Adaptive Packet Marking for Achieving Fairness in DiffServ Networks

Vasilios A. Siris* and Harris M. Marinakis

Institute of Computer Science, FORTH
and Department of Computer Science, University of Crete
P.O. Box 1385, GR 711 10 Heraklion, Crete, Greece

Abstract

Current DiffServ (Differentiated Services) networks supporting Assured Forwarding (AF) service succeed in providing minimum rate guarantees, but fail to distribute network capacity in a fair way. To address the unfairness problem we propose simple adaptive markers operating at the network edge, whose marking function adapts to changes of the traffic mix. Extensive simulation experiments indicate that adaptive marking can effectively address the unfairness problem, without decreasing network utilization. The proposed approach can be implemented using a Bandwidth Broker architecture, whereby a centralized Bandwidth Broker sets the control parameter of the adaptive markers, based on the current traffic contracts.

Keywords: Assured Forwarding, fairness, equal sharing

1 Introduction

The Differentiated Services [4] (DiffServ) architecture proposes a scalable solution for QoS provision in IP networks, based on aggregating flows to a small number of traffic classes. In addition to the basic best-effort service offered by the current Internet, DiffServ introduces two additional packet forwarding mechanisms (Per

*Corresponding author: V. A. Siris, vsiris@ics.forth.gr, tel./fax: +30 2810 391726/391601.

Hop Behaviors - PHBs) called *Expedited Forwarding* (EF) and *Assured Forwarding* (AF) [3]. EF services provide strict guarantees on throughput, delay, jitter, and packet loss probability. On the other hand, AF services are appropriate for applications that can adapt their throughput, but typically require an uncongested network to achieve satisfactory performance. Network providers that offer AF services aim to provide a minimum amount of bandwidth, or *target rate*, to each customer; this target rate is defined in the traffic contract, which is part of the Service Level Agreement (SLA) between the customer and the provider. When the network load is low, AF customers can utilize the excess capacity, and achieve a throughput higher than their target rate.

Implementations of AF services are based on differential packet marking at the edge of the network, depending on the conformance to the SLA parameters, and differential dropping at the core of the network, with a higher probability of dropping non-conforming packets. *Traffic Conditioners* based on the *Token Bucket* algorithm or the *Time Sliding Window* (TSW) [1] average rate estimator are used for metering, to check for conformance and mark packets appropriately. The work in [5] discusses problems with Token Bucket marking for supporting throughput guarantees of TCP flows. At the network core, differential dropping is achieved through Active Queue Management (AQM) techniques like RIO (RED with IN and OUT) [1], which defines two sets of RED parameters, one for in-profile (IN) packets and another one for out-of-profile (OUT) packets.

Although, current AF implementations succeed in providing minimum rate guarantees, they do not achieve fair distribution of the network capacity among all users. In under-subscribed networks, excess bandwidth should be distributed in a fair manner. On the other hand, due to network failures or high bandwidth demand periods, there are cases where the network capacity is not sufficient to provide the target rate to all users. In such over-subscription cases, throughput degradation should also be done in a fair way. Fairness may be specified in various ways. A system aiming to provide equal share fairness should guarantee the target rate of each user, and distribute any excess capacity equally among them, irrespective of their target rates. In over-subscribed networks, users should experience equal degradations of throughput below their target rates. Another definition of fairness is to distribute excess bandwidth proportionally to each user's target rate. According to this definition of fairness, in over-subscribed networks degradation of throughput should be proportional to each user's target rate. Studies [11] [8], show that unfairness phenomena can be caused by differences in TCP parameters, such as the packet size, the round trip time (RTT), the TCP protocol stack, the number of TCP connections in an aggregate flow, and the target rates. Unfairness phenomena also occur among UDP flows with different sending rates. Finally, when TCP and UDP traffic share the same traffic queue, unfairness is observed in favour of the non-responsive UDP flows.

To address the unfairness problem, we propose *adaptive markers* that operate at the network edge, and consist of simple marking algorithms that adapt their marking function to changes of the network traffic mix. Extensive simulation experiments indicate that adaptive marking can successfully tackle the unfairness problem in DiffServ networks. Moreover, because of the simplicity of the packet marking algorithm, the proposed approach is less complex compared to other approaches involving modifications to the transport protocols running at the end-systems or to the dropping/marking procedures in the network core. Finally, the proposed approach can be implemented using the Bandwidth Broker architecture, whereby a centralized Bandwidth Broker adjusts the control parameter of the adaptive markers, based on the current set of traffic contracts.

The rest of this paper is structured as follows. In Section 2 we summarize representative work on packet marking algorithms for dealing with unfairness in DiffServ networks. In Section 3 we motivate and discuss the proposed adaptive marking approach, describing three algorithms for achieving equal sharing of excess capacity. The algorithms are investigated in Section 4 for a wide range of scenarios, which include different causes of unfairness and different traffic mixes. In Section 5 we discuss how the proposed approach can be implemented using a Bandwidth Broker architecture. Finally, in Section 6 we conclude the paper.

2 Related Work

The work in [1] proposes the TSW (Time Sliding Window) two color marker (TSW2CM), which operates as follows:

With each packet arrival calculate the AverageRate. If AverageRate is below

TargetRate then mark packet as IN, else mark packet as OUT with probability

$$P_{TSW2CM}(AverageRate) = \frac{AverageRate - TargetRate}{AverageRate} \qquad (1)$$

The TSW2CM marker can guarantee the target rates in under-subscribed networks, but suffers from the unfairness problems identified in the previous section.

Next we summarize representative work that proposes new marking algorithms for dealing with the unfairness problem. The Memory-Based Marker (MBM) [7] proposes a simple marking algorithm for TCP aggregates, that adjusts the marking probabilities based on rate changes, and comparison of the average rate with the target rate. MBM improves fairness in some cases, but in other cases its performance remains far from optimum. The work in [2] proposes different intelligent traffic conditioners to deal with different causes of unfairness, hence is not a general solution. Equation Based Packet marking (EBPM) [12] proposes an effective but quite complex mechanism based on packet loss rate estimation to deal with unfairness among heterogeneous TCP sources. Furthermore, EBPM needs to identify packet losses, which is not always easy to achieve, and its performance is not evaluated for aggregates of individual TCP connections. Finally, the Rate Adaptive Marking (RAM) [6] algorithm marks IN packets proportional to the target rate and inversely proportional to a factor reflecting how much higher the average throughput is compared to the target rate. In a scenario where a single flow consists of many individual TCP connections, the RAM algorithm improves the fairness of the TSW2CM algorithm by a factor of 5% (93% fairness with TSW2CM, 98% fairness with RAM), which is lower than the improvement achieved by the adaptive marking scheme proposed in this paper.

The adaptive markers proposed in this paper differ from the above proposals in that they entail simple marking algorithms at the network edge that adapt their marking probability function to changes of the traffic mix. Moreover, they address the unfairness problem in a wide variety of scenarios, which capture different causes of unfairness.

3 Adaptive marking for equal sharing of excess capacity

The majority of studies related to the unfairness problems in DiffServ networks propose solutions that follow the general marking rule that the rate of IN packets should be close to the target rate. In this paper, as in [6], we propose an approach that changes this marking rule to deal with the unfairness problem.

A key observation is that fairness depends on the rate threshold above which packets begin to be marked as OUT. The TSW2CM marking algorithm, given by (1), begins marking packets as OUT when the average rate exceeds the target rate. As our experimental evidence shows, if marking packets as OUT begins at a different threshold than the target rate, then fairness can be improved. Based on this observation we propose *adaptive markers*, which consist of simple parameterized marking functions, where the control parameter determines the threshold rate above which packets begin to be marked as OUT. The optimal value of this parameter, which leads to fair sharing of bandwidth, depends on the network traffic mix. Based on the above discussion, we define the following three simple marking algorithms:

A (Step) - *With each packet arrival estimate the AverageRate. If AverageRate is below (TargetRate + h) then mark packet as IN, else mark packet as OUT.*

B (Linear) - *With each packet arrival estimate the AverageRate. If AverageRate is below (TargetRate + h) then mark packet as IN, else mark packet as OUT with probability*

$$P(AverageRate) = d \cdot (AverageRate - TargetRate - h)\,,$$

where the slope factor d is a constant.

C (Concave) - *With each packet arrival estimate the AverageRate. If AverageRate is below (TargetRate + h) then mark packet as IN, else mark packet as OUT with probability*

$$P(AverageRate) = \frac{AverageRate - TargetRate - h}{AverageRate - h}\,.$$

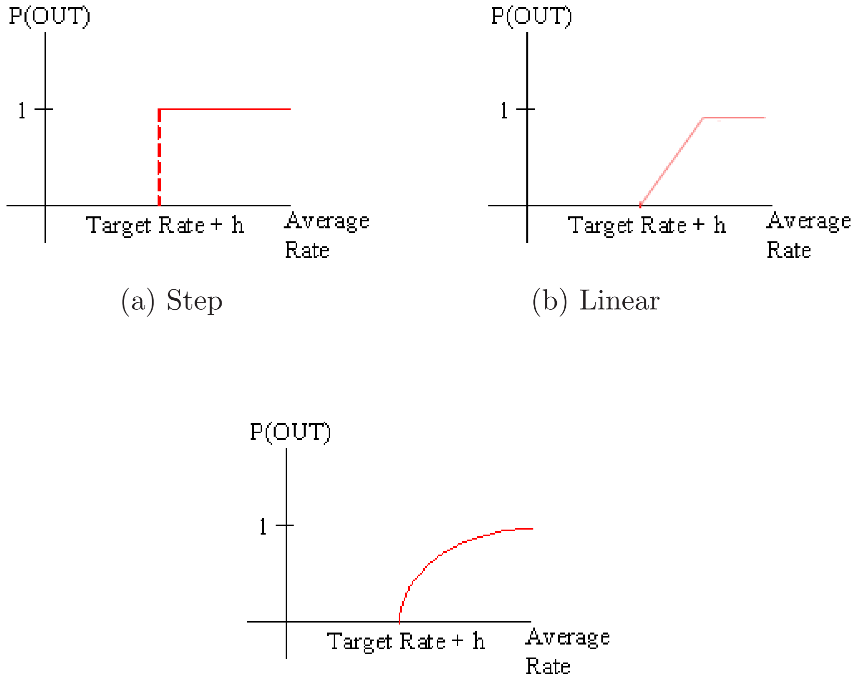

(a) Step (b) Linear (c) Concave

Figure 1: Marking probabilities of algorithms A,B,C

Note that the three marking algorithms, as the TSW2CM algorithm, use the TSW rate estimator for estimating the average rate. Fig. 1 shows the probability of marking packets as OUT, as a function of the average rate, for the three marking algorithms.

The above algorithms are adaptive in the sense that the control parameter h can be adjusted based on the current traffic mix. Indeed, the sum $TargetRate+h$ is the threshold above which marking of packets as OUT begins. Hence, by changing the value of h, the marking function shifts along the average rate axis. Note that the marking function of algorithm C (concave) is equivalent to $P_{TSW2CM}(AverageRate - h)$, which is given by (1), hence can be considered an adaptive version of the TSW2CM algorithm. In the investigations of Section 4, we first consider that the parameter h is optimally set. The optimal value is found through experimentation, by adjusting h in the direction where fairness increases. Latter, in Section 4.3 we discuss a simple and effective approach for setting its value based on the target rate of current traffic contracts.

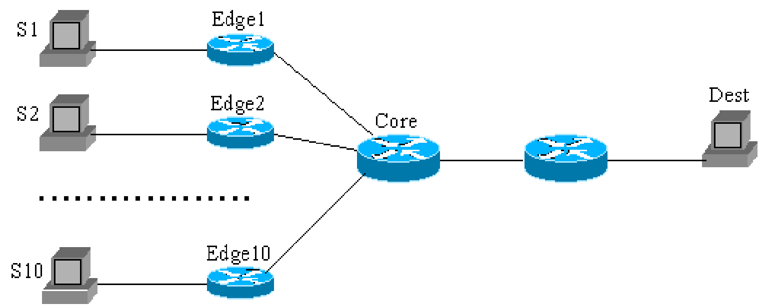

Figure 2: Network Topology

4 Simulation experiments

In this section we investigate the effectiveness of the adaptive marking algorithms proposed in the previous section using the ns-2 simulator [13]. Fig. 2 shows the network topology used in the experiments. Sources S1 to S10 generate traffic (aggregates of TCP or UDP flows), which enters the network through their corresponding edge routers. Traffic in edge routers is metered using the TSW average rate estimator, and marked as IN or OUT according to the particular marking algorithm. Inside the network, traffic traverses the core router, which implements the RIO queue management algorithm. RIO's queue size was set to 70 packets, and its parameters were $(Min_{th}, Max_{th}, Max_p) = (20, 40, 1.0)$ for the OUT packets and $(Min_{th}, Max_{th}, Max_p) = (40, 60, 1.0)$ for the IN packets. Our experiments showed that the specific values of the RIO's parameters did not have a significant influence on fairness.

All links have capacity 5 Mbps and sources generate FTP/TCP or CBR/UDP traffic. Unless otherwise stated, all links have latency 1 ms, and sources producing FTP traffic generate an aggregate of 5 TCP flows, with packet size 500 bytes and an aggregate target rate 200 Kbps. The duration of each experiment was 80 seconds, and the throughput was measured at the destination node. The slope factor d for the linear marking function (algorithm B) was set to 10^{-5}.

Fairness is evaluated using R. Jain's fairness index:

$$Fairness\ Index = \frac{(\sum x_i)^2}{n \cdot \sum x_i^2},$$

where x_i is the excess throughput of source i, and n is the number of sources.

The fairness index varies from $1/n$ (where one source obtains all the capacity) to 1 (where all sources obtain an equal capacity share).

4.1 Fairness improvement with adaptive marking

The experiments of this subsection determine the upper bound for the fairness achieved with the three marking algorithms presented in the previous section, when parameter h is optimally set for each algorithm. The optimal value of h was determined through experimentation. We consider a number of different scenarios, each corresponding to a different cause of unfairness. As a benchmark, for each experiment we also computed the fairness achieved with the TSW2CM algorithm.

Different TCP packet sizes This experiment involved sources generating TCP traffic with different packet sizes: 5 sources had packet size 100 bytes and 5 sources had packet size 1500 bytes. Fig. 3 shows the results for the three marking algorithms. Observe that the fairness achieved with the three adaptive marking algorithms approaches the optimal value, with the fairness of the linear algorithm being slightly less. The fairness index for the TSW2CM marker is much lower: 0.73; such a low value corresponds to sources with 100 byte packets achieving a throughput of approximately 320 Kbps and sources with 1500 byte packets achieving a throughput of 660 Kbps, i.e. more than double. On the other hand, the fairness index for the step and concave marking algorithm is 0.99, which corresponds to sources with 100 byte packets achieving a throughput of approximately 483 Kbps and sources with 1500 byte packets approximately 495 Kbps; the ideal case would be for all sources to obtain a throughput of 500 Kbps.

Differences in other TCP parameters Next we investigate the fairness when other parameters differ: Number of TCP flows in each aggregate (5 sources each with 5 TCP flows in the aggregate, and 5 sources each with 100 TCP flows), Round Trip Time - RTT (5 sources with propagation delay 7 msec, and 5 sources with 107 msec), and TCP protocol stacks (5 sources with the TCP Tahoe stack, and 5 sources with TCP Sack). Results of these experiments are presented in

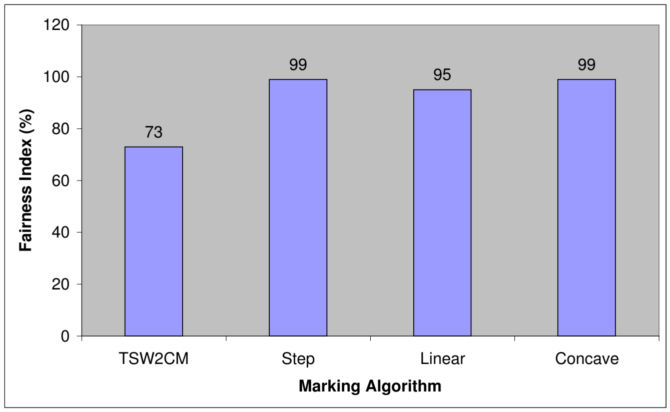

Figure 3: Fairness index for different TCP packet sizes: 5 sources with packet size 100 bytes and 5 sources with packet size 1500 bytes.

Fig. 4, which also presents the results of a more complicated scenario, where 5 sources each generate an aggregate of 5 TCP flows with 100 byte packets, 3 sources generate an aggregate of 5 TCP flows with 1500 byte packets, and 2 sources generate an aggregate of 100 TCP flows with 1500 byte packets. In all cases we observe that the three adaptive marking algorithms improve fairness. The improvements are higher, when the unfairness is due to the number of individual TCP flows, packet size, and RTT.

Different target rates In the majority of scenarios where sources have different target rates, the results are similar to the previous ones. In some cases, where the target rates differ significantly, which is the case in Fig. 5 (5 sources with target rate 700 Kbps and 5 sources with target rate 100 Kbps), the step and linear marking algorithms performed worse than then original TSW2CM algorithm. After some investigation, we conjecture that this is mainly because the average rate measured at the ingress edge node is different (higher) than the rate at the receiving node, which results in bias in favour of small target rates. This bias is overcome with the TSW2CM and the concave marking algorithms, because with these algorithms, for average rates higher than the marking threshold, the

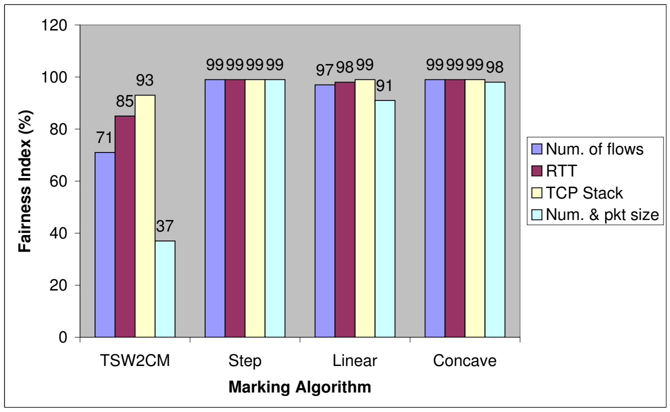

Figure 4: Fairness index for different number of TCP flows in aggregate, RTT, TCP stack, and combination of number of flows and packet size.

marking probability decreases as the target rate increases.

Over-subscribed network This experiment investigates the case where the sum of the target rates is greater than the total capacity. All target rates are set to 800 Kbps, when the network is only capable of providing 500 Kbps to each source. 5 sources generate 100 byte packets and 5 sources generate 1500 byte packets. In such an over-subscribed situation, to achieve fairness the throughput of all sources should be equally decreased. Fig. 6 shows that in this case the adaptive marking algorithms achieve a higher fairness compared to the TSW2CM algorithm.

UDP flows Fig. 6 shows the fairness results in the case of UDP traffic with a different sending rate. In this experiment, 5 sources have rate 600 Kbps, and 5 sources have rate 2 Mbps. Fig. 6 also shows the fairness results when TCP and UDP traffic coexist. In this experiment, 5 sources generate TCP traffic, and 5 sources generate UDP traffic each with rate 5 Mbps. In all the above cases, we observe that adaptive markers can effectively deal with the unfairness problem.

In all the previous experiments, adaptive marking achieved a very high fair-

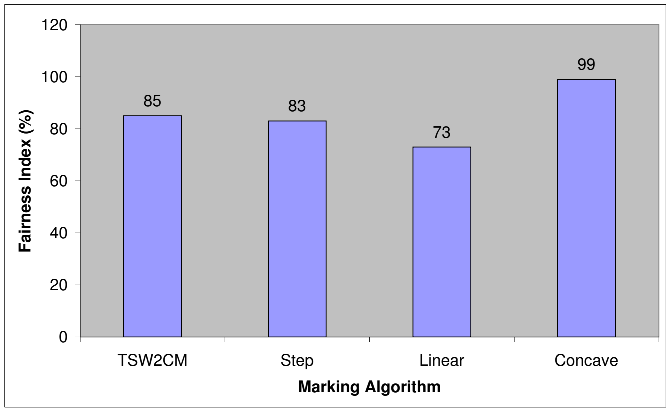

Figure 5: Fairness index for different Target Rates: 5 sources with Target Rate 700 Kbps and 5 sources with Target Rate 100 Kbps.

ness, independent of the causes for unfairness. This is achieved because the parameter h is set so as to protect each source's share of capacity, either their target rate plus some share of excess capacity in the case of under-subscribed networks, or some percentage of their target rate in the case of over-subscribed networks, by marking packets conforming to the fair share as IN. Such an approach is different than that taken by other approaches, which try to directly alleviate the causes of unfairness. Hence, our approach is more general and can address any cause of unfairness.

4.2 Sensitivity to the parameter h

In an actual network, the number of sources (users) changes. The optimal value of the control parameter h depends on the number of active sources. Fig. 7(a) and 7(b) show the fairness index for the three adaptive marking algorithms changes when the number of sources changes (decreases), while the value of h remains equal to its optimal value for the case of 10 sources. Observe that the fairness index does not decrease significantly, even when the number of sources is reduced from 10 to 6. Nevertheless, also observe that the linear marking algorithm is

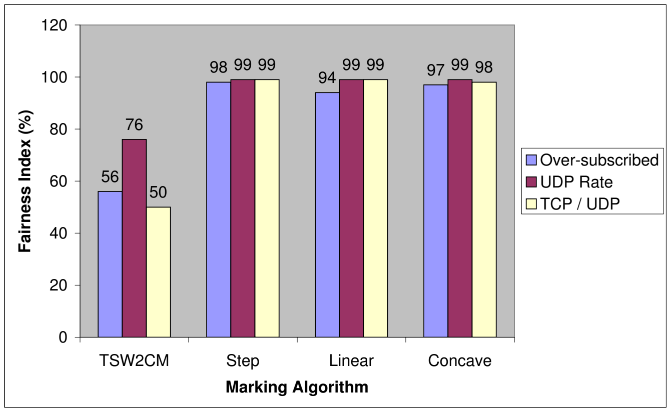

Figure 6: Fairness index in three experiments: over-subscribed network, UDP traffic, and mixed TCP/UDP traffic.

slightly less sensitive than the other two, and the step marking algorithm is the most sensitive. Indeed, the sensitivity of the linear algorithm depends on the slope parameter d: a larger slope increases sensitivity, and when the slope parameter obtains a large value the linear marking algorithm behaves identical to the step marking algorithm.

4.3 Selection of the parameter h

In the experiments up to now, the parameter h was set to its optimal value, i.e. the value for which the fairness index obtains its maximum value. Hence the results demonstrate the maximum gains achievable with adaptive marking. The optimal value of h was determined through experimentation, by estimating the fairness index and iteratively adjusting h in the direction where the fairness index increased. Although it would be possible to also apply such an approach online, in this subsection we describe an alternative, simpler, approach. According to the description of the adaptive marking algorithms in Section 3, the sum $TargetRate + h$ represents the threshold above which packets start to be marked as OUT. With equal sharing of excess capacity, each flow should claim an equal

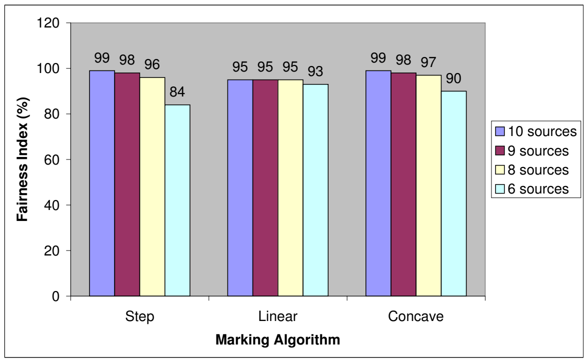

(a) TCP sources: half with 100 byte pkts and half with 1500 byte pkts

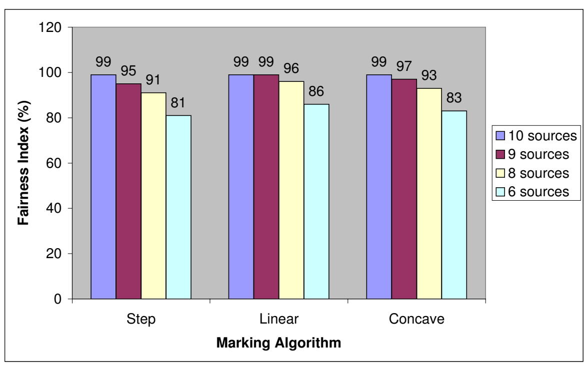

(b) UDP sources: half with rate 600 Kbps and half with rate 2 Mbps

Figure 7: Fairness index when the number of sources change, but the control parameter remains the same (equal to the optimal for 10 sources).

Table 1: Fairness and utilization for different TCP packet sizes: 5 sources with packet size 100 bytes and 5 sources with packet size 1500 bytes.

	TSW2CM	Step	Linear	Concave
Fairness (%)	73	99	95	99
Utilization (%) with optimum selection of h	97	99	99	99
Utilization (%) with h computed using (2)	-	96	99	97

share of the excess capacity, i.e. $ExcessRate = \frac{ExcessCapacity}{\#\ of\ Sources}$, hence a logical selection of h would be

$$h = ExcessRate = \frac{Capacity - \sum_i TargetRate_i}{\#\ of\ Sources}, \quad (2)$$

where $TargetRate_i$ is the target rate for source i. Such a selection, in the case of an over-subscribed network, allows the marker to "protect" each source's fair share of unreserved capacity in addition to its target rate.

Fig. 8(a) and 8(b) compares the fairness index for the optimal value of h, and the fairness index when h is set according to (2). The figures show that the performance of the concave adaptive marking algorithm for parameter h set according to (2) is very close to the performance when h is set to its optimal value; this is not the case with the other two marking algorithms (step and linear). Hence, the concave marking algorithm can address the unfairness problem, with a simple analytical computation of the adaptation parameter h.

4.4 Capacity utilization

The results in the previous section showed that adaptive marking can improve fairness. In general, procedures that address the unfairness problem can result in decreased capacity utilization. Table 1, which presents the results for the utilization corresponding to the experiment in Fig. 3 (which was for TCP connections with different packet size), shows that the proposed marking algorithms increases fairness, while maintaining a very high capacity utilization.

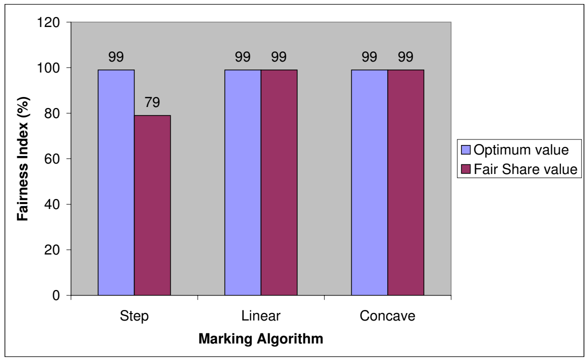

(a) 5 UDP sources with rate 600 Kbps and 5 UDP sources with rate 2 Mbps

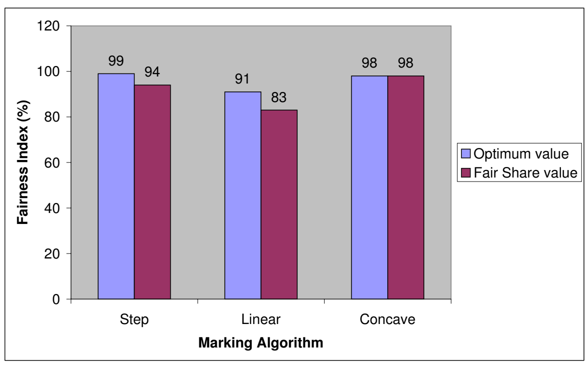

(b) 5 TCP sources (each with 5 TCP flows) with 100 byte pkts, 3 TCP sources (each with 5 TCP flows) with 1500 byte pkts, 2 TCP sources (each with 100 TCP flows) with 1500 byte pkts

Figure 8: Fairness index with optimal value of h and the value determined by their fair share (2).

5 Implementation

Our proposal for adaptive marking algorithms involves setting the control parameter h, according to (2); this requires knowledge of the target rates of all sources, and of the capacity of the bottleneck link. After computing parameter h, its value needs to be distributed to the adaptive markers at the edge nodes. An approach to implement the above is to use the Bandwidth Broker architecture [9], since Bandwidth Brokers are centralized entities having knowledge of the resources and the users of a particular domain, and already communicate with edge nodes to inform them of admission and policy decisions. The Bandwidth Broker would periodically update parameter h, when the traffic mix changes. All the information required to compute h is available in the SLAs, since they contain the target rate for each source (organization). Note that the computation does not require knowledge of the number of individual TCP connections that each source has, since the fair share takes into account only the target rate of each source; indeed, unfairness that can be caused due to the different number of TCP connections is an issue that adaptive marking successfully addresses.

The cost of the adaptive marking approach is the increased signalling overhead for the communication between the Bandwidth Broker and the edge devices. However, our results in Section 4.2 show that the performance of adaptive marking is not very sensitive to the optimal selection of h. Moreover, in the case of a network provider offering services to organizations, SLAs are not expected to change frequently, hence the update of parameter h does not need to be frequent.

6 Conclusions

The main theme of this work is that simple yet adaptive packet markers, where the marking probability adapts to changes of the traffic mix, can significantly improve fairness. Based on this, we propose and evaluate adaptive marking algorithms in a variety of scenarios containing different causes of unfairness. From the results we can conclude that the concave marking algorithm has the best overall performance, including the case where the target rates differ significantly. The step marking algorithm has good overall performance, for cases where the target rates did not differ significantly, but is more sensitive to the optimal setting of the adaptation parameter. The linear marking function also has similar good performance, and is less sensitive to the optimal setting of the adaptation parameter. Moreover, the concave marking algorithm achieves near optimal performance when the control parameter is set according to its theoretical value; this is not that case with the linear and step algorithms.

Adaptive packet marking algorithms for achieving proportional fairness and for equal sharing over multiple congested links is discussed in [10]. Further work in the direction set in this paper includes investigating the transient performance of the proposed adaptive marking algorithms in dynamic scenarios, where sources enter and leave the network.

References

[1] D. Clark and W. Fang. Explicit allocation of best-effort packet delivery service. *IEEE/ACM Trans. on Networking*, 6(4):362–372, August 1998.

[2] B. Nandy et al. Intelligent traffic conditioners for assured forwarding based differentiated services networks. In *Proc. of IFIP High Performance Networking*, June 2000.

[3] J. Heinanen et al. Assured Forwarding PHB Group . RFC 2597, 1999.

[4] S. Blake et al. An architecture for differentiated services. RFC 2475, December 1998.

[5] S. Sahu et al. On achievable service differentiation with token bucket marking for tcp. In *Proc. of ACM SIGMETRICS'00*, 2000.

[6] Seung-Joon Seok et al. Unfairness of assured service and a rate adaptive marking strategy. In *Proc. of QofIS 2002*, 2002.

[7] K. R. R. Kumar, A. L. Ananda, and L. Jacob. A memory-based approach for a TCP-friendly traffic conditioner in DiffServ networks. In *Proc. of 9th International Conference on Network Protocols*, November 2001.

[8] B. Nandy, N. Seddigh, and P. Pieda. Diffserv's Assured Forwarding PHB: What assurance does the customer have? In *Proc. of NOSSDAV'99*, 1999.

[9] K. Nichols, V. Jacobson, and L. Zang. A two-bit differentiated service architecture for the internet. RFC 2638, July 1999.

[10] V. A. Siris and H. M. Marinakis. Adaptive Packet Marking for Achieving Fairness in DiffServ Networks. Technical Report # 355, FORTH-ICS, March 2005.

[11] N. Seddigh, B. Nandy, and P. Pieda. Bandwidth assurance issues for TCP flows in a differentiated services network. In *Proc. of IEEE GLOBECOM*, 1999.

[12] M. A. El-Gendy K. G. Shin. Equation-based packet marking for assured forwarding services. In *Proc. of IEEE INFOCOM'02*, 2002.

[13] UCB/LBNL/VINT. Network Simulator - ns (version 2). `http://www-mash.cs.berkeley.edu/ns.html`.