

Effect of Number of Drop Precedences in Assured Forwarding

draft-goyal-dpstdy-diffserv-02.txt

Mukul Goyal, Arian Duresi, Raj Jain

The Ohio State University

Columbus, OH 43210-1277

Jain@cis.ohio-state.edu

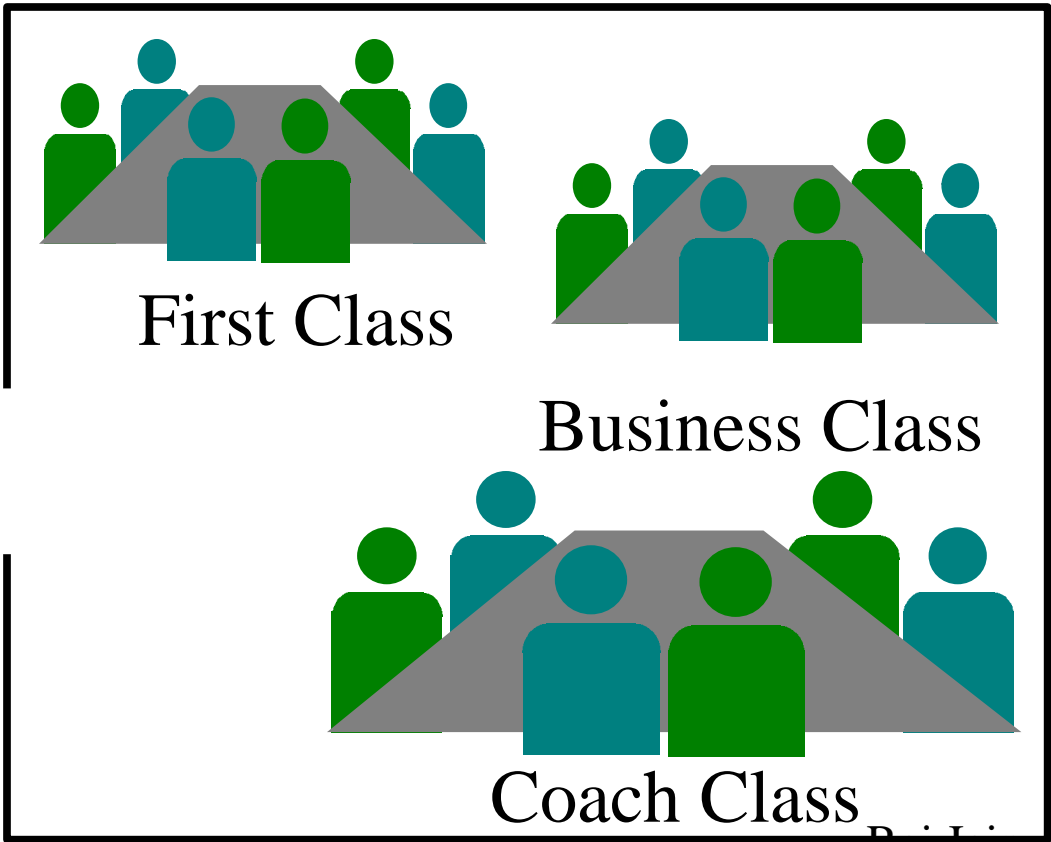
These slides, ID, and a paper are available on-line at

<http://www.cis.ohio-state.edu/~jain/ietf/dpstd2.htm>



- ❑ Key Variables
- ❑ Buffer Management Classification: Types of RED
- ❑ Traffic Types and Treatment
- ❑ Level of Reserved Traffic
- ❑ Two vs Three: Best Results
- ❑ Summary

Key Variables



Key Variables

□ **Bandwidth Management:**

- Number of colors: One, Two, or Three
- Percentage of green (reserved) traffic: Low, high, oversubscribed

□ **Buffer Management:**

- Tail drop or RED
- RED parameters, implementations


□ **Traffic Types and their treatment:**

- Congestion Sensitivity: TCP vs UDP
- Excess TCP vs Excess UDP

□ **Network Configuration:**

Our goal is to identify results that apply to all configs.


Buffer Management Classification



- ❑ Accounting (queued packets):
Per-color, per-VC, per-flow, or Global
Multiple or Single
- ❑ Threshold: Single or Multiple
- ❑ Four Types:
 - Single Accounting, Single threshold (SAST)
 - Single Accounting, Multiple threshold (SAMT)
 - Multiple Accounting, Single threshold (MAST)
 - Multiple Accounting, Multiple threshold (MAMT)

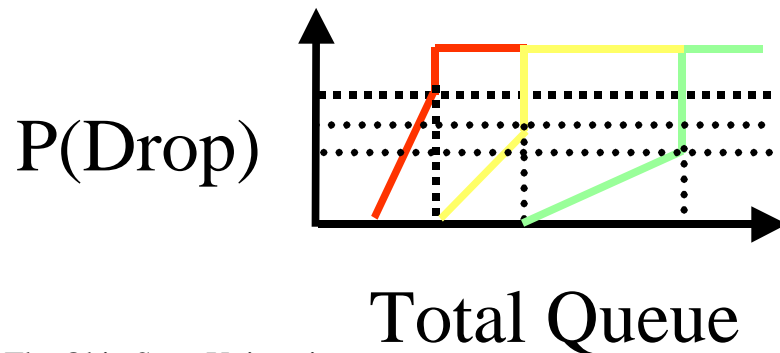
Types of RED

- Single Accounting Single Threshold (SAST):
Color-blind Random Early Discard (RED)



Used in present
diffserv-unaware routers


- Single Accounting Multiple Threshold (SAMT):
Color-Aware RED as implemented in some products



Used in this study

Types of RED (Cont)


- Multiple Accounting Single Threshold (MAST):



G, G+Y, G+Y+R Queue

Used in our
previous study

- Multiple Accounting Multiple Threshold (MAMT):



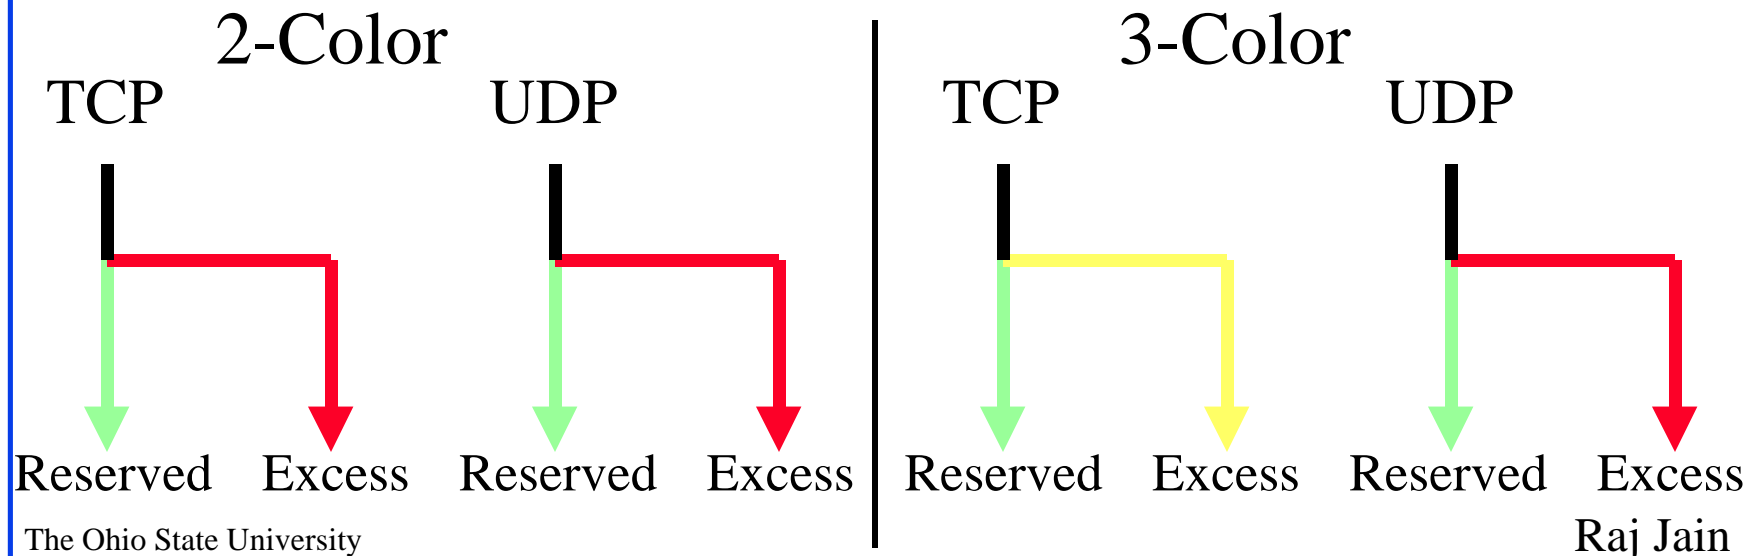
R, Y, G Queue

Conclusion:

More Complexity
 \Rightarrow More Fairness

Traffic Types and Treatment


- Both TCP and UDP get their reserved (green) rates
- Excess TCP competes with excess UDP
- UDP is aggressive
 - ⇒ UDP takes over all the excess bandwidth
 - ⇒ Give excess TCP better treatment than excess UDP



Level of Reserved Traffic

- ❑ Percentage of reserved (green) traffic is the most important parameter
- ❑ If the green traffic is high
 - ⇒ No or little excess capacity
 - ⇒ Two or three colors perform similarly
- ❑ If the green traffic is low
 - ⇒ Lots of excess capacity
 - ⇒ Behavior of TCP vs UDP impacts who gets excess
 - ⇒ Need 3 colors + Need to give excess TCP yellow
 - + Need to give excess UDP red colors

Simulation Configuration



1 μ s	5 μ s	30 ms	5 μ s
10 Mbps	1.5 Mbps	1.5 Mbps	1.5 Mbps

Link Parameters

	Link B/W	Link Delay	Drop Policy
Between TCP _i /UDP & Customer _i	10 Mbps	1 μ s	DropTail
From Customer _i to R ₁	1.5 Mbps	5 μ s	DropTail w marker
From R ₁ to Customer _i	1.5 Mbps	5 μ s	DropTail
From R ₁ to R ₂	1.5 Mbps	30 ms	RED _n
From R ₂ to R ₁	1.5 Mbps	30 ms	DropTail
Between R ₂ & Snk _i	1.5 Mbps	5 μ s	DropTail

Simulation Parameters

- ❑ *Single Accounting Multiple Threshold RED*
- ❑ RED Queue Weight for All Colors: $w = 0.002$
$$Q_{avg} = (1-w)Q_{avg} + w Q$$
- ❑ Maximum Queue Length (For All Queues): 60 packets
- ❑ TCP flavor: Reno
- ❑ TCP Maximum Window: 64 packets
- ❑ TCP Packet Size: 576 bytes
- ❑ UDP Packet Size: 576 bytes
- ❑ UDP Data Rate: 1.28Mbps

Two Color Simulations

Simulation Configuration	Green Token Generation Rate [kbps]	Green Token Bucket Size (in Packets)	Maximum Drop Probability {Green, Red}	Drop Thresholds {Green, Red}
1 Through 1152	12.8, 25.6, 38.4, 76.8, 102.4, 128, 153.6, 179.2	1, 2, 4, 8, 16, 32	{0.1,0.1} {0.1,0.5} {0.1,1} {0.5,0.5} {0.5,1} {1,1}	{40/60,0/10} {40/60,0/20} {40/60,0/5} {40/60,20/40}

Three Color Simulations

Simulation Config.	Green Token Gener. Rate [kbps]	Green Token Bucket Size in Packets	Yellow Token Bucket Size in Packets	Max Drop Probability {Green, Yellow, Red}	Drop Thresholds {Green, Yellow, Red}	Yellow Token Gener. Rate [kbps]
1 Through 2880	12.8, 25.6, 38.4, 76.8	1, 2, 4, 8, 16, 32	1, 2, 4, 8, 16, 32	{0.1,0.5,1} {0.1,1,1} {0.5,0.5,1} {0.5,1,1} {1,1,1}	{40/60,20/ 40, 0/10} {40/60,20/ 40, 0/20}	128, 12.8


Fairness Index

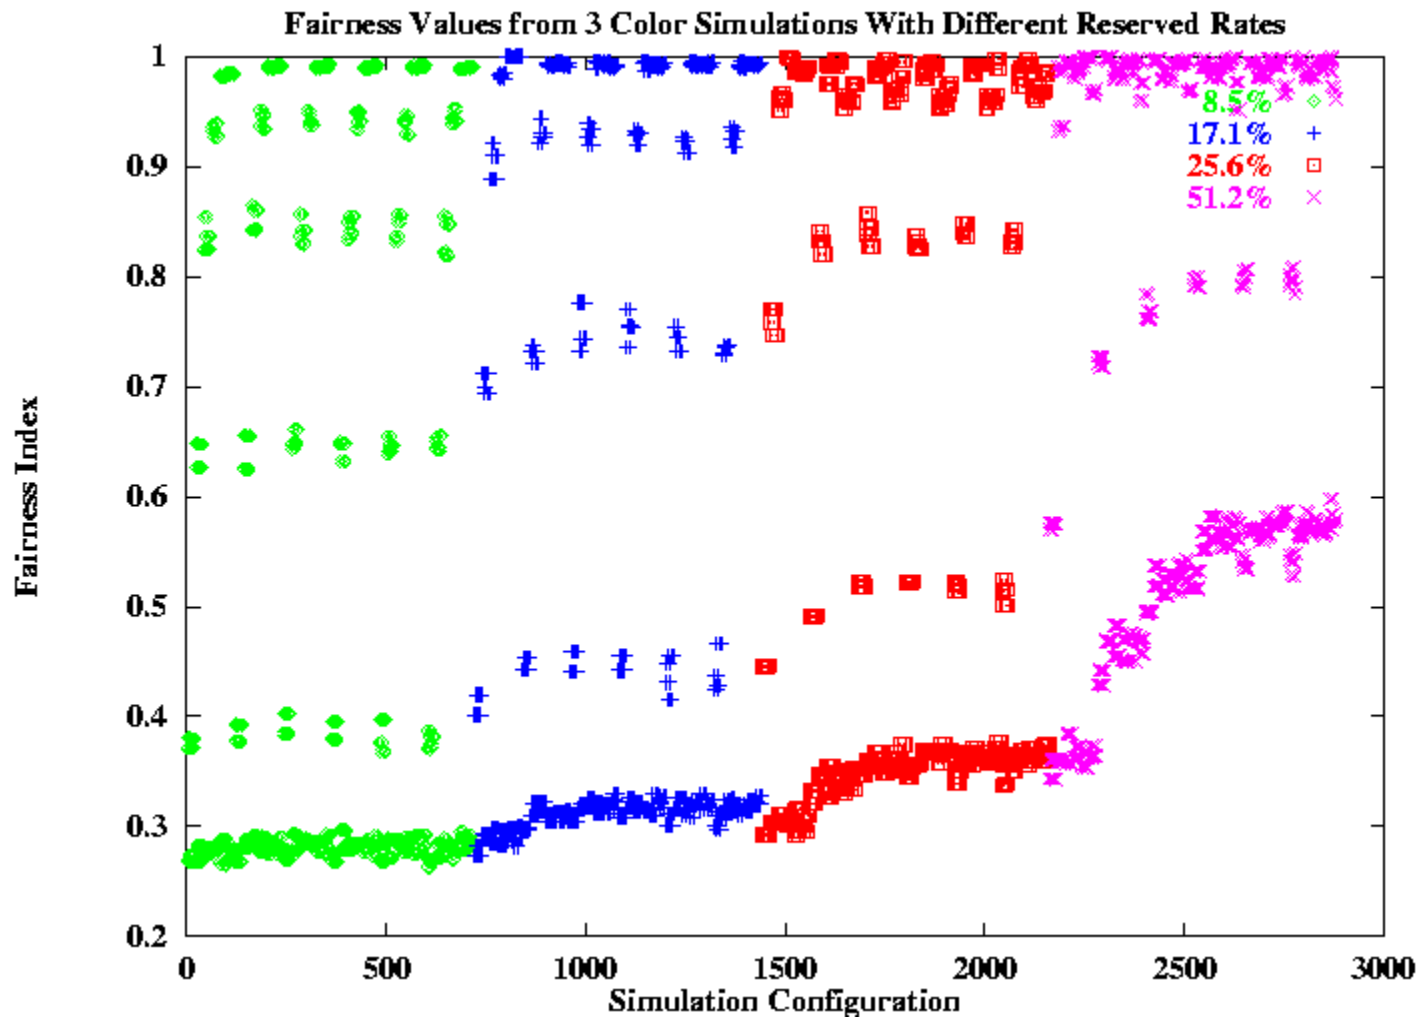
- Measured Throughput: (T_1, T_2, \dots, T_n)
- Use any criterion (e.g., max-min optimality) to find the Fair Throughput (O_1, O_2, \dots, O_n)
- Normalized Throughput: $x_i = T_i/O_i$

$$\text{Fairness Index} = \frac{(\sum x_i)^2}{n\sum x_i^2}$$

Example: 50/50, 30/10, 50/10 \Rightarrow 1, 3, 5


$$\text{Fairness Index} = \frac{(1+3+5)^2}{3(1^2+3^2+5^2)} = \frac{9^2}{3(1+9+25)} = 0.81$$





Two vs Three: Best Results

Fairness



ANOVA For 2 Color Simulations

- Most Important Factors Affecting Fairness:
 - Green Rate (Explains 65.6% of the Variation)
 - Bucket Size (Explains 19.2% of the Variation)
 - Interaction between Green Rate and Bucket Size (Explains 14.8% of the Variation)

ANOVA For 3 Color Simulations

- Most Important Factors Affecting Fairness:
 - Yellow Rate (Explains 74% of the Variation)
 - Yellow Bucket Size (Explains 8.9% of the Variation)
 - Interaction Between Yellow Rate And Yellow Bucket Size (Explains 7.7% of the Variation)
 - Green Rate (Explains 5.6% of the Variation)



Summary

1. The key performance parameter is the level of green (reserved) traffic
2. If reserved traffic level is high or if there is any overbooking, two and three colors give the same throughput and fairness
3. If the reserved traffic is low, three colors give better fairness than two colors
4. Classifiers have to distinguish TCP and UDP:
Reserved TCP/UDP \Rightarrow Green, Excess TCP \Rightarrow Yellow,
Excess UDP \Rightarrow Red
5. RED parameters and implementations have significant impact.

Thank You!

