q_j denotes the backlog at link j buffer. For simplicity of analysis we assume that the buffer sizes are infinite. The first condition in (6) represents the capacity constraint. The second constraint says that there is backlogged fluid at a resource only if the total rate through it equals the capacity. The third constraint follows from that the window size of connection i should equal the sum of the amount of fluid in transmit and the backlogged fluid in the buffers, i.e.,

$$w_i = x_i \cdot \bar{d}_i + q^i,$$

where q^i denotes connection i's total backlog in the buffers.

Let $W = [0, w_{1,max}] \times \cdots \times [0, w_{I,max}]$, where $w_{i,max}$ is connection i's maximum congestion window size announced by the receiver and is assumed to be sufficiently large, and $\mathcal{X} = [0, \min_{j \in \mathcal{J}_i} C_j] \times \cdots \times [0, \min_{j \in \mathcal{J}_i} C_j]$. It has been shown in [18] that the rate vector x is a well defined function of the window sizes w, and we can define a function $\mathcal{W}$ that maps a window size vector $w \in W$ to a rate vector $x \in \mathcal{X}$.

Under the $(p,1)$-proportionally fair algorithm by Mo and Walrand [18] each connection has a target queue size $p_i > 0$ and attempts to keep p_i packets at the switch buffers. Let $d_i(t), w_i(t)$, and $x_i(t)$[10] denote the round trip delay, the congestion window size, and the rate of connection i at time t, respectively. Suppose that each connection i has a fixed target queue size p_i. Connection i updates its window size $w_i(t)$ according to the following differential equation

$$\frac{d}{dt} w_i(t) = -\kappa \frac{\bar{d}_i}{d_i(t)} \frac{s_i(t)}{w_i(t)},$$

---

[10]We use $x_i(t)$ to denote $x_i(w(t))$ when there is no confusion.

where κ is some positive constant, and $s_i(t) = w_i(t) - x_i(t) \cdot \bar{d}_i - p_i$. Under this algorithm the window sizes converge to a unique point w^* such that for all $i \in \mathcal{I}$

$$w_i^* - x_i(w^*) \cdot \bar{d}_i = p_i,$$

where $x_i(w^*)$ is connection i's throughput when the window sizes are w^*. They show that at the unique stable point w^* of the algorithm the resulting allocation $x(w^*)$ is weighted proportionally fair.

Suppose that users update their window sizes according to the following system of differential equations

$$\frac{d}{dt}w_i(t) = -\kappa x_i(t)s_i(t)u_i(t),$$

where

$$s_i(t) = w_i(t) - x_i(t) \cdot \bar{d}_i - \frac{p_i}{(x_i(t) + 1)^{\alpha-1}}$$

and

$$u_i(t) = \bar{d}_i - (\alpha - 1)\frac{p_i}{(x_i(t) + 1)^{\alpha}}.$$

This algorithm is called a (p, α)-proportionally fair algorithm. They prove that the above algorithm converges to a unique stable point of the system for all fixed α and the max-min fair allocation is achieved as a limit as $\alpha \to \infty$. However, since their work is motivated by the fundamental question of the existence of a fair end-to-end congestion control mechanism, they have not considered the problem of maximizing the aggregate utility of the users, while maintaining fairness.

4   Pricing Scheme for Charge-Sensitive TCP

The algorithms suggested in [10, 7] are based on the assumptions that the network can provide the necessary feedback to the users, and users adjust their *rates* based on the feedback information. However, in the current Internet, many, if not most, connections use TCP, which is a window-based congestion control mechanism, to control their rates. Thus, users control the rates only through the window sizes.

There are several arguments for a window-based congestion control algorithm in the Internet over a rate-based algorithm. One of the arguments is the stability of the Internet. Suppose that users use a rate-based congestion control algorithm. If they have incorrect estimates of the available bandwidth and release packets into the network at a rate that is much higher than they should, the

9

network could temporarily experience an extremely high packet loss rate due to buffer overflows and may take a long time to recover from it. A window-based congestion control algorithm not only controls the transmission rate, but also limits the maximum number of outstanding packets according to the congestion window size. This helps alleviate the long term effect of the inaccurate estimation of the available bandwidth by the users. This is an important advantage of a window-based algorithm. In an open system such as the Internet, it is important to control the number of packets the connections can keep within the network for stability and to ensure that no users can arbitrarily penalize other users by increasing their rates during congestion periods due to incorrect estimates.

4.1 Pricing Scheme

Section 3.2 tells us that, given $(p_i, i \in \mathcal{I})$, the users can reach a solution to *NETWORK(A,C;p)* using a window-based congestion control mechanism, namely the $(p, 1)$-proportionally fair algorithm of Mo and Walrand. Throughout this subsection, we refer to this $(p, 1)$-proportionally fair algorithm when we say TCP. The challenge now is to design a mechanism that drives the users to the right p^* where the resulting rate allocation solves *SYSTEM(U, A, C)*. In this subsection we describe a simple pricing mechanism that can achieve this. We show that the price a user pays can be directly computed by the user without any feedback from the network, by using the same mechanism already built into TCP.

When the total rate through a link is strictly smaller than its capacity there is no congestion, as each user receives its desired rate without contention from other users. However, when the total rate reaches the link capacity, any increase in a congestion window size by a user i in an attempt to increase its own rate results in increased backlog at the resource. This leads to higher queueing delay at the resource. If the users are delay sensitive this increase in queueing delay represents an increase in the implicit[11] cost the users pay through larger delay. From the perspective of the network, this increase in queueing delay may be interpreted as an increase in undesirable delay cost for the system. In other words, the queueing delay caused by a user can be interpreted as the delay cost it imposes on other users [15].

Suppose that the network attempts to recover the increase in system cost due to queueing delay through a pricing mechanism as follows. Let $q_j, j \in \mathcal{J}$, denote the backlog at resource j. When the total rate through the resource is strictly smaller than its capacity, we assume that there is no backlog, from (7). When resource $j \in \mathcal{J}$ is congested, i.e., the total rate through it equals its

---

[11]We say the cost is implicit because the users do not necessarily have to pay in monetary form.

capacity, the resource charges a price, where the price per unit flow g_j is the queueing delay at the resource, i.e., $g_j = \frac{q_j}{C_j}$. Hence, the total price per unit flow for a user with route $\mathcal{J}_i \subset \mathcal{J}$ is $\sum_{j \in \mathcal{J}_i} g_j$, and the total price per unit time user i pays is $x_i \cdot \sum_{j \in \mathcal{J}_i} g_j$. Then, the net benefit or the objective function of the user is given by

$$U_i(x_i) - x_i \cdot \sum_{j \in \mathcal{J}_i} g_j = U_i(x_i) - x_i \cdot \sum_{j \in \mathcal{J}_i} \frac{q_j}{C_j} \tag{10}$$

We have, however, claimed that no information is explicitly fed back to the end hosts from the network. Hence, the switches are not allowed to send any information regarding the current price per unit flow back to the end users.

In order to maximize the objective function in (10) given the *current prices per unit flow*, a user only needs to know the total price per unit flow of its route, but not the price at each resource. We now show that users using TCP can compute their prices per unit flow without any help from the network. Suppose that each connection knows the propagation delay of its route. In practice this is done by keeping track of the minimum round trip time of the packets [2, 17]. Suppose that the target queue sizes of the users are given by $p = (p_i, i \in \mathcal{I})$, and the users' window sizes converge to the stable point of TCP, where the resulting rates are weighted proportionally fair with the weight vector p. Then, the price per unit time user i pays, h_i, can be computed as follows:

$$h_i = x_i \cdot \sum_{j \in \mathcal{J}_i} \frac{q_j}{C_j} = \sum_{j \in \mathcal{J}_i} q_j \cdot \frac{x_i}{C_j} = \sum_{j \in \mathcal{J}_i} q_j^i = p_i, \tag{11}$$

where q_j^i is connection i's queue size at resource j. Here we have implicitly assumed that the queue size of each connection at a congested resource is proportional to its rate, which is a consequence of the assumption that the switches exercise FIFO service discipline. Therefore, at the stable point of TCP for a fixed p, the price of a user equals its target queue size, and the user can infer its own price from its target queue size. This eliminates the need for any feedback from the network. One thing to note is that the rate of a user depends not only on its own target queue size, but also on those of other users. Hence, the prices per unit flow at the resources depend on the congestion level in the network.

One important aspect of a pricing scheme is fairness. The price a connection pays for using resource $j \in \mathcal{J}$ should be proportional to its rate. In other words, the price per unit flow for each connection at a resource, should be the same. This is obviously the case with the above pricing scheme. Moreover, it is easy to see that when a new user comes into the network, the price the new user pays is exactly the increase in the total system cost, i.e., the increase in the total queueing delay experienced by the packets. This can be seen from that the price user i pays per unit time

equals its target queue size p_i and the total system cost per unit time is given by

$$\sum_{j \in \mathcal{J}} C_j \cdot g_j \;=\; \sum_{j \in \mathcal{J}} C_j \cdot \frac{q_j}{C_j} = \sum_{j \in \mathcal{J}} q_j = \sum_{i \in \mathcal{I}} p_i.$$

Therefore, the fairness of the pricing scheme is automatically guaranteed.

In the following two sections, based on this pricing mechanism, we propose two algorithms that can be deployed over the current Internet without an extensive modification inside the network. All that is required is a modification of the already existing TCP at the end hosts. We demonstrate, using a fluid model, that at an equilibrium of the algorithm the resulting rates are the optimal rates that solve $SYSTEM(U, A, C)$.

5   Algorithm I

In the previous section we have analyzed the case where users have fixed their target queue sizes $p = (p_i, i \in \mathcal{I})$. However, as more intelligence is embedded in end systems in the future, the users may be able to vary the parameters $p_i, i \in \mathcal{I}$, to maximize their net utility given by (10). In this section we propose a user algorithm and show that this algorithm has a unique equilibrium, at which it yields the optimal rates that solve $SYSTEM(U, A, C)$.

For the rest of this paper we assume that the utility functions satisfy the technical assumptions of Appendix A. These assumptions are satisfied by the most commonly used utility functions such as $U(x) = a \cdot \log(x + b)$ and $U(x) = c \cdot x^d$, where $a > 0$, $0 \le b \le 1$, $c > 0$, and $0 < d < 1$ are arbitrary constants.

Suppose that ϵ is some small positive constant and users update their willingness to pay or price $p_i, i \in \mathcal{I}$, at time t according to

$$p_{i,\epsilon}(t) = \begin{cases} \hat{p}_i(\lambda_i(t^-)) & \text{if } \hat{p}_i(\lambda_i(t^-)) \ge \epsilon \\ \epsilon & \text{otherwise} \end{cases} \tag{12}$$

where $\lambda_i(t^-)$ is the price per unit flow along user i's route at time t^- and $\hat{p}_i(\lambda) = \arg\max_{p_i} U_i(\frac{p_i}{\lambda}) - p_i$. We assume that the price updates take place on a much larger time scale, while users allow their window sizes to converge to a point close to the unique stable point of TCP for fixed $(p_i(t), i \in \mathcal{I})$. This is a natural assumption since the window sizes typically take only seconds to converge and users are not likely to keep changing their parameters before estimating the current throughput. The intuition behind the updating rule is as follows. At time t, based on the price per unit flow at time t^-, user i computes its optimal price that maximizes its net utility. If the current price per unit flow is too high user i prefers to wait till the price per unit flow is lower. In such a case user

i needs to probe the network for the available bandwidth and price per unit flow along its route. In order to measure the residual bandwidth not used by the other users, if there is any, and the price per unit flow, user i needs to set its window size large enough so that it utilizes all residual capacity not taken by the other users. Hence, we assume that user i sets the target queue size, which is its willingness to pay in our model, to some small positive constant ϵ that is arbitrarily small, so that each user can estimate its well defined price per unit flow along the route, which is strictly positive. We now consider the limit as $\epsilon \downarrow 0$, where (12) can be written as

$$
\begin{aligned}
p_i(t) &= \arg\max_{p_i} U_i\left(\frac{p_i}{\lambda_i(t^-)}\right) - p_i \\
&= \begin{cases} 0, & \text{if } \lambda_i(t^-) \geq U_i'(0) \\ p_i \text{ such that } U_i'\left(\frac{p_i}{\lambda_i(t^-)}\right) = \lambda_i(t^-), & \text{if } 0 < \lambda_i(t^-) < U_i'(0) \end{cases}
\end{aligned} \tag{13}
$$

where $U_i'(\cdot) = \frac{\partial U_i(\cdot)}{\partial x_i}$. The updating rule in (13) can be motivated as follows. If the price per unit flow λ_i is larger than or equal to $U_i'(0)$, then user i receives a negative net utility from any positive p_i. Thus, user i should wait till the price per unit flow is smaller than $U_i'(0)$. If the price per unit flow λ_i is strictly smaller than $U_i'(0)$, then there exists a unique solution to the $USER_i(U_i; \lambda_i)$ problem in (2). This solution is the unique p_i such that $U_i'(\frac{p_i}{\lambda_i}) = \lambda_i$, which maximizes user i's net utility. Hence, user i should set its price to such p_i.

Equation (13) implicitly assumes that a user does not anticipate the effect of its own action on prices [10] and updates its parameter in a myopic manner. This is a reasonable assumption when the size of the network is large and each user occupies at most a small fraction of a resource in the network or when a user cannot correctly compute the price per unit flow as a function of its own p_i. In some simple cases, however, users may be able to correctly estimate the effect of their own actions on the prices. This is discussed in [12].

Let us define a mapping $\mathcal{T} : \mathcal{P} \to \mathcal{P}$ to be

$$
\mathcal{T}_i(p) = \arg\max_{\tilde{p}_i} U_i\left(\frac{\tilde{p}_i}{\lambda_i(p)}\right) - \tilde{p}_i \tag{14}
$$

where $\mathcal{P} = \Re_+^I$, and $\lambda(p)$ is the price per unit flow vector at the unique stable point of TCP with price vector p. A fixed point p^* of the mapping $\mathcal{T}$ is a vector in $\mathcal{P}$ such that $\mathcal{T}(p^*) = p^*$.[12]

---

[12]The fixed point of this mapping should be interpreted as the limit of a sequence of fixed points $p(n)$ of the mappings $T_{\epsilon(n)}$, where

$$
T_{i,\epsilon(n)}(p) = \begin{cases} \hat{p}_{i,\epsilon(n)}(\lambda_i(p)) & \text{if } \hat{p}_i(\lambda_i(p)) \geq \epsilon(n) \\ \epsilon(n) & \text{otherwise} \end{cases}
$$

and $\epsilon(n) \downarrow 0$.

Theorem 1 *There exists a unique fixed point p^* of the mapping $\mathcal{T}$, and the resulting rate allocation from p^* is the optimal rate allocation x^* that solves $SYSTEM(U, A, C)$.*

Proof: The proof is given in Appendix B. ∎

Theorem 5 tells us that when the users adopt the algorithm given in this section, at an equilibrium where no user i changes its parameter p_i the resulting rate allocation coincides with the system optimal rate allocation. This demonstrates that solving $SYSTEM(U, A, C)$ could be accomplished using a simple window-based TCP algorithm in a distributed environment with no additional feedback from the network.


Figure 2: Single bottleneck case.

A natural question that arises now is whether or not the user prices $p_i, i \in \mathcal{I}$, converge to the unique fixed point of the mapping $\mathcal{T}$. Due to the complex relationship between the price vector and the resulting rate allocation, showing the convergence of users prices in a general network is a challenging problem. In this section we only show the convergence of user prices in the case of a single bottleneck link, i.e., every user has only one and the same bottleneck link. In such a single bottleneck case the price per unit flow is same for all users. Hence, each user updates its price based on the current price per unit flow of the system. Suppose that the current price is $p \neq 0$ and $p^t = \sum_i p_i$. Let

$$\check{p}_i(p^t) = \arg\max_{p_i \in \mathcal{P}_i} U_i\left(\frac{p_i}{\lambda}\right) - p_i,$$

where $\lambda = \frac{p^t}{C}$, C is the capacity of the bottleneck link, and $\mathcal{P}_i = [0, P^i_{max}]$. Here P^i_{max} represents user i's budget constraint, which is assumed to be sufficiently large so that the constraint is not active at the equilibrium. We assume that there exists at least one $i \in \mathcal{I}$ such that $U'_i(0) > \lambda_{max} = \frac{\sum_i P^i_{max}}{C}$ and that the initial price vector is such that $p^t(0) = \sum_i p_i(0) > 0$. This ensures that $p(n) \neq 0$ for all $n \geq 1$.

Consider the following update scheme. We first assume that all users are synchronized and model the user updates with a discrete-time model. At each period $n \geq 1$, every user i updates its price based on its rate $x_i(n-1)$ and price $p_i(n-1)$ as follows:

$$p_i(n) = p_i(n-1) + \frac{\check{p}_i(n) - p_i(n-1)}{M+1}, \tag{15}$$

14

where $\breve{p}_i(n) = \arg\max_{p_i \in \mathcal{P}_i} U_i(\frac{p_i}{\lambda(n-1)}) - p_i, \lambda(n-1) = \frac{p^t(n-1)}{C}$, and $p^t(n-1) = \sum_i p_i(n-1)$. Once every user finishes updating its price, they wait long enough till the window sizes converge. After window sizes converge, the users repeat the above update procedure, based on the new rate allocation and prices $p_i(n)$. This is called the Jacobi update scheme.

The following theorem tells us that the user prices $p(n)$ converge to p^* under the Jacobi update scheme.

Theorem 2 *In a single bottleneck case the user prices $p(n) = (p_1(n), ..., p_I(n))$ converge to $p^* = (p_1^*, ..., p_I^*)$ as $n \to \infty$ under the Jacobi update scheme, i.e.,*

$$\lim_{n \to \infty} p(n) = p^*.$$

Proof: The proof is given in Appendix C.   ∎

In practice, however, the network users are almost never synchronized. Hence, it is important to show the convergence of user prices under an asynchronous update scheme with possibly delayed information. In other words, users do not necessarily update their prices simultaneously and some users may not have an access to the most recent value of the price per unit flow (possibly due to still fluctuating window size) and may decide to use an old value instead.

Let T_i be the set of periods at which user i updates its price, and

$$p_i(n+1) = p_i(n) + \frac{\breve{p}_i(\lambda(\tau_i(n))) - p_i(n)}{M+1}, \text{ for all } n \in T^i,$$

where $0 \le \tau_i(n) \le n$. We assume that the sets $T^i, i \in \mathcal{I}$, are infinite and if $\{n_k\}$ is a sequence of elements in T^i that tends to infinity, then

$$\lim_{k \to \infty} \tau_i(n_k) = \infty.$$

The update scheme described here is called a *totally asynchronous update scheme* [1].

Theorem 3 *In a single bottleneck case the user prices $p(n) = (p_1(n), ..., p_I(n))$ converge to $p^* = (p_1^*, ..., p_I^*)$ as $n \to \infty$ under the totally asynchronous update scheme.*

Proof: The proof is given in Appendix D.   ∎

A numerical example of the convergence of user prices to the unique fixed point of the mapping $\mathcal{T}$ with utility functions of $U(x_i) = a_i \cdot \log(x_i + b_i), a_i > 0$ and $0 \le b_i \le 1$ is given in section 7.

6 Algorithm II

The algorithm described in section 5 assumes that the users[13] explicitly compute the optimal prices based on the current prices per unit flow and use the $(p, 1)$-proportionally fair algorithm to solve the *NETWORK(A, C; p)* problem with the given prices $p = (p_i, i \in \mathcal{I})$. In other words, the problem is formulated as a discrete model where, given the prices per unit flow from the previous period $n-1$, users solve *USER(U_i; $\lambda_i(n-1)$)* for period n. Then the $(p, 1)$-proportionally fair algorithm is used to drive the users to the solution of *NETWORK(A, C; p(n))*. However, there are a few implementation issues that need to be addressed. First, since $(p, 1)$-proportionally fair algorithm converges asymptotically, the users need to know how long they should wait before updating their prices again or how often they should update their prices. Second, even with increasing computational power, it may still require non-negligible amount of CPU time to solve *USER(U_i; λ_i)*.

In this section we introduce an algorithm that does not require a computation of the optimal prices, but the users' preferences are implicitly reflected in the window size updating rule. Suppose that the users update their window sizes according to the following system of differential equations:

$$\frac{dw_i(t)}{dt} = -\kappa \cdot M_i(t) \cdot r_i(t) \tag{16}$$

where

$$M_i(t) = \frac{d_i + U_i'(x_i(t)) + x_i(t) \cdot U_i''(x_i(t))}{d_i(t)}, \tag{17}$$

$U_i'(\cdot) = \frac{d}{dx_i}U_i(\cdot), U_i''(\cdot) = \frac{d^2}{dx_i^2}U_i(\cdot)$, $d_i(t) = \frac{w_i(t)}{x_i(t)} = d_i + \sum_{j \in \mathcal{J}_i} \frac{q_j(t)}{C_j}$, $q_j(t)$ is the backlog at resource j at time t, and

$$
\begin{aligned}
r_i(t) &= \frac{w_i(t) - x_i(t)d_i - x_i(t)U_i'(x_i(t))}{w_i(t)} \\
&= 1 - \frac{x_i(t)d_i}{w_i(t)} - \frac{x_i(t)U_i'(x_i(t))}{w_i(t)}. \tag{18}
\end{aligned}
$$

Note that user i's utility function appears in both $M_i(t)$ and $r_i(t)$.

The following theorem states that the algorithm given by (16) - (18) converges to the unique stable point, where the resulting rates solve the *SYSTEM(U, A, C)* problem.

Theorem 4 *Let $V(w) = \sum_i (r_i(w))^2$. Then V is a Lyapunov function for the system of differential equations (16) - (18). The unique value minimizing V is a stable point of this system, to which all trajectories converge.*

---

[13]In practice an agent at the end host may carry out the computation on behalf of the user.

Proof: The proof is given in Appendix D.   ∎

In sections 5 and 6 for the purpose of analysis we have assumed that users are delay insensitive. However, if the users are sensitive to delay and the cost due to queueing delay is given by h_i in (11), then the algorithm does not require any pricing mechanism. Recall that the purpose of the pricing mechanism is to impose the system cost due to queueing delay on the users in an incentive-compatible way. In this case when users maximize their utility functions with *explicit* delay cost, the resulting rate allocation at an equilibrium is the optimal rate allocation.

7   Numerical Examples

In this section we give a numerical example of a simple network and demonstrate the convergence of users parameters through simulations. The simulations are run using the Network Simulator (NS) developed at Lawrence Berkeley Laboratory (LBL) and UC Berkeley.

7.1   Algorithm I

Although the convergence results for the first algorithm have been proved only for single bottleneck cases, the simulation results indicate that the user rates converge to the system optimal rates even in general networks with the utility functions satisfying the assumption 1 in the previous section, convexity of $\hat{p}_i(\lambda_i)$ in λ_i, and an appropriate damping constant M.

1 : e2-B2-B8-B4-E4
2 : e2-B2-B7-B5-e5
3 : e1-B1-B2-B8-B4-e4
4 : e1-B1-B7-B5-e5
5 : e3-B3-B8-B7-B6-e6
6 : e2-B7-B6-e6
7 : e1-B1-B2-e2
8 : e3-B3-B4-e4
9 : e2-B2-B8-B4-e4
10 : e2-B2-B7-B5-e5
11 : e1-B1-B2-B8-B4-e4

Figure 3: Topology of the simulated network and the users' routes.

The topology of the simulated network is shown in Figure 3. There are 11 users that share the network. The source-destination pairs of the users are given in the figure, and the capacities and

17

user	$U_i(x_i)$	x_i^* (Mbps)	p_i^*
1	$5 \log(x + 1)$	0.584	4.932
2	$7 \log(x + 1)$	0.283	6.808
3	$9 \log(x + 1)$	0.521	8.864
4	$6 \log(x + 1)$	0.724	5.934
5	$8 \log(x + 1)$	1.076	7.941
6	$10 \log(x + 1)$	0.424	9.815
7	$4 \log(x + 1)$	0.459	3.931
8	$7 \log(x + 1)$	1.500	6.963
9	$10 \log(x + 1)$	1.175	9.932
10	$12 \log(x + 1)$	0.492	11.808
11	$9 \log(x + 1)$	0.521	8.864

Table 1: Utility functions of the users and the optimal rates and prices.

the delays of the links are indicated by the numbers next to the links. The utility functions of the users are given in Table 1. Each connection updates its price after it receives the acknowledgement for the k_i-th packet after the previous update. Hence, the price updates are not synchronized due to different round trip delays. For the simulation we have set $k_i = 25$, $\kappa = 0.1$, and the damping constant $M = 5$ for all connections. Packet sizes are fixed at 1,000 bytes.


Figure 4: Convergence of user prices and rates.

The prices and rates of some of the users are plotted in Figure 4. The dotted lines in the plots represent the optimal rates and unique equilibrium prices, respectively. These plots clearly

demonstrate the convergence of the users' prices and, thus, the rates to the system optimum.

7.2 Algorithm II

```
1: e2-B2-B3-B4-B5-e5          10: e2-B2-B9-B10-B6-e6
2: e2-B2-B9-B10-B6-e6         11: e1-B1-B9-B11-B5-e5
3: e1-B1-B9-B11-B5-e5         12: e4-B4-B5-B6-B7-e7
4: e1-B1-B8-B7-B6-e6          13: e4-B4-B11-B10-B8-e8
5: e3-B3-B9-B10-B7-e7         14: e2-B2-B1-B8-B7-e7
6: e2-B2-B1-B8-e8             15: e1-B1-B8-B7-e7
7: e1-B1-B2-e2                16: e3-B3-B2-B1-B8-e8
8: e3-B3-B4-5-e5              17: e3-B3-B4-B5-B6-e6
9: e2-B2-B3-4-B5-e5           18: e4-B4-B11-B10-B8-e8
```
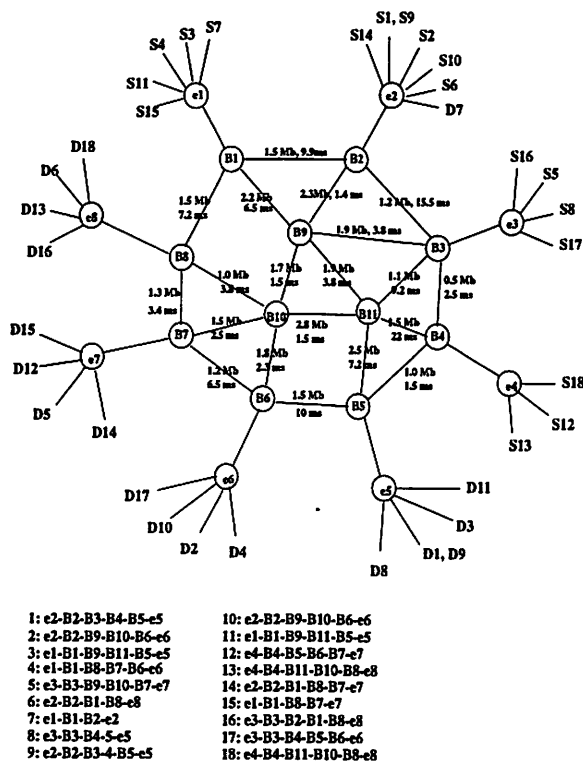
Figure 5: Topology of the simulated network and the users' routes.

The topology of the second simulated network is given in Figure 5. There are 18 users that share the network, and the routes of the users are as listed in Figure 5. The capacities and delays of the links are given next to the links in the figure. Table 2 has the utilities functions of the users. In this simulation each connection computes $M_i(t)$ and $r_i(t)$ after receiving the acknowledgement for each packet and updates its window size according to (16). The parameter κ is set to 0.1 for all connections. Packet sizes are fixed at 1,000 bytes. The simulation is run for 500 seconds.

The evolution of the prices and the rates of users 1, 3, 5, 7, 9, and 11 are plotted in Figure 6. Although the optimal rates and the equilibrium prices are not plotted, one can see that users' rates converge to a region around the optimal rates and the prices converge to the unique equilibrium. Since the users use instantaneous values of the rates and the queueing delays, their rates oscillate slightly around the optimal rates.

The convergence rates of the algorithms obviously depend on the parameters such as κ and the damping constant M. The examples in this section indicate that it could take several minutes before the users converge to the equilibrium. This slow convergence is, however, due to the fact

user	$U_i(x_i)$	user	$U_i(x_i)$
1	$0.7x^{0.35}$	10	$0.7x^{0.32}$
2	$0.7x^{0.28}$	11	$0.7x^{0.29}$
3	$0.7x^{0.39}$	12	$0.7x^{0.12}$
4	$0.7x^{0.65}$	13	$0.7x^{0.39}$
5	$0.7x^{0.50}$	14	$0.7x^{0.82}$
6	$0.7x^{0.36}$	15	$0.7x^{0.79}$
7	$0.7x^{0.42}$	16	$0.7x^{0.50}$
8	$0.7x^{0.59}$	17	$0.7x^{0.49}$
9	$0.7x^{0.40}$	18	$0.7x^{0.49}$

Table 2: Utility functions of the users.

Figure 6: Convergence of user prices and rates.

that all users start from the initial condition together. Suppose that a user enters a network with many users that is already in an equilibrium with the existing users. Then, the simulation results suggest that the system takes only a few seconds to converge to the new equilibrium because the arrival of a new user does not perturb the existing users too much. This allows the new user to quickly estimate the price per unit flow of its route and reach the equilibrium. Due to the size of the Internet and a large number of users, if the arrivals and departures of the users are reasonably frequent, we expect the system to remain close to an equilibrium at any given time.

8 Controlling the Backlog and Queueing Delay in the Network

In the previous sections for simplicity of analysis we have assumed that the buffer sizes at the resources are infinite. In the real network, however, the buffer sizes are finite and there may be packet losses due to temporary buffer overflows. In this section we discuss how our pricing scheme and algorithms can be extended to cope with these packet losses.

In our pricing scheme we have assumed that the price per unit flow charged by a resource j is the queueing delay at the resource. Suppose that the price per unit flow at a resource is proportional to the queueing delay, i.e., $g_j = \alpha \cdot \frac{q_j}{C_j}$, where $\alpha > 0$, but not necessarily equal to the queueing delay. In this case the users face the same user optimization problem, except that now the target queue size should be set to its willingness to pay divided by α. More specifically, in the first algorithm the target queue size p_i of the $(p,1)$-proportionally fair algorithm should be set to $\alpha^{-1} \cdot \hat{p}_i(\lambda)$, where $\hat{p}_i(\lambda) = \arg\max_{p_i} U_i(\frac{p_i}{\lambda}) - p_i$, and in the second algorithm $U_i'(x_i(t))$ and $U_i''(x_i(t))$ in (17) and (18) need to be replaced by $\alpha^{-1} \cdot U_i'(x_i(t))$ and $\alpha^{-1} \cdot U_i''(x_i(t))$, respectively. One can show that under this more general pricing scheme and modified algorithms, there exists a unique equilibrium of the algorithms, and at the equilibrium the resulting rates are the optimal rates that solve $SYSTEM(U, A, C)$. In other words, the equilibrium prices and resulting rates are invariant under change of α. One consequence of this is the following. Because the equilibrium prices are the same for all $\alpha > 0$, the target queue sizes of the users are inversely proportional to α. Therefore, by increasing α we can arbitrarily reduce the target queue sizes of the users and, hence, the backlog inside the network. Further, as α goes to ∞, one can see that the queueing delay goes to 0 for all users because the target queue sizes of the users go to 0. What this means in practice is that α can be chosen to ensure that packet losses are negligible even when the buffer sizes are finite.

9 Conclusions

In this paper we have investigated the fundamental problem of the existence of an algorithm that achieves the system optimum in a distributed network without any explicit feedback from the network elements to the end hosts. We have described a pricing scheme and two algorithms that solve the system problem at the unique equilibrium point of the algorithms. We have demonstrated that our algorithms are incentive compatible and do not require any feedback from the network, making them easy to deploy. Moreover, they achieve weighted proportional fairness. We have proved the convergence of the first algorithm to the optimal rates in a single bottleneck case under both synchronous and asynchronous update schemes under some mild technical assumptions on the

utility functions. The second algorithm is proven to converge to the system optimum in any general network. A few numerical examples are given to demonstrate the convergence of the algorithms in simple networks.

In this paper we have assumed that the routes of the users are fixed. However, if the network charges the users based on the congestion level along the route of the connection, the network should ensure that no users get preferential treatment or penalized through the selection of the routes. For instance, the routing has to be done in such a way that no user is routed through a path that has the highest price per unit flow among the available paths, while other connections are routed through the cheapest path.

We are currently attempting to design a mechanism that uses measured packet loss rates to solve the *SYSTEM(U, A, C)* problem, instead of the queueing delays. We are also working on how one can solve a system problem with a different objective function, such as the total revenue, using a distributed window-based algorithm. Another interesting question that remains is how the algorithm described in this paper can be extended to provide different quality of service (QoS).

References

[1] D. Bertsekas and J. Tsitsiklis. *Parallel and Distributed Computation.* Prentice-Hall, Englwood Cliffs,N.J., 1989.

[2] L.S. Brakmo and L.L. Peterson. Tcp vegas: end to end congestion avoidance on a global internet. *IEEE Journal on Selected Areas in Communications*, 13(8):1465–80, October 1995.

[3] P. Dubey. Inefficiency of Nash equilibria. *Mathematics of Operations Research*, 11:1–8, 1986.

[4] R. Edell and P. Varaiya. INFOCOM 1999 Keynote Talk. available at http://www.INDEX.berkeley.edu/reports/99-009S.

[5] S. Floyd and V. Jacobson. Connection with multiple Congested Gateways in packet-Switched Networks,Part1: One-way Traffic". *ACM Computer Communication Review*, 21(5):30–47, August 1991.

[6] S. Floyd and V. Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.

[7] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. Available at http://www.statslab.cam.ac.uk/~ frank, June 1998.

[8] V. Jacobson. Congestion avoidance and control . *Computer Communication Review*, 18(4):314–29, August 1988.

[9] F. Kelly. Charging and rate control for elastic traffic. *European Transactions on Telecommunications*, vol.8(1):33–7, January 1997.

[10] F. Kelly, A. Maulloo, and D. Tan. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, vol.49(3):237–52, March 1998.

[11] F.P. Kelly. On tariffs, policing, and admission control for multiservice networks. *Operations Research Letters*, vol.15(1):1–20, Feb. 1994.

[12] R. J. La and V. Ananthram. Charge-Sensitive TCP and Rate Control in the Internet. In *Proc. IEEE INFOCOM '00*, March 2000.

[13] T. V. Lakshman and U. Madhow. The performance of TCP/IP for networks with high bandwidth-delay products and random loss. *IEEE/ACM Transactions on Networking*, 5(3):336–50, June 1997.

[14] S. H. Low and D. E. Lapsley. Optimization flow control. *IEEE/ACM Transactions on Networking*, 7(6):861–74, Dec. 1999.

[15] J.K. Mackie-Mason and H. R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 13(7):1141–9, September 1995.

[16] L. Massoulie and J. Roberts. Bandwidth sharing : objectives and algorithms. In *Proceedings of IEEE INFOCOM '99*, New York, NY, March 1999.

[17] J. Mo, R. J. La, V. Ananthram, and J. Walrand. Analysis and Comparison of TCP Reno and Vegas. In *Proc. INFOCOM '99*, March 1999.

[18] J. Mo and J. Walrand. Fair end-to-end window-based congestion control. Technical Report M98/46, UCB/ERL, July 1998. available at http://www.path.berkeley.edu/~ jhmo.

[19] W. Rudin. *Principles of Mathematical Analysis, 3rd Edition*. McGraw-Hill, 1976.

[20] H.R. Varian and J.K. Mackie-Mason. Pricing the internet. *Public Access to the Internet*, 1994.

A Assumptions on the Utility Functions

Assumption 1 $\hat{p}_i(\lambda_i)$ *is decreasing in* $\lambda_i > 0$, *where* $\hat{p}_i(\lambda_i) = \arg\max_{p_i} U_i(\frac{p_i}{\lambda_i}) - p_i$.

Figure 7: Assumption 2.

Assumption 2 *There exists $M > 0$ such that*

(i) for all $p^t < p^{*t}, \check{p}^t(p^t) < 2Mp^{*t} - (2M - 1)p^t$

(ii) for all $p^t > p^{*t}, \check{p}^t(p^t) > 2Mp^{*t} - (2M - 1)p^t$

where $p^t = \sum_i p_i, p^{*t} = \sum_i p_i^*, \check{p}^t(p^t) = \sum_i \check{p}_i(p^t)$, *and* $\check{p}_i(p^t) = \arg\max_{p_i'} U_i(\frac{p_i' \cdot C}{p^t}) - p_i'$.

Assumption 3 *Utility functions of the users satisfy the following:*

$$U_i'(x_i) + x_i \cdot U_i''(x_i) \geq 0, \quad x_i \in [0, C^i],$$

where $C^i = \min_{j \in \mathcal{J}_i} C_j$.

B Proof of Theorem 1

At an equilibrium p^* of the algorithm, where user prices stay constant, we have the following:

1. $p_i^* = x_i^* \cdot \lambda_i^*$, where x^* and λ^* are the resulting weighted proportionally fair allocation with weight vector p^* and the price per unit flow vector, which is strictly positive.

2. x^* solves $NETWORK(A, C; p^*)$.

3. p_i^* solves $USER_i(U_i; \lambda_i^*)$ for all $i \in \mathcal{I}$.

4. p_i^* solves $USER_i(U_i; \lambda_i^*)$ for all $i \in \mathcal{I}$.

(1) follows trivially from that $p_i^* = x_i^* \cdot \lambda_i^*$ by definition. (2) holds at the unique stable point of TCP with fixed p^* [18]. (3) is true by the construction of the algorithm. Furthermore, any p^* with its corresponding λ^* and x^*, that satisfies (1)-(3) is an equilibrium of the algorithm. By Theorem 2 in [9], the existence of such p^* is guaranteed and the resulting rate allocation x^* is the unique solution to $SYSTEM(U, A, C)$.

The uniqueness of equilibrium point can be shown as follows. First, recall that there exists unique x^* that solves $SYSTEM(U, A, C)$, from the strict concavity of the utility functions. From the Kuhn-Tucker conditions, which are the necessary and sufficient conditions, $\hat{p}_i$ solves $USER_i(U_i; \lambda_i^*)$ if and only if it satisfies

$$\begin{aligned} \hat{p}_i > 0 & \quad \Rightarrow \quad U_i'(\frac{\hat{p}_i}{\lambda_i^*}) = \lambda_i^* \\ \hat{p}_i = 0 & \quad \Rightarrow \quad U_i'(0) \le \lambda_i^*. \end{aligned}$$ (19)

Hence, at any equilibrium p^* one can see that

$$\begin{aligned} x_i^* > 0 & \quad \Rightarrow \quad U_i'(\frac{p_i^*}{\lambda_i^*}) = U_i'(x_i^*) = \lambda_i^* \\ x_i^* = 0 & \quad \Rightarrow \quad U_i'(0) \le \lambda_i^*. \end{aligned}$$ (20)

This immediately tells us that there exists a unique fixed point from

$$p_i^* = x_i^* \cdot \lambda_i^* = \begin{cases} x_i^* \cdot U_i'(x_i^*) & \text{if } x_i^* > 0 \\ 0 & \text{if } x_i^* = 0 \end{cases}$$

C Proof of Theorem 2

Note that, in a single bottleneck case, updating scheme of the users depends only on the current price and price per unit flow $\lambda(n) = \frac{p_i(n)}{x_i(n)} = \frac{p^t(n)}{C}$, where C is the bottleneck link capacity. Hence, in order to show the convergence of $p(n)$, it suffices to show the convergence of $p^t(n) = \sum_i p_i(n)$ to $p^{*t} = \sum_i p_i^*$.

Since we have assumed that $U_i(\cdot), i \in \mathcal{I}$, are continuously differentiable, one can see that $\breve{p}^t(p^t)$ is a continuous function of p^t. Thus, it suffices to show that if $p(n-1) \ne p^*$,

$$\left| p^t(n) - p^{*t} \right| < \left| p^t(n-1) - p^{*t} \right| \text{ for all } n \ge 1.$$ (21)

This can be easily shown as follows. Let us discuss two separate cases.

1. $\lambda(n) < \lambda^*$, i.e., $p^t(n) < p^{*t}$:

   In this case we have from (15)

   $$p^t(n+1) = p^t(n) + \frac{\breve{p}^t(n+1) - p^t(n)}{M+1} > p^t(n),$$

where the last inequality follows from assumption 1. Now we show that

$$p^t(n) < p^t(n+1) < p^{*t}$$

as follows.

$$
\begin{aligned}
p^t(n+1) &= \sum_{i \in \mathcal{I}} p_i(n) + \frac{\breve{p}_i(n+1) - p_i(n)}{M+1} \\
&= \frac{M}{M+1} p^t(n) + \frac{1}{M+1} \breve{p}^t(n+1) \\
&< \frac{M}{M+1} p^t(n) + \frac{1}{M+1} (p^{*t} + M \cdot (p^{*t} - p^t(n))) \qquad (22) \\
&= p^{*t},
\end{aligned}
$$

where the inequality in (22) follows from assumption 2.

2. $\lambda(n) > \lambda^*$, i.e., $p^t(n) > p^{*t}$:

   Using a similar argument, one can show that

$$p^{*t} < p^t(n+1) < p^t(n).$$

The convergence of $p(n)$ follows from that $\breve{p}^t(p)$ is a continuous function of p and assumption 2.

D   Proof of Theorem 3

We first show that there is a sequence of nonempty sets $\{\mathcal{P}(n)\}$ with

$$\cdots \subset \mathcal{P}(n+1) \subset \mathcal{P}(n) \subset \cdots \mathcal{P}(0) \subseteq \mathcal{P},$$

where $\mathcal{P} = \mathcal{P}_1 \times \cdots \times \mathcal{P}_I$, satisfying the following two conditions:

1. Synchronous convergence condition : we have

$$\mathcal{T}(p) \in \mathcal{P}(n+1) \text{ for all } n \text{ and } p \in \mathcal{P}(n).$$

   Furthermore, if $\{p^k\}$ is a sequence such that $p^k \in \mathcal{P}(k)$ for every k, then every limit point of $\{p^k\}$ is the unique fixed point of the mapping $\mathcal{T}$.

2. Box condition : for every n, there exist sets $\mathcal{P}_i(n) \subset \mathcal{P}_i$ such that

$$\mathcal{P}(n) = \mathcal{P}_1(n) \times \cdots \times \mathcal{P}_I(n).$$

26

From the assumption there exists at least one $i \in \mathcal{I}$ such that $U_i'(0) > \lambda_{max} = \frac{\sum_i P_{max}^i}{C}$. Thus, we know that $p^t(k) > 0$ for all $k \geq 1$.

Let $\mathcal{P}(0) = [\min\{p_i(0), \check{p}_i(\lambda_{max})\}, \ p_{max}^i]$. Define for all $n \geq 1$, $\mathcal{P}'(n) = \{ T'(p) \mid p \in \mathcal{P}(n-1) \}$, where

$$T_i'(p) = p_i + \frac{\check{p}_i(p^t) - p_i}{M + 1}.$$

Take the projection for each $i \in \mathcal{I}$

$$\mathcal{P}_i(n) = \{p_i \mid p_i \text{ is the } i\text{-th element of } p \in \mathcal{P}'(n)\},$$

and define

$$\mathcal{P}(n) = \times_{i \in \mathcal{I}} \mathcal{P}_i(n).$$

Then, $\mathcal{P}'(n) \subset \mathcal{P}'(n-1)$ and, hence, $\mathcal{P}(n) \subset \mathcal{P}(n-1)$. Further, one can show that $\mathcal{P}(n)$ satisfies the *synchronous convergence condition*. From its construction $\mathcal{P}(n)$ satisfies the *box condition*. The theorem follows from [1].

E   Proof of Theorem 4

We first define interior and boundary points. An interior point is defined to be a window size vector, around which we can find an open ball such that all the points in the ball have the same set of bottlenecks. A window size vector, for which we cannot find such an open ball is said to be a boundary point. One can show that the boundary regions are linear and divide the window size space into a finite number of convex regions. Hence, if $w(t)$ is an interior point, then for sufficiently small $\epsilon > 0$, $w(t + \epsilon)$ is in the same region as $w(t)$. If $w(t)$ is a boundary point, then for any sufficiently small $\epsilon > 0$ $w(t + \epsilon)$ is in one of the neighboring regions of the boundary point.

Recall that $V(t) = \sum_i (r_i(t))^2$, where

$$r_i(t) = \frac{w_i(t) - x_i(t)(d_i + U_i'(x_i(t)))}{w_i(t)}.$$

We show that for small $\epsilon > 0$,

$$V(t + \epsilon) - V(t) = r^T(t) J_r(t)(w(t + \epsilon) - w(t)) + o(\epsilon), \tag{23}$$

where $J_r(t)$ is the partial derivative

$$J_r(t) = \left[\frac{\partial r_i}{\partial w_j}, i, j \in \mathcal{I} \right],$$

which depends on the region in which $w(t + \epsilon)$ is.

We first state a claim that will be used in the proof.

27

Claim 1 *At an interior point w the partial derivative $J_x = \left[\frac{\partial x_i}{\partial w_j}, i, j \in \mathcal{I}\right]$ is given by*

$$J_x = \bar{D}^{-1}(I - XA_B(A_B^T X \bar{D}^{-1} A_B)^{-1} A_B^T \bar{D}^{-1}). \tag{24}$$

Proof: We first show that the mapping $\mathcal{W}$ from the window sizes to the rates is differentiable at the interior points. Suppose that w is an interior point and $x = \mathcal{W}(w)$. Let $\mathcal{B}$ be the set of bottlenecks at w. We consider two cases.

<u>Case (1)</u> A_B has full rank, i.e., $rank(A_B) = B = |\mathcal{B}|$, where A_B is the submatrix of A with columns corresponding to the bottlenecks. Let

$$g_B(q_B', w) = A_B^T x(q_B', w) - C_B, \tag{25}$$

where

$$x_i(q_B', w) = \frac{w_i}{d_i + A_i \cdot q_B'} \tag{26}$$

and A_i. is the i-th row of A_B. We use q_B' to denote the queueing delays, not queue sizes that are denoted by q_B. We define

$$J_{g_q} = \left[\frac{\partial g_i}{\partial q_j'}, i, j = 1, \cdots, B\right],$$

where

$$
\begin{aligned}
\frac{\partial g_i}{\partial q_j'} &= \frac{\partial}{\partial q_j'}(\sum_m A_{m,i} x_m(q_B', w) - C_i) \\
&= \frac{\partial}{\partial q_j'}(\sum_m A_{m,i} \frac{w_m}{d_m + \sum_k A_{m,k} q_k'}) \\
&= \sum_m A_{m,i}(\frac{\partial}{\partial q_j'} \frac{w_m}{d_m + \sum_k A_{m,k} q_k'}) \\
&= \sum_m A_{m,i}(-\frac{w_m}{(d_m + \sum_k A_{m,k} q_k')^2}) A_{m,j} \\
&= -\sum_m A_{m,i} A_{m,j} \frac{w_m}{(d_m + \sum_k A_{m,k} q_k')^2} \\
&= -\sum_m A_{m,i} A_{m,j} \frac{x_m(q_B', w)}{\bar{d}_m(q_B')},
\end{aligned}
$$

where $\bar{d}_m(q_B') = d_m + A_m \cdot q_B'$. This can be written in a matrix form as

$$J_{g_q} = -A_B^T X \bar{D}^{-1} A_B, \tag{27}$$

where $\bar{D} = \text{diag}(\bar{d}_1, \cdots, \bar{d}_I)$.

28

One can easily show that J_{g_q} is nonsingular because $A_B^T X \bar{D}^{-1} A_B$ is positive definite. Therefore, by the implicit function theorem, $q'_B(w)$ is differentiable [19], and so is $x(q'_B, w)$ from (26). Further, from the implicit function theorem,

$$
\begin{aligned}
J_q(w) &= \left[\frac{\partial q'_j}{\partial w_i}, j = 1, \cdots, \text{B}, i \in \mathcal{I} \right] \\
&= -(J_{g_q})^{-1} J_{g_w},
\end{aligned}
$$

where

$$
\begin{aligned}
J_{g_w} &= \left[\frac{\partial g_j}{\partial w_i}, j = 1, \cdots, \text{B}, i \in \mathcal{I} \right] \\
&= A_B^T \bar{D}^{-1} \qquad \text{from (25)}.
\end{aligned}
$$

Thus,

$$
J_q(w) = (A_B^T X \bar{D}^{-1} A_B)^{-1} A_B^T \bar{D}^{-1}. \tag{28}
$$

<u>Case (2)</u> $rank(A_B) = \gamma < \text{B}$. Let $\bar{\mathcal{B}} \subset \mathcal{B}$ such that $rank(A_{\bar{B}}) = \gamma$ and $|\bar{\mathcal{B}}| = \gamma$. Then, applying the implicit function theorem to $g_{\bar{B}}$ as defined in (25) yields the differentiability of $x(q'_{\bar{B}}, w)$. One can show that $x(q'_{\bar{B}}, w)$ is the unique solution of g_B as follows. Since we delete dependent rows from A_B^T, if x is a solution of $A_{\bar{B}}^T x = C_{\bar{B}}$, then it is a also a solution of $A_B^T x = C_B$. Since the solution is unique, $x(q'_{\bar{B}}, w)$ is the solution of g_B. Therefore, the differentiability follows.

We proceed to prove the claim with the assumption that A_B has full rank. If not, we can take $A_{\bar{B}}$ as described before and prove the claim. Since A_B has full rank, we know that the queueing delays are well defined from [18]. First, from (8) for all $i \in \mathcal{I}$ we have

$$
w_i = x_i(w) \cdot d_i + x_i(w) A_{i\cdot} q'_B(w). \tag{29}
$$

If we differentiate (29) with respect to $w_j, j \in \mathcal{I}$,

$$
\delta_{i,j} = (A_{i\cdot} q'_B(w) + d_i)(J_x)_{i,j} + x_i(w)(A_{i\cdot}(J_q)_{\cdot j}).
$$

Again, note that J_q is well defined from (28). In a matrix form this is given by

$$
I = \bar{D} J_x + X A_B J_q,
$$

which yields

$$
J_x = \bar{D}^{-1}(I - X A_B J_q).
$$

Since $J_q = (A_B^T \bar{D}^{-1} X A_B)^{-1} A_B^T \bar{D}^{-1}$ from (28), we get

$$
J_x = \bar{D}^{-1}(I - X A_B (A_B^T \bar{D}^{-1} X A_B)^{-1} A_B^T \bar{D}^{-1}).
$$

29

We continue the proof of the theorem. If $w(t)$ is an interior point, then from the algorithm and that $x(w)$ is differentiable at interior points from the above claim, we have

$$
\begin{aligned}
J_r &= DXW^{-2} - DW^{-1}J_x + XU'W^{-2} - W^{-1}(U' + XU'')J_x \\
&= DXW^{-2} - DW^{-1}\bar{D}^{-1} \\
&\quad + DW^{-1}\bar{D}^{-1}XA_B(A_B^T X\bar{D}^{-1}A_B)^{-1}A_B^T\bar{D}^{-1} \\
&\quad + XU'W^{-2} - W^{-1}(U' + XU'')\bar{D}^{-1} \\
&\quad + W^{-1}(U' + XU'')\bar{D}^{-1}XA_B(A_B^T X\bar{D}^{-1}A_B)^{-1}A_B^T\bar{D}^{-1} \qquad (30) \\
&= -W^{-1}XU''\bar{D}^{-1} + (D + U' + XU'')\bar{D}^{-2}A_B(A_B^T X\bar{D}^{-1}A_B)^{-1}A_B^T\bar{D}^{-1}, \qquad (31)
\end{aligned}
$$

where $W = \mathrm{diag}(w(t))$, $X = \mathrm{diag}(x(t))$, $\bar{D} = \mathrm{diag}(d_1(t), \cdots, d_I(t))$, $D = \mathrm{diag}(d_1, \cdots, d_I)$, $U' = \mathrm{diag}(U_1'(x_1(t)), \cdots, U_I'(x_I(t)))$, $U'' = \mathrm{diag}(U_1''(x_1(t)), \cdots, U_I''(x_I(t)))$, and A_B is the submatrix of A with columns that correspond to bottlenecks at $w(t)$[14]. If $w(t)$ is a boundary point, then A_B is the submatrix of A with columns that correspond to bottlenecks at $w(t + \epsilon)$. The second inequality in (30) follows from that $J_x = \bar{D}^{-1}(I - XA_B(A_B^T X\bar{D}^{-1}A_B)^{-1}A_B^T\bar{D}^{-1})$.

Fix $\epsilon > 0$. Define $\delta w(t, \epsilon) = \frac{w(t+\epsilon) - w(t)}{\epsilon}$. Then, since $J_r(\cdot)$ is well defined in $[t, t + \epsilon]$ if the path $w(t')$ follows the straight line between $w(t + \epsilon)$ and $w(t)$ and $V(\cdot)$ is a continuous function of window sizes, one can see that

$$
\begin{aligned}
V(t + \epsilon) - V(t) &= \int_t^{t+\epsilon} r^T(s) J_r(s) \delta w(t, \epsilon) ds \\
&= \int_t^{t+\epsilon} (r(t) + \Delta r(s))^T [J_r(t) + \Delta J_r(s)] \delta w(t, \epsilon) ds,
\end{aligned}
$$

where $\Delta r(s) = r(s) - r(t)$, $\Delta J_r(s) = J_r(s) - J_r(t)$, $s \in [t, t + \epsilon]$, and $J_r(\cdot)$ is defined as described above. Then,

$$
\begin{aligned}
V(t + \epsilon) - V(t) &= \int_t^{t+\epsilon} r^T(t) \, J_r(t) \, \delta w(t, \epsilon) \, ds \\
&\quad + \int_t^{t+\epsilon} r^T(t) \, \Delta J_r(s) \, \delta w(t, \epsilon) \, ds \\
&\quad + \int_t^{t+\epsilon} \Delta r^T(s) \, J_r(t) \, \delta w(t, \epsilon) \, ds
\end{aligned}
$$

[14]For sufficiently small $\epsilon > 0$, $w(t + \epsilon)$ has the same bottlenecks as $w(t)$.

30

$$+ \int_t^{t+\epsilon} \Delta r^T(s) \, \Delta J_r(s) \, \delta w(t, \epsilon) \, ds$$

$$= r^T(t) \, J_r(t) \, (w(t+\epsilon) - w(t))$$

$$+ r^T(t) \cdot \int_t^{t+\epsilon} \Delta J_r(s) ds \cdot \delta w(t, \epsilon)$$

$$+ \int_t^{t+\epsilon} \Delta r^T(s) ds \cdot J_r(t) \cdot \delta w(t, \epsilon)$$

$$+ \int_t^{t+\epsilon} \Delta r^T(s) \, \Delta J_r(s) ds \cdot \delta w(t, \epsilon)$$

$$= r^T(t) \, J_r(t) \, (w(t+\epsilon) - w(t)) + o(\epsilon). \tag{32}$$

Hence, from (32) if $r(t) \neq 0$,

$$\lim_{\epsilon \downarrow 0} \frac{V(t+\epsilon) - V(t)}{\epsilon} = \lim_{\epsilon \downarrow 0} \frac{r^T(t) J_r(t) (w(t+\epsilon) - w(t))}{\epsilon} + \lim_{\epsilon \downarrow 0} \frac{o(\epsilon)}{\epsilon}$$

$$= r^T(t) J_r(t) \dot{w}(t)$$

$$= -\kappa r^T(t) J_r(t) M(t) r(t) < 0, \tag{33}$$

where $M(t) = \bar{D}^{-1}(D + U' + X \cdot U'')$. The last inequality follows from that $J_r(t) M(t)$ is a positive definite matrix for (30). Note that for a boundary point $w(t)$ such that $\dot{w}(t)$ lies along a boundary region $r^T(t) J_r(t) \dot{w}(t)$ is the same for all $J_r(t)$ with different A_B of the neighboring regions. Hence, the limit is well defined. Since (33) is true for all t unless $r_i(t) = 0$ for all $i \in \mathcal{I}$, the convergence follows.