Fairness Comparisons of Per-flow and Aggregate Marking Schemes in DiffServ Networks

VIRPI LAATU[1], JARMO HARJU[2] AND PEKKA LOULA[3]

[1]Tampere University of Technology, P.O. Box 300, FIN-28101, Pori, Finland, email virpi.laatu@tut.fi

[2]Tampere University of Technology, P.O. Box 553, FIN-33101, Tampere, Finland, email jarmo.harju@tut.fi

[3]Tampere University of Technology, P.O. Box 300, FIN-28101, Pori, Finland, email pekka.loula@tut.fi

Abstract--***This paper examines the fairness issues among TCP flows originated from a single customer network in a Differentiated Services (DS) network. In a DS network, performance commitments for customer traffic are often specified at aggregate levels, rather than on the level of individual flows. 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. Target 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 Services class. Since aggregation is a key factor in any DS capable network to achieve better link utilization, it is essential that an improvement of QoS should be seen also by individual flows.***

I. INTRODUCTION

It is widely accepted that traditional best effort service model does not meet the delivery requirements for new real time and performance-critical services. Thus a wide range of QoS (Quality of Service) support is needed to meet application requirements. QoS measures are necessary not just for real time services but also for the transfer of documents of various data types that may have certain bandwidth demands.

The Differentiated Services (DS) architecture [1] is a respectable approach for providing preferential treatment to certain packets inside the IP network. Although not yet widely deployed, DS architecture will be the basis for future tiered service levels. In Differentiated Services, service discrimination is based on the value of the DS codepoint field (DSCP). DSCP mark indicates that a packet should receive a particular forwarding treatment (Per-Hop Behavior, PHB) at the network nodes. PHBs are implemented in nodes by means of some buffer management and packet scheduling mechanisms. After the packets have been marked, policed and shaped according to a service level agreement (SLA) at the boundary of the network, they are forwarded through the core network according to the PHB associated with the DSCP value. [2]

Providing Assured Services for TCP flows has been an active research issue. There have been a number of studies [e.g. 3, 4, 5, 6, 7, 8, 9] addressing fairness issues among flows with different characteristics. Most of these studies have concluded that the throughput attained by a customer is affected, not only by the marking strategies at the edge router, but by the presence of other flows in the same bottleneck link. The authors in [9] pointed out that UDP flows should be penalised at the network congestion, mapping UDP flows to higher drop precedence level of the same AF class than TCP flows. Seddigh et al. in [8] propose the use of two separate AF class queues, one for TCP and one for UDP flows. In [3] the authors studied the impact of five different factors on offering predictable bandwidth assurance services on customer: round trip time (RTT), target rate, TCP/UDP interaction, number of flows in an aggregate, and packet size.

However, most of these above studies suffer from the same problem. They look at fairness issues on a per-aggregation basis, not within an aggregate, as addressed in this study. In addition, many studies (e.g. [10], [11]) deal with per-flow marking, or the studies dealt with only individual sources. Yeom and Reddy in [7] look at the fairness issues, but they assume that the individual flows within the aggregate have already been pre-allocated an individual contract rate from the aggregate marker tokens. Usually, in a differentiated services network, for scalability reasons service contracts (SLAs) will not typically be on a per end-user (per-connection) basis. One common scenario will be the case where a company contracts a target rate with an ISP, i.e. the agreements cover the aggregate rate sent by the company, so that at any one time a large number of source flows originating from the company would then compete for the aggregate target rate. Many current aggregate traffic markers mark individual packets of an aggregate using token-bucket based markers [e.g. 5, 9, 12, 13]. These markers do not look at fairness issues within an aggregate such as token distribution among individual flows. One way of doing this is in advance divide up the aggregate SLA token distribution into smaller token traffic specifications for individual flows at the edge marker, handled by policy-driven distribution of tokens (e.g. as proposed in [14]). However, the edge marker should have knowledge of individual token bucket specifications for every incoming flows. This will cause scalability problem.

There are various issues to be addressed before a reasonable level of service differentiation between TCP users can be offered. We focused our research on services built on top of the AF PHB. We investigated the unfairness problem between individual TCP flows within an aggregation. In this paper, 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 in assigning applications to the different AF classes.

The rest of this paper is organized as follows: Section 2 describes briefly the Assured Forwarding scheme. In Section 3, we describe the measurement scenario used in the study, i.e. what things were measured, how, and in what kind of conditions. In Section 4, we present the performed tests and the results of these tests. Finally, conclusions follow in Section 5.

II. ASSURED SERVICE SCHEME

In this section we describe the components that compose the Assured Service used in this study, i.e. traffic conditioning mechanisms, such as a policer, a traffic meter and a packet marker, and an active queue management technique that implements the AF PHB group itself.

The AF [15] defined by IETF provides the delivery of IP packets in four independent traffic classes (AF classes), each with three level of drop precedence (DP0-green, DP1-yellow and DP2-red). IP packets assigned to different AF classes are forwarded independent of each other. 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 node, and queue management mechanism at the core nodes. 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 routers at the edge of the network 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 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) [13] 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, CBS (Committed Burst Size) and PBS (Peak Burst Size) respectively. CBR is used as the green token bucket size, whereas 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. When traffic exceeds its CIR, but falls below its maximum contracted sending rate (PIR), packets are marked as yellow. A packet is marked as red if it exceeds the PIR.

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 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 to give, such as Random Early Detection (RED) [16], RIO (RED with IN and OUT) [17], and Core-stateless Fair Queuing (CSFQ) [14]. Following the AF specification [16], 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 (Generalised 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.

III. MEASUREMENT ARRANGEMENT

The details of the measurements are presented in this section. First, we give a brief review of the goal of the measurements. After that, we present the measurement methodologies and the test network used in the measurements. We also describe the software tools and traffic workloads used in the study and explain how these were used in each experiment.

A. Goals of the Measurements

The goal of the measurements was to investigate the fairness issues among TCP flows originating from a single customer's network and using the Assured Service. We studied two marking strategies in the edge nodes, per-flow marking and 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 four bulk TCP connections sharing one aggregated CIR rate with that of four bulk TCP connections, each having its own individual CIR rate.

B. 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. the per-flow goodput).

In these experiments, we consider the following performance metrics: (1) utilization of the committed rate (CIR) by the customer (individual or aggregated flows), (2) the average goodput[1] and throughput obtained by a 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. We call this the green throughput. Utilization of the committed rate by a flow is measured as the ratio of 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 by a flow is also measured. The excess throughputs of the flows are determined by the number of yellow and red packets received at the destination(s).

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 among N flows sharing a link was computed using the following formula (1) [18]:

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

where N is the number of flows and x_i, i=1,..N is the bandwidth obtained by flow i. The fairness index always lies between 0 and 1. If fairness is perfectly achieved, the fairness index should be equal to one.

C. 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 Edge1, Edge2 and Edge3, and two core routers labeled as Core1 and Core2. 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 marking but not for dropping. Both core routers implement the AF PHB using a non-overlapping GRED buffer management mechanism. All the network nodes were Pentium III PCs running Linux RedHat 8.0 distribution, including kernel version 2.4.18 (kernel distribution has built-in diffserv support).

The background traffic used during the measurements was generated by the Adtech AX/4000 (equipped with two 10/100 Ethernet interfaces) traffic generator and looped back to the same equipment through the bottleneck link. The monitored traffic consisted of four FTP flows originating from TCP traffic sources connected to Edge 1, crossing the DS domain and destined to TCP sink nodes that are connected to Edge2, as indicated by the arrows in Fig. 1. Each monitored TCP flow uses the same AF class. The bandwidth of all the links was set to 10 Mbps. All links were bi-directional, but the traffic is unidirectional with only ACKs on the return path.

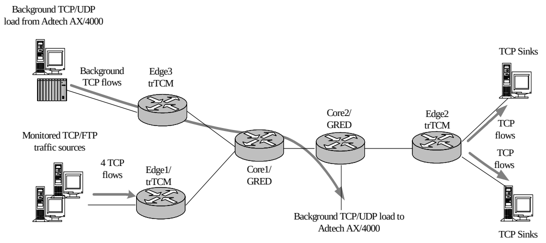

Figure 1. The test network topology

D. Parameters 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 performance. Min and max thresholds were chosen so that they do not overlap. The queue weight[2] used to calculate RED average queue length was always set to 0.002. The parameters setting for GRED are listed in Table 1. These settings have been validated carefully with the help from the studies in [16, 19, 20].

---

[1] Goodput measures the rate of successfully transmitted packets.

[2] An exponentially weighted moving average (EWMA) constant was used for calculating average queue size.

Table 1. GRED parameters settings used for experiments

	Min_th [pkts]	*Max_th [pkts]*	*P_max*	Q_w
DP0 (green)	30	60	0.02	0.002
DP1 (yellow)	15	30	0.05	0.002
DP2 (red)	7	15	0.1	0.002

The amount of buffer in bytes allocated to all queues (three virtual AF queues and the BE traffic queue) was set to 120 kB. 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. In the underprovisioned network the sum of CIRs exceeds the bandwidth allocated (by CBQ) for AF packets in the bottleneck link.

E. Test Methodology

Several different scenarios were created in order to evaluate the characteristics of a TCP flow with bulk data contents. TCP traffic injected into the network were generated by a publicly available benchmarking tool, TCP sender/receiver pttcp [21] utility. The traffic was captured using tcpdump near the receiver and the traces we analyzed off-line with the help of tcptrace utility [22]. The TCP version used in our measurements was New Reno with SACK option. The segment size was fixed at 1460 bytes (not headers included) and the TCP maximum sending window size was set to 64kB. We used FTP bulk data flows with infinite file transfers for the monitored TCP traffic in all our experiments.

The scenarios were composed of four measured TCP flows with a mixture of TCP and UDP flows as background traffic. There were several reasons for using background traffic. Background TCP flows were used to make the AF aggregate contain 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 Core1 and Core2 nodes. The total sum of green colored background TCP packets was limited to 800 kbps. PIR for aggregated background TCP flows was set to 1.4 Mbps. In addition, we added unbounded best effort UDP traffic to fill up the rest of the available link bandwidth on the 10 Mbps bottleneck link to see how it is competing with assured traffic. During the measurements (except the parameters validating process) the background load was always turned on.

For simplicity, we used one AF class for all traffic generated during the experiments. However, more than one class can be considered, allowing more flexibility in assigning applications to the different AF classes.

Each of the tests was carried out five times with the same parameters to gain confidence in our results. We assumed that after ten seconds the background traffic get stabilized and the monitored sources could start their transmission. The monitored four TCP flows were started simultaneously. All tests lasted for 120 seconds.

IV. TEST CASES AND RESULTS

In this section, we present the measurements that we performed on our test network, and discuss the results obtained from the experiments. Several different scenarios were conducted in order to compare the performance of four flows sharing one aggregated CIR rate with that of four flows, each with an individual (per-flow) CIR rate.

Before the actual experiments, we tuned both edge and core routers to fulfil 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. The parameters were configured so that green packets from Edge1 enter the buffer of Core1 node at the rate of 800 kbps.

The measurements were made both in the under- and overprovisioned network conditions in order to study the effect that the AF bandwidth available in core router had on the performance of FTP flows.

A. 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 flow was individually marked with a trTCM, so that all the monitored AF flows have the same CIR of 200 kbps and PIR of 350 kbps.

Two sets of experiments were conducted varying the total amount of offered AF bandwidth in core router. Table 2 and 3 summarize the results of the measurements in both the under- and the overprovisioned network respectively. The first and second rows in Table 2 and 3 tell the utilization of the committed rate and excess rate for individual flows. The third row tells the achieved average goodput, and the achieved average throughputs for individual flows are depicted in the fourth row. The fairness indexes for the committed bandwidth sharing (green packets), the excess (yellow and red packets) bandwidth sharing and the total bandwidth sharing are indicated in the next three rows. The last row tells the average aggregated throughput. The results presented in tables have been collected from the results of one test run.

The results show that in both the under- and the overprovisioned case, all the monitored TCP flows share the bandwidth in a relative fair manner in terms of committed, excess and total achieved bandwidth. Moreover, all the flows are able to close to fully utilize their committed rate in both cases. Only the utilization of excess bandwidth is better in the overprovisioned case.

Table 2. Measurement results for per-flow marking case in underprovisioned network

	Monitored TCP flows			
	Flow1	Flow2	Flow3	Flow4
Utilization of committed rate	0.94	0.97	0.91	0.96
Utilization of excess rate	0.09	0.12	0.09	0.09
Goodput [kbps]	209.0	222.7	200.4	209.3
Throughput [kbps]	220.8	236.7	214.6	223.9
$\text{Fairness}_{\text{committed}}$	0.999			
$\text{Fairness}_{\text{excess}}$	0.978			
$\text{Fairness}_{\text{total}}$	0.999			
Avg. aggregate throughput [kbps]	896.0			

Table 3. Measurement results for per-flow marking case in overprovisioned network

	Monitored TCP flows			
	Flow1	Flow2	Flow3	Flow4
Utilization of committed rate	0.98	0.93	0.97	0.96
Utilization of excess rate	0.14	0.13	0.14	0.14
Goodput [kbps]	229.4	215.0	225.1	229.2
Throughput [kbps]	248.7	229.5	245.2	244.6
$\text{Fairness}_{\text{committed}}$	0.999			
$\text{Fairness}_{\text{excess}}$	0.999			
$\text{Fairness}_{\text{total}}$	0.993			
Avg. aggregate throughput [kbps]	968.0			

B. 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 four flows share the same service profile (CIR value of 800 kbps). The four monitored TCP flows share an aggregated contract rate (CIR= 800 kbps) metered in Edge1 node, whereas the background TCP flows share the same aggregated rate of 800 kbps, metered in Edge3 node. PIR for aggregated flows was set to 1.4 Mbps.

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.

Table 4 summarizes the results of the measurements in the underprovisioned case. The results show that there is a wide variation in fairness in allocation of committed bandwidth with the best result being 1.17, i.e. Flow 4 obtains more than its fair share of the committed bandwidth. The results also show that the overall network utilization for monitored TCP flows (aggregated throughput) with aggregate marking is smaller than with per-flow marking.

Table 4. Measurement results for aggregate marking case in underprovisioned network

	Monitored TCP flows			
	Flow1	Flow2	Flow3	Flow4
Utilization of committed rate	0.70	0.64	0.63	1.17
Utilization of excess rate	0.05	0.06	0.05	0.08
Goodput [kbps]	157.0	146.2	144.0	255.3
Throughput [kbps]	161.7	155.4	147.7	267.7
$\text{Fairness}_{\text{committed}}$	0.926			
$\text{Fairness}_{\text{excess}}$	0.961			
$\text{Fairness}_{\text{total}}$	0.933			
Avg. aggregate throughput [kbps]	732.5			

Table 5 summarizes the results of the measurements in the overprovisioned case. The results show that the committed rates achieved by individual TCP flows increases slightly compared to the underprovisioned case. The results show that there is almost even distribution of the excess bandwidth between all of four competing TCP flows. Instead, the distributions of committed and total achieved bandwidth in the overprovisioned network are quite close to the results achieved in the underprovisioned case.

Table 5. Measurement results for aggregate marking case in overprovisioned network

	Monitored TCP flows			
	Flow1	Flow2	Flow3	Flow4
Utilization of committed rate	0.49	1.07	0.72	0.98
Utilization of excess rate	0.11	0.13	0.13	0.12
Goodput [kbps]	124.9	249.6	176.9	224.3
Throughput [kbps]	134.3	264.1	189.0	240.5
Fairness$_{committed}$	0.927			
Fairness$_{excess}$	0.993			
Fairness$_{total}$	0.944			
Avg. aggregate throughput [kbps]	827.9			

V. 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.

The main benefit of using an experimental testbed instead of a simulated one is that we can evaluate actual implementations of algorithms, and avoid any assumptions or simplifications on network behavior. In this study we compared the fairness in bandwidth allocation among homogenous TCP flows 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 aggregate marking scheme.

Since aggregation is a key factor in any Diffserv network, an improvement of QoS should be seen by individual flows. Also a fair distribution of committed bandwidth is one of the main concerns. A relevant and important issue is how an ISP can configure and provision its network. There is a clear lack of understanding of how the several marking parameter, such as buffer space and burst size provisioning influence the achieved rate.

REFERENCES

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss. An Architecture for Differentiated Services, *Network Working Group, RFC2475,* December 1998.

[2] K. Nichols, S. Blake, F. Baker, D. Black, Definition of the Differentiated Services Field (DS Filed) in the IPv4 and IPv6 Headers, *Network Working Group, RFC2474,* December 1998.

[3] N. Seddigh, B. Nandy, P. Pieda, Bandwidth Assurance Issues for TCP Flows in a Differentiated Services Network, *Proceedings of IEEE Globecom'99*, Rio De Janeiro, December 1999.

[4] I. B. H. de Alves, J. F. de Rezende, L. M. de Moraes, Evaluating Fairness in Aggregated Traffic Marking, *Proceedings of the IEEE Globecom'00,* November 2000.

[5] B. Nandy, J. Ethridge, A. Lakas, A. Chapman, Aggregate Flow Control: Improving Assurances for Differentiated Services Network, *Proceedings of IEEE Infocom'01,* Alaska, USA, April 2001.

[6] Y. Chait, C. V. Hollot, V. Misra, D. Towsley, H. Zhang, J. C. S. Lui, Providing Throughput Differentiation for TCP Flows Using Adaptive Two-Color Marking and Two-Level AQM, *Proceedings of IEEE Infocom'02,* New York, June 2002.

[7] I. Yeom, A. L. N. Reddy, Adaptive Marking for Aggregated Flows. *Proceedings of IEEE Globecom'01,* San Antonio, Texas, November 2001.

[8] N. Seddigh, B. Nandy, P. Pieda, Study of TCP and UDP Interaction for the AF PHB, *Internet Draft,* Expired February 2000.

[9] M. Goyal, A. Dusseri, R. Jain, C. Liu, Effect of Number of Drop Precedences in Assured Forwarding, *Proceedings of IEEE Globecom'99,* Vol. 1(A), pp. 188-193, Rio De Janeiro, December 1999.

[10] J. Rezehde, Assured Service Evaluation, *Proceedings of IEEE Globecom'99,* Rio De Janeiro, December 1999.

[11] H. Shimonishi, I. Maki, T. Murase, M. Murata, Dynamic Fair Bandwidth Allocation for Diffserv Classes, *Proceedings of International Conference on Communications (ICC'02),* New York, April 2002.

[12] H. Su, M. Atiquzzaman, ItswTCM: A New Aggregate Marker to Improve Fairness in DiffServ, *Proceedings of IEEE Globecom'01,* San Antonio, Texas, November 2001.

[13] J. Heinanen, R. Guerin, A Two Rate Three Color Marker, *Network Working Group, RFC2698,* September 1999.

[14] A. Das, D. Dutta, A. Helmy, Fair Stateless Aggregate Traffic Marking using Active Queue Management Techniques, *Proceedings of 5th IFIP/IEEE International Conference on Management of Multimedia Networks and Services (MMNS'02),* pp. 211-223, Santa Barbara, USA, October 2002.

[15] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, Assured Forwarding PHB Group, *Network Working Group, RFC2597,* June 1999.

[16] S. Floyd, V. Jacobson, Random Early Detection Gateway for Congestion Avoidance, *IEEE/ACM Transactions on Networking, Vol. 1, No. 4,* pp. 397-413, August 1993.

[17] D. David, C. Fang, W. Fang, Explicit Allocation of Best Effort Packet Delivery Service, *IEEE/ACM Transactions on Networking, Vol. 6, No. 4,* pp. 397-413, August 1998.

[18] R. Jain, The Art of Computer Systems Performance Analysis: Techniques for Experimental Design, Measurement, Simulation and Modelling, *John Wiley and Sons, Inc.,* 1991, pp. 36-37.

[19] W. Feng, D. Kandlur, D. Saha, K. Shin, A self-configuring RED gateway, *Proceedings of ACM SIGCOMM'02*, Stockholm, August 2000.

[20] V. Firoui, M. Borden, A study of active queue management for congestion control, *Proceedings of IEEE Infocom'00*, Tel-Aviv, Israel, March 2000.

[21] I. Pratt, Pratt's Test TCP (pttcp), http:/www.cl.cam.ac.uk/Research/SRG/netos/netx

[22] S. Ostermann, TCP dump file analysis tool (tcptrace), http://www.tcptrace.org