Improving the Performance of TCP/IP over ATM UBR+ Service

Raj Jain
The Ohio State University
Columbus, OH 43210
Jain@CIS.Ohio-State.Edu
http://www.cis.ohio-state.edu/~jain/

Overview

- TCP/IP over Plain UBR
- Slow Start, FRR, SACK, New Reno
- PPD
- EPD
- Fair Buffer Allocation, Selective Drop
- Guaranteed Rate

TCP Over Plain UBR

- Low throughput
- Unfair
- Anomalies: More receiver buffer
  $\Rightarrow$ Lower throughput
  Due to Silly window avoidance + Delayed Ack
- Solution: Min sender buffer size should be 3 × MSS

Ref: Comer

Improving Performance of TCP over UBR

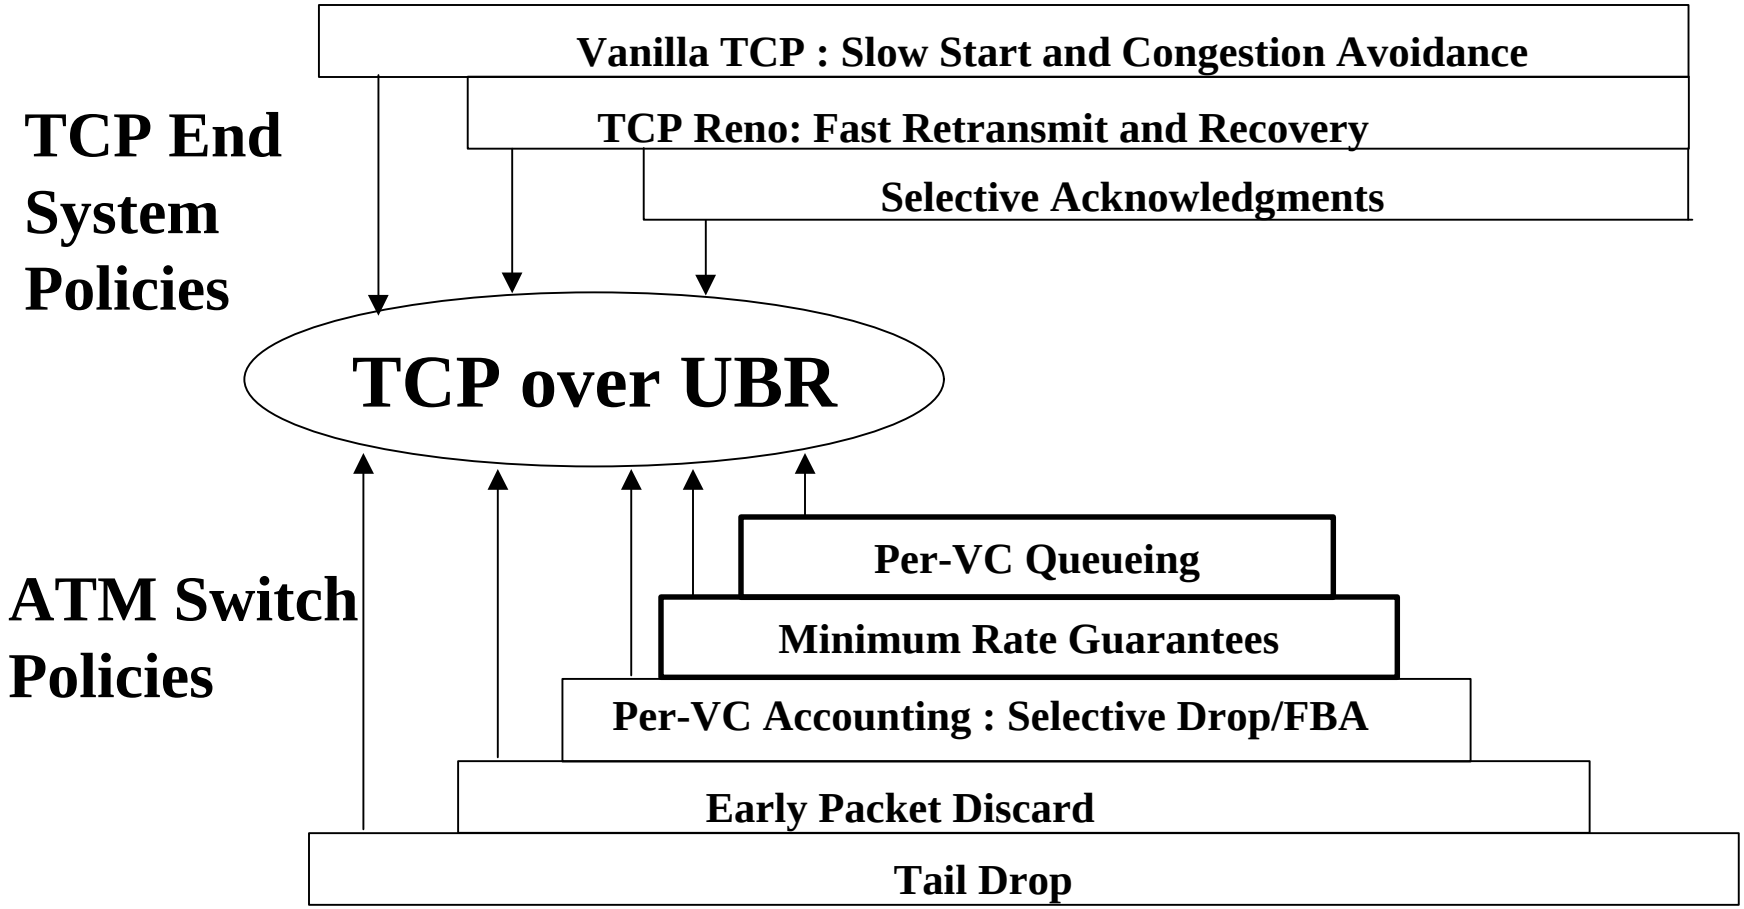

TCP/IP over UBR: Improvements

- Switch Based Mechanisms:
  - PPD
  - EPD
  - EPD + per-VC queueing
  - EPD + per-VC Accounting
- Source Based Mechanisms:
  - Fast Retransmit and Recovery
  - New Reno
  - Selective Acknowledgement

PPD and EPD

- Plain ATM: Discard all cells if Q > threshold
- Partial Packet Discard:
  Discard all cells of a packet if one cell dropped
  Q > threshold
- Early Packet Discard:
  Discard all cells of the next packet if Q > threshold

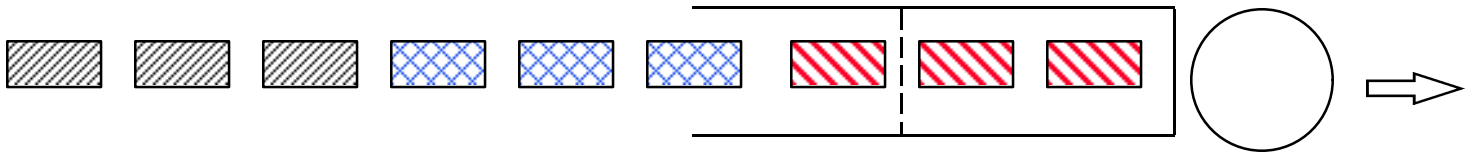

EPD
Threshold

PPD vs EPD

- Plain ATM ⇒ Many packets dropped
- Dropping all cells of a packet is better than dropping randomly
  ⇒ PPD is better than plain UBR
- Never drop the EOM cell of a packet unless the first cell has also been dropped.
  Otherwise two packets are lost.
- EPD ⇒ Even fewer packets dropped
  ⇒ better throughput
- Plain ATM << PPD << EPD
- EPD improves efficiency but not fairness

EPD + Per-VC Queueing

- Accept the next packet if Xi/(X/N) < Z
- Round-robin scheduling ⇒ Fairness improved
- However, more VC's have packets dropped
  ⇒ Lower total throughput

Ref: Siu

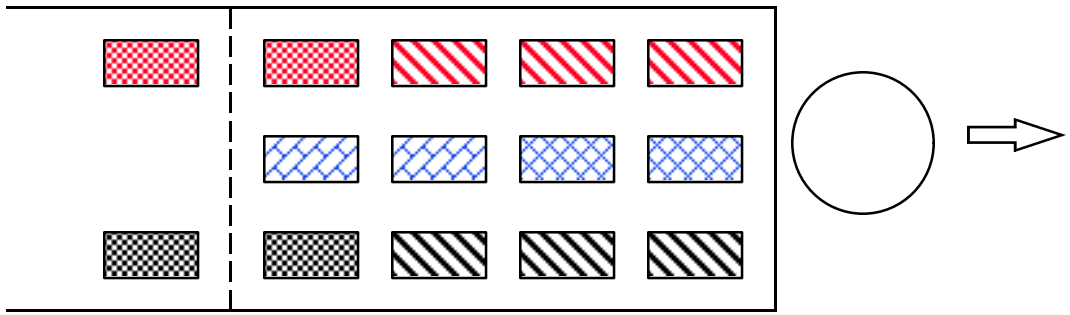

EPD
Threshold

Fair Buffer Allocation

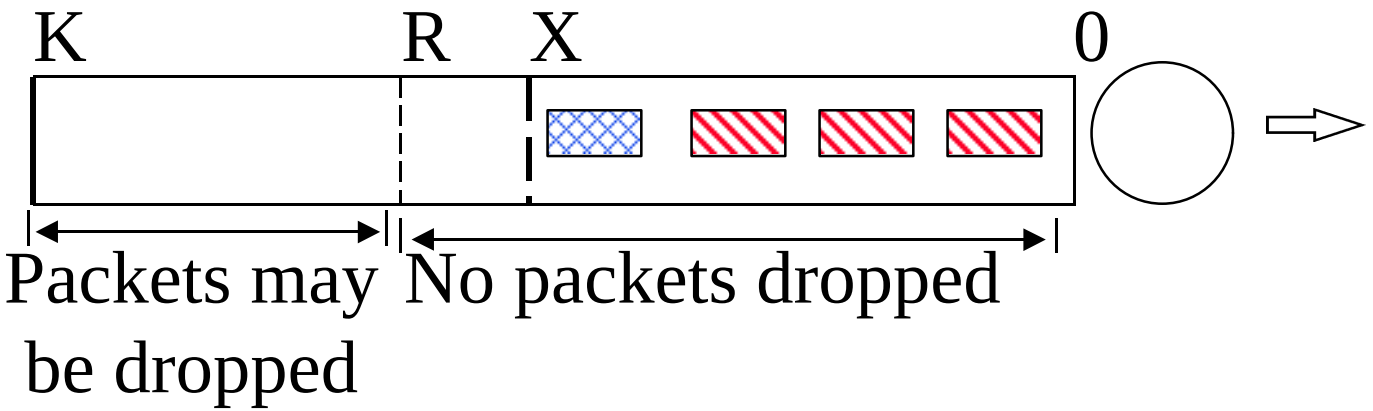

- Drop packets of only high rate VCs
- No per-VC queueing ⇒ All VCs share a single FIFO queue
- per-VC accounting ⇒ track per-VC cell count
- Decrease per-VC buffer allowance as total occupancy increases

FBA (Cont)

- Drop complete packet of VC_i if
  $(X > R)$ AND $(X_i * N_a / X > W(X)$
  $W(X) = Z*((K - R)/ (X - R))$
- X_i = Cells of ith VC, X= Total Cells = ΣX_i
- N_a = Number of active VCs (i.e., $X_i > 0$)
- K = Total buffers, R = Threshold
- Z = parameter between 0.5 and 1.

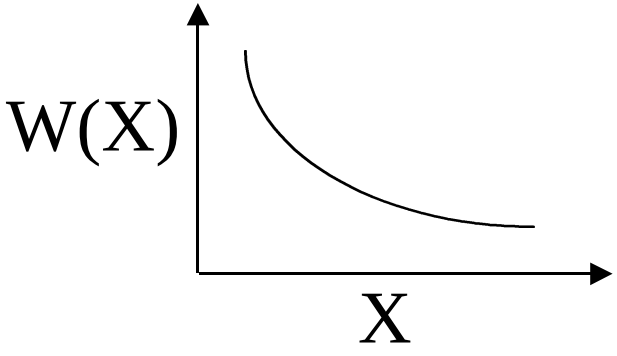

- Note that packets from more and more flows are dropped as queue X increases
- FBA improves fairness and efficiency
- Can we make it simpler?

Selective Packet Discard

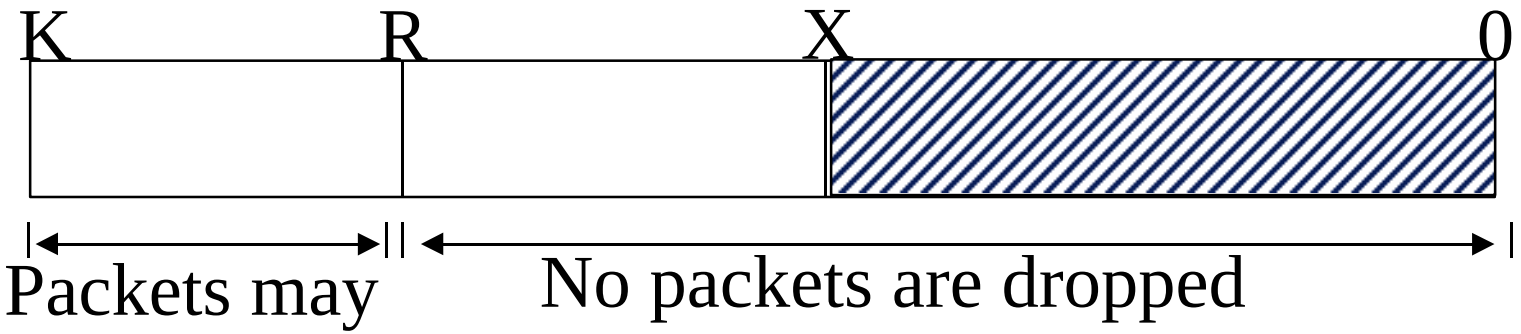

- A simplification of FBA
- Drop complete packet of VCi if:

  $(X > R)$ AND $(X_i/(X/N_a) > Z)$
- Selective drop also improves fairness and is less sensitive to parameters than FBA

Drop Policies: Other Ideas

- Do not drop successive packets
- Drop from front of queues not tails $\Rightarrow$ earlier effect

Policies

Switch Policies		End-System Policies: No FRR	FRR	New Reno	SACK + New Reno
No EPD					
	Plain EPD				
EPD	Selective Drop				
	Fair Buffer Allocation				

Policies: Results

- In LANs, switch improvements (PPD, EPD, SD, FBA) have more impact than end-system improvements (Slow start, FRR, New Reno, SACK). Different variations of increase/decrease have little impact due to small window sizes.
- In satellite networks, end-system improvements have more impact than switch-based improvements
- FRR hurts in satellite networks.
- Fairness depends upon the switch drop policies and not on end-system policies

Policies (Continued)

- In Satellite networks:
  - SACK helps significantly
  - Switch-based improvements have relatively less impact than end-system improvements
  - Fairness is not affected by SACK
- In LANs:
  - Previously retransmitted holes may have to be retransmitted on a timeout
    $\Rightarrow$ SACK can hurt under extreme congestion.

Guaranteed Rate Service

- Guaranteed Rate (GR): Reserve a small fraction of bandwidth for UBR class.

GR	GFR
per-class reservation	per-VC reservation
per-class scheduling	per-VC accounting/scheduling
No new signaling	Need new signaling
Can be done now	In TM4+

Guaranteed Rate: Results

- Guaranteed rate is helpful in WANs.
- For WANs, the effect of reserving 10% bandwidth for UBR is more than that obtained by EPD, SD, or FBA
- For LANs, guaranteed rate is not so helpful. Drop policies are more important.
- For Satellites, end-system policies seem more important.

Summary

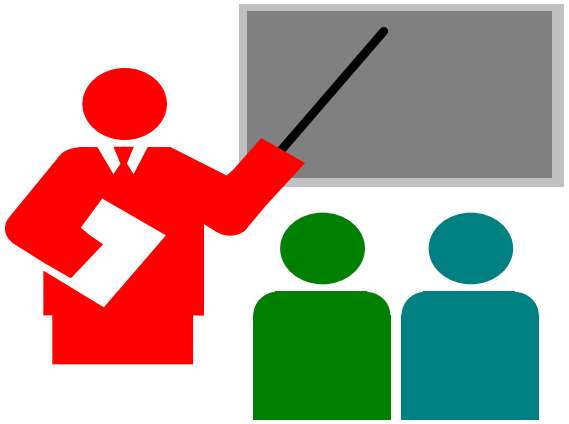

- End system policies are more important than switch policies in WAN. Opposite is true in LANs
- Selective drop and Fair Buffer Allocation improve fairness and efficiencyiy
- FBA is more sensitive to parameters than SD
- In WANs, reserving a small amount of bandwidth helps UBR more than other switch policies

TCP over UBR: References

- A. Romanow and S. Floyd, "Dynamics of TCP Traffic over ATM Networks," IEEE Journal on Selected Areas in Communications, Vol. 13, No. 4, May 1995, pp. 633-641, ftp://ftp.ee.lbl.gov/papers/tcp_atm.ps.Z
- J. Heinanen and K. Kilkki, "A Fair Buffer Allocation Scheme," Telecom Finland Draft 17 March 1995.
- H. Li, K-Y Siu, and H-Y Tzeng, "TCP Performance over ABR and UBR Services in ATM," Proc. IPCCC'96, March 1996.

References (Cont)

- D. E. Comer and J. C. Lin, "TCP Buffering and Performance over an ATM Network," Internetworking: Research and Experience, Vol. 6, 1995, pp. 1-13.
- M. Perloff and K. Reiss, "Improvements to TCP Performance in High-Speed ATM Networks," Communications of ACM, February 1995, pp. 90-100.
- B.j. Ewy, et al, "TCP/ATM Experiences in the MAGIC Testbed,"

References (Cont)

- L. Kalampoukas and A. Varma, "Performance of TCP over Multi-Hop ATM Networks: A Comparative Study of ATM-Layer Congestion Control Schemes," Technical Report, UCSC-CRL-95-13, in ftp://ftp.cse.ucsc.edu/

Our Contributions & Papers

- R. Goyal et al, "Buffer Management for the GFR Service," ATM_Forum/98-0405, July 1998, http://www.cis.ohio-state.edu/~jain/atmf/a98-0405.htm
- R. Goyal et al, "GFR -- Providing Rate Guarantees with FIFO Buffers to TCP Traffic" ATM Forum/97-0831, Sep 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0831.htm
- R. Goyal et al, "Simulation Experiments with Guaranteed Frame Rate for TCP/IP traffic," ATM Forum/97-0607, July 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0607.htm

- R. Goyal et al, "UBR Buffer Requirements for TCP/IP over Satellite Networks," ATM Forum/97-0616, July 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0616.htm
- R. Goyal et al, "Selective Acknowledgements and UBR+ Drop Policies to Improve TCP/UBR Performance over Terrestrial and Satellite Networks, ATM Forum/97-0423, April 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0423.htm

- R. Goyal et al, "Guaranteed Rate for Improving TCP Performance on UBR+ over Terrestrial and Satellite Networks," ATM Forum/97-0424, April 1997, http://www.cis.ohio-state.edu/~jain/atmf/a97-0424.htm
- R. Goyal et al, "Further Results on UBR+: Effect of Fast Retransmit and Recovery," ATM Forum/96-1761, December 1996, http://www.cis.ohio-state.edu/~jain/atmf/a96-1761.htm

- R. Goyal et al, "Performance of TCP over UBR+," ATM Forum/96-1269, October 1996, http://www.cis.ohio-state.edu/~jain/atmf/a96-1269.htm
- R. Jain et al, " Performance of TCP over UBR and buffer requirements," ATM Forum/96-0518, April 1996, http://www.cis.ohio-state.edu/~jain/atmf/af_ubr22.htm
- R. Jain et al, "TCP/IP over UBR," ATM Forum/96-0179, February 1996, http://www.cis.ohio-state.edu/~jain/atmf/af9602-3.htm

- R. Goyal et al, "Improving the Performance of TCP over the ATM-UBR service", To appear in Computer Communications, http://www.cis.ohio-state.edu/~jain/papers/cc.htm
- R. Goyal et al, "Design Issues for providing Minimum Rate Guarantees to the ATM Unspecified Bit Rate Service", Proceedings of ATM'98, May 1998, http://www.cis.ohio-state.edu/~jain/papers/atm98.htm
- R. Goyal et al, "Providing Rate Guarantees to TCP over the ATM GFR Service," Proceedings of LCN'98, October 1998, http://www.cis.ohio-state.edu/~jain/papers/atm98.htm

- S. Kota et al, "Satellite ATM Network Architectural Considerations and TCP/IP Performance", Proceedings of the 3rd Ka Band Utilization Converence, Italy, 1997, http://www.cis.ohio-state.edu/~jain/papers/kaband.htm
- R. Goyal et al, "TCP Selective Acknowledgments and UBR Drop Policies to Improve ATM-UBR Performance over Terrestrial and Satellite Networks", Proc. ICCCN97, Las Vegas, September 1997, http://www.cis.ohio-state.edu/~jain/papers/ic3n97.htm

- R. Goyal et al, "UBR+: Improving Performance of TCP over ATM-UBR service," Proc. ICC'97, Montreal, June 1997, http://www.cis.ohio-state.edu/~jain/papers/icc97.htm
- S. Kalyanaraman et al, "Performance and Buffering Requirements of Internet Protocols over ATM ABR and UBR Services," IEEE Communications Magazine, June 1998, http://www.cis.ohio-state.edu/~jain/papers/ieee-mag.htm

- S. Kalyanaraman et al, "Performance of TCP/IP Using ATM ABR and UBR Services over Satellite Networks," Proc. IEEE Communication Society Workshop on Computer-Aided Modeling, Analysis and Design of Communication Links and Networks, Mclean, VA, October 20, 1996, http://www.cis.ohio-state.edu/~jain/papers/satellit.htm