98-0876R1: Performance Analysis of TCP Enhancements for WWW Traffic using UBR+ with Limited Buffers over Satellite Links

Mukul Goyal, Rohit Goyal,
Raj Jain, Bobby Vandalore,
Sonia Fahmy

Tom vonDeak, Kul Bhasin
NASA Lewis Research Center
Sastri Kota, Norm Butts
Lockheed Martin Telecommunications

Raj Jain is now at
Washington University in Saint Louis
Jain@cse.wustl.edu
http://www.cse.wustl.edu/~jain/

io-state.edu/~jain/

Overview

- Goals
- TCP over UBR+
- Previous Work
- WWW Model
- Full Factorial Experimental Design and Analysis
- Simulation Results

Goals

- Analyze the effect of three factors:

1. TCP Flavors
   - Vanilla: Slow start and congestion avoidance
   - Fast retransmit and recovery (Reno)
   - New Reno
   - Selective Acknowledgements
2. Switch Drop Policies
   - EPD
   - Per-VC accounting
3. Satellite: WAN, MEO, GEO latencies.

TCP over UBR+

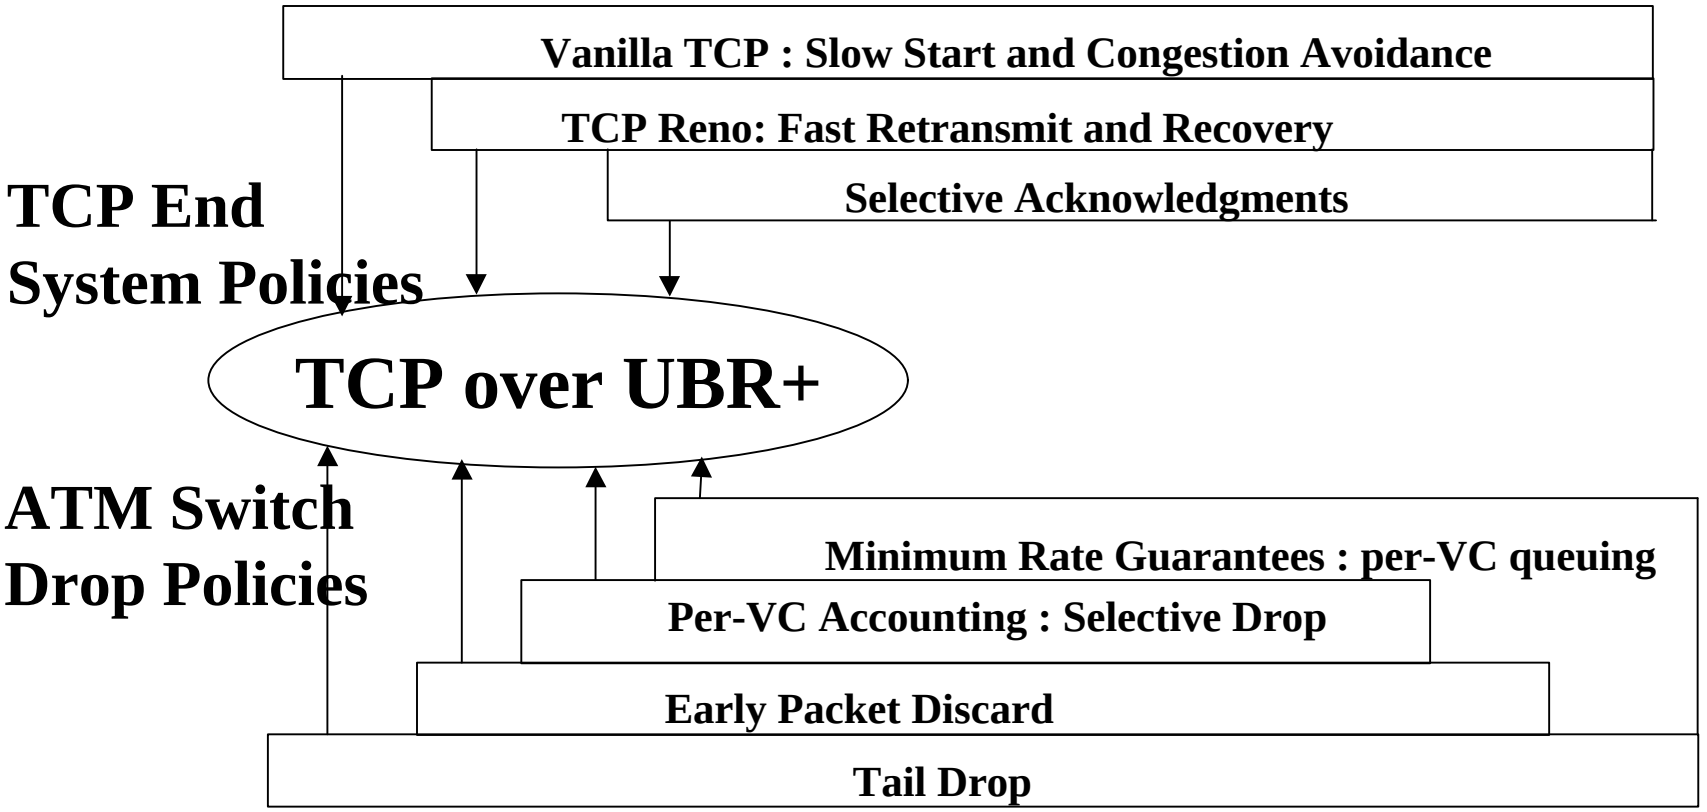
Vanilla TCP : Slow Start and Congestion Avoidance
TCP Reno: Fast Retransmit and Recovery
Selective Acknowledgments
TCP End System Policies
TCP over UBR+
ATM Switch Drop Policies
Minimum Rate Guarantees : per-VC queuing
Per-VC Accounting : Selective Drop
Early Packet Discard
Tail Drop

TCP Mechanisms

- Vanilla TCP
  - Slow start and congestion avoidance
- TCP Reno
  - Fast retransmit and recovery (FRR)
- TCP New Reno
  - Fast recovery phase
- TCP SACK
  - Fast recovery phase
  - Selective acknowledgements

TCP NewReno

- **Receive 3 duplicate acks**: Enter fast recovery phase
- **All lost packets acked**: Exit fast recovery phase
- **Each partial ack**: Send next lost segment
- **Every 2 duplicate acks**: Send 1 new segment (flywheel)
- Recovers from N packet losses in N RTTs
- Implementation based completely on ns simulator (ns2-l b3).
- [FLOYD98] has additional mechanism to avoid multiple retransmits. NOT IMPLEMENTED.

UBR+ Buffer Management

- X_i = Per-VC buffer occupancy. $X = \sum X_i$
- Na = Number of active connections
- **Early Packet Discard**
  - Drop threshold (R) = 0.8 * Buffer size
  - Packet is dropped if X > R
- **Selective Drop**
  - Drop threshold (R) = 0.8 * Buffer size
  - Fairness threshold (Z) = 0.8
  - Packet is dropped if
    - $X > R$ and $X_i > Z*X/Na$

Previous Results

- **Persistent TCP over UBR+**
- **Low delay**: Switch improvements (PPD, EPD, SD, FBA) have more impact than end-system improvements (Slow start, FRR, New Reno, SACK). SACK can hurt under extreme congestion.
- **Satellite networks**: End-system improvements have more impact than switch-based improvements. SACK helps significantly.
- Fairness depends upon the switch drop policies and not on end-system policies

SPECWeb 96 WWW Model

- Majority of traffic on the Internet is WWW
- Developed by Standard Performance Evaluation Corporation (SPEC), a consortium similar to the ATM Forum for performance benchmarking
- SPECMark, SPEC CPU95, SPECInt95, SPEC SFS
- SPECWeb96 is for benchmarking WWW servers
- Ref: http://www.specbench.org/ost/web96/webpaper.html

Modified SPECWeb96

Class 0	Class 1	Class 2	Class 3	Class 4
p = 0.2	p = 0.28	p = 0.40	p = 0.112	p = 0.008
0.1 kB	1 kB	10 kB	100 kb	1 MB
0.2 kB	2 kB	20 kB	200 kB	2 MB
...	...	...	...	...
0.9 kB	9 kB	90 kB	900 kB	9 MB

- Each web page consists of one index page and 4 images.
- First column: Index page (p = 1/5)
- Other columns: p = 0.8

Modified SPECWeb 96

- Average file size = 120.3 KB
- Bandwidth per client = 0.48 Mbps
- HTTP 1.1 ⇒ All components of a web page are fetched in one TCP connection.
- A client makes on average 5 requests every 10s.

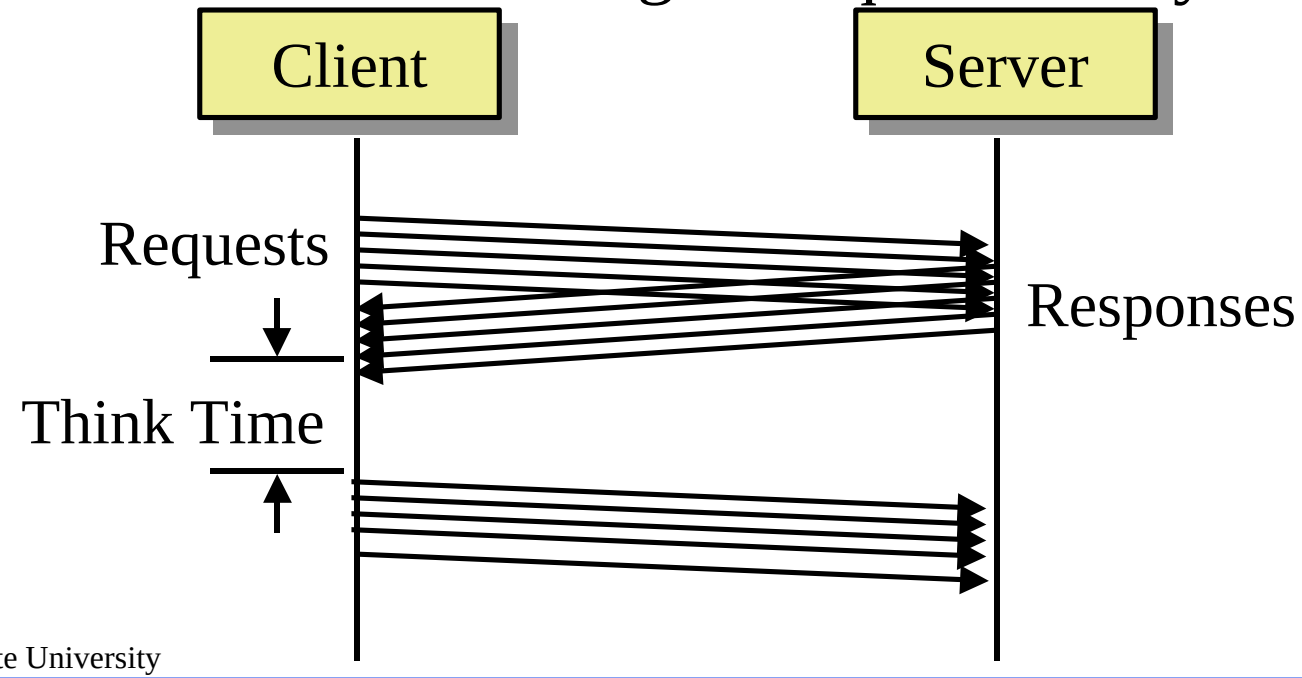

N Client-Server Configuration

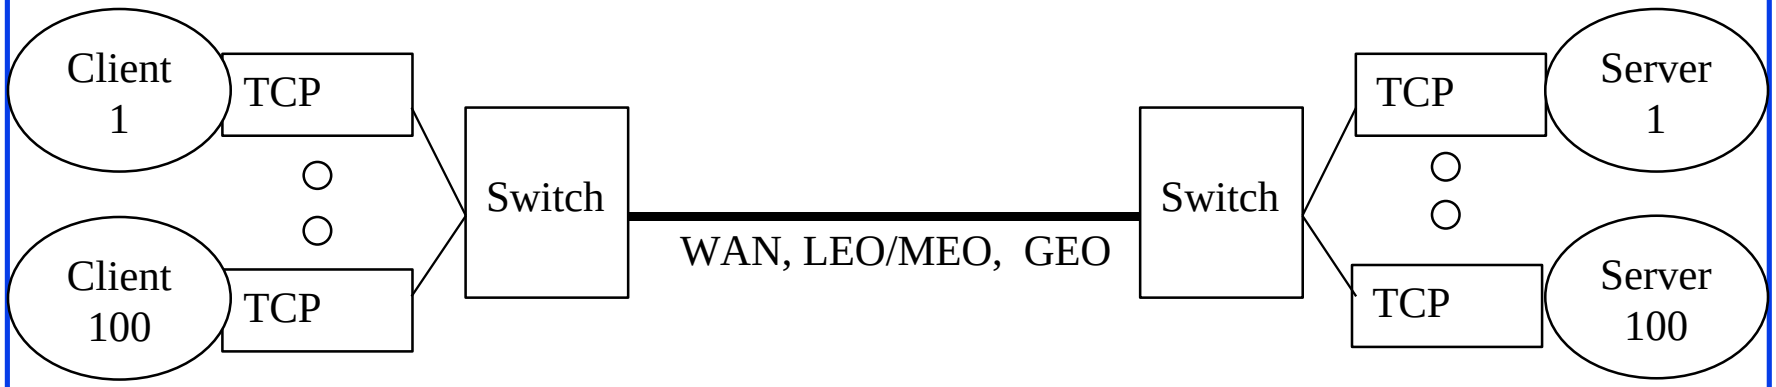

- 1 client per server, N clients and servers, N=100
- RTTs for WAN,multiple-hop LEO/Single-hop MEO and GEO link: 10ms, 200ms and 550ms
- Inter-switch link Bandwidth: 45 Mbps (T3)
- Simulation Time = 100secs i.e. 10 cycles of client requests

TCP Parameters

- MSS = 1024 (WAN), 9180 (LEO/MEO, GEO) bytes
- RCV_WND > RTT × Bandwidth
- "Silly Window Syndrome Avoidance" disabled, since WWW requests must be sent right away.
- Initial SS_THRESH = RTT × Bandwidth [HOE96]
- TCP delay ACK timer is NOT set ⇒ No ack delay
- TCP max window scaled using window scaling option
- TCP timer granularity = 100 ms

Switch Parameters

Link Type (RTT)	RTT-bandwidth product (cells)	Switch Buffer Sizes (cells)
WAN (10 ms)	1062	531, 1062, 2300
Multiple-Hop LEO/Single-Hop MEO (200 ms)	21230	10615, 21230, 42460
Single-Hop GEO (550 ms)	58380	29190, 58380, 116760

Performance Metrics

- Efficiency

$$E = \frac{\sum_{i=1}^{i=N} x_i}{C}$$

  - x_i = achieved TCP throughput. N = 100
  - C = max possible TCP throughput

- Fairness

$$F = \frac{\left(\sum_{i=1}^{i=N} x_i / e_i\right)^2}{N \times \sum_{i=1}^{i=N} \left(x_i / e_i\right)^2}$$

  - e_i = max-min fair throughput for server i

Analysis Technique

Factors	Levels
TCP Flavor	Vanilla, Reno, NewReno, SACK
Buffer Size	0.5 RTT, 1 RTT, 2 RTT
Switch Drop Policy	EPD, SD

- Separate analysis for Efficiency and Fairness results.
- $y_{ijk} = \mu + \underbrace{\alpha_i + \beta_j + \chi_k}_{} + \underbrace{\delta_{ij} + \gamma_{jk} + \phi_{ik}}_{} + \varepsilon_{ijk}$

Observation = Mean + Main Effects + Interaction + Error

Analysis Technique (contd.)

- $\Sigma y_i^2 = n\mu^2 + \Sigma\alpha_i^2 + \Sigma\beta_j^2 + \Sigma\chi_k^2 + \Sigma\delta_{ij}^2 + \Sigma\gamma_{jk}^2 + \Sigma\phi_{ik}^2 + \Sigma\varepsilon_i^2$
  $SS_Y = SS_{Mean} + SS_{Main\ Effects} + SS_{Interaction} + SS_{Error}$
- **Overall Mean** μ: Mean of all values
- **Overall Variation**: Sum of squares of Y
- **Main Effects**: Means of a particular level and factor
- **First Order Interactions**: Interactions between 2 levels of any two factors.
- **Allocation of Variations**: % of the overall variation explained by each effect
- **Confidence Intervals of Effects**: Is the main effect statistically significant?

Results: WAN Efficiency

- TCP flavor is most important factor (57% of variation)
  - NewReno and SACK show best performance
  - SACK is worse for low buffer (high congestion)
- Buffer size is next important factor (30% of variation)
  - Increase in buffer size increases efficiency
  - More room for improvement for Vanilla and Reno
  - Buffer size of 1 RTT is sufficient. This may be related to the number of TCP connections.
- Drop policies have little effect
  - For small buffer, SD is better than EPD

Results: WAN Fairness

- Buffer size most important (53 % of variation)
  - Fairness increase significant for 1 RTT.
- TCP flavor is also important (21% of variation)
  - NewReno has best fairness
  - SACK is very aggressive. Can reduce fairness.
- Drop policy not important for WANs unless buffer size is small

Results: MEO Efficiency

- TCP flavor explains 57% of variation
  - SACK clearly gives best performance
- Buffer size is next important factor (22% of variation)
  - Increase in buffer size increases efficiency
  - More room for improvement for Vanilla and Reno
  - Buffer size of 0.5 RTT is sufficient
- Drop policies have little effect

Results: MEO Fairness

- Fairness values are high for buffer sizes of 0.5 RTT or more.
- TCP flavor, and drop policy do not have much effect.

Results: GEO Efficiency

- TCP flavor explains 61% of variation
  - SACK clearly gives the best performance
- Buffer size is the next important factor (14% of variation)
  - Increase in buffer size increases efficiency
  - More room for improvement for Vanilla and Reno
  - Buffer size of 0.5 RTT is sufficient
- Drop policies have little effect

Results: GEO Fairness

- Fairness values are high for buffer sizes of 0.5 RTT or more.
- TCP flavor, and drop policy do not have much effect.

Overall Results: Efficiency

- End system policies have more effect as delay increases
  - SACK is generally best esp. for long delay
  - NewReno may be better for lower delay and severe congestion
- Drop policies have more effect on lower delays.
- Buffer size: Larger buffers improve performance.
  0.5 RTT to 1 RTT buffers sufficient.
  Optimal buffer size may be related to number of TCPs.

Overall Results: Fairness

- End system policies:
  - SACK hurts fairness for lower delay and extreme congestion
- Drop policies do not have much effect unless delay is lower and buffers are small.
- Buffer size has more effect on longer delays
  - Increase in buffer size increases fairness.

Summary

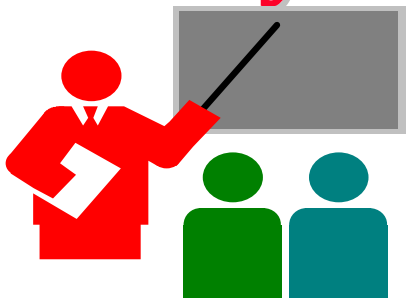

- WWW TCP over UBR+ for WAN and satellite delays
- TCP: Vanilla, Reno, NewReno, SACK
- UBR+: EPD, SD
- Buffer Size: 0.5 RTT, 1 RTT, 2 RTT
- RTT: 10 ms (WAN), 200 ms (MEO), 550 ms (GEO)
- WWW model using modified SpecWeb96

Summary (contd.)

- As delay increases, end system policies have more effect than drop policies or larger buffers.
- SACK is generally most effective
  - Exception: Lower delay and high congestion -- NewReno is best in these cases.
- Drop policies only have an effect for low delays and small buffers.
- Buffer size of 0.5 RTT to 1 RTT is sufficient for the experiments performed. Buffer size may be related to the number of TCP connections.

WAN Efficiency/Fairness

Drop Policy	TCP Flavor	Buffer = 0.5 RTT		Buffer = 1 RTT		Buffer = 2 RTT	
		Efficiency	Fairness	Efficiency	Fairness	Efficiency	Fairness
EPD	Vanilla	0.4245	0.5993	0.5741	0.9171	0.7234	0.9516
	Reno	0.6056	0.8031	0.7337	0.9373	0.8373	0.9666
	NewReno	0.8488	0.8928	0.8866	0.9323	0.8932	0.9720
	SACK	0.8144	0.7937	0.8948	0.8760	0.9080	0.8238
SD	Vanilla	0.4719	0.6996	0.6380	0.9296	0.8125	0.9688
	Reno	0.6474	0.8230	0.8043	0.9462	0.8674	0.9698
	NewReno	0.8101	0.9089	0.8645	0.9181	0.8808	0.9709
	SACK	0.7384	0.6536	0.8951	0.8508	0.9075	0.8989

WAN Allocation of Variation

Component	Sum of Squares		%age of Variation	
	Efficiency	Fairness	Efficiency	Fairness
Individual Values	14.6897	18.6266		
Overall Mean	14.2331	18.3816		
Total Variation	0.4565	0.2450	100	100
Main Effects:				
TCP Flavor	0.2625	0.0526	57.50	21.49
Buffer Size	0.1381	0.1312	30.24	53.55
Drop Policy	0.0016	0.0002	0.34	0.09
First-order Interactions:				
TCP Flavor-Buffer Size	0.0411	0.0424	8.99	17.32
TCP Flavor-Drop Policy	0.0104	0.0041	2.27	1.68
Buffer Size-Drop Policy	0.0015	0.0009	0.33	0.38
Standard Error, s_e = 0.0156(For Efficiency), 0.0472(For Fairness)				

WAN Confidence Intervals

Factor	Main Effect		Confidence Interval	
	Efficiency	Fairness	Efficiency	Fairness
TCP Flavor:				
Vanilla	-0.1627	-0.0308	(-0.1734,-0.1520)	(-0.0632,0.0016)
Reno	-0.0208	0.0325	(-0.0315,-0.0101)	(0.0000, 0.0649)
NewReno	0.0939	0.0573	(0.0832,0.1046)	(0.0248, 0.0898)
SACK	0.0896	-0.0590	(0.0789,0.1003)	(-0.0914, -0.0265)
Buffer Size:				
0.5 RTT	-0.1000	-0.1034	(-0.1087,-0.0912)	(-0.1299,-0.0769)
1 RTT	0.0163	0.0382	(0.0076,0.0250)	(0.0117, 0.0647)
2 RTT cells	0.0837	0.0651	(0.0749,0.0924)	(0.0386, 0.0916)
Drop Policy:				
EPD	-0.0081	-0.0030	(-0.0142, -0.0019)	(-0.0217,0.0157)
SD	0.0081	0.0030	(0.0019,0.0142)	(-0.0157, 0.0217)

MEO Efficiency/Fairness

Drop Policy	TCP Flavor	Buffer = 0.5 RTT		Buffer = 1 RTT		Buffer = 2 RTT	
		Efficiency	Fairness	Efficiency	Fairness	Efficiency	Fairness
EPD	Vanilla	0.8476	0.9656	0.8788	0.9646	0.8995	0.9594
	Reno	0.8937	0.9659	0.9032	0.9518	0.9091	0.9634
	NewReno	0.9028	0.9658	0.9105	0.9625	0.9122	0.9616
	SACK	0.9080	0.9517	0.9123	0.9429	0.9165	0.9487
SD	Vanilla	0.8358	0.9649	0.8719	0.9684	0.9009	0.9615
	Reno	0.8760	0.9688	0.8979	0.9686	0.9020	0.9580
	NewReno	0.8923	0.9665	0.8923	0.9504	0.8976	0.9560
	SACK	0.9167	0.9552	0.9258	0.9674	0.9373	0.9594

MEO Allocation of Variation

Component	Sum of Squares		%age of Variation	
	Efficiency	Fairness	Efficiency	Fairness
Individual Values	19.3453	22.1369		
Overall Mean	19.3334	22.1357		
Total Variation	0.0119	0.0012	100	100
Main Effects:				
TCP Flavor	0.0067	0.0003	56.75	29.20
Buffer Size	0.0026	0.0001	21.73	7.70
Drop Policy	0.0001	0.0001	0.80	6.02
First-order Interactions:				
TCP Flavor-Buffer Size	0.0016	0.0001	13.42	10.16
TCP Flavor-Drop Policy	0.0007	0.0003	6.11	22.60
Buffer Size-Drop Policy	0.0001	0.0001	0.53	6.03
Standard Error, s_e = 0.0036(For Efficiency), 0.0060(For Fairness)				

MEO Confidence Intervals

Factor	Mean Effect		Confidence Interval	
	Efficiency	Fairness	Efficiency	Fairness
TCP Flavor:				
Vanilla	-0.0251	0.0037	(-0.0276,-0.0226)	(-0.0004,0.0078)
Reno	-0.0005	0.0024	(-0.0030,0.0019)	(-0.0017,0.0065)
NewReno	0.0038	0.0001	(0.0013,0.0062)	(-0.0040,0.0042)
SACK	0.0219	-0.0062	(0.0194,0.0244)	(-0.0103,-0.0020)
Buffer Size:				
0.5 RTT	-0.0134	0.0027	(-0.0154,-0.0114)	(-0.0007,0.0060)
1 RTT	0.0016	-0.0008	(-0.0005,0.0036)	(-0.0042,0.0026)
2 RTT	0.0119	-0.0019	(0.0098,0.0139)	(-0.0052,0.0015)
Drop Policy:				
EPD	0.0020	-0.0017	(0.0006,0.0034)	(-0.0041,0.0007)
SD	-0.0020	0.0017	(-0.0034,-0.0006)	(-0.0007,0.0041)

GEO Efficiency/Fairness

Drop Policy	TCP Flavor	Buffer = 0.5 RTT		Buffer = 1 RTT		Buffer = 2 RTT	
		Efficiency	Fairness	Efficiency	Fairness	Efficiency	Fairness
EPD	Vanilla	0.7908	0.9518	0.7924	0.9365	0.8478	0.9496
	Reno	0.8050	0.9581	0.8172	0.9495	0.8736	0.9305
	NewReno	0.8663	0.9613	0.8587	0.9566	0.8455	0.9598
	SACK	0.9021	0.9192	0.9086	0.9514	0.9210	0.9032
SD	Vanilla	0.8080	0.9593	0.8161	0.9542	0.8685	0.9484
	Reno	0.8104	0.9671	0.7806	0.9488	0.8626	0.9398
	NewReno	0.7902	0.9257	0.8325	0.9477	0.8506	0.9464
	SACK	0.9177	0.9670	0.9161	0.9411	0.9207	0.9365

GEO Allocation of Variation

Component	Sum of Squares		% age of Variation	
	Efficiency	Fairness	Efficiency	Fairness
Individual Values	17.3948	21.4938		
Overall Mean	17.3451	21.4884		
Total Variation	0.0497	0.0054	100	100
Main Effects:				
TCP Flavor	0.0344	0.0008	69.16	14.47
Buffer Size	0.0068	0.0006	13.65	11.48
Drop Policy	0.0001	0.0001	0.25	2.31
First-order Interactions:				
TCP Flavor-Buffer Size	0.0037	0.0012	7.54	22.16
TCP Flavor-Drop Policy	0.0025	0.0014	4.96	26.44
Buffer Size-Drop Policy	0.0002	0.0001	0.41	1.45
Standard Error, s_e = 0.0182(For Efficiency), 0.0139(For Fairness)				

GEO Confidence Intervals

Factor	Mean Effect		Confidence Interval	
	Efficiency	Fairness	Efficiency	Fairness
TCP Flavor:				
Vanilla	-0.0295	0.0037	(-0.0420,-0.0170)	(-0.0058,0.0133)
Reno	-0.0252	0.0027	(-0.0377,-0.0127)	(-0.0068,0.0123)
NewReno	-0.0095	0.0034	(-0.0220,0.0030)	(-0.0062,0.0129)
SACK	0.0642	-0.0098	(0.0517,0.0768)	(-0.0194,-0.0003)
Buffer Size:				
0.5 RTT	-0.0138	0.0050	(-0.0240,-0.0036)	(-0.0029,0.0128)
1 RTT	-0.0099	0.0020	(-0.0201,0.0004)	(-0.0058,0.0098)
2 RTT	0.0237	-0.0070	(0.0134,0.0339)	(-0.0148,0.0009)
Drop Policy:				
EPD	0.0023	-0.0023	(-0.0049,0.0095)	(-0.0078,0.0033)
SD	-0.0023	0.0023	(-0.0095,0.0049)	(-0.0033,0.0078)

Thank You!