ItswTCM: A New Aggregate Marker to Improve Fairness in DiffServ

Hongjun Su
Department of Electrical & Computer Engineering
University of Dayton, Dayton, OH 45469-0226.
suhongju@flyernet.udayton.edu

Mohammed Atiquzzaman
School of Computer Science
University of Oklahoma, Norman, OK 73019-6151.
Tel: (405) 325 8077 Fax: (405) 325 4044
atiq@ou.edu

Abstract—**Recent demand for real time applications has given rise to a need for Quality of Service (QoS) in the Internet. Differentiated Services is one of such efforts currently pursued by IETF. Previous researchers found unfairness in the DiffServ network. To solve the unfairness problem, we propose a new *TSW* based three color marker (*ItswTCM*) which achieves proportional fair sharing of excess bandwidth among aggregates. We have compared the fairness of our proposed *ItswTCM* marker with *srTCM*, *trTCM*, and *tswTCM*. Results show that our proposed marker performs better than the other three schemes for low to middle network provision level ($20\% - 70\%$); we believe this is the region where all well provisioned network will operate. Results also show that our proposed marker is not as sensitive to the number of flows as previous markers. We point out that yellow packets play a significant role in achieving proportional fair sharing of excess bandwidth among aggregates. We conclude that in order to achieve proportional fair sharing of excess bandwidth, it is important to inject the right amount of yellow packets into the network.**

Keywords—**Differentiated Service Network (DiffServ), Fairness, Aggregate Marker, Performance Evaluation, Quality of Service, Traffic Marker, Time Slide Windows (TSW)**

I. Introduction

The Internet, based on the TCP protocol, has succeeded in providing world wide data communication service for the past few decades. However, Internet does not provide any Quality of Service (QoS) guarantee to applications. With increasing emergence of new service types, such as real-time audio/video applications, there is an increasing demand for providing QoS support in the Internet. Differentiated Services (DiffServ) network [1], [2] is currently being developed by IETF to add QoS to Internet without fundamental change to the current Internet protocol. DiffServ relies on using the type-of-service bit (TOS/DS) bit in the IP header to provide coarse-grained QoS to applications. Unlike Integrated Service (IntServ) [3] which has scalability problem in the core network, DiffServ avoids the scalability issue by providing QoS guarantees to aggregates rather than individual connections/flows.

DiffServ provides statistical QoS to a few predefined service classes instead of providing guarantees to individual flows. It provides service differentiation among traffic aggregates over a long time scale. A DiffServ network achieves its service goals by distinguishing between the edge and core network. It pushes all complex tasks, such as administration, control, traffic classification, traffic monitoring, traffic marking etc., to the edge network where per flow based schemes may be used. Traffic passing through the edge network will be classified into different service classes and marked with different drop priorities (such as In/Out packets). Core routers implement active queue management schemes [4], [5], such as RED with In and Out (RIO) [6], and provide service differentiation to the traffic according to preassigned service classes and drop priorities carried in the packet header. RIO-like schemes achieve this objective by dropping low priority packets earlier and with a much higher probability than dropping high priority packets.

Many recent research results, both simulation based [7], [8], [9] and modeling based [10], [11], [12], have shown that the current DiffServ proposal has three types of *fairness* problems: (a) fairness between TCP-friendly traffic and UDP like traffic, (b) fair sharing of bandwidth among flows in an aggregate, (c) fair sharing of excess network bandwidth among aggregates. The fairness between TCP and UDP flows can be alleviated by mapping them to different drop precedences [13], [8], by using separate queues, or by using per-flow marking schemes [8], [14]. The fair sharing of flows in an aggregate can be achieved using per-flow marking [14]. However, *the authors are not aware of any significant effort to solve the problem of fair sharing of excess bandwidth among aggregates*. The *objective* of this paper is to propose techniques to solve the unfair sharing of excess bandwidth among aggregates.

In this paper, we *focus* on the fair sharing of excess bandwidth among aggregates. We first analyze the effect of aggregate based marking schemes on the fair sharing of excess bandwidth. We then *propose* an Improved Time Sliding Window (TSW) based Three Color Marker (*ItswTCM*). Our proposed scheme is *based* on the following principles: (a) Protecting well behaved senders (sending traffic according their service contract), (b) punishing aggressive senders by assigning their traffic to higher dropping precedence having high dropping probability, (c) improving the proportional fair sharing of the excess network bandwidth among aggregates. With proportional fair sharing, we mean that the excess network bandwidth should be shared among aggregates in proportion to their service subscribing rate. This is important for larger service subscribers.

Our work *differs* from previous work in the sense that our new marking scheme improves the proportional fair sharing of the excess network bandwidth among aggregates. The *contributions* of this paper are as follows:

- We *propose* an *ItswTCM* marker which can improve the proportional fair sharing of the excess network bandwidth among aggregates.
- We discuss *design issues* to address the fairness problems in DiffServ networks.
- Using simulations, we show that our proposed marker can alleviate the unfairness problem among aggregates for a large range of network load, especially for middle load ($40\% - 70\%$ provision level).
- We *compare* fairness of our marker with other well known marker schemes. Simulation results shown that our new scheme achieves better fairness for most cases.

The rest of the paper is organized as follows. Section II gives a brief introduction to some commonly used traffic markers. Section III presents our proposed *ItswTCM* marking scheme. In section IV, we present simulation results comparing fairness of some previous proposed marking schemes with our new scheme. Section V discusses design issues of fair aggregate based markers, and the effect of parameter settings on our proposed marker. Section VI concludes this paper.

II. Traffic Marker

In this section, we give a brief introduction to some commonly used traffic markers which are the most important building blocks of a Diffserv edge router. A marker marks the senders's traffic according to its service profiles. Traffic that conforms to the service profile will be marked with low dropping priority and will receive better service, while nonconformant part of the traffic will be marked with high drop priority and receive a best effort service. Based on the mechanism used to check the conformity of the traffic, traffic markers can be classified into two broad categories: *token bucket based markers* and *average rate estimator based markers*.

Token bucket based markers have been used in many studies [7], [15], [16], [17], [12]. These use one or more token buckets to measure the sending traffic. A single token bucket marker is very easy to implement; an arriving packet is marked as *In-profile* if there is token in the bucket, otherwise, it is marked as *Out-profile*. To improve the fair sharing of excess bandwidth between adaptive and non-adaptive traffic inside an Assured Forwarding (AF) class [13], single rate three color marker (*srTCM*) and two rate three color marker (*trTCM*) have been proposed [15], [16]. For the three color marking scheme, two token buckets are used: one is to hold green tokens, the other to hold yellow tokens. When a packet arrives, it is marked as green if both green and yellow tokens are available; it is marked as yellow if only yellow tokens are available, otherwise, it is marked as red [15], [16]. The number of tokens in the bucket is reduced by the number of used tokens. These schemes are easy to implement.

Average rate estimator based markers use average rate estimating algorithm to estimate the rate of individual flows or aggregated flows. The most common rate estimating algorithm used in such markers is the Time Sliding Window (*TSW*) algorithm [6]. To take advantage of three color marking, *TSW* based three color marker (*tswTCM*) has been proposed recently [18]. For *tswTCM*, whenever a packet arrives, the marker calculates the estimated arrival rate. If the estimated arrival rate is less than the Committed Information Rate (*CIR*), arriving packets are marked as green; otherwise, they are marked as green, yellow or red according to a calculated probability as shown in Fig. 1. We will discuss fairness of this marker in the next section.

avg-rate = Estimated Avg Traffic Sending Rate
if (avg-rate <= CIR)
    the packet is marked as green;
else if (avg-rate <= PIR) AND (avg-rate > CIR)
    calculate P0 = $\frac{\text{(avg-rate - CIR)}}{\text{avg-rate}}$
    with probability P0 the packet is marked as yellow;
    with probability (1-P0) the packet is marked as green;
else
    calculate P1 = $\frac{\text{(avg-rate - PIR)}}{\text{avg-rate}}$
    calculate P2 = $\frac{\text{(PIR - CIR)}}{\text{avg-rate}}$
    with probability P1 the packet is marked as red;
    with probability P2 the packet is marked as yellow;
    with probability (1-(P1+P2)) the packet is marked as green;

Fig. 1. Marking algorithm for the tswTCM marker

III. Proposed Marker

In this section, we list the notations used in this paper, analyze the shortcomings of previous marker algorithms, and then present our proposed marking algorithm along with its advantages over previous algorithms.

A. Notations

1. *MSS* = TCP's *Maximum Segment Size*.
2. *RTT* = *Round Trip Time* of a TCP connection.
3. *CIR* = *Committed Information Rate*, is the contracted average sending rate. In *srTCM* and *trTCM*, it is used as green token arriving rate.
4. *PIR* = *Peak Information Rate*, is the maximum contracted sending rate. In *trTCM* scheme, it is used as the yellow token bucket size.
5. *CBS* = *Committed Burst Size*, is the contracted traffic burst size. In *srTCM* and *trTCM*, it is used as the green token bucket size.
6. *EBS* = *Excess Burst Size*, is the excess token bucket (used to hold excess tokens) size in the *srTCM* scheme.
7. *PBS* = *Peak Burst Size*, is the yellow token bucket size in the *trTCM* scheme.

B. Analysis of Unfairness of Previous Algorithms

As thoroughly analyzed in our previous work [12] (and to some extent in [7], [8], [9]), unfairness exists in the sharing of excess network bandwidth among traffic aggregates. It was found that the excess bandwidth in the core network is shared among flows instead of aggregates, and DiffServ tends to favor small service subscribers. There are mainly two reasons for this: One is due to the TCP congestion algorithm, and the other is that the marker only marks the green packets according to their committed information rate (*CIR*), while the marking of yellow and red packets has little to do with their *CIR*. Green packets correspond to the traffic that commit to their service profile, and enjoy high priority in the core network; yellow packets have lower priority; and red packets have the lowest priority. While the service rate is ensured by the green packets, the *excess network bandwidth at the core network is acquired by yellow or red packets*. Marking the yellow and red packets of aggregates without considering their service profiles leads to unfairness among traffic aggregates (i.e., service subscribers).

From the above observatrions, let's look at the marking of yellow and red packets in the *tswTCM* scheme, and discuss its unfairness problem in sharing of excess bandwidth among aggregates. The *tswTCM* algorithm is shown in Fig. 1. For traffic rate less than CIR, all packets are marked as green packets, no packets are marked yellow or red. For traffic rate greater than CIR, portion of the traffic less than CIR will be marked as green, the remaining portion of the traffic will be marked as yellow or red according to a calculated probability as given below.

For $avg\_rate > PIR$, yellow packets marking probability will decrease with an increase in the sending rate, and the red packets marking probability will increase with an increase in the sending rate; this will help to punish the aggressive sender. However, note that, for $CIR < avg\_rate \leq PIR$, the yellow packet marking probability is proportional to the excess bandwidth defined by this aggregate rather than the service rate (CIR) of this aggregate.

The above behavior is improper because *it encourages the aggressive senders in this region*, and small service subscribers or an aggregate with more flows usually benefit. The marking probability of this algorithm is shown in Fig. 2, which plots the marking probability versus avg_rate/CIR, where avg_rate is the average sending rate of an aggregate measured by the *TSW* algorithm. In this figure, the overall drop precedence is defined as: $d_1 * p_{green} + d_2 * p_{yellow} + d_3 * p_{red}$, where, d_1,d_2 and d_3 represent three drop precedence levels, and p_{green}, p_{yellow} and p_{red} represent the marking probabilities of the Green, Yellow and Red packets respectively. The figure shows that for $PIR < 2 * CIR$, marking of packets as yellow is rare.

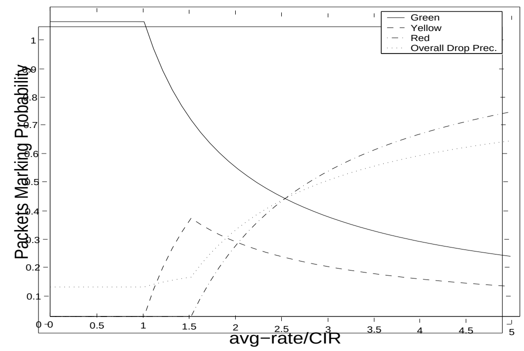

Fig. 2. Marking probability and overall drop precedence of the tswTCM algorithm

C. The Proposed Marker ItswTCM Algorithm

To alleviate unfair sharing of excess bandwidth among aggregates, in this paper, *we propose the Improved TSW based Three Color Marker scheme* (*ItswTCM*) as shown in Fig. 3. The basic idea is to allow an aggregate to inject yellow packets in proportion to their *CIR*. This means a large service subscriber will be able to inject more yellow packets than small service subscribers.

In the proposed scheme, if $avg\_rate < CIR$, all packets are marked as green with no yellow or red packets; this help to guarantee the service rate. When $CIR < avg\_rate < c{*}CIR$, all packets are marked as yellow. When $avg\_rate > c * CIR$, packets are marked as yellow with probability $\left(\frac{c*CIR}{avg\_rate}\right)^2$, and the rest of traffic are marked as red. Note that the probability of marking yellow packets is proportional to their service rate (CIR); this helps to achieve better proportional fair sharing of excess bandwidth. The quadratic term in the probability calculation punishes aggressive senders once their rates exceed $c * CIR$. The marking probability of our proposed scheme is shown in Fig. 4. It is seen that, in contrast to previous algorithms, yellow packets are marked in proportion to their *CIR*.

In our proposed marker, no packet from an aggregate is marked as green when the rate of that aggregate exceeds its CIR. Note that previous marking schemes mark a portion of the excess traffic to green when the traffic rate exceeds its CIR. However, $avg\_rate > CIR$ implies that the source is sending traffic in excess of its service rate and the network has excess bandwidth; it is therefore not very critical whether the throughput is realized by green or yellow packets. Furthermore, our proposed marker has following benefits over previous markers:

- It provides a strong incentive for sources to obey their service contracts. If they send traffic exceeding their service rate, all their traffic will be marked as yellow or red (no more green traffic as in previous algorithms). Senders have to face possible service quality degradation, such as excess traffic marked as yellow packets suffering more delay.
- It will help to protect TCP-friendly traffic from UDP-like traffic. In the proposed scheme, all UDP-like traffic will be marked as yellow if it does not obey its traffic profile. UDP-like traffic may therefore suffer higher drop probability in the case of network congestion. TCP-friendly sources can respond to congestion by reducing their rate and obtain their desired service by getting all their packets marked as green.

In the next section, we will compare our proposed marker with *srTCM*, *trTCM* and *tswTCM* in terms of fairness among aggregates in a DiffServ network.

```
avg-rate = Estimated Avg Traffic Sending Rate
c = a constant (c>1)
if (avg-rate <= CIR)
    the packet is marked as green;
else if (avg-rate <= c*CIR)
    the packet is marked as yellow;
else
    calculate P = (c*CIR / avg-rate)^2
    with probability P the packet is marked as yellow;
    with probability 1-P the packet is marked as red;
```

Fig. 3. Marking algorithm of the proposed ItswTCM marker.

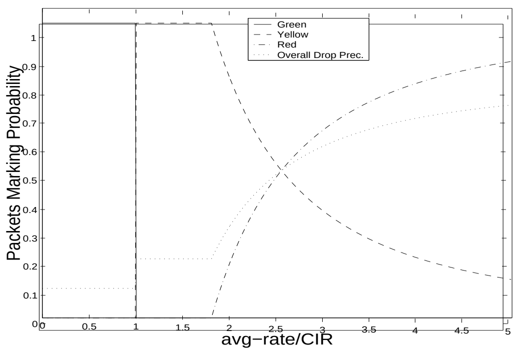

Fig. 4. Marking probability and overall drop precedence of improved tswTCM algorithm

IV. Simulation and Results

In this section, we will compare the fairness of our proposed Improved TSW based Three Color Marker (*ItswTCM*) with three of the previously proposed marking schemes: Single Rate Three Color Marker (*srTCM*) [16], two rate Three Color Marker(*trTCM*) [15], and TSW based Three Color Marker (*tswTCM*) [18]. Simulations have been performed using the ns-2 network simulator [19] together with DiffServ patch from Nortel [20].

Figure 5 shows the simulation topology used in this study. There are two groups of sources called *aggregate 0* and *aggregate 1*. *Aggregate 0* sends traffic through edge router E0 to destination D0, while *aggregate 1* sends traffic through edge router E1 to destination D1. Each aggregate contains a number of flows from different source nodes (S). E0 and E1 are the edge routers responsible for monitoring and marking traffic aggregates 0 and 1 respectively.

C1 is the core router which implements RIO. We use the notation $\{x, y, z\}$ to represent minimum threshold, maximum threshold, and weight parameter respectively of a RED queue [4]. The settings of the core router is $\{0, 40, 0.2\}$, $\{40, 80, 0.1\}$ and $\{80, 120, 0.02\}$ for Red, Yellow and Green packets respectively. We use TCP Reno in all the simulations, and all TCP flows are long lived FTP applications with $MSS = 1500$ bytes and $RTT = 40ms$ for both aggregates.

For *srTCM* and *trTCM*, token arrival rate is set to their *CIR*s respectively. *CBS* is set to 30000 bytes. For *srTCM*, *EBS* is

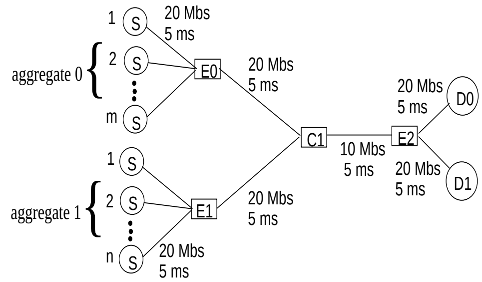

Fig. 5. Simulation topology.

set to $30000 * 120\% = 36000$ bytes. For *trTCM, PBS* is set to 36000 bytes, while *PIR* is set to 1.2 times *CIR*. The *PIR* for *tswTCM* is also set to 1.2 times of the *CIR*. These settings simulate the case where a source has negotiated an average sending rate of *CIR* with its peak sending rate which is at most 120% of their CIR. The factor c for *ItswTCM* is set to two, and the effect of c will be discussed in Section V.

We will study the fairness of four marking schemes in three different cases. In all the cases, the *CIR* for aggregate 1 is fixed at 1 Mbps, while the *CIR* for aggregate 0 increases from 1 Mbps to 8 Mbps. This corresponds to the network going from a provision level of 20% to 90% (we define the *provision level* to be the ratio of total provisioned bandwidth to the bandwidth of the bottleneck link).

In this study, the fairness index (*FI*) [21] is defined as follows:

$$FI = \frac{(\sum_i x_i)^2}{N * \sum_i {x_i}^2} \quad (1)$$

where $0 < FI < 1, x_i = \frac{excess\ bandwidth\ obtained\ by\ aggregate\ i}{CIR\ of aggregate\ i}$, and N is the total number of aggregates under consideration (in our case, $N = 2$). According to this definition, the closer the fairness index is to one, the fairer is the distribution of the excess bandwidth between the aggregates. Note that in this paper, we use proportional fair sharing rather than equal fair sharing.

A. Case 1: Aggregate 0 has 16 flows and aggregate 1 has 16 flows

In this case, aggregate 0 and aggregate 1 each has 16 flows. This scenario illustrates how a small service subscriber and a larger service subscriber, emitting the same number of flows, competes for the excess bandwidth. Figure 6 shows the fairness index versus provision level for the four marking schemes. We see that our *ItswTCM* performs significantly better than others for a large range of provision level, while the *trTCM* and *tswTCM* perform slightly better than *srTCM*. Note that at the high provision level (let's say above 75%), our scheme is worse than *trTCM* and *tswTCM*. The main reason for the improved performance of our marker is that both of the aggregates have only green and yellow packets. Although yellow packets compete for the excess bandwidth, the small service subscriber will grab more of the excess bandwidth due to the fact that TCP

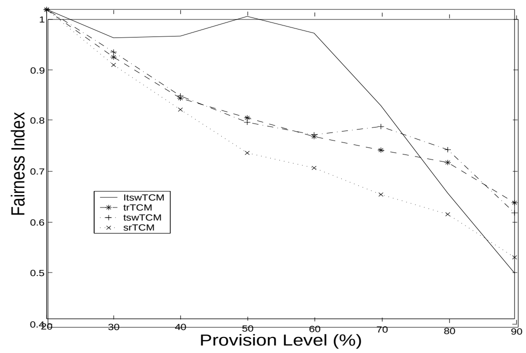

Fig. 6. Fairness Index versus provision level, for aggregate 0(16 flows) and aggregate 1(16 flows) case.

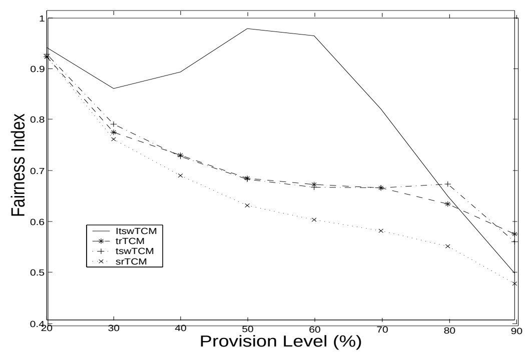

Fig. 7. Fairness Index versus provision level, for aggregate 0(16 flows) and aggregate 1(32 flows) case.

congestion control algorithm favors small flows [12], [7], [8], [9]. This also applies to the two case discussed in Sections IV-B and IV-C.

B. Case 2: Aggregate 0 has 16 flows and aggregate 1 has 32 flows

In this case, aggregate 0 and 1 has 16 and 32 flows respectively. This case illustrates the sharing of excess bandwidth between a small service subscriber emitting large number of flows and a larger service subscriber emitting few flows. Figure 7 shows the fairness index versus provision level for the four marking schemes. We see that our proposed marking scheme achieves better fair share of the excess bandwidth than other schemes for a large range of provision level. The worst one is *srTCM*.

C. Case 3: Aggregate 0 has 32 flows and aggregate 1 has 16 flows

In this case, aggregate 0 has 32 flows, while aggregate 1 has 16 flows. This case corresponds to the sharing of excess bandwidth between a small service subscriber emitting a small number of flows and a large service subscriber emitting a large number of flows. Figure 8 shows the fairness index versus provision level for the four marking schemes. We see that our scheme

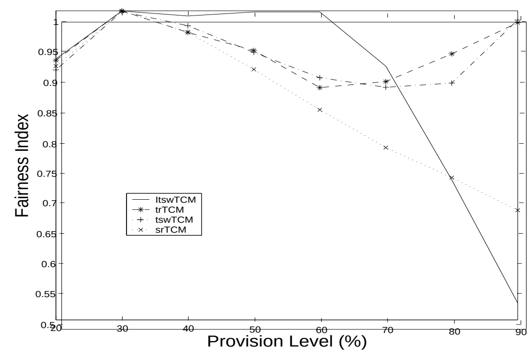

Fig. 8. Fairness Index versus provision level, for aggregate 0(32 flows) and aggregate 1(16 flows) case.

again performs better than other schemes in the $20\% - 70\%$ provision level. *srTCM* is still the worst choice. All schemes performs better in this case than their corresponding ones in the previous two cases given in Sections IV-A and IV-B. This is because the larger profile aggregate 0 has more flows, and therefore, it is easier to achieve the proportional sharing of the excess bandwidth. As in Case 1, previous schemes perform better than our proposed scheme for a high provision level. However, it is expected that a typical network will be provisioned below 80% provision level, in which case, *our scheme is better than all previous schemes in all ranges*.

V. DISCUSSION

From the simulation results, we see that our proposed *ItswTCM* scheme achieves better fairness than other schemes for low to middle level of provision (say from 20% to 75%). By comparing Figs. 6, 7 and 8, we see that the shape of the fairness curve for our proposed scheme does not change much, while the curves change significantly for the other three schemes. This means that our proposed scheme is not as *sensitive* as other schemes to a change in the number of flows per aggregate; its performance is mainly aggregate based.

Figure 9 shows the performance of our proposed *ItswTCM* marker for different values of c. It is seen that our proposed scheme still can not achieve proportional fair sharing of excess bandwidth over the whole region. It achieves better performance in terms of fair sharing in selected regions (depend on value of c). Except for $c = 1$, we see that as c decreases, the fair sharing region moves from low provision level toward high provision level. However, we see that the maximum value of the fairness index in the region decreases as it moves toward high provision level. This is because at a high provision level, the packet drop probability is significantly higher than in the low provision level. Because it costs significant amount of time for the large profile aggregate to recover its sending rate following a packet loss, the effect of dropping a packet is much more serious for large profile aggregates than for small profile aggregates.

It is interesting to have a closer look for $c = 1$. When $c = 1$, the performance of our proposed scheme is similar to the worst

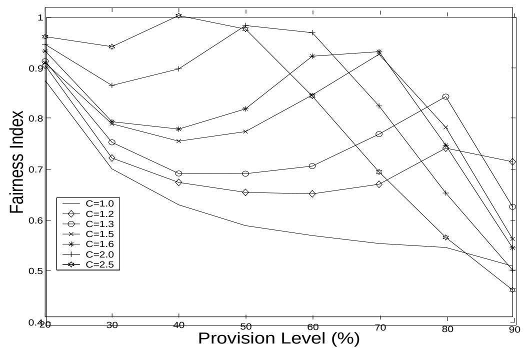

Fig. 9. Effect of factor c to ItswTCM

scheme, *srTCM*. From the marking probability in Fig. 4 and the *ItswTCM* algorithm in Fig. 3, we see that for $c = 1$, the probability of marking a packet as yellow is small, i.e., most of the traffic will be marked as either green or red. The result emphasizes our earlier observation that *yellow packets play a very important role* in realizing proportional fair sharing of excess bandwidth. This also explains why the *tswTCM* scheme performs poorly with respect to fairness. From Fig. 2, we see that for $PIR < 2 * CIR$, marking of packets as yellow is rare. Note that setting PIR to about $2 * CIR$ is not a reasonable service contract because PIR is the peak information rate.

VI. Conclusions

In this paper, we proposed a new aggregate based marker (*ItswTCM*) scheme to improve the fairness in the proportional sharing of excess bandwidth between aggregates in DiffServ network. The results from this work can be summarized as follows:

- We *proposed* a new *TSW* based three color marker, *ItswTCM*. It is an aggregate marker which is intended to achieve proportional fair sharing of the excess bandwidth among aggregates in a DiffServ network.
- We *compared the fairness* of our proposed *ItswTCM* marker with *srTCM*, *trTCM*, and *tswTCM*. Simulation results show that our proposed marker performs better than the other three markers for low to middle network provisioning level ($20\% - 70\%$). We believe that all well provisioned networks will operate in this region.
- Comparing the performance of our proposed marker under different network conditions, we find that the *ItswTCM* marker is *not as sensitive to the number of flows in an aggregate as the other three markers*. Our scheme is mainly aggregate based.
- By studying the effect of parameter settings on the performance of our marker, and the marking probability of our proposed marker, we find that yellow packets play a significant role in achieving proportional fair sharing of excess bandwidth among aggregates. We conclude that in order to achieve proportional fair sharing of excess bandwidth, it is important to *inject the right amount of yellow traffic into the network*.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, “An architecture for differentiated services.” RFC 2475, December 1998.

[2] K. Nichols, V. Jacobson, and L. Zhang, “A two-bit differentiated services architecture for the internet.” RFC 2638, Internet Engineering Task Force, July 1999.

[3] R. Braden, D. Clark, and S. Shenker, “Integrated services in the internet architecture: An overview.” RFC 1633, June 1994.

[4] S. Floyd and V. Jacobson, “Random early detection gateways for congestion avoidance,” *ACM/IEEE Transactions on Networking*, vol. 1, no. 4, pp. 397–413, August 1993.

[5] B. Braden, D. Clark, J. Crowcroft, S. Deering, D. Estrin, S. Floyd, V. Jacobason, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, and L. Zhang, “Recommendations on queue management and congestion avoidance in the internet.” RFC 2309, Internet Engineering Task Force, April 1998.

[6] D. Clark and W. Fang, “Explicit allocation of best effort packet delivery service,” *IEEE/ACM Transactions on Networking*, vol. 6, no. 4, pp. 362–373, August 1998.

[7] J. Ibanez and K. Nichols, “Preliminary simulation evalustion of an assured service.” draft-ibanez-diffserv-assured-eval-00.txt, Internet Draft, August 1998.

[8] B. Nandy, N. Seddigh, P. Pieda, and J. Ethridge, “Intelligent traffic conditioners for assured forwarding based differentiated services networks,” *IFIP High Performance Networking (HPN 2000)*, Paris, France, June 2000.

[9] I. Yeom and A. L. Narasimha Reddy, “Realizing throughput guarantees in a differentiated services network,” *Proceedings of ICMCS*, pp. 372–376, vol. 2, June 7-11, 1999.

[10] I. Yeom and A. L. Narasimha Reddy, “Modeling TCP behavior in a differentiated services network.” Texas A&M Technical Report, May 1999.

[11] M. Baines, B. Nandy, P. Pieda, N. Seddigh, and M. Devetsikiotis, “Using TCP models to understand bandwidth assurance in a differentiated services network.” Nortel Technical Report, July 2000.

[12] H. Su and M. Atiquzzaman, “Performance modeling of differentiated service network with a token bucket marker,” *11th IEEE Workshop on Local and Metropolitan Area Networks*, Boulder CO., USA, pp. 81–82, March 18-21, 2001.

[13] J. Heinanen, F. Blaker, W. Weiss, and J. Wroclawski, “Assured forwording PHB group.” RFC 2597, Internet Engineering Task Force, June 1999.

[14] I. Yeom and A. L. Narasimha Reddy, “Impact of marking strategy on aggregated flows in a differentiated-services network,” *International Workshop on QOS*, Landon, UK, pp. 156–158, May 1999.

[15] J. Heinanen and R. Guerin, “A two rate three color marker.” RFC 2698, September 1999.

[16] J. Heinanen and R. Guerin, “A single rate three color marker.” RFC 2697, September 1999.

[17] S. Sahu, P. Nain, D. Towsley, C. Diot, and V. Firoiu, “On achievable service differentiation with token bucket marking for TCP,” Technical Report 99-72, University of Massachusetts, Dept. of Computer Science, 1999.

[18] W. Fang, N. Seddigh, and B. Nandy, “A time sliding window three colour marker (tswtcm).” RFC 2859, June 2000.

[19] “Network simulator 2 (ns-2).” University of California at Berkeley, CA, Available via http://www.isi.edu/nsnam/ns/, 1997.

[20] “Diffserv model for the ns2 simulator.” Open IP Group at Nortel, available from http://www7.nortel.com:8080/CTL/.

[21] R. Jain, *The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation, and Modeling*, p. 36 & 387. New York, NY: John Wiley and Sons Inc., April 1991.