EFFECT OF NUMBER OF DROP PRECEDENCES IN ASSURED FORWARDING

Mukul Goyal, Arian Durresi, Padmini Misra, Chunlei Liu, Raj Jain

The Ohio State University

New Address: Raj Jain, Washington University in Saint Louis,
jain@cse.wustl.edu, http://www.cse.wustl.edu/~jain

Abstract

This paper explores the issue of fair allocation of excess network bandwidth between congestion sensitive and insensitive flows in an Assured Forwarding traffic class. In the absence of any mechanism to distinguish between out-of-profile traffic of congestion sensitive and insensitive flows, congestion insensitive flows will get most of the excess network bandwidth. However, if out-of-profile packets of congestion sensitive and insensitive flows are 'colored' differently, network can be configured so as to give better treatment to excess packets of congestion sensitive flows and achieve fair allocation of excess network bandwidth. With a view to clearly distinguish between out-of-profile packets of congestion sensitive and insensitive flows, three levels of drop precedence are required. However, if the network operates close to its capacity, three levels of drop precedence are redundant as there is not much excess bandwidth to be shared.

1 Introduction

Differentiated Services (DS) aims to provide *scalable service differentiation in the Internet* that can be used to permit *differentiated pricing of Internet service* [1]. The service to be received by a traffic is marked as a code point in the DS field in the IPv4 or IPv6 header. The DS code point in the header of an IP packet is used to determine the Per-Hop Behavior (PHB), i.e. the forwarding treatment it will receive at a network node. Currently, formal specification is available for two PHBs - Assured Forwarding [2] and Expedited Forwarding [3]. In Expedited Forwarding, a transit node uses policing and shaping mechanisms to ensure that the maximum arrival rate of a traffic aggreagate is less than its minimum departure rate. In Assured Forwarding (AF), IP packets are classified as belonging to one of four traffic classes. Within a traffic class, a packet is assigned one of three levels of drop precedence (green,yellow,red). In case of congestion, an AF-compliant DS node drops low precedence (red) packets in preference to higher precedence (green, yellow) packets. Multiple levels of drop precedence can be used to mitigate the effect of round-trip time on TCP flows [4] and to achieve fair allocation of excess network bandwidth among congestion sensitive TCP and insensitive UDP flows. In this study, we perform wide ranging simulations with two and three levels of drop precedence (or colors) in order to understand the factors influencing fair allocation of excess network resources among congestion sensitive and insensitive flows.

2 Simulation Configuration

The simulations performed in this study use the network configuration shown in Figure 1. Here, customers 1 through 10 send data over the link between Routers 1 and 2 using the same AF traffic class. Traffic is one-dimensional with only ACKs coming back from the other side. Customers 1 through 9 carry an aggregated traffic coming from 5 Reno TCP sources. Customer 10 gets its traffic from a single UDP source sending data at a rate of 1.28 Mbps. Common configuration parameters are detailed in Table 1. All TCP and UDP packets are marked green at the source before being 'recolored' by a traffic conditioner at the customer site. The traffic conditioner consists of two 'leaky' buckets (green and yellow) that mark packets according to their token generation rates (called reserved/green and yellow rate). In two color simulations, yellow rate of all customers is set to zero. Thus, in two color simulations, both UDP and TCP packets will be colored either green or red. In three color simulations, customer 10 (the UDP customer) always has a yellow rate of 0. Thus, in three color simulations, TCP packets coming from customers 1 through 9 can be colored green, yellow or red and UDP packets coming from customer 10 will be colored green or red. All the traffic coming to Router 1 passes through a Random Early Drop (RED) queue. The RED policy implemented at Router 1 can be classified as *Single Average Multiple Threshold* RED as explained in next section.

We have used NS simulator version 2.1b4a [5] for these simulations. The code has been modified to implement the traffic conditioner and multi-color RED (RED_n).

3 Calculating Average Queue Length for RED

In RED, the drop probability of a packet depends on the average queue length which is an exponential average of instantaneous queue length at the time of the packet's arrival [6]. The drop probability increases linearly from 0 to max_p as average queue length increases from min_th to max_th. With packets of multiple colors, one can calculate average queue length in many ways and have multiple sets of drop thresholds for packets of different colors. In general, with
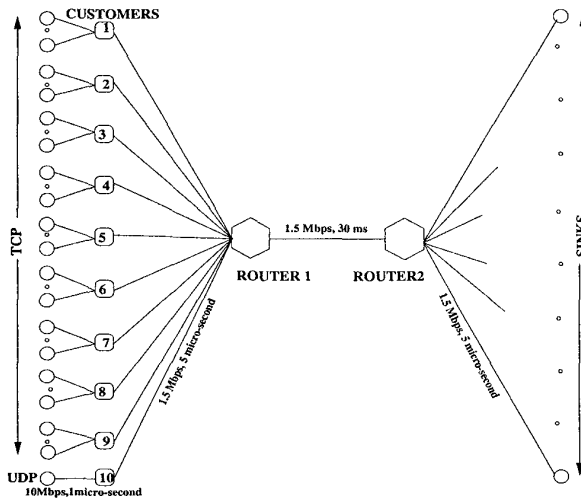
Figure 1: Simulation Configuration

Table 1: Simulation Configuration Parameters

Simulation	100 s	Customers-Router 1 Links:	
Time		Link B/W	1.5 Mbps
TCP Window	64 pkt	1 Way Delay	5 μs
IP Packet Size	576 bytes	Drop Policy	DropTail
UDP Rate	1.28 Mbps	Router 1-Router 2 Link:	
Queue Size	60 pkt	Link B/W	1.5 Mbps
(for each)		1 Way Delay	30 ms
Customers-UDP/TCPs Links:		Drop Policy	
Link B/W	10 Mbps	(At Router 1)	RED_n
1 Way Delay	1 μs	(At Router 2)	DropTail
Drop Policy	DropTail	Router 2-Sinks Links:	
		Link B/W	1.5 Mbps
		1 Way Delay	5 μs
		Drop Policy	DropTail

multiple colors, RED policy can be implemented as a variant of one of four general categories:

- *Single Average Single Threshold (SAST)*
- *Single Average Multiple Thresholds (SAMT)*
- *Multiple Average Single Threshold (MAST)*
- *Multiple Average Multiple Thresholds (MAMT)*

Single Average Single Threshold RED has a single average queue length and same *min_th* and *max_th* thresholds for packets of all colors. Such a policy does not distinguish between packets of different colors and can also be called *color blind* RED.

In Single Average Multiple Thresholds RED, average queue length is based on total number of packets in the queue irrespective of their color. However, packets of different colors have different drop thresholds. For example, if maximum queue size is 60 packets, the drop thresholds for green, yellow and red packets can be {40/60, 20/40, 0/10}. In these simulations, we use Single Average Multiple Thresholds RED.

In Multiple Average Single/Multiple Threshold RED, average queue length for packets of different colors is calculated differently. For example, average queue length for a color can be calculated using number of packets in the queue with same or better color [4]. In such a scheme, average queue length for green, yellow and red packets will be calculated using number of green, yellow + green, red + yellow + green packets in the queue respectively. Another possible scheme is where average queue length for a color is calculated using number of packets of that color in the queue [7]. In such a case, average queue length for green, yellow and red packets will be calculated using number of green, yellow

and red packets in the queue respectively. Multiple Average Single Threshold RED will have same drop thresholds for packets of all colors whereas Multiple Average Multiple Threshold RED will have different drop thresholds for packets of different colors.

4 Experimental Design

In this study, we perform full factorial simulations involving many factors:

- Green Traffic Rates: Green traffic rate is the token generation rate of green bucket in the traffic conditioner. We have experimented with green rates of 12.8, 25.6, 38.4 and 76.8 kbps per customer. These rates correspond to a total of 8.5%, 17.1%, 25.6% and 51.2% of network capacity (1.5 Mbps). In order to understand the effect of green traffic rate, we also conduct simulations with green rates of 102.4, 128, 153.6 and 179.2 kbps for two color cases. These rates correspond to 68.3%, 85.3%, 102.4% and 119.5% of network capacity respectively. Note that in last two cases, we have oversubscribed the available network bandwidth.
- Green Bucket Size: 1, 2, 4, 8, 16 and 32 packets of 576 bytes each.
- Yellow Traffic Rate (only for three color simulations): Yellow traffic rate is the token generation rate of yellow bucket in the traffic conditioner. We have experimented with yellow rates of 12.8 and 128 kbps per customer. These rates correspond to 7.7% and 77% of total capacity (1.5 Mbps) respectively. We used a high yellow rate of 128 kbps so that all excess (out of green rate) TCP packets are colored yellow and thus can be distinguished from excess UDP packets that are colored red.

- Yellow Bucket Size (only for three color simulations): 1, 2, 4, 8, 16, 32 packets of 576 bytes each.

Table 2: Two Color Simulation Parameters

Simulation ID	Green Rate	Max Drop Probability {Green,Red}	Drop Thresholds {Green,Red}	Green Bucket Size (in Packets)
1-144	12.8 kbps	{0.1,0.1}	{40/60,0/10}	1
201-344	25.6 kbps	{0.1,0.5}	{40/60,0/20}	16
401-544	38.4 kbps	{0.1,1}	{40/60,0/5}	2
601-744	76.8 kbps	{0.5,0.5}	{40/60,20/40}	32
801-944	102.4 kbps	{0.5,1}		4
1001-1144	128 kbps	{1,1}		8
1201-1344	153.6 kbps			
1401-1544	179.2 kbps			

- Maximum Drop Probability: Maximum drop probability values used in the simulations are listed in Tables 2 and 3.
- Drop Thresholds for red colored packets: The network resources allocated to red colored packets and hence the fairness results depend on the drop thresholds for red packets. We experiment with different values of drop thresholds for red colored packets so as to achieve close to best fairness possible. Drop thresholds for green packets have been fixed at {40,60} for both two and three color simulations. For three color simulations, yellow packet drop thresholds are {20,40}.

In these simulations, the queue weight used to calculate RED average queue length is 0.002. For easy reference, we have given an identification number to each simulation (Tables 2 and 3). The simulation results are analyzed using ANOVA techniques [8].

5 Performance Metrics

Simulation results have been evaluated based on utilization of reserved rates by the customers and the fairness achieved in allocation of excess bandwidth among different customers.

Utilization of reserved rate by a customer is measured as the ratio of *green* throughput of the customer and the reserved rate. Green throughput of a customer is determined by the number of green colored packets received at the traffic destination(s). Since in these simulations, the drop thresholds for green packets are kept very high in the RED queue at Router 1, chances of a green packet getting dropped are minimal and ideally green throughput of a customer should equal its reserved rate.

The fairness in allocation of excess bandwidth among n customers sharing a link can be computed using the following formula [8]:

$$Fairness\ Index = \frac{(\Sigma x_i)^2}{n \times \Sigma(x_i^2)}$$

Where x_i is the *excess* throughput of the ith customer. Excess throughput of a customer is determined by the number of yellow and red packets received at the traffic destination(s).

6 Simulation Results

Simulation results of two and three color simulations are shown in Figure 2. In this figure, a simulation is identified by its *Simulation ID* listed in Tables 2 and 3. Figures 2a and 2c show the fairness achieved in allocation of excess bandwidth among ten customers for each of the two and three color simulations. Figures 2b and 2d show the utilization of reserved rate by each of ten customers for each simulation.

It is clear from figure 2a that fairness is not good in two color simulations. With three colors, there is a wide variation in fairness results with best results being close to 1. Note that fairness is zero in some of the two color simulations. In these simulations, total reserved traffic uses all the bandwidth and there is no excess bandwidth available to share.

As shown in Figures 2b and 2d, there is a wide variation in reserved rate utilization by customers in two and three color simulations. Figure 3 shows the reserved rate utilization by TCP and UDP customers. For TCP customers, we have plotted the average reserved rate utilization in each simulation. Note that in some cases, reserved rate utilization is slightly more than one. This is because token buckets are initially full which results in all packets getting green color in the beginning. Figures 3b and 3d show that UDP customers have good reserved rate utilization in almost all cases. In contrast, TCP customers show a wide variation in reserved rate utilization.

In order to determine the influence of different simulation factors and their interactions on the reserved rate utilization and fairness achieved in excess bandwidth distribution, we analyze simulation results statistically using Analysis of Variation (ANOVA) technique. ANOVA involves calculating the *Total Variation* in simulation results around the *Over-*

Table 3: Three Color Simulation Parameters

Simulation ID	Green Rate	Max Drop Probability {Green,Yellow,Red}	Drop Thresholds {Green,Yellow,Red}	Yellow Rate	Bucket Size (in Packets)	
					Green	Yellow
1-720	12.8 kbps	{0.1,0.5,1}	{40/60,20/40,0/10}	128 kbps	16	1
1001-1720	25.6 kbps	{0.1,1,1}	{40/60,20/40,0/20}	12.8 kbps	1	16
2001-2720	38.4 kbps	{0.5,0.5,1}			2	2
3001-3720	76.8 kbps	{0.5,1,1}			32	32
		{1,1,1}			4	4
					8	8

Table 4: Main Factors Influencing Reserved Rate Utilization by TCP Customers

Factor/Interaction	Allocation of Variation	
	2 Colors	3 Colors
Green Rate	18.46%	10.36%
Green Bucket Size	77.14%	81.88%
Green Rate - Green Bucket Size	3.65%	3.34%

Table 5: Main Factors Influencing Fairness Results in Three Color Simulations

Factor/Interaction	Allocation of Variation
Yellow Rate	77.15%
Yellow Bucket Size	10.78%
Yellow Rate- Yellow Bucket Size	9.85%

all Mean and doing *Allocation of Variation* to contributing factors and their interactions. Details about ANOVA can be found in [8].

7 ANOVA Analysis for Reserved Rate Utilization Results

As shown in figure 3, reserved rate utilization of UDP customers is almost always good for both two and three color simulations. However, in spite of very low proability of a green packet getting dropped in the network, TCP customers are not able to fully utilize their reserved rate in all cases. Table 4 shows the Allocation of Variation to contributing factors for reserved rate utilization by TCP customers. For TCP customers, green bucket size is the main factor in determining reserved rate utilization. TCP traffic because of its bursty nature is not able to fully utilize its reserved rate unless bucket size is sufficiently high. In our simulations, UDP customer sends data at a uniform rate of 1.28 Mbps and hence is able to fully utilize its reserved rate even when bucket size is low. The minimum size of the leaky bucket required to fully utilize the token generation rate depends on the burstiness of the traffic.

8 ANOVA Analysis for Fairness Results

Fairness results shown in figure 2a indicate that fairness in allocation of excess network bandwidth is very poor in two color simulations. With two colors, excess traffic of TCP as well as UDP customers is marked red and hence is given same treatment in the network. Congestion sensitive TCP flows reduce their data rate in response to congestion created by UDP flow. However, UDP flow keeps on sending data at the same rate as before. Thus, UDP flow gets most of the excess bandwidth and the fairness is poor. In three color simulations, fairness results vary widely with fairness being good in many cases. Table 5 shows the important factors influencing fairness in three color simulations as determined by ANOVA analysis. Yellow rate is the most important factor in determining fairness in three color simulations. With three colors, excess TCP traffic can be colored yellow and thus distinguished from excess UDP traffic which is colored red. Network can protect congestion sensitive TCP traffic from congestion insensitive UDP traffic by giving better treatment to yellow packets than to red packets. Treatment given to yellow and red packets in the RED queues depends on RED parameters (drop thresholds and max drop probability values) for yellow and red packets. Fairness can be achieved by coloring excess TCP packets as yellow and setting the RED parameter values for packets of different colors correctly. In these simulations, we experiment with yellow rates of 12.8 kbps and 128 kbps. With a yellow rate of 12.8 kbps, only a fraction of excess TCP packets can be colored yellow at the traffic conditioner and thus resulting fairness in excess bandwidth distribution is not good. However with a yellow rate of 128 kbps, all excess TCP packets are colored
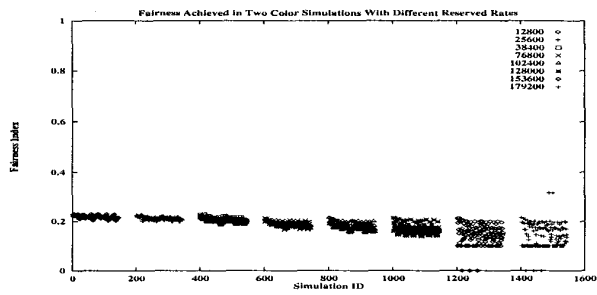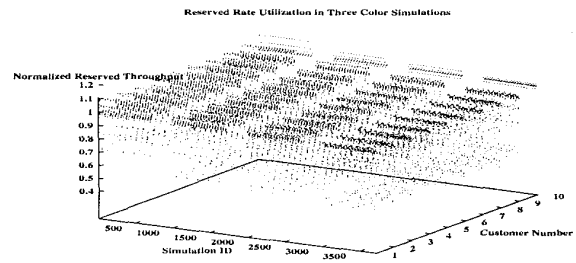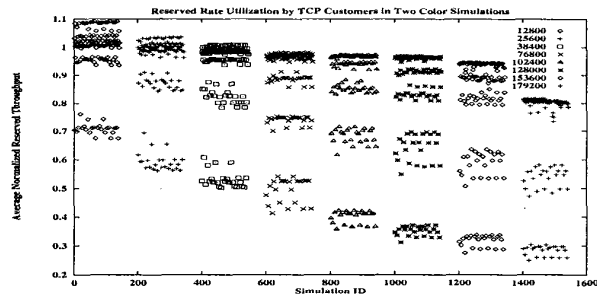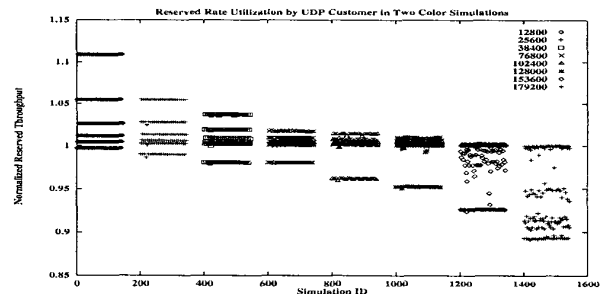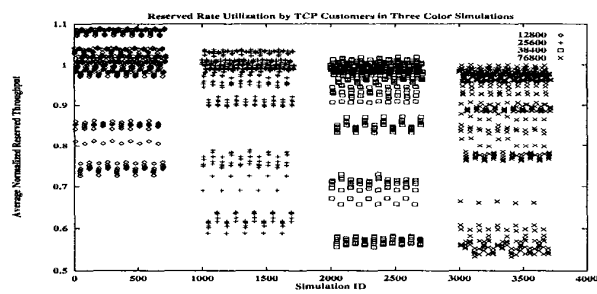
(a)
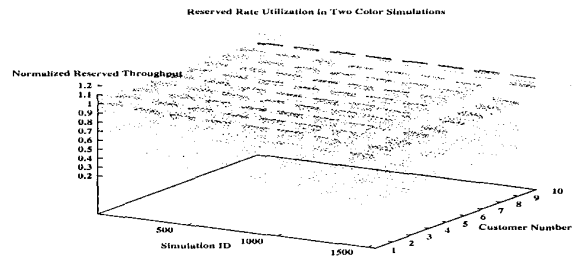
(b)
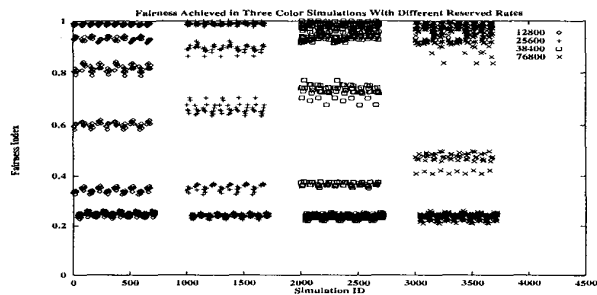
(c)

(d)

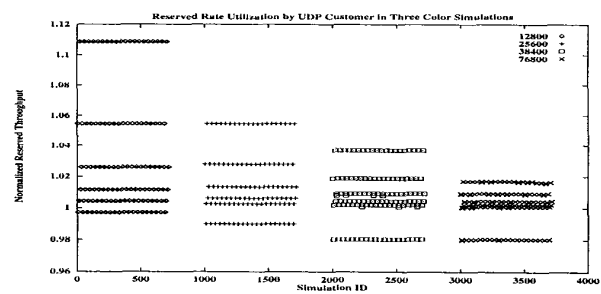Figure 2: Simulation Results

(a)

(b)

(c)

(d)

Figure 3: Reserved Rate Utilization by TCP and UDP Customers

yellow and good fairness is achieved with correct setting of RED parameters. Yellow bucket size also explains a substantial portion of variation in fairness results for three color simulations. This is because bursty TCP traffic can fully utilize its yellow rate only if yellow bucket size is sufficiently high. The interaction between yellow rate and yellow bucket size for three color fairness results is because of the fact that minimum size of the yellow bucket required for fully utilizing the yellow rate increases with yellow rate.

It is evident that three colors are required to enable TCP flows get a fairshare of excess network resources. Excess TCP and UDP packets should be colored differently and network should treat them in such a manner so as to achieve fairness. Also, size of token buckets should be sufficiently high so that bursty TCP traffic can fully utilize the token generation rates.

9 Conclusions

One of the goals of deploying multiple drop precedence levels in an Assured Forwarding traffic class is to ensure that all customers achieve their reserved rate and a fair share of excess bandwidth. It is assumed that combined reserved rate for all customers is less than the network capacity. Network should be configured in such a manner so that in-profile traffic (colored green) does not suffer any packet loss and is successfully delievered to the destination. The fair allocation of excess network bandwidth can be achieved only by giving different treatment to out-of-profile traffic of congestion sensitive and insensitive flows. The reason is that congestion sensitive flows reduce their data rate on detecting congestion however congestion insensitive flows keep on sending data as before. Thus, in order to prevent congestion insensitive flows from taking advantage of reduced data rate of congestion sensitive flows in case of congestion, excess congestion insensitive traffic should get much harsher treatment from the network than excess congestion sensitive traffic. Hence, it is important that excess congestion sensitive and insensitive traffic is colored differently so that network can distinguish between them. Clearly, three colors or levels of drop precedence are required for this purpose. However, if the total reserved traffic is close to network capacity, three levels of drop precedence are redundant as there is not much excess capacity to be shared. Thus, utility of three levels of drop precedence in a traffic class depends on the proportion of reserved traffic to the total capacity.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, W. Weiss, *An Architecture for Differentiated Services*, RFC 2475, December 1998.

[2] J. Heinanen, F. Baker, W. Weiss, J. Wroclawski, *Assured Forwarding PHB Group*, RFC 2597, June 1999.

[3] V. Jacobson, K. Nichols, K. Poduri, *An Expedited Forwarding PHB*, RFC 2598, June 1999.

[4] D. Clark, W. Fang, *Explicit Allocation of Best Effort Packet Delivery Service*, IEEE/ACM Transactions on Networking, August 1998.

[5] NS Simulator, Version 2.1 , Available from http://www-mash.cs.berkeley.edu/ns.

[6] S. Floyd, V. Jacobson, *Random Early Detection Gateways for Congestion Avoidance*, IEEE/ACM Transactions on Networking, August 1993.

[7] N. Seddigh, B. Nandy, P. Pieda, *Study of TCP and UDP Interactions for the AF PHB*, Internet Draft - Work in Progress, draft-nsbnpp-diffserv-tcpudpaf-00.pdf, June 1999.

[8] Raj Jain, *The Art of Computer Systems Performance Analysis*, John Wiley and Sons Inc., 1991.