

# Service differentiation and performance of weighted window-based congestion control and packet marking algorithms in ECN networks

Vasilios A. Siris<sup>a,\*</sup>, Costas Courcoubetis<sup>a,b</sup>, George Margetis<sup>a</sup>

<sup>a</sup>*Institute of Computer Science (ICS), Foundation for Research and Technology—Hellas (FORTH), P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece*

<sup>b</sup>*Department of Informatics, Athens University of Economics and Business, Athens, Greece*

Received 12 May 2002; accepted 14 May 2002

## Abstract

We investigate the service differentiation, in terms of average throughput, and the performance achieved using weighted window-based congestion control in networks supporting Explicit Congestion Notification (ECN). Our results show how service differentiation, queuing delay, and average throughput are affected by the increase and decrease rules of the end-system congestion control algorithms, and how they depend on the marking algorithms operating in the routers. The end-system algorithms we investigate include Willingness-To-Pay (willingness-to-pay) and MulTCP congestion control, and the packet marking algorithms include RED, virtual queue marking, and load-based marking. Our investigations consider both single and multiple link topologies, and connections with different round trip times.

© 2002 Elsevier Science B.V. All rights reserved.

*Keywords:* Service differentiation; Congestion control; Active queue management; Explicit congestion notification

## 1. Introduction

The Transmission Control Protocol (TCP) [1,2] has played an important role in the Internet's growth. With TCP, however, all connections with the same round trip time receive the same average throughput in the equilibrium; hence, it cannot support service differentiation. In the highly competitive telecommunications market, the ability to provide, in a flexible and efficient way, differentiated services will be extremely important. This push, along with the need to develop new algorithms better suited to the characteristics and requirements of the traffic being transferred, e.g. streaming video/audio, will continue to foster the appearance of new congestion control algorithms. Moreover, new active queue management algorithms [3] combined with Explicit Congestion Notification (ECN) [4], which has recently been approved as an IETF proposed standard, open up new possibilities for alternative approaches to congestion control.

One approach for supporting service differentiation, which is followed by the differentiated services (DiffServ) architecture, is to add mechanisms inside network routers. Drawbacks to this approach are the increased complexity,

compared to the Internet today, and the need for cooperation among the routers in order to provide coherent end-to-end services. An alternative approach has emerged [5–7], see also overview in Ref. [8], which suggests that service differentiation, as well as efficient and stable network operation and growth to meet increasing demand, can be achieved by a network with a simple feedback mechanism, such as ECN marking, that informs users of the congestion cost their traffic is incurring. The users react to these congestion signals using some rate control algorithm. Moreover, by charging a small fixed price per mark, the network provider can give users the incentive to react to the congestion signals, hence use the network, according to their actual needs. Note that the above approach places intelligence in the end-systems, while keeping the network simple. In this sense, it adheres to the end-to-end arguments put forward in Ref. [9], which have had a profound impact on the design of Internet protocols. The reasons for following such an approach are straightforward: it is at the end-systems where information on how to respond to congestion resides, since end-systems know the requirements of applications but also the valuations of users, in terms of the performance versus price tradeoff; both of these can be expressed in the form of utility functions, as discussed in Section 2. Indeed, the rate or congestion control

\* Corresponding author. Tel.: +30-81-391726; fax: +30-81-391601.  
E-mail address: vsiris@ics.forth.gr (V.A. Siris).

algorithms operating at the end-systems have the role of enforcing the policies expressed by these utility functions.

In this paper we investigate the latter approach, and consider a network where congestion feedback is in the form of ECN marks. For such a network we investigate the service differentiation achieved using weighted window-based congestion control, and how its performance is affected by the marking algorithms operating in the routers. With service differentiation we refer to the ability of the end-system algorithms, working in conjunction with the marking algorithms in routers, to offer different throughput to connections with different weights or willingness-to-pay values.

Our simulation experiments consider both a simple network with a single congested link, and a network with multiple links. The latter configuration allows us to study service differentiation when connections encounter different round trip times, and when marking occurs on more than one links.

The congestion control algorithms we investigate are Willingness-To-Pay (willingness-to-pay), which was first proposed in Ref. [5], and MulTCP [10], which is a modification of TCP that makes one connection approximate the behavior of multiple normal TCP flows. With willingness-to-pay, the congestion window follows a *single multiplicative decrease*: the rate of ECN marks is roughly<sup>1</sup> proportional to the sending rate and the congestion window decreases by some *fixed* amount for each ECN mark received. On the other hand, MulTCP, as TCP, follows a *double multiplicative decrease*: both the rate of ECN marks and/or losses is proportional to the sending rate and the congestion window is *halved* upon detection of congestion. willingness-to-pay achieves an average rate that, depending on the marking algorithm in the routers, is proportional to a weight or willingness-to-pay, and inversely proportional to the marking probability inside the network. On the other hand, MulTCP achieves a rate that is proportional to a weight and inversely proportional to the square root of the marking and/or loss probability.

The three marking algorithms we investigate are Random Early Detection (RED) [11], virtual queue marking [5], that marks packets depending on the overflow of a virtual queue whose buffer and capacity are some percentage of the link's actual buffer and capacity, and load-based marking, where the marking probability is a linear function of the link utilization measured over some time interval.

Our main results are summarized as follows:

- Weighted window-based congestion control is an effective way for an end-system to control its throughput. In cases where connections encounter different round trip times, these round trip times should be taken into account, in order to appropriately set weights.

- With the virtual queue algorithm, where the marking probability depends on the traffic burstiness, service (i.e. average throughput) differentiation depends on the characteristics of the congestion control algorithm, and in particular on the increase and decrease rules.
- With appropriate tuning, all three marking algorithms can exhibit the same marking probability for a range of average utilizations. As discussed in Ref. [6], the marking probability as a function of average utilization affects the convergence and stability behavior of the system comprising of the end-system congestion control algorithms and the router marking algorithms.
- Both RED and load-based marking, where packet marking is probabilistic, exhibit smaller average queueing delay and delay variation compared to virtual queue marking, which deterministically marks packets from the time the virtual queue overflows until the time it first becomes empty again.<sup>2</sup>
- Comparison of willingness-to-pay with MulTCP, which follows TCP's double multiplicative decrease but supports service differentiation, shows that MulTCP can be less fair and can achieve lower link utilization compared to willingness-to-pay.

The rest of this paper is structured as follows: In Section 2 we summarize some background work. In Section 3 we present the willingness-to-pay class of congestion control algorithms, and in Section 4 we discuss the three marking algorithms considered in our investigations. In Section 5 we present and discuss the results from our simulation experiments. Finally, in Section 6 we discuss related work, and in Section 7 we present some concluding remarks.

## 2. Theoretical background

In this section we summarize the theoretical background of our work as developed by Kelly and other researchers [6, 13,14].

Consider a network of  $J$  resources [13,14]. Each resource (link)  $j$  marks packets with some probability  $p_j(y_j)$ , where  $y_j$  is the aggregate arrival rate at resource  $j$ . Let  $R$  be the set of routes, or connections, active in the network. Each connection  $r$  updates its rate  $x_r$  according to the equation

$$\frac{dx_r(t)}{dt} = \kappa_r \left( w_r - x_r(t) \sum_{j \in r} \mu_j(t) \right), \quad (1)$$

where  $w_r$  is a weight,  $\kappa_r$  is a gain factor, and

$$\mu_j(t) = p_j \left( \sum_{s: j \in s} x_s(t) \right)$$

<sup>1</sup> As we discuss in more detail later, this depends on the marking algorithm operating in the router.

<sup>2</sup> Other variations of the virtual queue algorithm exist, such as the one considered in Ref. [12] where a packet is marked when it causes the virtual queue to overflow.

is the marking probability at resource  $j$ . With Eq. (1), connection  $r$  adjusts its sending rate  $x_r$  so that the rate of marks it receives  $x_r(t) \sum_{j \in r} \mu_j(t)$  becomes equal to the weight  $w_r$ . If the network charges a fixed amount for each mark returned to the end-system, then the sending rate is adjusted so that the rate of marks multiplied by the price per mark becomes equal to  $w_r$ , which thus represents the willingness-to-pay for connection  $r$ .

Assume that the marking probability  $p_j(y)$  depends on the cost  $C_j(y)$  incurred at resource  $j$  as follows

$$p_j(y) = \frac{dC_j(y)}{dy}.$$

It can be shown that, if  $C_j(y)$  is differentiable and feedback is instantaneous, then the above system converges to a point that maximizes

$$\sum_{r \in R} w_r \log x_r - \sum_{j \in J} C_j \left( \sum_{s: j \in s} x_s(t) \right).$$

If  $U_r(x_r) = w_r \log x_r$  is taken to be the utility for a connection  $r$ , then the last expression represents the social welfare of the system (resources and connections).

The above results can be generalized for the case where a connection  $r$  has a utility of a general form  $U_r(x_r)$ , if the willingness-to-pay changes smoothly according to

$$w_r(t) = x_r(t) U'_r(x_r(t)).$$

Up to now we have assumed instantaneous feedback. Next we summarize the approach of Kelly [6] for taking feedback delay into account, for the case of a single resource when all connections have the same gain parameter  $\kappa$  and under the assumption that the queueing delay represents a small fraction of the round trip time  $T$  (feedback delay), which is fixed and common for all connections. Let  $x(t) = \sum_r x_r(t)$  and  $w = \sum_r w_r$ . Summing Eq. (1) for all connections and taking the time lag into account we get

$$\frac{dx(t)}{dt} = \kappa(w - x(t - T)p(x(t - T))). \quad (2)$$

The linear delay equation

$$\frac{du(t)}{dt} = -\alpha u(t - T) \quad (3)$$

converges to zero as  $t$  increases if  $\alpha T < \pi/2$ . Moreover, the convergence is non-oscillatory if  $\alpha T < 1/e$ . Letting  $x(t) = x + u(t)$ , where  $x$  is the equilibrium of Eq. (2), and linearizing Eq. (2) about  $x$ , we get Eq. (3) with  $\alpha = \kappa(p + xp')$ , where  $p, p'$  are the marking probability and its derivative at rate  $x$ . Hence, the equilibrium of Eq. (2) is stable and the convergence is non-oscillatory if

$$\kappa T(p + xp') < e^{-1}.$$

The rate-based control algorithm given by Eq. (1) can be approximated using a window-based, self-clocking algorithm, similar to TCP's congestion avoidance algorithm, if a

connection's congestion window  $cwnd$  is updated with the reception of an acknowledgement using, see [5],

$$cwnd+ = \bar{\kappa} \left( \frac{\bar{w}}{cwnd} - f \right), \quad (4)$$

where  $\bar{\kappa} = \kappa T$  is the gain factor per round trip time,  $\bar{w} = wT$  is the willingness-to-pay per round trip time, and  $f$  equals 1 if the acknowledgement contains a mark or 0 if it does not. Note that Eq. (1) and its window-based version Eq. (4) follows a single multiplicative decrease, since the rate of marks is (roughly) proportional to the sending rate and the congestion window is decreased by some fixed amount for each ECN mark received.

### 3. Weighted window-based congestion control

In this section we describe a class of window-based congestion control algorithms that are based on the algorithm given by Eq. (4) [5]. The congestion window  $cwnd$  is updated according to

$$cwnd+ = \bar{\kappa} \left( \frac{\bar{w}_{inc}}{cwnd} - \frac{f}{\bar{w}_{dec}} \right),$$

hence the average change of the rate per unit time is

$$\frac{\bar{\kappa} \left( \frac{\bar{w}_{inc}}{cwnd} - \frac{p}{\bar{w}_{dec}} \right) / T}{T / cwnd} = \frac{\bar{\kappa}}{T} \left( \frac{\bar{w}_{inc}}{T} - \frac{xp}{\bar{w}_{dec}} \right), \quad (5)$$

where as before,  $\bar{\kappa}$  is the gain factor per round trip time and  $f = 1$  or 0 if the acknowledgement contains or does not contain an ECN mark, respectively. From the last equation we can see that for connections with the same round trip time, the average rate is proportional to  $\bar{w}_{inc} \bar{w}_{dec} = \bar{w}$ . For  $\bar{w}_{inc} = \bar{w}, \bar{w}_{dec} = 1$  we have the proportional increase algorithm given by Eq. (4). For  $\bar{w}_{inc} = 1, \bar{w}_{dec} = \bar{w}$  we have the inversely proportional decrease algorithm given by

$$cwnd+ = \bar{\kappa} \left( \frac{1}{cwnd} - \frac{f}{\bar{w}} \right). \quad (6)$$

The values of  $\bar{w}_{inc}, \bar{w}_{dec}$  represent a trade-off between the aggressiveness in probing for available bandwidth, hence in reaching the steady state, and the stability and size of the fluctuations around the average congestion window (equivalently, around the average throughput).

Substituting  $\kappa = \bar{\kappa}/T$  and  $w_{inc} = \bar{w}_{inc}/T$  in Eq. (5), summing for all connections, and taking the time lag into account we get

$$\frac{dx(t)}{dt} = \kappa \left( w_{inc} - \frac{x(t - T)p(x(t - T))}{\bar{w}_{dec}} \right),$$

which using the results of Kelly [6] that are summarized in Section 2, is stable and the convergence is non-oscillatory if




Fig. 1. For  $\bar{w}_{inc} = 4, \bar{w}_{dec} = 1$  the congestion window reaches the equilibrium faster, but fluctuations are larger. On the other hand, for  $\bar{w}_{inc} = 1, \bar{w}_{dec} = 4$  the congestion window reaches the equilibrium slower, but fluctuations are smaller. In both cases, the average congestion window is approximately the same.

$$\frac{\kappa T(p + xp')}{\bar{w}_{dec}} < e^{-1}. \quad (7)$$

The last equation is easier to satisfy for larger values of  $\bar{w}_{dec}$ .

The trade-off between rate of convergence and magnitude of fluctuations in the equilibrium is depicted in Fig. 1. As noted above, increasing  $\bar{w}_{dec}$ , while keeping the product  $\bar{w}_{inc}\bar{w}_{dec} = \bar{w}$  constant, results in smaller fluctuation around the equilibrium, but in the same average throughput. Smaller fluctuations can be advantageous for streaming applications, which typically perform better when their sending rate does not change abruptly.

As indicated by Eq. (7), convergence also depends on the marking probability  $p$  and its derivative  $p'$ : in general, a larger value of  $p$  and  $p'$  results in faster convergence, but, on the other hand, will also result in larger fluctuations around the equilibrium [14]. As we investigate in Section 5, the marking probability depends on both the end-system congestion control algorithms and the router marking algorithms.

## 4. Packet marking algorithms

In this section we describe the three packet marking algorithms that will be investigated in Section 5.

### 4.1. Random early detection

With RED [11], the marking probability is a piecewise linear function of the average queue length  $\bar{q}$ , which is estimated using exponential averaging with weight factor  $w_q$ . Specifically, if  $\bar{q}$  is smaller than a minimum queue length  $\min_{th}$ , then the marking probability is zero. If  $\bar{q}$  is between  $\min_{th}$  and a maximum queue length  $\max_{th}$ , then the marking probability ranges linearly from 0 to some

maximum value  $\max_p$ . Finally, if  $\bar{q}$  is above  $\max_{th}$ , then packets are dropped. The ‘gentle\_’ variation<sup>3</sup> of the original RED algorithm suggests to linearly vary the marking probability from  $\max_p$  to 1, when  $\bar{q}$  is between  $\max_{th}$  and  $2\max_{th}$ . In our simulation investigations we use RED with this variation, modified so that packets are marked rather than dropped when the average queue length is above  $\max_{th}$ , and when the buffer limit is exceeded.

### 4.2. Virtual queue marking

The virtual queue marking algorithm [5] presents an early warning of congestion. The algorithm maintains a virtual buffer of size  $\theta B$  serviced at rate  $\theta C$ , where  $B, C$  are the actual buffer and capacity of the output link, respectively. Note that  $B$  is not necessarily the total buffer of the output link, but can be some value that corresponds to a target maximum delay. The algorithm marks all packets that arrive at the link from the time a loss occurs in the virtual buffer until the first time the virtual buffer becomes empty; this interval is called the busy period of the virtual buffer. Note that there are other variations for when to mark packets: For example, the algorithm in Ref. [12] marks incoming packets when they cause the virtual queue to overflow.

As we will see in Section 5, a property of the virtual queue algorithm is that it differentiates flows based on their burstiness. Another property, demonstrated in experiments that are not presented in this paper, is that it does not avoid cases of synchronization of phases in closed-loop congestion control algorithms, such as those observed for drop-tail routers [15], which can result in some connections achieving very large average throughput, and other connections achieving very small throughput. Indeed, one motivation for introducing RED in routers was to avoid such effects [11,15].

### 4.3. Load-based marking

The third marking algorithm we investigate is load-based marking. According to this algorithm, the marking probability is a piecewise linear function of the average utilization measured over some time interval. In contrast, with RED the marking probability is a piecewise linear function of the average queue length. With the load-based algorithm, the marking probability is zero when the average load is less than a minimum value  $\rho_0$ . For utilization  $\rho$  larger than  $\rho_0$ , the marking probability is given by  $\min\{\alpha(\rho - \rho_0), 1\}$ . Hence, the algorithm has three parameters: the minimum utilization  $\rho_0$ , the parameter  $\alpha$  that controls the slope of the marking probability, and the averaging interval  $t_{avg}$ ; this interval determines how quickly the algorithm adjusts the marking probability to changes of the load.

<sup>3</sup> See ‘Recommendation on using the “gentle\_” variant of RED’, Floyd, 2000, <http://www.aciri.org/floyd/red/gentle.html>.

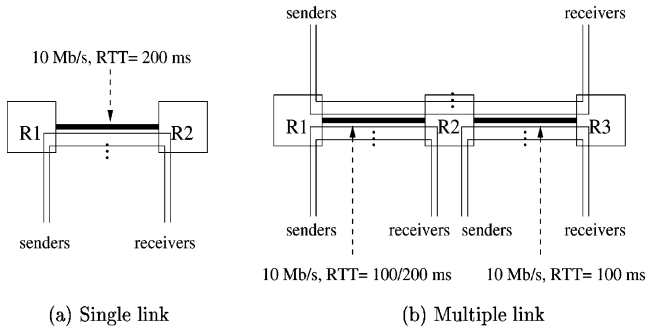


Fig. 2. Network topologies used in the simulation experiments. (a) Single link. (b) Multiple link.

Moreover, the interval determines the timescales over which congestion is detected, hence the timescales over which traffic burstiness affects the marking probability.

Load-based marking can be appropriate in cases where there is no shared buffer, such as the uplink in wide area wireless networks and Ethernet local area networks, since unlike the previous two marking algorithms, it does not rely on the queue length as a measure of congestion.

## 5. Simulation results and discussion

In this section we present and discuss simulation results investigating the interaction of willingness-to-pay and MulTCP with the three marking algorithms discussed in Section 4. The specific issues we investigate include the following:

- Service differentiation: Dependence of the average throughput achieved by different connections on their weight or willingness-to-pay, and how this is affected by the marking algorithm.
- Dependence of the marking probability on the average utilization: As discussed in Section 2, this dependence

affects the stability and convergence behavior of the whole system, i.e. the congestion control algorithms in the end-systems and the marking algorithms in the routers.

- Queueing delay and link utilization: In particular, we investigate the queueing delay for different marking algorithms, when the average link utilization remains the same.
- Comparison of willingness-to-pay and MulTCP: Recall that willingness-to-pay follows a single multiplicative decrease, whereas MulTCP follows a double multiplicative decrease, similar to normal TCP.

The simulations were performed using the ns-2 simulator [16]. The two network topologies considered in the simulation experiments are shown in Fig. 2. Since our focus is on investigating the interaction of the congestion control algorithms at the end-systems with the marking algorithms at the routers, we consider link parameters that do not result in packet loss, since this would generate retransmissions, which would affect the throughput measurements. Also, we focus on the congestion avoidance phase; hence our measurements begin after some time interval has elapsed from the start of each experiment.

### 5.1. Service differentiation

Fig. 3(a) shows the average ratio of throughput for different ratios of willingness-to-pay values. Also shown is the 90% confidence interval. The results were obtained from 10 independent runs of the experiment with the same parameters. Each connection carried data from a long ftp transfer, and the start time of each connection was selected randomly from the interval [0,5] seconds. Finally, the throughput was computed for the interval [60,180] seconds.

Fig. 3(a) indicates that a connection can control its throughput by adjusting its willingness-to-pay. Moreover,

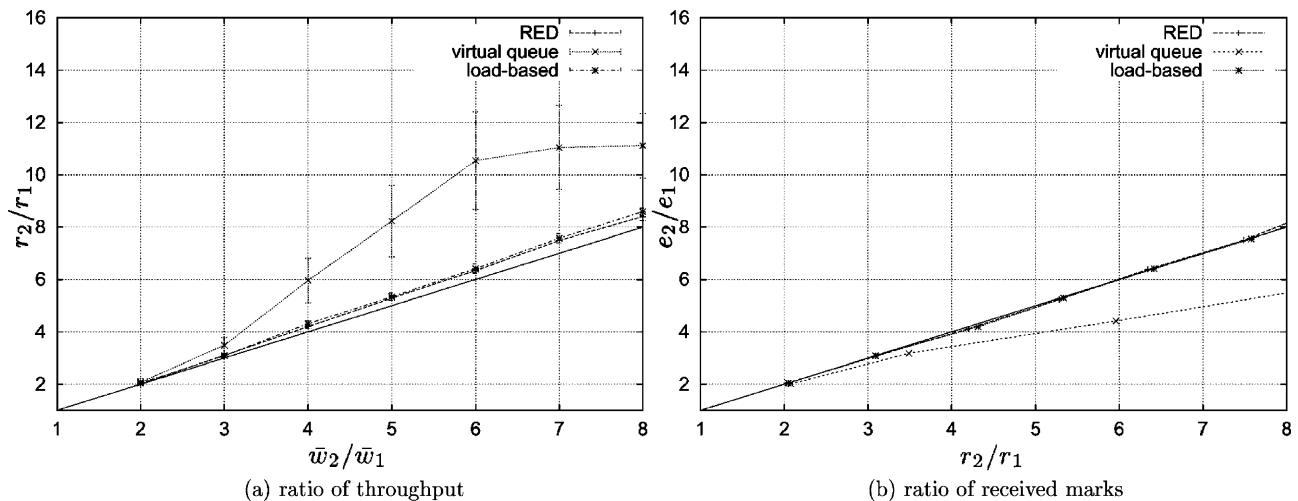


Fig. 3. Service differentiation using proportional increase willingness-to-pay (single link).  $\bar{\kappa} = 0.5$ .  $C = 10$  Mbps,  $B = 30$  pkts,  $RTT = 200$  ms,  $N = 10$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ . (a) Ratio of throughput, (b) ratio of received marks.

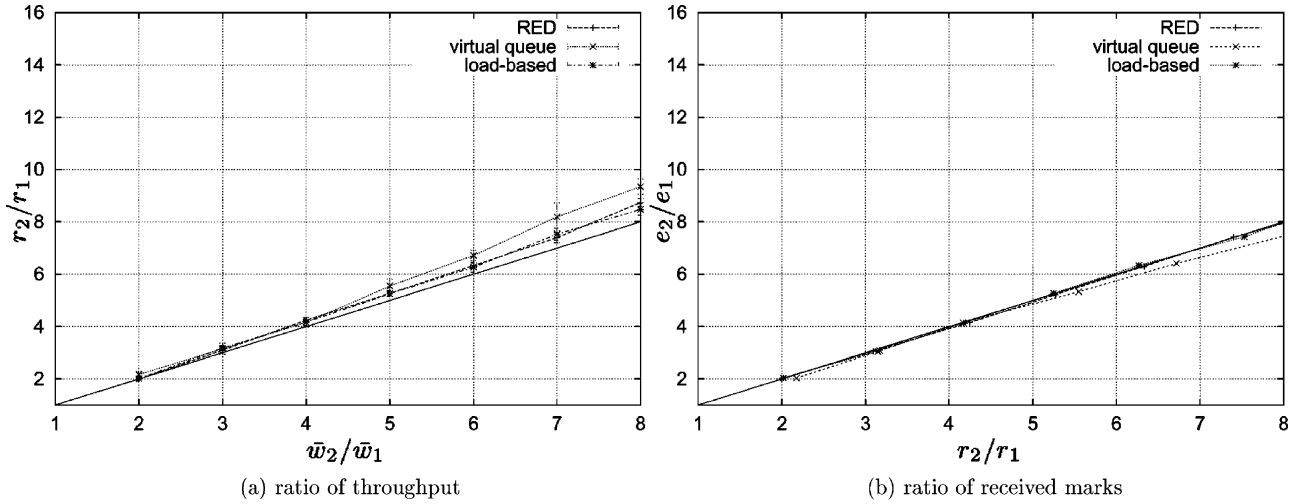


Fig. 4. Service differentiation using inversely proportional decrease willingness-to-pay (single link).  $\bar{\kappa} = 0.5$ .  $C = 10$  Mbps,  $B = 30$  pkts,  $RTT = 200$  ms,  $N = 10$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ . (a) Ratio of throughput, (b) ratio of received marks.

observe that the service differentiation achieved by both RED and load-based marking is roughly the same, and the ratio of the average throughput is very close to the ratio of willingness-to-pay. Note, however, that the two curves are slightly above the diagonal; this is attributed to the window-based nature of the congestion control algorithm since, in results not shown here, the ratio of the congestion windows is approximately equal to the ratio of willingness-to-pay.

On the other hand, in the case of virtual queue marking observe in Fig. 3(a) that for large values of the ratio of willingness-to-pay, the ratio of throughput is larger than the ratio of willingness-to-pay. Moreover, as indicated by the confidence interval, the fluctuations in service differentiation are also much larger compared to the fluctuations for the other two marking algorithms. Next we explain both of these observations.

The first observation can be explained with the help of Fig. 3(b), which shows the ratio of marks as a function of the ratio of average throughput. The figure shows that a smaller percentage of the packets belonging to a connection with a larger willingness-to-pay are marked, i.e. the marking probability is smaller for a connection with a larger willingness-to-pay. This is due to the combination of the following two factors: first, the connection with a smaller willingness-to-pay sends a smaller number of segments in one round trip time and second; the segments are typically sent back-to-back; the latter being a property of any window-based control mechanism. As a result, the connection with a smaller willingness-to-pay produces a burstier traffic stream compared to the connection with a larger willingness-to-pay. Burstier traffic streams, however, are more difficult for a multiplexer (link) to handle, hence require more bandwidth than their

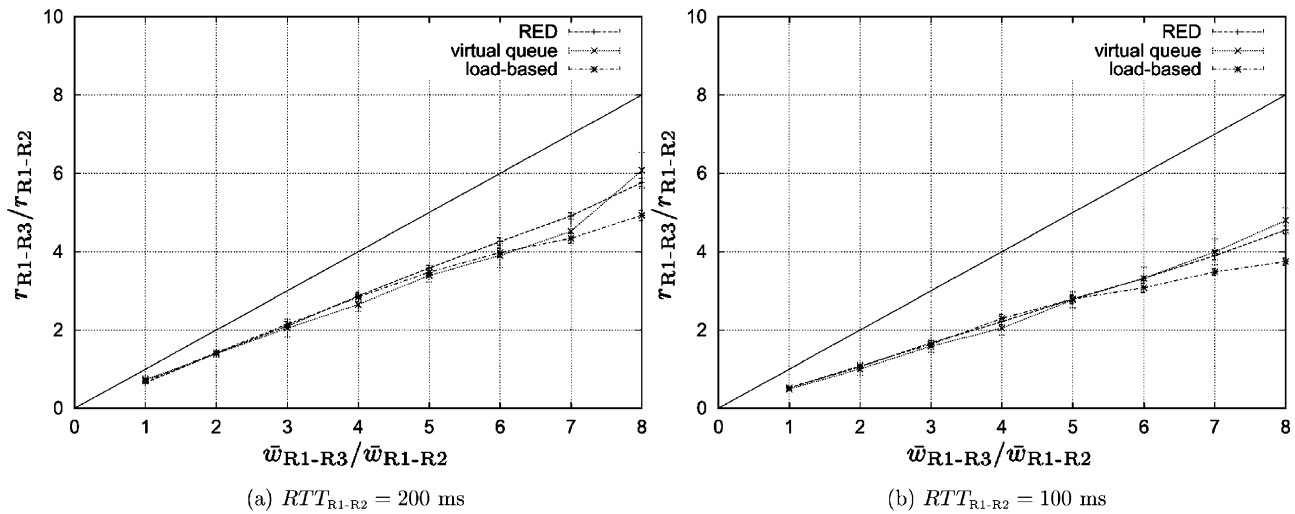


Fig. 5. Service differentiation for different RTT's (multiple link, proportional increase willingness-to-pay).  $\bar{\kappa} = 0.5$ .  $C_{R1-R2} = C_{R2-R3} = 10$  Mbps,  $B_{R1-R2} = B_{R2-R3} = 30$  pkts,  $RTT_{R2-R3} = 100$  ms,  $N_{R1-R2} = N_{R1-R3} = 5$ ,  $N_{R2-R3} = 0$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ . (a)  $RTT_{R1-R2} = 200$  ms, (b)  $RTT_{R1-R2} = 100$  ms.

average rate. The virtual queue marking algorithm has the property of differentiating streams based on their burstiness; see also discussion in Ref. [5]. Hence, the algorithm marks a higher percentage of packets belonging to the burstier stream, which, as explained above, is the stream with smaller willingness-to-pay.

The above is not the case when the end-systems implement inversely proportional decrease willingness-to-pay, as shown in Fig. 4(a). In this case, and for the timescales in which packet marking occurs, the traffic from willingness-to-pay connections with a different willingness-to-pay exhibit roughly the same burstiness, hence are not differentiated and the behavior of virtual queue marking is close to that of RED and load-based marking. This is also illustrated in Fig. 4(b), which shows that in this case the ratio of marks for all three algorithms is roughly proportional to the ratio of willingness-to-pay values.

Now we explain the second observation made from Fig. 3(a), namely the larger variations of service differentiation with virtual queue marking, compared to that of RED and load-based marking. The latter two algorithms mark packets probabilistically; hence marks tend to be spread out, which results in smoother changes of the congestion window. On the other hand, the virtual queue algorithm marks packets deterministically from the time the virtual queue overflows until the time it first becomes empty, hence tends to produce bursts of marks, which result in larger fluctuations of the congestion window.

Next we present simulation results for the multiple link topology in Fig. 2(b). Fig. 5(a) and (b) show the average ratio of throughput for different ratios of willingness-to-pay values for two connections: one that traverses only link R1–R2 in Fig. 2(b), and one that traverses both links R1–R2 and R2–R3. Moreover, in the results of Fig. 5(a) and (b) there are no connections that traverse only link R2–R3 (i.e.  $N_{R2-R3} = 0$ ).

Our first observation regarding Fig. 5(a) and (b) is that, as in the single link scenario, a connection's throughput can be controlled through its willingness-to-pay, and the performance of all three marking algorithms is roughly the same, at least for small ratios of willingness-to-pay. Indeed, the behavior of the virtual queue algorithm is similar to the other two because in the corresponding timescales of the virtual queue overflow, connections with small and large willingness-to-pay values exhibit the same burstiness.

The second observation regarding Fig. 5(a) and (b) is that the line giving the ratio of throughput is lower than the diagonal. The reason is that the horizontal axis in Fig. 5(a) and (b) depicts the ratio of willingness-to-pay values *per round trip time*,  $\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$ . In the single link case, all connections had the same round trip time, hence the ratio  $\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$  was equal to the ratio of willingness-to-pay values *per unit of time*,  $w_{R1-R3}/w_{R1-R2}$ . On the other hand, in the multiple link case, the two connections considered experience a different round trip time: 200 ms in Fig. 5(a) (100 ms in Fig. 5(b)) for the connection traversing only link




Fig. 6. Throughput ratio as a function of willingness-to-pay ratio (multiple link, proportional increase willingness-to-pay).  $\bar{\kappa} = 0.5$ .  $C_{R1-R2} = C_{R2-R3} = 10$  Mbps,  $B_{R1-R2} = B_{R2-R3} = 30$  pkts,  $RTT_{R1-R2} = 200$  ms,  $RTT_{R2-R3} = 100$  ms,  $N_{R1-R2} = N_{R1-R3} = 5$ ,  $N_{R2-R3} = 0$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ .

R1–R2, and 300 ms in Fig. 5(a) (200 ms in Fig. 5(b)) for the connection traversing both R1–R2 and R2–R3. Focusing on Fig. 5(a), the round trip time for the connection traversing only link R1–R2 is 200 ms, hence its willingness-to-pay per round trip time is  $\bar{w}_{R1-R2} = w_{R1-R2} \cdot 200$ . On the other hand, the round trip time for the connection traversing both R1–R2 and R2–R3 is 300 ms, hence its willingness-to-pay per round trip time is  $\bar{w}_{R1-R3} = w_{R1-R3} \cdot 300$ . From the above, it follows that the ratio of willingness-to-pay per unit of time is  $w_{R1-R3}/w_{R1-R2} = (2/3)\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$ . In the steady state, see Eq. (1), a connection's throughput is proportional to the willingness-to-pay per unit of time, hence if the connections experience the same marking probability, we expect the ratio of throughput to be two thirds of the ratio  $\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$ , a result that Fig. 5(a) confirms. For Fig. 5(b), we have  $w_{R1-R3}/w_{R1-R2} = (1/2)\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$ , hence as the figure confirms, the ratio of throughput is one half of the ratio  $\bar{w}_{R1-R3}/\bar{w}_{R1-R2}$ .

One approach to take into account the round trip time is to set the willingness-to-pay per round trip time such that  $\bar{w} \propto (RTT/RTT_{ref})w$ , where  $RTT_{ref}$  some reference round trip time. For example, if we consider the network scenario corresponding to Fig. 5(a), we can take  $RTT_{ref} = 200$  ms, and in order to achieve the same throughput on connections that traverse only link R1–R2

Table 1

Multiple link experiment.  $\bar{\kappa} = 0.5$ ,  $\bar{w}_{R1-R2} = \bar{w}_{R1-R3} = \bar{w}_{R2-R3} = 1$ .  $C_{R1-R2} = C_{R1-R3} = 10$  Mbps  $RTT_{R1-R2} = 100$  ms,  $N_{R1-R2} = N_{R1-R3} = N_{R2-R3} = 5$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $\omega_p = 0.002$

$RTT_{R1-R2}$ (ms)	$r_{R1-R3}/r_{R1-R2}$	$p_{R1-R2}$	$p_{R2-R3}$
200	0.231	0.023	0.042
100	0.251	0.042	0.042




Fig. 7. Service differentiation for different RTT's (multiple link, proportional increase willingness-to-pay).  $\bar{\kappa} = 0.5$ .  $C_{R1-R2} = C_{R2-R3} = 10$  Mbps,  $B_{R1-R2} = B_{R2-R3} = 30$  pkts,  $RTT_{R2-R3} = 100$  ms,  $N_{R1-R2} = N_{R1-R3} = N_{R2-R3} = 5$ . RED:  $\min_{th} = 5$ ,  $\max_{th} = 15$ ,  $\max_p = 0.1$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ . (a)  $RTT_{R1-R2} = 200$  ms, (b)  $RTT_{R1-R2} = 100$  ms.

and connections that traverse both links R1–R2 and R2–R3, we can set  $\bar{w}_{R1-R2} \propto w_{R1-R2}$  and  $\bar{w}_{R1-R3} \propto 1.5w_{R1-R3}$ , with the same proportionality factor in both cases. The results of such an approach are shown in Fig. 6. Observe in this figure that, for large ratios of willingness-to-pay, the ratio of throughput for load-based marking tends to be smaller than that for the other two marking algorithms. This is because with load-based marking, the marking probability depends on the utilization, and when the willingness-to-pay values for the connections traversing both R1–R2 and R2–R3 becomes large, packets are also marked on link R2–R3, hence the marking probability experienced by connections that traverse both R1–R2 and R2–R3 is higher than the marking probability experienced by connections traversing only link R1–R2. (Indeed, it is for the same reason that the ratio of throughput for load-based marking in Fig. 5(a) and (b) is lower than for the other two marking algorithms, for ratios of willingness-to-pay larger than approximately 6).

Next, we present simulation experiments where there are three types of connections: connections that traverse only link R1–R2, only link R2–R3, and both links R1–R2 and R2–R3. The results are shown in Fig. 7(a) for  $RTT_{R1-R2} = 200$  ms, and Fig. 7(b) for  $RTT_{R1-R2} = 100$  ms. Similar to the above, these figures confirm that a connection can control its throughput through its willingness-to-pay.

To illustrate the agreement between simulation results and those predicted by theory, consider a particular experiment where  $\bar{w}_{R1-R2} = \bar{w}_{R1-R3} = 1$ , and both links R1–R2 and R3–R4 implement RED. The ratio of the throughput of a connection traversing both link R1–R2 and R2–R3, and the throughput of a connection traversing only link R1–R2,  $r_{R1-R3}/r_{R1-R2}$ , and the marking probabilities at links R1–R2 and R2–R3 are shown in Table 1.

In the equilibrium, from Eq. (1), we have that a connection's rate  $r$  will satisfy  $r = w/p$  where  $w$  is the

willingness-to-pay and  $p$  is the aggregate marking probability the connection encounters. For  $RTT_{R1-R2} = 200$  ms, the connection traversing only link R1–R2 will have  $w_{R1-R2} = \bar{w}_{R1-R2}/200 = 1/200$ , whereas the connection traversing both links R1–R2 and R2–R3 will have  $w_{R1-R3} = \bar{w}_{R1-R3}/300 = 1/300$ . From Table 1, the aggregate marking probability encountered by each is  $p_{R1-R2} = 0.023$  and  $p_{R1-R3} = p_{R1-R2} + p_{R2-R3} = 0.065$ , respectively. Hence, the ratio of throughput predicted by theory is

$$\frac{r_{R1-R3}}{r_{R1-R2}} = \frac{\frac{w_{R1-R3}}{p_{R1-R3}}}{\frac{w_{R1-R2}}{p_{R1-R2}}} = \frac{\frac{1}{300 \cdot 0.065}}{\frac{1}{200 \cdot 0.023}} \approx 0.236,$$

whereas the ratio of throughput for the case  $RTT_{R1-R2} = 100$  ms is

$$\frac{r_{R1-R3}}{r_{R1-R2}} = \frac{\frac{w_{R1-R3}}{p_{R1-R3}}}{\frac{w_{R1-R2}}{p_{R1-R2}}} = \frac{\frac{1}{200 \cdot 0.084}}{\frac{1}{100 \cdot 0.042}} = 0.25,$$

both of which are in good agreement with the simulation results shown in the second column of Table 1.

Finally, we note that results from experiments with more complex network topologies, including cases with different link capacities, agree with the results and conclusions concerning service differentiation presented in this subsection.

## 5.2. Marking probability

Next we investigate the marking probability as a function of the average utilization, for the three marking algorithms we consider. Note that this probability depends, in addition to the load, also on the burstiness of traffic, which in turn is affected by both the congestion control algorithm in the end-



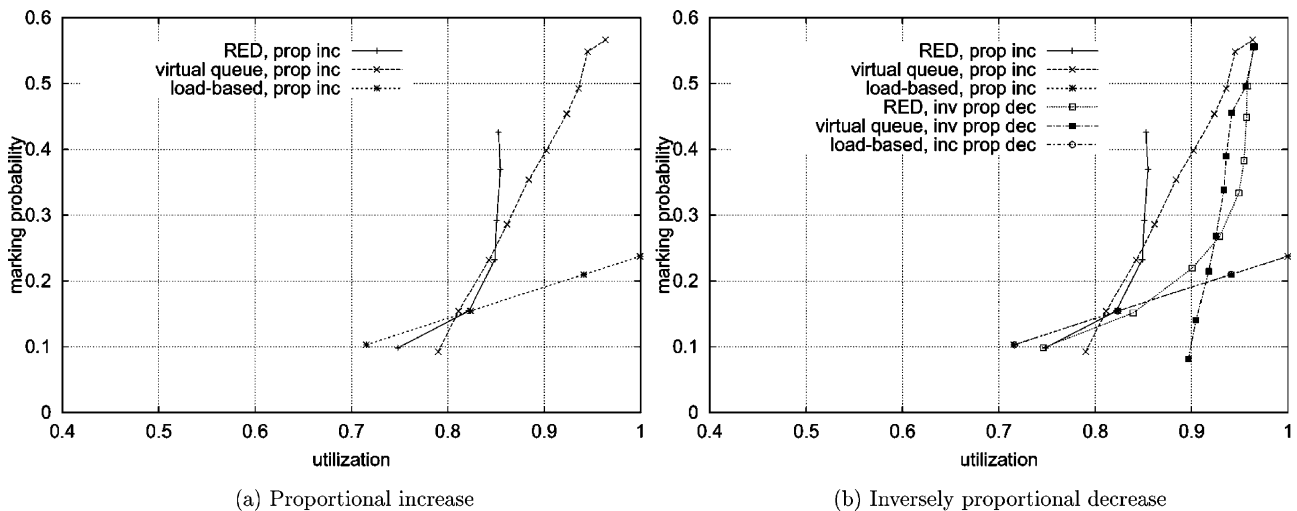


Fig. 8. Marking probability as a function of utilization for proportional increase willingness-to-pay:  $\bar{\kappa} = 0.5$ ,  $\bar{w}_{inc} = 1$  and 6,  $\bar{w}_{dec} = 1$ , and inversely proportional decrease willingness-to-pay:  $\bar{\kappa} = 0.5$ ,  $\bar{w}_{inc} = 1$ ,  $\bar{w}_{dec} = 1$  and 6.  $C = 10$  Mbps,  $B = 30$  pkts,  $RTT = 200$  ms. RED:  $\min_{th} = 2$ ,  $\max_{th} = 6$ ,  $\max_p = 0.1$ ,  $w_q = 0.02$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.6$ ,  $\alpha = 0.71$ . (single link), (a) proportional increase, (b) inversely proportional decrease.

systems and the marking algorithm in the routers. The shape of this curve, as discussed in Section 2 and Ref. [6], affects the convergence and stability behavior of the particular marking algorithm.

Fig. 8(a) and (b) show the marking probability as a function of average utilization for the three marking algorithms, and for the proportional increase Eq. (4) and proportional decrease Eq. (6) willingness-to-pay algorithms. Observe that for both virtual queue and load-based marking, there is a linear dependence of the marking probability on the link utilization. On the other hand, the marking probability for RED is convex and becomes steeper as the utilization increases (we will see later that this is not always the case). Using Fig. 8(a), we can assess the robustness of a marking algorithm with the load as follows: For RED configured with the specific parameters, the performance in terms of convergence and utilization would be similar to that of the

other two marking algorithms for utilization levels less than approximately 80%. However, for utilization levels above 85%, because the marking probability curve becomes very steep, RED would exhibit large fluctuations and stability can be compromised.

Fig. 8(b) shows that for both virtual queue marking and RED, the end-system congestion control algorithm has an affect on the marking probability. Indeed, for inversely proportional decrease willingness-to-pay, the same marking probability is achieved for a higher utilization, compared to proportional increase willingness-to-pay. This is expected, since inversely proportional decrease willingness-to-pay results in smoother traffic, as shown in Fig. 1, hence can achieve a higher utilization, for the same marking probability.

We have performed experiments for other round trip times, and when connections have heterogeneous round trip times. The results show that when  $\kappa$ ,  $w_{inc}$  are kept constant, the marking probability curve is slightly higher for larger round trip times: the reason for this is that a larger round trip time,  $RTT$ , results in larger values for  $\bar{\kappa} = \kappa \cdot RTT$  and  $\bar{w}_{inc} = w_{inc} \cdot RTT$ , hence the window-based congestion control algorithm produces burstier traffic, resulting in a higher marking probability for the same utilization.

Fig. 8(a) and (b) were obtained for some arbitrary<sup>4</sup> selection of parameters for each marking algorithm. Fig. 9 shows that these parameters can be chosen such that the marking probability curves for the three algorithms become very close. This indicates that, at least in terms of the macroscopic convergence and stability properties, all marking algorithms have similar behavior when their parameters are tuned appropriately. Note, however, that these curves are averages hence, as we will see next, do not give a complete picture of the queuing delay.




Fig. 9. Marking probability as a function of utilization (single link).  $\bar{\kappa} = 0.5$ ,  $\bar{w}_{inc} = 1$  and 6,  $\bar{w}_{dec} = 1$ .  $C = 10$  Mbps,  $B = 30$  pkts,  $RTT = 200$  ms. RED:  $\min_{th} = 3$ ,  $\max_{th} = 9$ ,  $\max_p = 0.95$ ,  $w_q = 0.002$ . VQ:  $\theta = 0.95$ . LB:  $t_{avg} = 0.5$  s,  $\rho_0 = 0.7$ ,  $\alpha = 2$ .

<sup>4</sup> For RED we followed the guidelines suggested in the literature for  $\max_p$ , and for the relative values of  $\min_{th}$  and  $\max_{th}$ .

Table 2  
Queueing delay (in ms). Long ftp flows,  $\rho \approx 0.90$

Marking algorithm	Average	Maximum	Std. dev.
RED ( $\min_{th} = 5, \max_{th} = 15, \max_p = 0.10, w_q = 0.002$ )	26.9	89.6	26.4
RED ( $\min_{th} = 3, \max_{th} = 9, \max_p = 0.95, w_q = 0.002$ )	5.3	22.4	3.6
Virtual queue ( $\theta = 0.95$ )	8.6	39.4	6.2
Load-based ( $t_{avg} = 0.5 \text{ s}, \rho_0 = 0.7, \alpha = 2$ )	6.3	38.4	4.8




Fig. 10. Effect of the virtual queue factor  $\theta$  and the buffer size (single link).  $C = 10 \text{ Mbps}$ ,  $RTT = 200 \text{ ms}$ ,  $\bar{w}_{inc} = 1$  and  $6$ ,  $\bar{w}_{dec} = 1$ ,  $\bar{\kappa} = 0.5$ . (a) factor  $\theta$ , (b) buffer  $B$ .

5.3. Queueing delay

Table 2 shows the queueing delay for the three marking algorithms when the average utilization is kept the same. The bottom three rows in the table show that the queueing delay is smaller and exhibits smaller fluctuations for RED and load-based marking, compared to virtual queue marking. This is

expected since, as discussed in Section 5.1, due to probabilistic marking, both RED and load-based marking smooth out marks, which results in smoother traffic; on the other hand, virtual queue marking produces bursts of marks.

The first two rows in Table 2 show that for the same marking algorithm the queueing delay depends on the parameters of the marking algorithm, even when the




Fig. 11. Comparison of willingness-to-pay and MultTCP (single link). WTP:  $\bar{\kappa} = 0.5$ .  $C = 10 \text{ Mbps}$ ,  $RTT = 200 \text{ ms}$ ,  $N = 10$ . RED Marking:  $\min_{th} = 5, \max_{th} = 15, \max_p = 0.1, w_q = 0.002$ . (a) Service differentiation, (b) average utilization.

utilization is kept the same. Indeed, the first row corresponds to an operating point located at the steep region of the marking probability curve for RED, similar to the one shown in Fig. 8(a).

#### 5.4. Effect of marking algorithm parameters

In addition to their performance, how easy or difficult it is to tune the parameters of a marking algorithm is equally, if not more, important. Towards this end, in this subsection we provide some insight on how the parameters of the marking algorithm affect the marking probability function.

Fig. 10(a) shows how the marking probability for the virtual queue algorithm is affected by factor  $\theta$ , which determines the buffer and capacity of the virtual queue. Observe that as  $\theta$  increases, the marking probability shifts to the right, i.e. for larger  $\theta$ , the same utilization corresponds to a lower marking probability. Also, observe that the factor has a very small effect on the slope of the curve. Finally, Fig. 10(b) shows that a larger buffer results in a steeper curve for the marking probability.<sup>5</sup> Indeed, for small utilizations, the marking probability is higher for a smaller buffer, since the buffer overflows more frequently. On the other hand, for large utilizations the effect is the opposite, i.e. the marking probability is higher for a larger buffer, because now the dominating effect is the larger overflow (marking) periods, since a larger buffer needs more time to empty.

For load-based marking the effect of its parameters is straightforward: The parameter  $\alpha$  affects the slope of the marking probability curve,  $\rho_0$  (for constant  $\alpha$ ) affects the marking probability for a given average utilization, and the averaging interval affects how fast the algorithm responds to changes of the load, and the timescales of traffic burstiness that affect the marking probability. Moreover, the averaging interval affects the maximum queue backlog that appears when the load changes.

Finally, for RED we have observed that when  $\max_p$  is small, then the marking probability is convex, as in Fig. 8. On the other hand, when  $\max_p$  approaches 1, then the marking probability tends to be linear, as in Fig. 9. In general, the rules for tuning the parameters of RED in the case of willingness-to-pay congestion control are not the same as the rules in the case of TCP.

#### 5.5. Comparison of willingness-to-pay and MulTCP

Fig. 11 compares willingness-to-pay with MulTCP [10]. The latter follows TCP's double multiplicative decrease, but supports service differentiation. Fig. 11(a) shows that the ratio of average throughput is higher than the ratio of weights for MulTCP, indicating that MulTCP favors connections with larger weights.

Fig. 11(b) compares the utilization achieved with willingness-to-pay and MulTCP, when the same number of connections are multiplexed in a link with RED marking. Observe that for small weights, willingness-to-pay achieves a higher utilization than MulTCP. This is expected since MulTCP decreases its window more aggressively when congestion is detected. As the weights increase, however, the difference between the two algorithms is smaller. We conjecture that this is due to the following: For MulTCP, connections are bursty even for small weights, due to the proportional decrease of the congestion window when an ECN mark is received; moreover, as weights increase, the resulting increase of aggressiveness helps connections achieve a higher average throughput. On the other hand, for small weights willingness-to-pay connections produce smooth traffic, hence the average utilization is high; when weights increase, however, connections become burstier, and as a result the average utilization decreases.

## 6. Related work

Next we discuss some related work. Note that this is not an exhaustive survey of all the work in the area.

The authors of Ref. [17] investigate the limitations of using double proportional decrease algorithms for achieving service differentiation. Their results are in agreement with the results of Section 5.5. They also investigate an algorithm where the congestion window is adjusted based on the percentage of lost packets. The loss adaptive algorithm in Ref. [17] achieves in the steady state an average throughput inversely proportional to the loss probability. This is similar to the dependence of the average throughput on the marking probability with the willingness-to-pay algorithms investigated in this paper. Comparison of the two algorithms in terms of throughput and delay is an area for further investigation.

The authors of Ref. [18] investigate the convergence and steady state behavior, using both simulation and dynamical system modeling, of the rate-based congestion control algorithm given by Eq. (1), when routers implement threshold-based marking, i.e. all packets are marked when the queue length exceeds a threshold. Additionally, they discuss issues regarding the timescales for a connection's rate to reach equilibrium in relation to the timescales of the inter-arrival time between new connections and the duration of each connection.

The author of Ref. [19] discusses the fairness of marking algorithms, using ideas from large deviations theory and economics. A number of marking algorithms that have been proposed in the literature are investigated, and a threshold-based marking algorithm that marks all packets in the queue when the queue length exceeds a threshold, which is adaptively set, is proposed. The authors of Ref. [20] propose an active queue management algorithm, called Random Exponential Marking (REM), where a link's marking

<sup>5</sup> Varying the buffer size, while keeping the same factor for both the capacity and buffer of the virtual queue is equivalent to having a different factor for each of these two resources.

probability is a concave and exponential function of the link's price, which is updated based on the rate mismatch (difference between input rate and link capacity) and the queue mismatch (difference between current and target queue length) at the link.

The authors of Ref. [21] investigate, using a fluid model, the decentralized selection of the marking rate at each router of a network in order to achieve loss-free operation. The convergence of the scheme is investigated using a timescale decomposition of the system into a slow and fast system model. Finally, the works of Refs. [22–24] present insightful results on the behavior of TCP and RED-like algorithms that compute the packet marking probability based on the queue length.

## 7. Concluding remarks

We have investigated, using simulation, the service differentiation and performance achieved with end-to-end weighted window-based congestion control algorithms, and how this is affected by the marking algorithms operating in the routers. Our simulation experiments included both single and multiple link topologies, where different connections had different round trip times. The congestion control algorithms<sup>6</sup> investigated include willingness-to-pay, which follows a single multiplicative decrease of the congestion window upon congestion. Our results show that, for networks with no losses supporting ECN, willingness-to-pay can be fairer and lead to higher utilization compared to MulTCP, which supports service differentiation and follows a double multiplicative decrease, similar to TCP. Although our results were for experiments where no losses occurred, we expect that they will hold when the loss probability is much smaller than the marking probability; the latter is likely to be the case in real networks, in order to avoid the overhead due to packet retransmissions. In cases where the loss rate tends to be large, algorithms that decrease the congestion window by a large amount, e.g. by half as in the case of TCP and MulTCP, might be more appropriate, since a large decrease of the congestion window can result in a lower loss rate, hence a smaller number of retransmissions. The marking algorithms investigated include RED, virtual queue marking, and load-based marking. Our results show that service differentiation and queueing delay can be worse for the virtual queue algorithm, where packet marking is deterministic and occurs at the timescales of the virtual queue overflow, compared to RED and load-based marking, where marking is probabilistic. On the other hand, the virtual

queue marking algorithm can differentiate flows according to their burstiness.

An issue for further investigation is how to automate the tuning of the parameters (adaptation) of the marking algorithms. Motivation for such automated adaptation is that, in agreement with the results regarding the tuning of RED parameters for TCP connections, no set of parameters are optimal for all possible traffic mixes and characteristics. Another issue with important engineering implications, which can assist in tuning the parameters of a marking algorithm, is the estimation of the equilibrium given the congestion control and marking algorithms; a related issue, if each mark is charged by a fixed price, is how the price per mark affects the link utilization, and subsequently, given the aggregate demand, how to determine this price in order to achieve a target utilization. Results in this direction are reported in Ref. [26]. Another important direction for further work is the application of ideas and results investigated in this paper to wireless networks. This will be particularly important since there is a limited capability, compared to fixed networks, to increase the capacity of wireless networks and the emerging multimedia services and applications will increase the demand for bandwidth in such networks, hence the need to develop efficient and robust mechanisms for controlling wireless network resources. The framework underlying the work presented in this paper is a strong candidate for constructing such mechanisms. Moreover, ECN marking for improving the performance of TCP over wireless networks has been suggested in Ref. [27], and for 3G networks in Ref. [28].

## Acknowledgments

This work was supported in part by the European Commission under IST project M3I 'Market Managed Multi-service Internet' (IST-1999-11429).

## References

- [1] M. Allman, V. Paxson, W. Stevens, TCP Congestion Control, RFC 2581, 1999, April.
- [2] V. Jacobson, M.J. Karels, Congestion avoidance and control, Proceedings of ACM SIGCOMM'88.
- [3] B. Braden, et al., Recommendations on Queue Management and Congestion Avoidance in the Internet, RFC 2309, 1998, April.
- [4] K.K. Ramakrishnan, S. Floyd, D. Black, The Addition of Explicit Congestion Notification (ECN) to IP, RFC 3168, 2001, September.
- [5] R.J. Gibbens, F.P. Kelly, Resource pricing and congestion control, *Automatica* 35 (1999) 1969–1985.
- [6] F.P. Kelly, Models for a self-managed Internet, *Philosophical Transactions of the Royal Society A358* (2000) 2335–2348.
- [7] P.B. Key, D.R. McAuley, Differential QoS and pricing in networks: Where flow control meets game theory, *IEE Proceedings Software* 146 (2) (1999) 39–43.
- [8] R.J. Gibbens, P. Key, Distributed control and resource marking using best-effort routers, *IEEE Network* (2001) 54–59.

<sup>6</sup> The willingness-to-pay and packet marking algorithms have also been implemented in a real test-bed comprising of FreeBSD workstations. The implementations are available at [http://www.ics.forth.gr/netgroup/publications/wtp\\_vq.html](http://www.ics.forth.gr/netgroup/publications/wtp_vq.html). Results from experiments conducted over the test-bed are reported in Ref. [25].

- [9] J. Saltzer, D. Reed, D. Clark, End-to-end arguments in system design, *ACM Transactions on Computer Systems* 2 (4) (1984) 277–288.
- [10] J. Crowcroft, P. Oechslin, Differentiated end-to-end Internet services using a weighted proportional fair sharing TCP, *Computer Communications Review* 28 (3) (1998) 53–67.
- [11] S. Floyd, V. Jacobson, Random early detection gateways for congestion avoidance, *IEEE/ACM Transactions on Networking* 1 (1993) 397–413.
- [12] R.J. Gibbens, P.B. Key, S.R.E. Turner, Properties of the Virtual Queue Marking Algorithm, *Proceedings of Seventeenth UK Teletraffic Symposium*, 2001.
- [13] F.P. Kelly, Charging and rate control for elastic traffic, *European Transactions on Telecommunications* 8 (1997) 33–37.
- [14] F.P. Kelly, A. Maulloo, D. Tan, Rate control in communication networks: shadow prices, proportional fairness and stability, *Journal of the Operational Research Society* 49 (1998) 237–252.
- [15] S. Floyd, V. Jacobson, On traffic phase effects in packet-switched gateways, *Internetworking: Research and Experience* 3 (1992) 115–156.
- [16] UCB/LBNL/VINT, Network Simulator—ns (version 2), <http://www-mash.cs.berkeley.edu/ns.html>.
- [17] T. Nandagopal, K.-W. Lee, J.-R. Li, V. Bharghavan, Scalable service differentiation using purely end-to-end mechanisms: features and limitations, *Proceedings of IEEE IWQoS'00*.
- [18] K. Laevens, P. Key, D. McAuley, An ECN-based end-to-end congestion control framework: experiments and evaluation, *Technical Report MSR-TR-2000-104*, Microsoft Research, October 2000.
- [19] D. Wischik, How to Mark Fairly, *Proceedings of Workshop on Internet Service Quality Economics*, MIT, 1999.
- [20] S. Athuraliya, S.H. Low, V.H. Li, Q. Yin, REM: Active Queue Management, *IEEE Network* (2001) 48–53.
- [21] S. Kunniyur, R. Srikant, A time scale decomposition approach to adaptive ECN marking, *Proceedings of IEEE INFOCOM'01*.
- [22] V. Firoiu, M. Borden, A study of active queue management for congestion control, *Proceedings of IEEE INFOCOM'00*.
- [23] C.V. Hollot, V. Misra, D. Towsley, W.-B. Gong, On designing improved controllers for AQM routers supporting TCP flows, *Proceedings of IEEE INFOCOM'01*.
- [24] P. Kuusela, P. Lassila, J. Virtamo, Stability of TCP-RED Congestion Control, *Proceedings of the Seventeenth International Teletraffic Congress (ITC-17)*, 2001.
- [25] P. Antoniadis, C. Courcoubetis, G. Margetis, V.A. Siris, G.D. Stamoulis, Efficient Adaptation to Dynamic Pricing Communicated by ECN Marks: Scenarios for Experimental Assessment, *Proceedings of SPIE International Symposium on Information Technologies*, 2000.
- [26] V.A. Siris, C. Courcoubetis, Interaction of Congestion Control and Packet Marking Algorithms in ECN Networks, *Technical Report No. 296*, ICS-FORTH, 2001, November.
- [27] G. Montenegro, S. Dawkins, M. Kojo, V. Magret, N. Vaidya, Long Thin Networks, *RFC 2757*, 2000, January.
- [28] H. Inamura, G. Montenegro, M. Hata, J. James, A. Hamed, R. Garces, F. Wills, TCP over 2.5g and 3g wireless networks, *Internet Draft—draft-ietf-pilc-2.5g3g-00*, 2001, February.