

Fairness Measurements about TCP Flows in DS Networks: Comparison of Per-flow and Aggregate Marking Schemes

VIRPI LAATU, JARMO HARJU AND PEKKA LOULA

Department of the information technology

Tampere University of Technology

Pohjoisranta 11 P.O Box 300, FIN-28101 Pori

FINLAND

virpi.laatu@tut.fi, jarmo.harju@tut.fi, pekka.loula@tut.fi

Abstract:- Traffic marking is an essential component in any Differentiated Services (DS) architecture. Traffic marking can be done with respect to a single flow or a flow aggregate. In a DS network, traffic marking is typically performed at aggregate levels, rather than on the level of individual flows. The use of aggregation raises a number of concerns, such as fairness concern. When we deal with aggregated sources in such networks, we need to consider not only the fairness issues among aggregates but also the fairness issues among individual flows within an aggregate.

We have carried out measurements in a test network to examine the effect of DS mechanisms on the fairness among individual TCP flows. The aim of the measurements was to evaluate how per-flow marking and aggregate marking affect the performance of individual TCP connections when competing with each other within the same Assured Service class. In order to eliminate the effects of different RTTs (Round Trip Times) and lengths of the flows, we used homogeneous TCP flows, i.e., flows with equal RTT and duration.

Our measurements showed that per-flow marking not only provides much fairer access to the assured bandwidth, but also provides higher throughputs seen by individual flows.

Key-Words: -. QoS and service differentiation, network measurements, network access, traffic control

1 Introduction

The Differentiated Services (DS) architecture [1] is a respectable way to provide preferential treatment to certain packets inside the IP network, but it also brings new problems, especially to TCP users. There are many problems when running TCP over DS networks. Most of DS router functionalities, such as traffic markers and scheduling mechanisms are designed mainly to support service differentiation for aggregated flows and not tailored for TCP traffic, which constitutes the majority of the traffic on the Internet today.

Our study focuses on services built on top of the AF PHB [2]. The fair sharing of the available bandwidth is a major issue in the context of the AF service. Many recent studies have tackled fairness problems in the AF based networks. However, these research efforts have almost exclusively looked at fairness on a per-aggregation basis, not within an aggregate, as addressed in this study. Some studies have dealt with fairness between individual flows within an aggregate, but they applied to an aggregate consisting of TCP and UDP traffic or the traffic marking is performed on a per-flow basis.

We are not aware of any previous study that evaluated the problems of the fair bandwidth allocation among individual TCP flows of an aggregate when using an aggregate marking scheme. Thus, the goal of this study is to investigate how significantly the aggregate marking mechanism affects the bandwidth distribution among the TCP flows within a single AF class. The evaluation is carried out in the context of both individual and aggregate service models. The performance is mainly evaluated in terms of throughput and fairness indices for TCP flows.

For sake of simplicity, we only consider bulk TCP traffic traversing through AF enabled DS domains. However, diversity of traffic can be considered if we use more than one AF class and assign different applications to the different AF classes.

The rest of this paper is organized as follows: Section 2 gives an overview of the current work on the problem of fairness within the context of the AF service. Section 3 describes the components that compose the Assured Service used in this study, i.e., traffic conditioning mechanisms that implement the AF PHB group itself. Section 4 describes the measurement configuration used in the study and

corresponding network parameter settings. In Section 5, we present the performed tests and the results of these tests. Finally, Section 6 summarizes our findings and concludes the paper.

2 Related Work

Providing fair services to users with diverse QoS requirements is one of the main performance concerns in AF based networks. In a DS architecture, traffic marking at edge routers can be done either with respect to a single microflow or an aggregate of flows, while core routers operate on aggregated flows only. However, due to the scalability, traffic marking at an edge router is mainly performed on the aggregated traffic (i.e., service contract (SLA) covers the aggregate rate sent by the customer). Per-aggregation marking is also easier to manage. Furthermore, giving each flow a fraction of the contracted rate of the aggregated traffic can lead to an inefficient utilization of the bandwidth in the presence of idle flows. These idle flows waste their shares, while preventing active flows from using unusable network resources. However, per-aggregation marking can adversely affect fairness within an aggregate and between aggregates that are entering an Internet service provider's (ISP's) network. When using aggregate marking, all the flows of the same aggregate are treated to be a part of a big single traffic source. In this case, a meter determines whether the aggregate is in-profile or not.

The unfairness can be caused by a number of factors. The edge router is the most appropriate place that can tackle the fairness among the competing flows. There are two kinds of unfairness, the intra-class and the inter-class fairness in an edge network. Our study focuses on the intra-class fairness problems. The fairness among competing flows within the same service class can be addressed in both queue management and scheduling schemes.

Several alleviations have been proposed to overcome the unfairness problems in AF based networks. Many sophisticated markers have been proposed to improve the fair share of the available bandwidth between different AF flows. Most of these current aggregate traffic markers mark individual packets of an aggregate using token-bucket based markers [e.g., 3, 4, 5, 6] or using average rate estimating algorithm (e.g., used in [7]). These markers do not look at the fairness issues within an aggregate such as token distribution among individual flows.

Many recent studies have addressed fairness issues when competing flows differ in traffic

characteristics, and concluded that the throughput attained by a customer is affected, not only by marking strategies at edge routers but by the presence of other flows in the same bottleneck link. The unfairness can be caused by different round trip times (RTTs), different target rates, differently sized flows, UDP/TCP interaction, and number of flows in an aggregate. In [8] the authors evaluated bandwidth assurance issues for TCP flows in a RIO (Random Early Detection (RED) with IN and OUT)-like DS network, and found that aforementioned factors can cause different throughput rates for end-users in spite of having identical service contracts. The authors in [9] and [10] observed that marking with the same parameters does not assure fair bandwidth sharing between TCP flows with different RTTs. They showed that flows with longer RTTs receive less than their fair share of the available bandwidth. The same unfairness phenomenon occurs for TCP flows with different target rates [4, 11].

There is also one major unfairness problem concerning the interaction among short-lived and long-lived TCP connections. Due to TCP congestion control mechanisms (slow start) short-lived TCP connections (e.g., Web browsing) have much different characteristics and behavior than bulk FTP transfers. Short-lived TCP flows seldom need to operate beyond the slow start phase to finish their transmission, while long-lived TCP flows mostly operate in congestion avoidance phase producing larger packet burst than generated by short flows. Some studies [12, 13, 14] have investigated unfairness problems between short-lived and long-lived TCP flows. Matta and Guo in [12] showed that fairness could be much improved by isolating TCP flows based on their size. They proposed to put TCP flows into two classes: one of short flows and another of long flows.

One source of unfairness comes from the difference between UDP and TCP flows in the way they experience packet drops. UDP flows usually get more than their fair share of bandwidth leaving TCP flows not to get their fair share. The interaction of TCP and UDP has been studied e.g., in [3, 15, 16].

Most of these aforementioned studies look at fairness issues on a per-aggregation basis, not within an aggregate, as addressed in this study. Moreover, most of these studies (e.g., [15], [17]) deal with per-flow marking, or they deal with only individual sources. Yeom and Reddy in [18] looked at the fairness issues between individual flows within an aggregate, but they assumed that the individual flows

have already been pre-allocated an individual contract rate from the aggregate marker tokens.

One way to alleviate the problem of unfair bandwidth shares between individual flows within an aggregate is in advance divide up the aggregate SLA token distribution into smaller tokens for individual flows at the edge marker. The authors in [19] used this scheme and proposed two new marking algorithms that improve the fairness among the individual flows within an aggregation. The major drawback of their scheme is that the edge marker should have knowledge of individual token bucket specifications for every incoming flow, thus causing the scalability problem.

3 Assured Service Scheme

The AF PHB group as specified in [2] provides the delivery of IP packets in four independent traffic classes (AF classes), each with three level of drop precedence (green, yellow and red). The AF class cannot offer absolute service bounds. It just offers relative service differentiation between classes. Some kind of bandwidth provisioning exists between the AF classes. For each of the AF classes, there is a certain amount of buffer space and bandwidth allocated by each DS node. The unused bandwidth can be configured so that it can be divided between the other AF classes or the other PHB groups. In the Assured Service, each packet has a codepoint encoded in the DS field, which identifies the AF PHB. In all, there are twelve DSCPs for AF PHB group.

The Assured Service relies on packet monitoring and marking mechanisms, performed by the traffic conditioner at the edge nodes of a DS network, and queue management mechanism at the core nodes. An ISP ensures that the aggregate traffic generated complies with the traffic profiles specified in the SLAs between the users and the network. In this study we considered the SLAs that are made on a per-customer (e.g., a small company) basis rather than on a per-connection basis.

In an AF-compliant domain, the edge routers meter and mark packets of flows based on agreed-upon profiles. The traffic meter tracks the rate of the customer's aggregated traffic at the edge of the DS network. Using this rate information, packets of a flow are marked with different colors (two or three). We used a dual token-bucket based mechanism called Two Rate Three Color Marker (trTCM) [5] to check the traffic conformity at the edge routers, and to mark packets in agreement with service profile.

A trTCM measures incoming traffic from the customer and marks the packets based on two rates, Committed Information Rate (CIR) and Peak Information Rate (PIR), and their associated burst sizes, Committed Burst Size (CBS) and Peak Burst Size (PBS), respectively. The CBS is used as the green token bucket size, whereas the PBS is used as the yellow token bucket size. When the customer's measured traffic is within the contracted average sending rate (CIR), the packets are marked as green. The packets that exceed the contracted rate are not discarded immediately, but marked as yellow instead when the rate falls below the PIR. All packets that exceed the PIR are marked as red.

In case of congestion within the AF class, the drop precedence of each packet determines the relative importance of the packets. At the time of congestion, the core router at the DS network uses active queue management (AQM) technique to provide preferential treatment to in-profile packets at the cost of out-profile packets. There are many alternative AQM policies to be used at the core routers, such as RED [20], RIO [21], and Core-Stateless Fair Queuing (CSFQ) [19]. Following the AF specification [2], we chose RIO-like AQM policy, which implements the AF PHB using the three-priority (color) version of the RED. This mechanism is known as GRED (Generalized RED). GRED allows multiple drop precedence levels within an AF class. In our experiments, we used three sets of RED parameters, one for each color. Link sharing between the AF class and best effort (BE) traffic was implemented using a CBQ (Class Based Queuing) mechanism.

4 Measurement Arrangement

The details of the measurements are presented in this section. First, we give a brief review of the goal of the measurements. Then we describe the performance criteria that we have used in Section 5 to compare the fairness of a per-flow marking and an aggregate marking mechanisms. After that, we present the measurement methodologies and the test network used in the study. We also describe the software tools and traffic workloads used in the study and explain how these were used in each experiment.

4.1 Goals of the measurements

The goal of the measurements was to investigate fairness among AF TCP flows originating from a single customer's network. We studied two marking strategies in the edge nodes, a per-flow marking and

an aggregate marking. In the per-flow marking the traffic is classified into TCP flows, and metering and marking takes place individually for these flows, i.e., each flow has its own trTCM. In the aggregate marking only one trTCM is used for all incoming traffic from the customer. The focus of our experiments was to determine how significantly the aggregate marking mechanism affects the bandwidth sharing among TCP flows within a single AF class from an end-user's perspective. For that purpose, we compared the performance of bulk TCP connections sharing one aggregated CIR rate with that of bulk TCP connections, each having its own individual CIR rate.

4.2 Performance metrics

In a DS capable network, service agreements for customer traffic are often specified at aggregate levels (a fixed bandwidth assurance for traffic originating from a single company), rather than on the level of individual flows. However, the performance measures of real interest are usually the level of performance that individual users and applications experience (e.g., per-flow goodput).

In these experiments, we consider the following performance metrics: (1) utilization of the committed rate (CIR) and the excess rate by the customer (individual or aggregated flows), (2) the average goodput¹ and throughput obtained by an FTP flow at the receiver, (3) the fairness achieved in the allocation of committed, excess and total bandwidth among different flows, (4) the throughput achieved by the aggregated source.

For each flow, the number of green colored packets delivered to the corresponding destination is calculated (the green throughput). Utilization of the committed rate by a flow is measured as the ratio of the green throughput of a flow and the expected fair share of committed rate (CIR of aggregate rate shared equally by each flow within an aggregate). The utilization of the excess rate (yellow and red packet rate) by a flow is also measured.

To quantify the level of the fairness, we used the fairness index and the standard deviation. The fairness achieved in the allocation of committed, excess and total achieved bandwidth among flows was evaluated on the basis of throughput measurements. The fairness index f among N flows sharing a link was computed using the following formula (1):

$$f(x_1, x_2, \dots, x_N) = \frac{\left(\sum_{i=1}^N x_i\right)^2}{N \times \sum_{i=1}^N x_i^2} \quad (1)$$

$$x_i = \frac{t_i}{T_i}$$

where N represents the number of flows, x_i , $i=1 \dots N$, is the relative allocation of bandwidth of the i^{th} flow, i.e., t_i is the measured throughput by flow i and T_i is the expected ideal bandwidth for flow i . The fairness index always lies between 0 and 1. The closer the fairness index is to one, the fairer is the distribution of the bandwidth between individual flows.

4.3 Test network

The performance was evaluated using a single-bottleneck link topology illustrated in Fig. 1. The network setup consists of five router elements: three edge routers labeled as E1, E2 and E3, and two core routers labeled as C1 and C2. Other network nodes are acting as sources and sinks of TCP traffic. All edge routers do the metering and packet marking by means of a trTCM mechanism. Edge routers were configured so that they are not congested, i.e., they are responsible for packet marking but not for packet dropping. Both core routers implement the AF PHB using a GRED buffer management mechanism. Our DS testbed was set up using the Linux QoS support (Linux RedHat 8.0 distribution with kernel version 2.4.18).

The scenarios were composed of measured TCP flows with a mixture of TCP and UDP flows as background traffic. The monitored traffic consists of four traffic sources of TCP Reno from nodes $S_0 \dots S_3$ to nodes $D_0 \dots D_3$. Background TCP flows were used to make the AF aggregate contains more traffic in the core than just the single customer's (measured) traffic. They were also used to create bottleneck for the AF traffic on the link between C1 and C2 nodes. The source nodes from $B_0 \dots B_6$ are used to generate background traffic in form of seven AF TCP flows carrying bulk traffic. The background traffic was generated by the Adtech AX/4000 traffic generator and looped back to the same equipment through the bottleneck link. In all experiments, the background TCP flow aggregate has a committed rate (CIR) of 800 kbps and a peak rate (PIR) of 1.4 Mbps. In all experiments, both monitored TCP flows and background TCP flows go to the same AF class. In addition, we added a single BE UDP flow carrying CBR traffic to fill up the rest of the available link bandwidth on the 10 Mbps bottleneck link. During the measurements the background load was always

¹ Goodput measures the rate of successfully transmitted packets.

turned on. The bandwidth of all the links was set to 10 Mbps. All links were bidirectional, but all the traffic sources are transmitting in the same direction (only ACKs on the return path). The bandwidth for AF packets was limited to 2 Mbps.

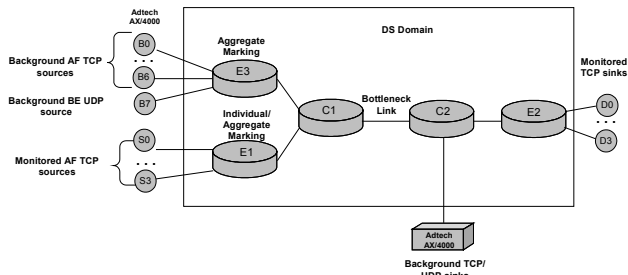


Fig. 1. The test network topology

4.4 Parameter configurations

Three sets of RED threshold parameters are maintained in the core node, one for each color. It is well known that the choice of different RED parameter values may have an important impact on the TCP performance. The general guidelines as given in [21, 22] are followed in setting the GRED parameters. Min and max thresholds were chosen so that they do not overlap. In our experiments, we used two physical queues, one for AF service and another for BE service. A single queue for AF traffic was made up of three virtual queues, one for each precedence level. Different GRED parameters used for the virtual queues are listed in Table 1.

Table 1. GRED parameters used for experiments

Parameters	Green	Yellow	Red
Virtual queue thresholds in kB			
Min _{th}	45	22.5	10.5
Max _{th}	90	45	22.5
Max _p	0.02	0.05	0.10
W _q	0.002	0.002	0.002

For our traffic markers, we used in aggregate and per-flow marking scenarios a CIR of 800 kbps and 200 kbps, respectively with a bucket size of 15 packets (CBS). PIR was set 1.4 Mbps and 350 kbps in aggregate and per-flow marking scenarios, respectively, while PBS was set to 18 packets. During the measurements, the network was provisioned so that the total rate of green colored packets generated by all the traffic sources does not exceed the bandwidth allocated (by CBQ) for AF packets in the bottleneck link.

4.5 Test methodology

Several different scenarios were created in order to evaluate the characteristics of a TCP flow with bulk data contents. The monitored TCP traffic was generated by pttcp [23] utility. The traffic was captured using tcpdump near the receiver and the traces we analyzed off-line with the help of tcptrace utility [24]. To eliminate other factors (e.g., flows with different RTTs and durations) that could cause unfairness problems, we used homogeneous TCP flows in all our experiments. All monitored flows were infinite FTP-like bulk data flow.

We assumed that after ten seconds the background traffic gets stabilized and the monitored sources could start their transmission. The monitored TCP flows were started simultaneously. The tests lasted for 120 sec., and each test was repeated five times to gain statistical confidence in our results. Unless otherwise stated, all presented results are averaged over five runs.

5 Test Cases and Results

Several different scenarios were conducted in order to compare the performance of a number of TCP flows sharing one aggregated CIR rate with that of a number of TCP flows, each with an individual (per-flow) CIR rate.

Before the actual experiments, we tuned both edge and core routers to fulfill the behavior expectations of the AF class, i.e., the green throughput of a customer should equal its committed rate. The GRED parameters and buffer sizes were optimized empirically in order to avoid any kind of unintentional bottleneck effects.

5.1 Each TCP flow has its own target rate

The main objective of this experiment was to see how AF distributes bandwidth among flows when marking is performed on a per-flow basis. In this section, we consider the case where each of the four TCP flows were individually marked with a trTCM, so that all the four monitored AF flows have the same CIR of 200 kbps and PIR of 350 kbps.

The obtained per-flow throughputs of the four TCP flows, along with their ideal fair shares are plotted in Fig. 2. The results show that all the monitored TCP flows are close to achieve their target rates. The results show also that all the flows share approximately equal amount of the available bandwidth.

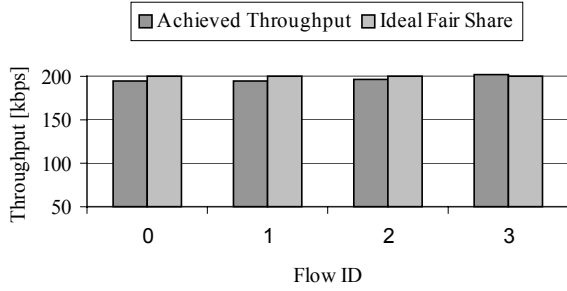


Fig. 2. Achieved per-flow throughputs in per-flow marking case (results gathered from one test run)

The graph in Fig. 3 shows the color distribution for each one of the monitored TCP flows. From the graph, we can see that the numbers of green and excess (yellow/red) colored packets allocated to each flow are shared quite fairly between the four TCP flows.

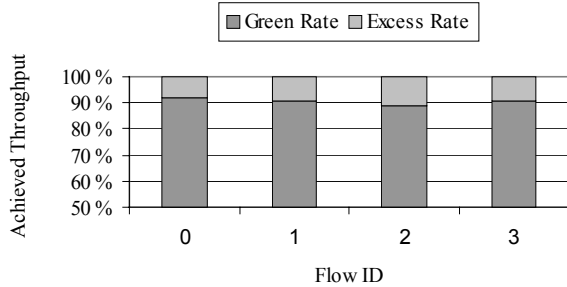


Fig. 3. Bandwidth distribution for green and excess traffic (results gathered from one test run)

Table 2 summarizes the statistical results of this experiment. The first and second rows in the table tell the utilization of the committed rate and the excess rate for monitored TCP flows aggregate. The third and the fourth rows tell the achieved average per-flow goodput and throughput, respectively. The average aggregate throughput is depicted in the last row.

The results show that all the flows are able to close to fully utilize their committed rate, i.e., the average green packet rate is about 90 % of the committed rate. We observed that on average in all the cases each monitored TCP flow is close to achieve its target rate. We also observed that standard deviation remains low (8.6 kbps), which indicates that the expected throughputs for each TCP flow are predictable, thus yielding better predictability of the assured service.

Table 2. Measurement results for per-flow marking case

	Mean	Sdev	Min/Max
Utilization of CIR	0.87	0.02	0.84/0.89
Utilization of excess rate	0.06	0.01	0.05/0.06
Per-flow goodput [kbps]	182.3	8.0	166.9/195.2
Per-flow throughput [kbps]	192.0	8.6	172.9/202.7
Aggr. throughput [kbps]	757.8	33.9	713.1/787.8

Table 3 shows statistical fairness details of the data we gathered in this experiment. The mean and standard deviation of fairness indices for the committed, excess and total bandwidth sharing are indicated in this table. The results show that all the monitored TCP flows share the bandwidth in a relative fair manner in terms of committed and excess bandwidth. Moreover, the overall fairness (total) is close to the ideal fairness and the standard deviation remains low.

Table 3. Fairness indices for per-flow marking case

<i>Committed</i>		<i>Excess</i>		<i>Total</i>	
Mean	Sdev	Mean	Sdev	Mean	Sdev
0.998	0.001	0.979	0.007	0.992	0.001

5.2 TCP flows share the same aggregated target rate

The objective of this experiment was to evaluate how aggregate marking affects the performance of individual TCP flows. In this experiment, we consider the same measurement setup as in the earlier experiment, but in this case the monitored TCP flows share the same service profile. In this case the marking behavior differs from the marking of individual flows. In the per-flow marking, the CIR rate for an individual flow is fixed. In the marking of aggregated flows, however, the CIR rate consumed by an individual flow is not fixed even though the aggregated CIR rate is fixed. In this experiment, the monitored TCP flows share an aggregated contract rate (CIR= 800 kbps). The PIR rate for aggregated TCP flows was set to 1.4 Mbps.

The achieved per-flow throughputs from this experiment, along with their ideal fair shares are shown in Fig. 4. As Fig. 4 shows, the monitored TCP connections do not show fair sharing of the available bandwidth. From this example run, we can see that

Flow 0 receives much less bandwidth compared to the other flows. In this example run, none of the four TCP flows could achieve their target rates either.

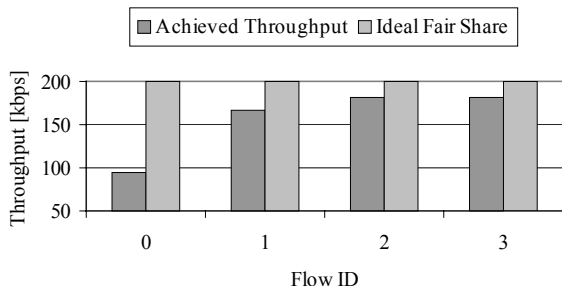


Fig. 4. Achieved per-flow throughputs in aggregate marking case (results gathered from one test run)

The graph in Fig. 5 shows the color distribution for each one of the monitored TCP flows within an aggregate. The results show that the tokens are more unevenly distributed among the flows than in the per-flow marking case.

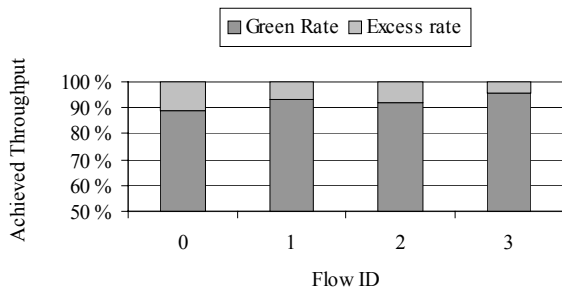


Fig. 5. Bandwidth distribution for green and excess traffic (results gathered from one test run)

Tables 4 and 5 summarize the statistical results of this experiment. We observed that deployment of the aggregate marking causes impairment in most of our metrics. Specifically, we observed degraded per-flow throughputs and predictability of service (as indicated by a higher standard deviation) when compared to the per-flow marking case. As seen from the Table 4, none of the monitored TCP flows could achieve their target rates. This is not the case when using per-flow marking, where almost all the monitored TCP flows could achieve their target rates. Moreover, when using aggregate marking, the aggregate marking can only get a total bandwidth of 616.1 kbps as compared to 757.8 kbps got by the individual marking. The utilization of committed rate is also degraded when flows are aggregated (71 % as compared to 87 %).

Table 4. Measurement results for aggregate marking case

	Mean	Sdev	Min/Max
Utilization of CIR	0.71	0.04	0.67/0.77
Utilization of excess rate	0.05	0.01	0.04/0.05
Per-flow goodput [kbps]	147.6	33.8	78.0/182.6
Per-flow throughput [kbps]	154.0	34.5	83.4/187.1
Aggr. throughput [kbps]	616.1	22.3	586.5/639.4

The results in Table 5 show that the overall fairness decreases in both fairness index and standard deviation when using aggregate marking. The committed, excess and total achieved bandwidths are more unfairly shared between the flows than in the per-flow marking case. For example, for the committed bandwidth the aggregate marking can only achieve a fairness index of 0.957 compared to 0.992 achieved by per-flow marking.

Table 5. Fairness indices for aggregate marking case

<i>Committed</i>		<i>Excess</i>		<i>Total</i>	
Mean	Sdev	Mean	Sdev	Mean	Sdev
0.953	0.027	0.962	0.024	0.957	0.026

6 Conclusions

In this study, we examined the effects that the Assured Services mechanism has on the behavior of TCP flows. We studied how per-flow and aggregated markers affect the bandwidth sharing for individual TCP flows forwarded in a single AF class. The behavior was mainly evaluated by observing the distribution of bandwidth among individual TCP flows generated by the FTP application. In this study, we compared the fairness in bandwidth allocation among homogenous TCP flows (i.e., flows with equal RTT and equal duration) of an AF aggregate when marking is performed either on the aggregated traffic or carried out on a per-flow basis.

The results showed that there is significant variation in the performance seen by individual end-users when using the aggregate marking scheme. Our measurements showed that per-flow marking not only provides much fairer access to the assured bandwidth, but also provides higher throughputs seen by individual flows. Moreover, standard deviation in per-flow marking case remains smaller, which indicates that the expected throughputs are

more predictable than with aggregate marking. While per-flow marking provides better fairness and higher utilization of the network bandwidth, when compared with aggregate marking, these benefits come with the increased cost of more elaborate configuration of the edge nodes.

References:

- [1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, An Architecture for Differentiated Services, *RFC2475*, December 1998.
- [2] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, Assured Forwarding PHB Group, *RFC2597*, June 1999.
- [3] M. Goyal, A. Dusseri, R. Jain, C. Liu, Effect of Number of Drop Precedences in Assured Forwarding, *Proceedings of the IEEE Globecom'99*, Vol.1(A), December 1999, pp. 188-193.
- [4] H. Su, M. Atiquzzaman, ItswTCM: A New Aggregate Marker to Improve Fairness in DiffServ, *Computer Communications Journal*, Vol.26, 2003, pp. 1018-1027.
- [5] J. Heinanen, R. Guerin, A Two Rate Three Color Marker, *RFC2698*, September 1999.
- [6] J. Heinanen, R. Guerin, A Single Rate Three Color Marker, *RFC2697*, September 1999.
- [7] W. Fang, N. Seddigh, B. Nandy, A Time Sliding Window Three Colour Marker (tswTCM), *RFC2859*, June 2000.
- [8] N. Seddigh, B. Nandy, P. Piedad, Bandwidth Assurance Issues for TCP Flows in a Differentiated Services Network, *Proceedings of the IEEE Globecom'99*, December 1999.
- [9] A. Sang, H. Zhu, S. Li, Weighted Fairness Guarantee for Scalable Differentiated Services (DiffServ), *Proceedings of the IEEE ICC'01*, June 2001.
- [10] O. Medina, J. Orozco, D. Ross, Bandwidth Sharing under the Assured Forwarding PHB, *Minutes of the 5th TF-NGN meeting*, irisa.fr/bibli/public/pi/2001/1478/1478.html, October 2001.
- [11] B. Nandy, N. Seddigh, P. Piedad, J. Ethridge, Intelligent Traffic Conditioners for Assured Forwarding Based Differentiated Services Networks, *Proceedings of the IFIP High Performance Networking*, June 2000.
- [12] I. Matta, L. Guo, Differentiated Predictive Fair Service for TCP Flows, *Proceedings of the 8th IEEE ICNP'00*, October 2000.
- [13] K. Tokuda, G. Hasegawa, M. Murata, Analysis and Improvement of the Fairness between Long-lived and Short-lived TCP Connections, *Proceedings of the 7th PfHSN'02*, April 2002.
- [14] M. Mellia, I. Stoica, H. Zhang, TCP-aware packet marking with DiffServ support, *Computer Networks Journal*, Vol.42, 2003, pp. 81-100.
- [15] M.A. El-Gendy, K.G. Shin, Assured forwarding fairness using equation-based packet marking and packet separation, *Computer Networks Journal*, Vol.41, No.3, 2003, pp. 435-450.
- [16] N. Seddigh, B. Nandy, P. Piedad, Study of TCP and UDP Interaction for the AF PHB, *Internet Draft*, Expired February 2000.
- [17] S. Yi, X. Deng, G. Kesidas, C.R. Das, Providing Fairness in DiffServ Architecture, *Proceedings of the IEEE Globecom'02*, November 2002.
- [18] I. Yeom, A. L. N. Reddy, Adaptive Marking for Aggregated Flows, *Proceedings of the IEEE Globecom'01*, November 2001.
- [19] A. Das, D. Dutta, A. Helmy, Fair Stateless Aggregate Traffic Marking using Active Queue Management Techniques, *Proceedings of the 5th IFIP/IEEE MMNS'02*, October 2002, pp. 211-223.
- [20] S. Floyd, V. Jacobson, Random Early Detection Gateway for Congestion Avoidance, *IEEE/ACM Transactions on Networking*, Vol.1, No.4, 1993, pp. 397-413.
- [21] D. David, C. Fang, W. Fang, Explicit Allocation of Best Effort Packet Delivery Service, *IEEE/ACM Transactions on Networking*, Vol.6, No.4, August 1998, pp. 362-373.
- [22] V. Firoui, M. Borden, A study of active queue management for congestion control, *Proceedings of the IEEE Infocom'00*, March 2000.
- [23] I. Pratt, Pratt's Test TCP (pttcp), <http://www.cl.cam.ac.uk/Research/SRG/netos/netx>
- [24] S. Ostermann, TCP dump file analysis tool (tcptrace), <http://www.tcptrace.org>