THE DYNAMICS OF TCP AND UDP INTERACTION IN IP-QOS DIFFERENTIATED SERVICES NETWORKS*

Peter Pieda, Nabil Seddigh, and Biswajit Nandy.

Nortel Networks
3500 Carling Avenue
Neapean, ON. Canada
K1Y 4H7
Fax: (613) 763-4222

ABSTRACT

The results of a study using different drop precedence assignments to address fairness issues when UDP and TCP traffic share the same Assured Forwarding (AF) PHB class in Differentiated Services are presented. Six different combinations of drop-precedence assignments were explored using two different models of RED parameter settings. Results indicate that the type of RED model utilized can play a role in the nature of bandwidth sharing between TCP and UDP flows. The results also show that with the current four Class, three Drop Precedence AF specification, complete fairness between TCP and UDP cannot completely be achieved using separate drop-precedence assignments. This is true for both under-provisioned networks and over-provisioned networks. Certain drop-precedence mapping schemes benefit TCP, while others give advantage to UDP.

Keywords: IP-QoS, Diffserv, Assured Forwarding PHB, TCP and UDP.

1.0 INTRODUCTION

Currently, all user packets compete equally for IP network resources. Recently, the surge in Internet use, coupled with multi-media heavy Web pages and new applications, such as voice and video, has fuelled research to improve the Quality of Service delivered by today's best-effort networks. The underlying concept in IP Quality of Service (IP-QoS) is the ability of network operators to offer different levels of treatment for traffic based on user requirements.

The Differentiated Services (Diffserv) architecture [1] has become the preferred method to address QoS issues in IP networks. This packet-marking approach to IP-QoS is attractive because of its simplicity and scalability. Diffserv is based on locating complicated functionality at the edge of the network and very simple functionality at its core.

Edge devices in this architecture ensure that individual user traffic conforms to traffic profiles, and aggregate flows into a small number of classes. Core devices perform differentiated aggregate treatment of these classes based on the marking performed by the edge devices. For example, core devices treat packet aggregates with Per-Hop-Behavior (PHB) according to their markings. PHB is the forwarding treatment that a packet receives at a network node.

An end-to-end differentiated service is obtained by linking per-domain services with Service Level Agreements (SLAs) between adjoining domains along the traffic's path from source to destination. Per-domain services are realized by traffic conditioning at the edge of the network and simple

* Presented at the 3rd Canadian Conference on Broadband Research, November 1999

differentiated forwarding mechanisms at its core. Two of the more popular proposed forwarding mechanisms are Expedited Forwarding (EF) [2] and Assured Forwarding (AF) [3] PHB.

.

2.0 ASSURED FORWARDING PHB

The host sending protocol must be notified of congestion so that it can react to lessen congestion and limit packet delay incurred in congested nodes. One method of notification is to have the congested node drop packets. The proposed Assured Forwarding PHB provides differentiated drop treatment to users during periods of congestion to realize IP QoS. The AF PHB approach delivers better than best-effort service by controlling the drop preference of packets at the time of congestion. AF PHB is an interesting alternative that may enable service offerings at less cost for audio, video, Web and other applications.

AF PHB is an extension of the RIO scheme [4], which uses a single FIFO queue and two-drop preferences. Most of the current studies of differentiated drop mechanisms are based on the RIO approach. RIO is RED [5] with in-profile and out-of-profile traffic differentiation. The AF PHB RFC [3] builds on this scheme by having four classes and three-drop preferences per class.
A congested device protects packets of low drop precedence (high assurance level) from being lost by discarding packets with a higher drop precedence value. The higher assurance level of a packet has, the less likely statistically that it will be discarded because of node congestion.

Random Early Detection (RED) is a packet discarding algorithm that operates on average queue size. It is designed to detect upcoming congestion and provide feedback to adaptive applications by dropping their packets. When the queue size is below a minimum threshold (min_{th}), RED admits all packets. When the queue is above a maximum threshold (max_{th}), RED drops all packets. When the average queue size is between min_{th} and max_{th}, RED drops packets with an ever-increasing probability up to max_p.

In AF PHB, a different set of RED parameters corresponds to each drop precedence value. The RED parameters of a particular queue cause drop differentiation. The drop precedence sets the relative importance of the packets within the AF class. Each set of RED parameters only acts on packets marked for that drop precedence, thus, appearing to create virtual queues within the physical queues. If a device treats packets with three-drop precedence values in three different ways, it can be considered to have three virtual queues.

The essential components for service based on AF PHB are a Traffic Meter, a Policer, and an Active Queue Management technique, like RED. The Traffic Meter tracks the rate of a user’s traffic at the edge of the network. Using this rate information, the Policer gauges whether the user’s IP flow has exceeded its contracted target rate or profile as setout in its policy. Traffic not exceeding its profile is marked in-profile as policy dictates, while traffic exceeding its profile is marked out-of-profile. During times of congestion, user packets that have been marked in-profile are preferentially treated via the Active Queue Management technique at the core of the network.

3.0 MOTIVATION

Recent studies of the Assured Forwarding PHB have identified end-system-related and network-related factors that affect fair bandwidth distribution for traffic aggregates having equal target rates [4] [7] [8]. End-system factors include variation in TCP stack and in packet size. Network factors encompass impact of Round Trip Time (RTT), large number of active flows, impact of unresponsive flows, number of flows in an aggregate, and size of target rate. These factors cause an unfair distribution of excess bandwidth in an over-provisioned network as well as unfair degradation in an under-provisioned or over-subscribed network.

The effect of unresponsive flows, such as UDP, is important when they share the same AF class as TCP flows. In recent Diffserv IETF discussions on whether the AF PHB required two or three drop precedences, it has periodically been suggested that TCP packets can be protected from non-responsive UDP packets by assigning UDP to a different drop-precedence value than TCP. Similarly, proposed solutions for the UDP/TCP fairness issues for the AF PHB seem to have focused on penalizing the UDP flows. There clearly, is a need to ensure that responsive TCP flows are protected from non-responsive flows in the same class, but at the same time certain UDP flows will require the same fair treatment as TCP due to multimedia demands.

Essentially, the drop-precedence-mapping scheme must ensure fairness for both TCP and UDP. In this case, fairness means the following:

(a) In an over-provisioned network, both UDP and TCP target rates should be achieved with protection for in-profile traffic.

(b) In an over-provisioned network, UDP out-of-profile and TCP out-of-profile packets should have a reasonable share of the excess bandwidth. Neither TCP nor UDP should be denied access to the excess bandwidth.

(c) In an under-provisioned network, TCP and UDP flows should experience degradation in proportion to their target bandwidth.

This study examined different drop-precedence assignments to address fairness issues when UDP and TCP traffic share the same Assured Forwarding (AF) PHB class. In particular, six possible combinations of drop-precedence assignment were evaluated with two different models of RED parameter settings.

4.0 EXPERIMENTAL DETAIL

This study is based on the matrix of options for mapping TCP and UDP to different drop precedences shown in Table 1. (The designations in Table 1 are as follows: DP = drop precedence, IN = in-profile, OUT = out-of-profile.)

Table 1: Possible scenarios for mapping TCP and UDP to different drop precedences

	Scenario					
	1	**2**	**3**	**4**	**5**	**6**
TCP–IN	DP0	DP0	DP0	DP0	DP0	DP0
TCP–OUT	DP1	DP1	DP1	DP2	DP1	DP1
UDP–IN	DP0	DP1	DP1	DP1	DP2	DP0
UDP–OUT	DP1	DP1*	DP2	DP2	DP2 *	DP2

* No distinction is made between UDP–IN and UDP–OUT packets

Scenario 1 is the baseline case that has been used in various studies of fairness issues between UDP and TCP flows. The UDP and TCP flows all have target rates and are mapped to the same drop precedence in a single AF class.

Scenario 2 explores the possibility of mapping TCP in-profile packets to DP0 marking, while mapping TCP out-of-profile and all UDP packets to DP1. This is essentially a similar test case as the one performed by [8]. However, in this study a different RED model is used where the min–max thresholds do not overlap. As results show, this is an important factor.

In Scenario 3, TCP in-profile packets are mapped to DP0; TCP out-of-profile packets are mapped to DP1; UDP in-profile packets are mapped to DP1; and UDP out-of-profile packets are mapped to DP2. This test case differs from Scenario 2 in that the UDP out-of-profile traffic does not share the same drop precedence as the UDP in-profile traffic.

Scenario 4 is the same as Scenario 3 except that TCP out-of-profile packets are put in DP2. Thus, out-of-profile packets for both TCP and UDP are put in DP2 while in-profile traffic is mapped to DP0 and DP1, respectively.

Scenario 5 completely isolates TCP from UDP traffic. TCP in-profile traffic is mapped to DP0, TCP out-of-profile traffic is mapped to DP1, and UDP traffic is mapped to DP2. This mapping is similar to that performed in [8]

In Scenario 6 both TCP and UDP in-profile packets are mapped to DP0. However, TCP out-of-profile packets are mapped to DP1 while UDP out-of-profile packets are mapped to DP2.

Experiments for the above six scenarios were carried out with two different models for the RED [5] parameter settings as shown in Figure 1.

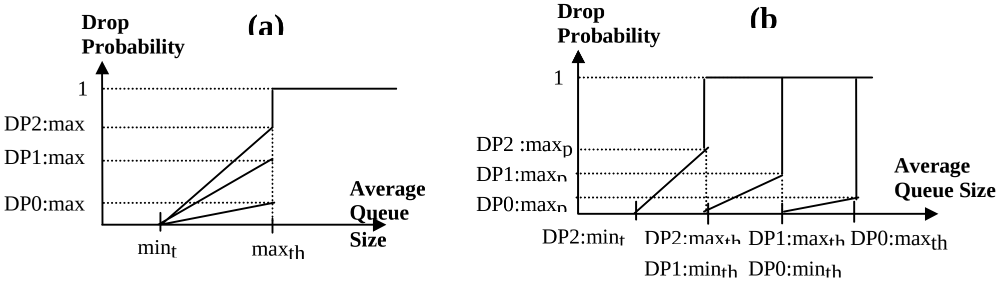

Figure 1 Models for RED parameter settings

The first model is the Overlap RED model and the second is the Non-overlap RED model. In the Overlap RED model, the min–max thresholds are the same for DP0, DP1, and DP2 packets. Consequently, the only factor causing differentiation is max_p – the drop probability. In the Non-overlap RED model (Figure 1b), the min–max thresholds for the different drop precedence decisions do not overlap at all. This model therefore, allows a greater opportunity for higher drop precedence (i.e. DP0) packets to reach their end destinations than the Overlap RED model. The values used for the experiments are listed in Table 2.

Table 2: RED parameter settings for experiments

Drop Precedence	**Model**					
	(a)			**(b)**		
	Min_{th}	**Max_{th}**	**Max_p**	**Min_{th}**	**Max_{th}**	**Max_p**
DP0	10	40	0.02	40	55	0.02
DP1	10	40	0.05	25	40	0.05
DP2	10	40	0.1	10	25	0.1

The experimental network configuration is depicted in Figure 2. Netperf tool [9] was used to generate all the competing TCP traffic and the UDPBLAST tool produced all the UDP non-responsive traffic. The competing TCP flows were all long lived. The link between E1 and the core as well as between E2 and core was 10 Mbps. The link between core and E3 is the bottleneck link and has a bandwidth of 5 Mbps.

A total of 24 TCP flows are generated from the sources that enter the network via edge devices E1 and E2. The 24 flows are divided among the source machines so that each machine sources six flows. A target rate is assigned for each group of six flows. All flows terminate in the sink machines that connect to edge device E3. UDP flows are started from two of the source machines.

The UDP flows source traffic at the rate of 1 Mbps. The target bandwidth for each UDP flow and TCP aggregate group is listed in the figures in Section 5.

The goal of the experiments is to explore six different combinations of drop-precedence mappings for TCP and UDP. All of these scenarios only considered TCP flows with single target rates. Packets from a single policy aggregate can be marked either in-profile or out-of-profile (a maximum of two different drop precedences). However, depending on the particular mapping scheme used, these packets may be assigned any one of the three drop-precedence markings for that class.

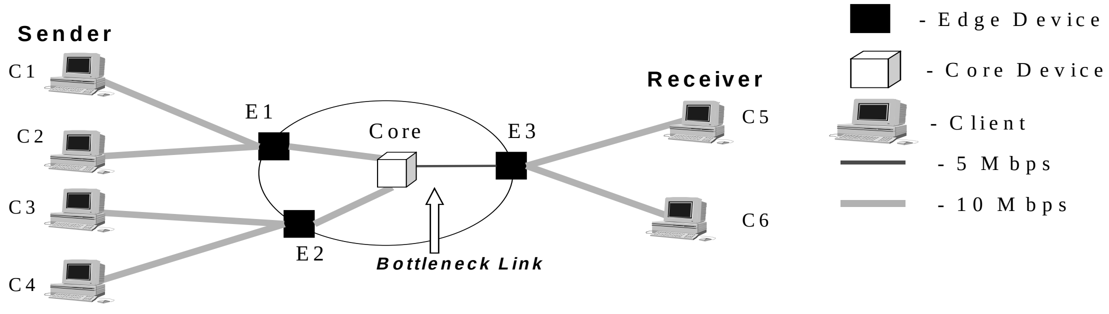

Figure 2 Experimental network setup

The study was carried out using VxWorks-based Diffserv Edge and Core device prototypes developed at the Computing Technology Lab, Nortel Networks. Figure 3 illustrates the functional blocks of a Diffserv Edge device. As a flow enters this device, the Micro Flow Classifier determines if this flow has a user policy associated with it. The Traffic Meter tracks the rate of a user's traffic and the Policer uses this information to decide whether the user's packets should be marked 'out-of-profile' or 'in-profile'. The marker changes the packet marking accordingly. The Buffer Management block enacts multiple-RED on the queues of the Diffserv Edge device. The user's packets then are sent into the network.

Figure 3 Edge device functionality

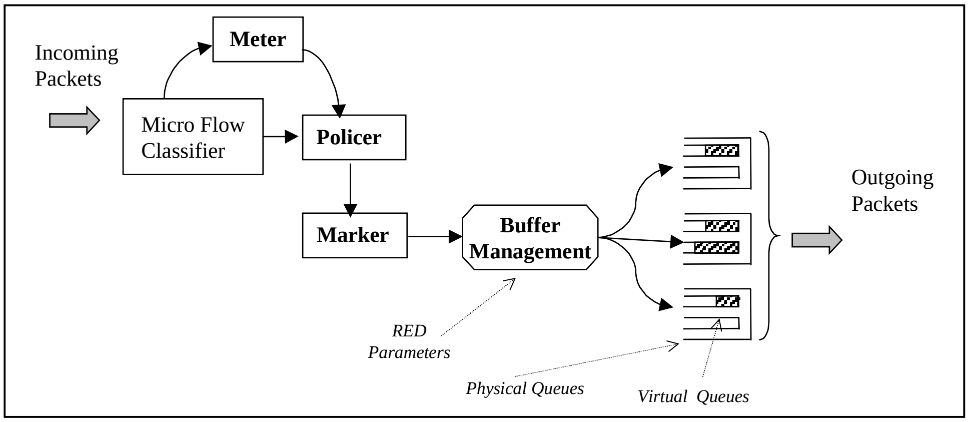

The core device functionality includes a Behavior Aggregate (BA) Classifier, a Scheduler, and the Active Queue Management Scheme MRED. The BA Classifier examines the packet marking to identify to which queue and drop precedence this packet should be mapped. The Scheduler manages the packet-output service each of the queues receives. The prototypes used a single AF class queue with three-drop precedences and a queue solely for Best Effort traffic. The Active Queue Management Scheme MRED controls the differential dropping of packets during congestion.

The devices implement the AF PHB using the Multiple-RED (MRED) algorithm. This document labels DP0 to specify the drop precedence value with lowest drop probability and DP2 to specify the drop precedence with highest drop probability. The MRED algorithm operates as specified in the RIO scheme of [4]. The possibility of dropping DP0 packets depends on the buffer occupancy of DP0 packets. The possibility of dropping DP1 packets depends on the buffer occupancy of DP0 and DP1 packets. The possibility of dropping DP2 packets depends on the buffer occupancy of DP0, DP1, and DP2 packets. The Policer used is the TSW (Time Sliding Window) tagger described in [4].

5.0 RESULTS OF DROP-PRECEDENCE MAPPING

This section presents the results for each RED model experiment for Scenarios 1 to 6. Figure 4 shows the results for tests with Scenario 1 in which TCP and UDP traffic are mapped to the same drop precedence. As the figure shows, TCP flows achieve their target rates in an over-provisioned network. The UDP flows not only achieve their targets but also get a share of the bandwidth for their out-of-profile packets. However, as the network approaches an under-provisioned state, the TCP flows suffer more degradation than the UDP flows. UDP gains unfairly at the advantage of TCP flows. This holds for both Overlap RED and Non-overlap RED models.

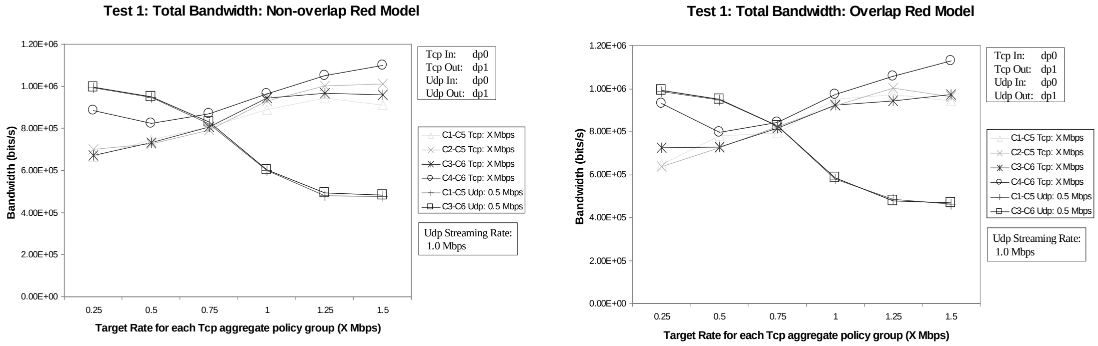

(a) (b)

Figure 4 Scenario 1

The Scenario 2 experiment is essentially similar to that performed in [8]. The results of this experiment are presented in Figure 5. As the network approaches an under-provisioned state, the TCP flows suffer minimal degradation compared to UDP and approach their specified traffic profile. The results for the Non-overlap RED model (a) appear to be slightly different than for the Overlap RED model (b). In Overlap RED model, UDP flows achieve some measure of their bandwidth – although not much. However, in Non-overlap RED model, as the network approaches an under-provisioned state, the UDP flows are severely punished and finally starved. The TCP gain is at the expense of the UDP in-profile and out-of-profile traffic.

Scenario 3 staggers the mapping of drop precedences. The results, as shown in Figure 6, are similar to those of Scenario 2. The only difference is that UDP traffic beyond its target profile is discarded even in an over-provisioned network. This is because UDP out-of-profile is mapped to DP2. In an under-provisioned situation, the DP1 and DP2 queue averages are close enough to the maximum RED threshold that most of their packets are discarded.

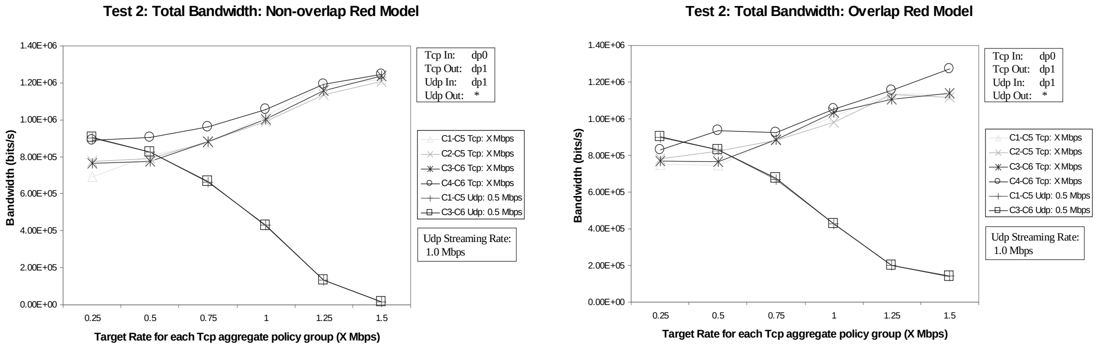

(a (b

Figure 5 Scenario 2

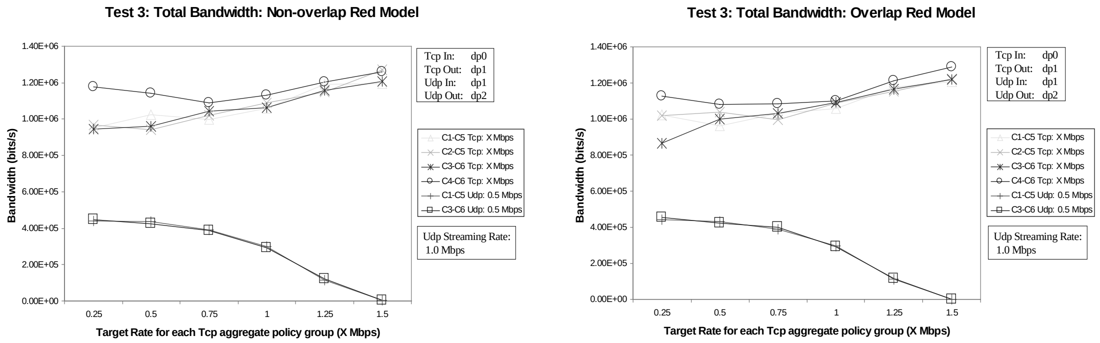

(a (b

Figure 6 Scenario 3

Scenario 4 is a slight variation of Scenario 3 with TCP out-of-profile packets mapped to DP2 instead of DP1. The purpose of this experiment was to protect UDP in-profile traffic by putting TCP-out-of-profile traffic in DP2. The results are illustrated in Figure 7. Moving the TCP out-of-profile packets to DP2 allows UDP to capture a greater share of the bandwidth in an over-provisioned network. However, UDP is still starved in the under-provisioned case because the DP1 average queue size is quite close to the maximum threshold in an under-provisioned network and all of its packets are dropped.

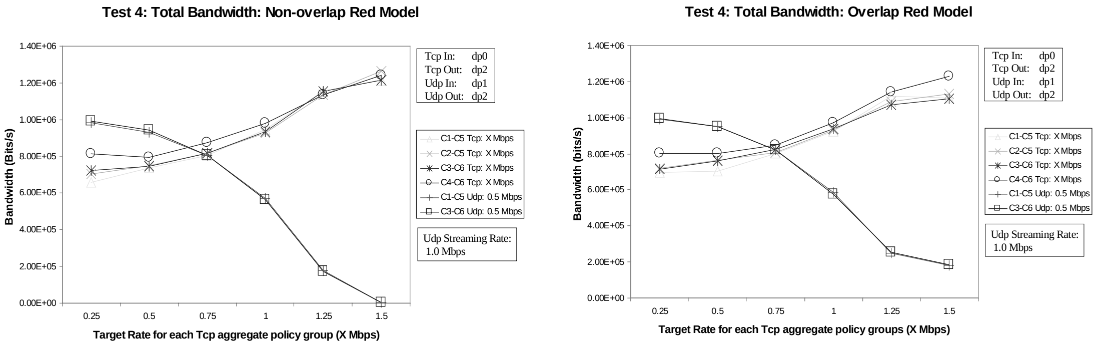

(a (b

Figure 7 Scenario 4

In Scenario 5 the TCP and UDP packets are totally isolated. All the UDP traffic is mapped to DP2 while the TCP traffic is mapped to DP0 and DP1, depending on whether it is in-profile or out-of-profile. The results in Figure 8 show that in the Non-overlap RED model, UDP flows are completely starved. In the Overlap RED model, UDP receives some but very minimal bandwidth. In either situation, the TCP flows are well protected from UDP.

In Scenario 6, TCP and UDP in-profile packets share the same drop preference while their out-of-profile packets are mapped to different drop preferences. The results are captured in Figure 9. In the over-provisioned case, both TCP and UDP achieve their target bandwidth. In the Overlap RED model, the UDP flows achieve a share of the excess bandwidth while the Non-overlap RED model does not allow the UDP flows to pick up any excess bandwidth. In an under-provisioned network, the TCP flows suffer maximal degradation from their target bandwidth, while the UDP flows experience little degradation.

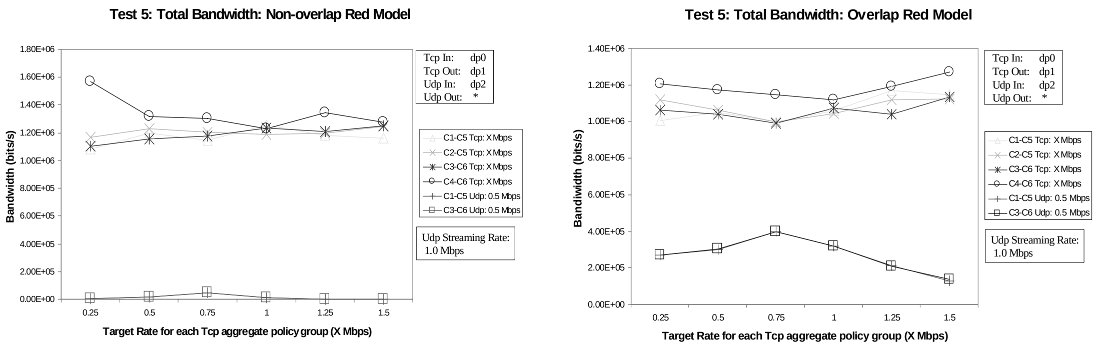

(a (b

Figure 8 Scenario 5

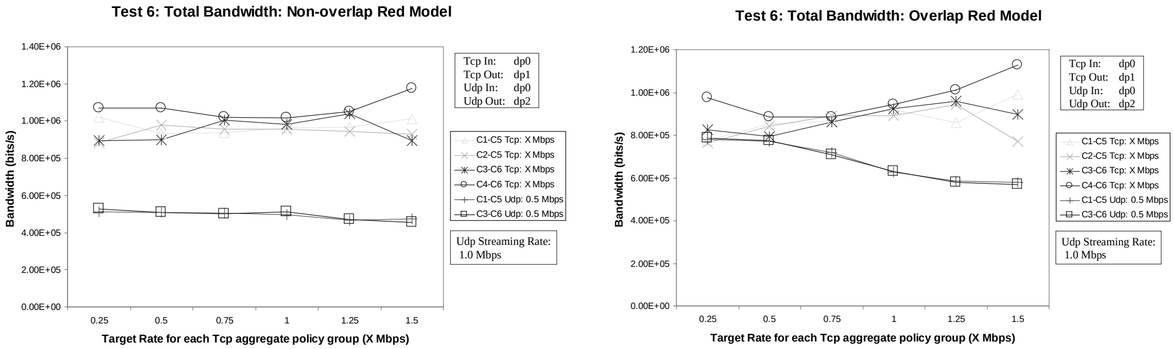

(a (b

Figure 9 Scenario 6

6.0 DISCUSSION OF RESULTS

Table 3 summarizes the results of the six scenarios and puts them in the context of the objectives of this study. None of the scenarios achieves all three objectives for both TCP and UDP. The three objectives translate into six criteria (two each for TCP and UDP) that can be applied to evaluate any drop precedence-mapping scheme.

The results for an over-provisioned network show, that if TCP is mapped to DP0, it mostly achieves its target rate irrespective of the drop preference to which UDP in-profile packets are mapped. UDP achieves its target rate if mapping them to DP0 (with the exception of Scenarios 3 and 5) protects its in-profile packets. However, the manner in which the excess bandwidth is shared remains dependent on the drop preference assigned to UDP out-of-profile packets – none of the scenarios are fair to both UDP and TCP. In an under-provisioned network, mapping TCP in-profile to DP0 and UDP in-profile to either DP1 or DP2 causes TCP to suffer less degradation from its target bandwidth than UDP. Mapping both UDP and TCP in-profile to DP0 results in unfairness to TCP as it experiences severe degradation from its target bandwidth in comparison to UDP.

Table 3: Summary of test results for Scenarios 1 to 6

Scenario	**Over-provisioned Network**				**Under-provisioned Network**	
	Achieves Target Rate		Gets Share of the Excess BW		Achieves Fair Degradation of Target Rate	
	TCP	UDP	TCP	UDP	TCP	UDP
1	✓	✓	✗	✓	✗	✓
2	✓	✓	✗	✓	✓	✗
3	✓	✗	✓	✗	✓	✗
4	✓	✓	✗	✓	✓	✗
5	✓	✗	✓	✗	✓	✗
6	✓	✓	✓	✗✻	✗	✓

✻ Depending on the RED model used, UDP captures some or none of the excess bandwidth.

To further compare the results against the original objectives defined in Section 3, quantitative analysis is performed. The analysis compared the expected versus the actual bandwidth obtained for each of the six scenarios.

For the analysis, each TCP and UDP group is considered a customer. Based on the objectives of Section 3, each customer should obtain bandwidth proportional to the target rate specified in their policy and the Link Bandwidth. The percentage deviation from expected share of bandwidth is calculated. These values give two sets of averages, one for the four TCP customers and one for the two UDP customers. A single average percent deviation is also obtained per scenario.

Equation (1) calculates the expected fair share of the bandwidth for UDP customers. The maximum sending Rate of UDP is considered by taking the minimum between the UDP stream maximum sending rate and the fair share bandwidth.

$$\text{Expected UDP Customer BW} = \text{Min} \left\{ \left(\frac{X_{udpi}}{\sum X_i} \right) * LBW, UR_i \right\} \qquad (1)$$

X_{udpi} = UDP customer target rate

X_i = target rate of each customer irrespective of traffic type

LBW= link bandwidth
UR_i = UDP maximum sending rate.

Equation (2) determines the expected fair share of the bandwidth for TCP customers. Unused bandwidth from the UDP customer(s) is divided between the TCP customers, proportional to their target rate. UDP customers have excess bandwidth when their sending rate is below the expected bandwidth.

$$\text{Expected TCP Customer BW} = \left(\frac{X_{tcpi}}{\sum X_i}\right) * LBW + EUDPBW \qquad (2)$$

$$\text{EUDPBW} = \left(\frac{X_{tcpi}}{\sum X_{tcpi}}\right) * Max\left(\left[\sum\left(\left(\frac{X_{udpi}}{\sum X_i}\right) * LBW\right) - \sum UR_i\right], 0\right) \qquad (3)$$

X_{tcpi} = TCP customer target rate
EUDPBW = unused UDP bandwidth.

$$\text{Percent Deviation} = \frac{Abs(ExpectedBandwidth - ActualBandwidth)}{ExpectedBandwidth} \qquad (4)$$

The graph in Figure 10 illustrates the results of the quantitative analysis. Each column consists of two parts. The total average deviation for the UDP customer group and the TCP customer group. From the graph we can see that the test 4 with the Overlapped Red Model had the least deviation from expected BW. Also note that the UDP customers received a lower quality of service in every test with respect to their expected share of link bandwidth. As well, Figure 10 indicates that all scenarios have at least 20% deviation.

Figure 10 Quantitative analysis results

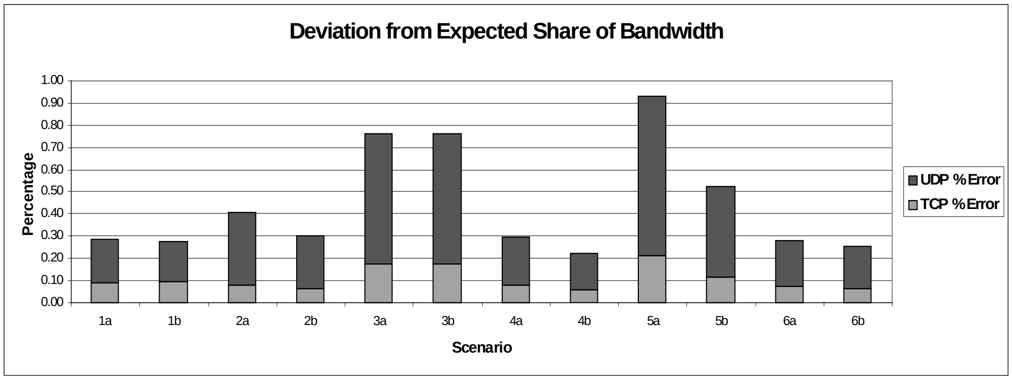

7.0 EXPERIMENTS WITH DE-COUPLED DROP DECISION

Based on the results in Section 6, it appears that the current RIO-based scheme fails to resolve fairness issues between TCP and UDP in a single AF class. It is evident that TCP and UDP packets must be isolated from each other in terms of their drop precedence to achieve fairness. However, as Figure 7 shows, even this is not sufficient. In this scenario, UDP packets in DP2 receive unfairly degraded service because their drop probability is dependent on the buffer occupancy of packets from DP1 and DP0. Additional experiments were carried out where the drop decision for packets of each drop precedence marking, was dependent only on the buffer occupancy of packets with its own marking. Thus, DP2 packet drop decision is dependent on the buffer occupancy of DP2 packets. This scheme is labeled as the de-coupled drop decision.

These tests utilized the same network setup as in the previous scenarios. Scenarios 1, 2 and 5 were repeated with the de-coupled drop decision algorithm as Scenario 1*, 2* and 5*. Table 4 shows the results when the network approaches an under-provisioned state. In these scenarios, each of the TCP groups has a target rate of 1 Mbps and each of the UDP groups has a target rate of 0.5 Mbps. The UDP flows source traffic at the rate of 1 Mbps.

For Scenarios 1* and 2*, the results are beneficial for UDP. All the UDP in-profile and out-of-profile traffic is protected at the expense of TCP in-profile traffic. This is in contrast to Scenario 2 in the previous section. In that experiment, UDP achieved limited or no bandwidth. Neither of these cases was fair.

In Scenario 5*, the UDP flows appear to receive their target bandwidths of 0.5 Mbps. However, on closer examination, the reason for this can be explained. The RED parameter settings and not their target rates govern the bandwidth of the UDP flows. Measurements on queue size revealed a queue size of 70 packets. Fifteen of these packets were DP2. Assuming an equal split among the UDP flows, each UDP flow contributed 7 packets to the queue. The rate of UDP traffic service is thus, (7/70)*5 Mbps = 0.5 Mbps. If the maxth for DP2 was increased, the service rate increased accordingly irrespective of the actual traffic rate. Thus, it seems that even de-coupling drop decisions will not solve the UDP/TCP fairness problem. It could be argued that the real test of isolation and de-coupling can only be done with 4 drop precedences where TCP in-profile and out-of-profile are mapped to DP0 and DP1 while UDP is mapped to DP2 and DP3. This is an area that needs further study.

Table 4: Results of tests with de-coupled drop decision

Test		**DP**	**Min$_{th}$**	**Max$_{th}$**	**Max$_{p}$**	**Bandwidth (Mbps)**					
						C1-C5 TCP	**C2-C5 TCP**	**C3-C6 TCP**	**C4-C6 TCP**	**C1-C5 UDP**	**C3-C6 UDP**
Scenario 1*	i)	0	10	40	0.02	.69	.55	.72	.91	1.02	1.03
		1	10	40	0.05						
	ii)	0	20	40	0.02	.65	.78	.77	.69	1.02	1.02
		1	10	20	0.05						
Scenario 2*	i)	0	10	30	0.02	.66	.77	.52	.88	1.05	1.05
		1	10	25	0.05						
		2	10	20	0.1						
	ii)	0	10	45	0.02	.74	.77	.74	.87	.91	.91
		1	10	20	0.05						
		2	7	15	0.1						
Scenario 5*		0	20	40	0.02	1.0	.9	.93	1.15	.47	.48
		1	10	20	0.05						
		2	7	15	0.1						

8.0 CONCLUSIONS

In summary, the key observations of this study are:

1. UDP and TCP interaction issues cannot be resolved completely using drop-preference mapping:
    - In an over-provisioned network, UDP and TCP traffic achieve their target rates if the in-profile traffic is protected.
    - In an over-provisioned network, the share of excess bandwidth is dependent on the mapping of out-of-profile packets.
    - In an under-provisioned network, if TCP and UDP share the same AF class, fair degradation for both traffic types cannot be achieved by different drop precedence assignments.
2. The choice of RED model impacts how TCP and UDP traffic interact. Some form of guideline would be useful.
3. A better solution that allows UDP and TCP to co-exist fairly is to put them in separate queues or AF classes. Such a scheme will not only address the fairness issues, but also will restrict the delay and jitter that UDP packets experience in the queue. This would make AF better suited for real-time services.

References

[1] Blake, S. Et al, "An Architecture for Differentiated Services", RFC 2475, December 1998.

[2] Jacobson, V. Et al, “An Expedited Forwarding PHB”, Internet Draft, RFC 2598, June 1999.

[3] Heinanen J., Et al, *“Assured Forwarding PHB Group”*, Internet Draft, RFC 2597, June 1999.

[4] Clark D. and Fang W., “Explicit Allocation of Best Effort Packet Delivery Service”, ACM Transactions on Networking, Aug 1998

[5] Floyd, S., and Jacobson, V., “*Random Early Detection gateways for Congestion Avoidance* “, IEEE/ACM Transactions on Networking, V.1 N.4, August 1993, p. 397-413.

[6] Seddigh N, Et al, *"Bandwidth Assurance Issues for TCP flows in a Differentiated Services Network"*, to be presented at GLOBECOM’99, Rio de Janeiro, December 1999. http://www7.nortel.com:8080/CTL/globe9806.pdf

[7] Nandy B, Et al, “Diffserv’s Assured Forwarding PHB: What Assurance does the Customer Have?”, NOSSDAV’99, New Jersey, June 99.

[8] Goyal M, Durresi A, and Jain R, *"Effect of Number of Drop Precedences in Assured Forwarding"*, Internet Draft, <draft-goyal-dpstdy-diffserv-02.txt>, June 1999.

[9] http://www.netperf.org/netperf/NetperfPage.html