

Figure 1: Variants of RED - MRED


the name WRED (Weighted RED). Two simple approaches to the MAMT variant of MRED are RIO-C (RED with In/Out and Coupled average Queues) and RIO-DC (RED with In/Out and Decoupled average Queues). Detailed descriptions of these different MRED variants are provided in the next section.

2.2 WRED – Weighted RED

Weighted RED calculates a single average queue that includes arriving packets of all colors[9]. For an arrival or departure of green, yellow or red packets, WRED updates a single average queue based on total number of packets of green, yellow or red color. However, multiple RED threshold parameters are maintained - one for each color. The parameters for packets of different colors can generally be set in three ways: (i) partially overlapped (Fig.2a) (ii) overlapped (Fig.2c) and (iii) staggered (Fig.2b). Guidance for setting of WRED parameters can be found in [9]. [9] recommends that the partially overlapped (Fig 2a) setting be used with WRED.

2.3 RIO-C (RED with In/Out and Coupled virtual Queues)

The traditional RIO [2] (RED with in/out) uses the same mechanism as in RED, but is configured with two sets of Parameters, one for IN packets and one for OUT packets. The packets of a traffic stream are marked in-profile (IN) when the stream is within the limits of specified policy. When the stream crosses the limits of the specified policy, its packets are marked as out-of-profile (OUT) packets. IN and OUT correspond to DP0 and DP1 or green and yellow packets. We refer to the RIO algorithm described in [2] as RIO-C. RIO-C is easily extended to three-drop precedence or colors [13]. We call RIO-C with three DP's as Generalized RIO-C. In GRIO-C, the average queue for packets of different



a) Partially Overlapped




Figure 2: Parameter settings for MRED

colors can be calculated by adding its average queue to the average queues of colors of lower drop precedence. For example, for red packets, the average queue will be calculated using red, yellow and green packets. With GRIO-C, parameters can be set in any of the way shown in Fig. 2.

2.4 RIO-DC (RED with In/Out and De-Coupled Queues)

In RIO-DC (RIO with Decoupled Queues)[11], the average queue for packets of each color is calculated using number of packets of that color in the queue. The average queue length for green, yellow and red packets will be calculated using the number of green, yellow and red packets in the queue respectively. The RED parameter settings can be chosen depending on the treatment to be given to different colors.

3.0 RELATED WORK

In [6] May et al perform some analytical modeling of Differentiated Service Architecture schemes. Firstly, using throughput as the performance indicator, they find “it does not matter whether the acceptance probability of tagged [marked] packets depends on the total number of packets or on the number of tagged packets in the queue”. Secondly, based on analytic evaluation of the loss probability, they conclude: “choice of different [RIO] parameter values can have a clear impact on performance”. This conclusion is investigated in detail as a part of this work.

In [8], Park et al develop a RI+O buffer management scheme. In RI+O, 3 different average queue calculations are maintained. The drop decision of IN packets is made based on the buffer occupancy of IN packets alone. The drop decision of OUT packets is made based on the buffer occupancy of IN + OUT packets. In addition, there is a further check of the average queue of OUT packets alone against a third set of thresholds to further punish OUT packets. It is not clear as to the significance of this third calculation. RI+O can be considered as a MAMT variant. However it is not clear how the scheme can be extended to include three-drop precedence levels as specified by the AF PHB RFC [4].

Many of the studies reporting experimental and simulation evaluation of Differentiated Services AF PHB are based on RIO-C [2][7][10][13]. However, the authors are unaware of any publications reporting studies with the WRED scheme.

4.0 EXPERIMENTAL CONFIGURATION

The experiments were done with two colors in order to isolate and study the specific buffer management algorithms: WRED & RIO-C. However the results obtained for two-color study can be extended for cases where three colors are used. The experimental setup is depicted in Figure 3. The setup consists of three edge devices E1, E2, E3, a core device C1 and end-hosts, H1 - H4. The edge devices are based on a Pentium-II platform and utilized a VxWorks-based RTOS (Real Time Operating System) implementation of Diffserv. The core device is based on a Pentium-II PC and utilized a Linux 2.2-12 implementation of Diffserv. The end hosts consisted of Pentium-II PC's running Linux-2.0.34 (TCP Reno stack) as the operating systems. There is a 10Mb link between edges and end hosts. The bottleneck link between C1 and E3 has a capacity of 1.7Mb. A delay emulator device is used in the testbed to introduce one-way transmission delay of 30ms. In carrying out the experiments, efforts were made to eliminate any non-homogeneity in hardware and software. Edge devices in the testbed classify, police and mark packets based on source/destination IP address. Token Bucket (single rate two color) meter/policer is used for traffic conditioning[3].

At the core RIO-C and WRED were used for differentiated dropping of packets during congestion. NETPERF[12] and TCPBLAST tools were used during the experiment to generate TCP and UDP traffic from the end-hosts. The RED settings used for tests are shown in Table 1.

5.0 EXPERIMENTAL DESCRIPTIONS AND RESULTS

Three sets of experiments were performed with three different types of Internet traffic: (i) ON-OFF bursty traffic (ii) transactional transfers to (iii) bulk transfers. All the tests were run over a ten minute duration. Each set of experiments was done using staggered, partially overlapped and fully overlapped RED parameter settings (Figure 2). The experim-




Figure 3: Experimental Network Configuration

TABLE 1: PARAMETER SETTINGS FOR DIFFERENT RED MODELS (FIGURE 1)

	Max_p	W_q	Staggered RED Parameters		Partially Overlapped RED Parameters		Fully Overlapped RED Parameters	
			Min_{th} (KB)	Max_{th} (KB)	Min_{th} (KB)	Max_{th} (KB)	Min_{th} (KB)	Max_{th} (KB)
Green (DP0)	0.02	0.002	25	50	25	50	25	50
Yellow (DP1)	0.1	0.002	10	25	10	50	25	50

ents were repeated a sufficient number of times to gain confidence in the results

5.1 Experiment 1: ON-OFF bursty traffic

This section reports results of experimentation when aggregated traffic arrives at an AF queue in ON-OFF fashion. The tests were repeated using the three sets of RED parameter models shown in Table 1.

a) Staggered RED

Staggered RED parameter settings shown in Table 1 were used. Continuous UDP traffic at the rate of 9Mb/s. was sent between H2 and H4 as background traffic to create congestion at the core. The entire traffic stream was marked as yellow (DP1). From H1 to H3, the test traffic consisted of short duration UDP bursts with 120ms 'ON' time and 2s 'OFF' time and all of it marked as green (DP0). A fixed packet size of 1024bytes was used for all traffic. The total number of packet drops was observed at the core. Table 2 summarizes the results. As seen from the table, RIO-C completely protects the bursty ON-OFF traffic even in the presence of congestion. For RIO-C, the DP0 queue is empty in the absence of any green traffic but the DP1 queue average operates around max_{th} (Figure 4a). When a small burst of green packets arrives, the DP0 average queue still remains below min_{th} , thus protecting them. In the case of WRED, this is not the case. Since there is a single average queue, even in the absence of green traffic, the average queue hovers around max_{th} of DP1 i.e. min_{th} of DP0 (Fig 4b). So when bursts of green traffic arrive, the DP0 queue average rises to a level between its min_{th} and max_{th} , thereby dropping some green packets.

b) Partially Overlapped RED

Experiment 5.1a was repeated with partially overlapped

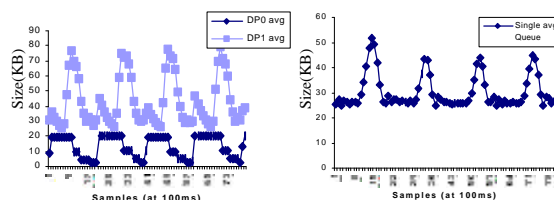


Figure 4: Queue variations for ON-OFF traffic a)RIO-C b)WRED

TABLE 2: ON-OFF BURSTY TRAFFIC-STAGGERED RED

Buffer Mngmt Scheme	Marking Policy	Total Packet (x'd(core))	Dropped Packets	% of total dropped	Observation
RIO-C					
ON-OFF CBR UDP	DP0 (Green)	42000	0	0	Green pkts Totally Protected
	DP1 (Yellow)	750000	656504	87.53%	
WRED					
ON-OFF CBR UDP	DP0 (Green)	42000	2878	6.85%	Green pkts Unprotected
	DP1 (Yellow)	750000	655720	87.43%	

(Fig. 1a) RED settings as depicted in Table 1. These settings are similar to those suggested in [9]. Congestion creating UDP traffic was reduced to 2Mb/s. UDP traffic at 9Mb/s is found to generate heavy drops for both RIO-C as well as WRED under these RED settings hence fair evaluation is not possible. Table 3 summarizes the results. Comparing Tables 2 and 3, we observe that under this new choice of RED parameter configuration, WRED performance in terms of protecting DP0 packets has decreased. It is because of increase of max_{th} for DP1 i.e. fewer forced drops of DP1 packets, cause more of the available buffer to be consumed by yellow packets. This causes an increase in DP0 packet drops. In contrast, RIO-C is still able to protect DP0 packets.

C) Overlapped RED

The third test for this type of ON-OFF traffic used overlapped RED parameters shown in Table 1. Experiment 5.1b was repeated with these new RED parameters. Table 4 shows the observed results at the core device. Examining the results for WRED, we observe that the number of dropped DP0 packets has increased. This is because the increase of min_{th} for DP1 packets provides more protection to yellow packets as compared to previous experiments, thereby increasing more dropped green packets.

5.2 Experiment 2: Transactional transfers

Netperf's TCP_CRR (connect request/response) was used to imitate transactional transfers. TCP_CRR opens a TCP connection. It sends a request of a given size and tears down the connection after receiving the response of prescribed size. A new connection is started immediately with zero delay. A request size of 1KB and response size of 5KB was used for transactions. The tests were done using the three different sets of RED parameter used in previous tests.

TABLE 3: ON-OFF BURSTY TRAFFIC- PARTIALLY OVERLAPPED RED

Buffer Mng't Scheme	Marking Policy	Total Packet rx'd (core)	Dropped Packets	% of total dropped	Observation
RIO-C					
ON-OFF	DP0 (Green)	42000	0	0	Green pkts totally
CBR UDP	DP1 (Yellow)	165000	70492	42.72%	Protected
WRED					
ON-OFF	DP0 (Green)	42000	4394	10.46%	Green pkts
CBR UDP	DP1 (Yellow)	165000	77272	46.83%	Unprotected

TABLE 4: ON- OFF BURSTY TRAFFIC- OVERLAPPED RED

Buffer Mng't Scheme	Marking Policy	Total Packet rx'd (core)	Dropped Packets	% of total dropped	Observation
RIO-C					
ON-OFF	DP0 (Green)	42000	0	0	Green pkts Totally
CBR UDP	DP1 (Yellow)	165000	70508	42.73%	protected
WRED					
ON-OFF	DP0 (Green)	42000	5254	12.5%	Green pkts
CBR UDP	DP1 (Yellow)	165000	75887	46.0%	Unprotected

a) Staggered RED

The staggered RED parameters are set as given in Table 1. Background traffic CBR UDP traffic at 9Mb/s was sent between H2 and H4, all marked yellow. Transactions were done with H3 sending a request and H1 giving a response. Forty TCP_CRR flows were started with zero delay between the start times. The TCP_CRR traffic was marked green as long as it conformed to target rate of 1.7Mb/s. Table 5 provides the results. From this table, we observe that with RIO-C, the TCP_CRR was able to complete 304 transactions, while with WRED, 267 transactions were completed during the same ten minute test period. In addition, while the WRED test showed some packet drops, there were no packet drops during the tests with RIO-C. The 13% downfall in transaction rate of WRED can be attributed to dropping of TCP_CRR packets marked as green and yellow.

b) Partially Overlapped RED

In this test, the RED settings are changed to use the partially overlapped as shown in Table 1. The background UDP CBR traffic between H2 and H4 is sent at the rate of 2Mb/s, all of which is marked yellow. Experiment 5.2a is then repeated under these new network conditions. The results in Table 6 show that the total number of transactions for both the WRED and RIO-C is reduced as compared against the results in Table 5. Again, the scenario utilizing RIO-C provides a higher transaction rate than the scenario with the WRED scheme. A large number of packet drops occur for WRED at the core because WRED's single average queue operate near max_{th} hence dropping DP0 packets, thus affecting Transaction flows. In the case of RIO-C, its

TABLE 5: TRANSACTIONAL TRANSFERS: STAGGERED RED

Buffer Mng't Scheme	Marking Policy	Total Packet rx'd (core)	Drop'd Packets	% of total dropped	Transactions/sec	Observation
RIO-C						
Trans. Xfer	DP0 (Green)	85374	0	0	0.52	Green pkts protected
UDP CBR + Excess Trans. Xfer	DP1 (Yellow)	751158	668645	89.0%		
WRED						
Trans. Xfer	DP0 (Green)	75095	56	0.07%	0.45	Green pkts not protected & lower transact. Rate
UDP CBR+ Excess Trans. Xfer	DP1 (Yellow)	750843	660738	88%		

TABLE 6: TRANSACTIONAL TRANSFERS: PARTIALLY OVERLAPPED RED

Buffer Mng't Scheme	Marking Policy	Total Packet rx'd (core)	Drop'd Packets	% of total dropped	Transactions/sec	Observation
RIO-C						
Trans. Xfer	DP0 (Green)	59382	0		0.39	Green pkts protected
UDP CBR+ Excess Trans. Xfer	DP1 (Yellow)	166434	70283	42.22%		
WRED						
Trans Xfer	DP0 (Green)	34233	6062	17.7%	0.16	Green pkts not protected & lower trans. Rate
UDP CBR	DP1 (Yellow)	165011	46815	28.37%		

transaction rate is reduced from the results of 5.1b because the DP1 packets are allowed a greater portion of the buffer.

c) Overlapped RED

In this experiment, the RED settings are changed to the overlapped with parameters as shown in Table 1. We repeat the experiment as in 5.2b. The observed results can be found in Table 7. RIO-C and WRED performance appears to improve slightly over the test results in 5.2b. Due to the increase of min_{th} , we observe a greater number of DP1 packets in the queue. As RED drop packets on the basis of bandwidth occupied by a flow, it results in increase in earlier drops thereby increasing the overall drop percentage of DP1 packets. This results in some improvement over partial RED settings.

5.3 Experiment 3: Bulk transfers

Experiments with bulk data transfer were also performed to understand the behavior of RIO-C and WRED with long-lived TCP flows such as FTP etc. The test was performed for same time duration and RED settings. All the UDP traffic is marked DP1 as in earlier tests. Netperf's TCP_RR(request/response) is used for generating continuous TCP flows between H1 and H3. TCP_RR opens one connection for the whole test period. It then keeps on generating requests and responses of the prescribed size without any delays. Request size of 64KB and 0B for the response size was used. Three TCP flows are started simultaneously. The three flows share the target rate of 1.7Mb/s – thus, causing green packets to arrive at the buffer at this rate.

a) Staggered RED

This experiment uses the staggered RED settings as used in previous experiments. UDP CBR traffic is sent at 9Mb/s from H2 to H4. Table 8 provides the results. RIO-C and WRED are found to give comparable performance under these settings. RIO-C is also unable to achieve its target rate. This can be attributed to some packet drops of TCP flows marked as DP1 and due to high transmission rate of UDP traffic, even a slight backing off by TCP traffic is rapidly taken by it.

TABLE 7: TRANSACTIONAL TRANSFERS: OVERLAPPED RED

Buffer Mng'ment Scheme	Marking Policy	Total Packet rx'd (core)	Drop'd Packets	% of total dropped	Transactions/sec	Observation
RIO-C						
Trans Xfer	DP0 (Green)	66187	0	0	0.405	Green Pkts protected
UDP CBR + Excess Trans. Xfer	DP1 (Yellow)	166399	75058	45.1%		
WRED					0.17	Green pkts not protected & lower trans Rate
Trans Xfer	DP0 (Green)	38282	6648	17.36%		
UDP CBR	DP1 (Yellow)	165007	49158	29.8%		

TABLE 8: BULK TRANSFERS: STAGGERED RED

Buffer Management Scheme	Marking Policy	Total Packet rx'd (core)	Drop'd Packets	% of total dropped	Throu Ghput (Mb/s)	Observation
RIO-C						
Bulk	DP0 (Green)	77113	0	0	1.5	Green pkts protected
UDP CBR + Excess Bulk	DP1 (Yellow)	750683	716140	95.4%		
WRED					1.3	Green pkts not protected & lower thro'put
Bulk	DP0 (Green)	67709	51	0.07%		
UDP CBR+ Excess Bulk	DP1 (Yellow)	750720	698066	93%		

b) Partially Overlapped RED

With partial overlapped RED settings same as in earlier test cases, background UDP traffic is similarly kept at 2Mb/s. Table 9 gives the observed readings. Performance of both RIO-C and WRED deteriorates. RIO-C performance is affected by increase in packet drops. However it vastly outperforms WRED. In case of WRED, single average queue results in high percentage of dropping of packets of TCP. As a result TCP flows never gets a chance to fully open its window, thereby affecting the throughput. Due to unresponsive UDP flow in the background it takes Full usage of the available bandwidth.

c) Overlapped RED

Experiment 5.3b is repeated with overlapped RED settings of Table 1. Table 10 gives the observed values. Performance of WRED improves slightly from partially overlapped test case. This is similarly explained as in test 5.2c.

5.4 Experiment 4: Varying RED parameters for WRED

In [6] the authors have concluded that for Differentiated

TABLE 9: BULK TRANSFERS: PARTIALLY OVERLAPPED RED

Buffer Management Scheme	Marking Policy	Total Packet rx'd (core)	Drop'd Packets	% of total dropped	Throu gput (Mb/s)	Observation
RIO-C						
Bulk	DP0 (Green)	55328	364	0.65%	1.1	Small % of green pkts not protect'd
UDP CBR	DP1 (Yellow)	165010	103934	63%		
WRED					0.15	High % of green pkts not protect'd & low thr'put
Bulk	DP0 (Green)	8371	1008	12.04%		
UDP CBR	DP1 (Yellow)	165000	38457	23.3%		

TABLE 10: BULK TRANSFERS: OVERLAPPED RED

Buffer Management Scheme	Marking Policy	Total Packet rx'd at core	Drop'd Packets	% of total dropped	Throu gput (Mb/s)	Observation
RIO-C						
Bulk	DP0 (Green)	54744	357	0.65%	1.1	Small % of green pkts not protected
UDP CBR	DP1 (Yellow)	165035	103162	62.5%		
WRED					0.2	High % of green pkts Not protected & low thr'put
Bulk	DP0 (Green)	10008	1010	10.1%		
UDP CBR	DP1 (Yellow)	165000	40638	24.63%		

Services Assured Forwarding based services, similar performance can be achieved with either a SAMT or MAMT scheme via proper selection of RED parameters. We decided to evaluate this conclusion by experimenting this with grossly staggered RED parameters i.e. increase min_{th} for DP0 well beyond max_{th} of DP1 to see at which point, WRED performed comparably to RIO-C. The test case with ON-OFF traffic is chosen. Thus, experiment 5.1a was repeated and RED parameters (min_{th} , max_{th}) were varied. The same max_p and w_q parameters were used as in all other tests. Table 11 illustrates the results. For the protection of DP0 packets for ON-OFF traffic, we observe from table that WRED is able to achieve the same level of performance as RIO-C, if we move the thresholds up at least by 25KB.

6.0 DISCUSSION

For WRED, in case of congestion, the staggered parameter setting provides the greatest assurance that packets with DP0 marking will not be dropped. For the overlapped and partially overlapped cases, WRED provides even less assurance of packets with DP0 marking. For RIO-C, regardless of parameters used, it always provide definite assurance for packets of lower drop precedence not being dropped unless all the packets of higher drop precedence are dropped. Thus lower drop precedence packets are always protected. As observed for ON-OFF traffic, for WRED to obtain same results as RIO-C, WRED needs to be grossly staggered. This improvement in WRED performance with different RED parameters has to be considered vis-à-vis buffer availability, acceptable latency and type of assured service required.

The specific service model created using AF PHB also affects type of buffer management technique to be chosen. If the drop precedence levels map to In-profile/out-profile and out-profile is dispensable at the cost of in-profile traffic, RIO-C scheme seems to be best suited. However, if packets in DP0 and DP1 are unrelated and there's no requirement to only drop DP0 packets after started dropping all DP1 packets, then WRED type scheme can also be used.

7.0 CONCLUSION

In this paper we have undertaken an evaluation of two Widely followed variants of MRED used in AF PHB for

Differentiated Services – WRED and RIO-C. Performance indicators used in the study included drop count of lower drop precedence packets (DP0), transaction rate, throughput and number of retransmissions. Our results find that:

- For ON-OFF traffic, RIO-C is better able to protect DP0 packets than WRED regardless of the RED model used.
- For traffic with characteristics of Transactional transfer, we also find that RIO-C is able to offer better transactional rate per second regardless of the RED model used.
- Both RIO-C and WRED offer greatest protection for DP0 packets when the staggered RED parameter model is utilized.
- WRED needs to have larger thresholds than RIO-C before it can protect DP0 packets. In a heterogeneous Internet, this has implications on delay for traffic as well as product cost.

8.0 REFERENCES

- [1] 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.
- [2] Clark D and Fang W, "Explicit Allocation of Best Effort Packet Delivery Service", ACM Transactions on Networking, Vol. 6, No 4, August 1998, pp 362-373
- [3] Blake S, Black D, Carlson M, Davies E, Wang Z, Weiss W, "An Architecture for Differentiated Services", RFC 2475, December 1998
- [4] Heinanen J, Baker F, Weiss W, and Wroclawski J, "Assured Forwarding PHB Group", Internet RFC 2597, June 1999
- [5] Jacobson V, Nichols K, and Poduri K, "An Expedited Forwarding PHB", Internet RFC 2598, June 1999
- [6] May M, Bolot J.C, Jean-Marie A, and Diot C, "Simple performance Models of differentiated services schemes for the Internet", Proceedings of INFOCOM'99, New York, March 1999
- [7] Goyal M, Durressi A, Jain R and Liu C, Misra P "Effect of number of Drop Precedences in Assured Forwarding", GLOBECOM 99, Rio De Janeiro, December 99.
- [8] Park W. H, Bahk S, Kim H, "A modified RIO algorithm that alleviates the bandwidth skew problem in Internet Differentiated Services" Proceedings of ICC2000, Vol 3, New Orleans, June 2000.
- [9] http://www.cisco.com/univercd/cc/td/doc/product/software/ios120/12cgcr/qos_c/qcpart3/qcwred.htm.
- [10] Seddigh N, Nandy B, and Pieda P, "Bandwidth Assurance issues for TCP flows in a Differentiated Services Network" Proceedings of Global Internet symposium, GLOBECOM' 99, Rio De Janeiro, December 99.
- [11] Seddigh N, Nandy B, Pieda P, Hadi Salim J, and Chapman A, "An experimental study of Assured services in a Diffserv IP QoS Network", Proceedings of SPIE symposium on QoS issues related to The internet, Boston, November 1998.
- [12] <http://www.netperf.org/netperf/NetperfPage.html>
- [13] Elloumi O, Snodder D and Pauwels K, "Usefulness of three drop precedence in Assured Forwarding service", IETF Draft, <draft-elloumi-diffserv-threestwo-00.txt>, July 1999

TABLE 11: WRED WITH GROSSLY STAGGERED PARAMETERS

Test Performed	Min_{th} (KB)	max_{th} (KB)	Marking Policy	Total Packet rx'd(core)	Dropped Packets	% of total
ON-OFF	30	55	DP0 (Green)	42000	714	1.7%
UDP-CBR	10	25	DP1 (Yellow)	750000	659898	88%
ON-OFF	40	65	DP0 (Green)	42000	85	0.2%
UDP-CBR	10	25	DP1 (Yellow)	750000	660811	88.1%
ON-OFF	45	70	DP0 (Green)	42000	36	0.08%
UDP-CBR	10	25	DP1 (Yellow)	750000	660762	88.1%
ON-OFF	50	75	DP0 (Green)	42000	4	
UDP-CBR	10	25	DP1 (Yellow)	750000	660913	88.12%