Which Service for TCP/IP Traffic on ATM: ABR or UBR?

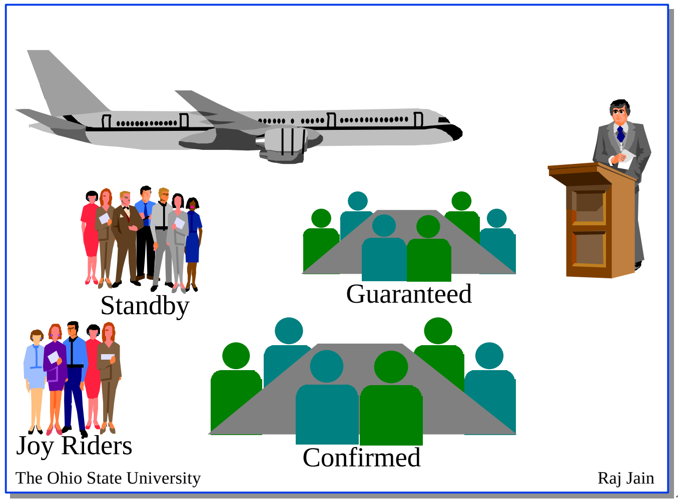

Raj Jain

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

- Service classes in ATM
- Seven Facts about TCP
- Performance on ABR
- Performance on UBR
- ABR or UBR?

Service Classes

- ABR (Available bit rate): Follows feedback
  Network gives max throughput with minimum loss.
- UBR (Unspecified bit rate):
  User sends whenever it wants. No feedback. No guarantee. Cells may be dropped during congestion.
- CBR (Constant bit rate): User declares required rate. Throughput, delay and delay variation guaranteed.
- VBR (Variable bit rate): Declare avg and max rate.
  - rt-VBR (Real-time): Conferencing.
    Max delay and delay variation guaranteed.
  - nrt-VBR (non-real time): Stored video. Mean delay

ABR: The Explicit Rate Scheme

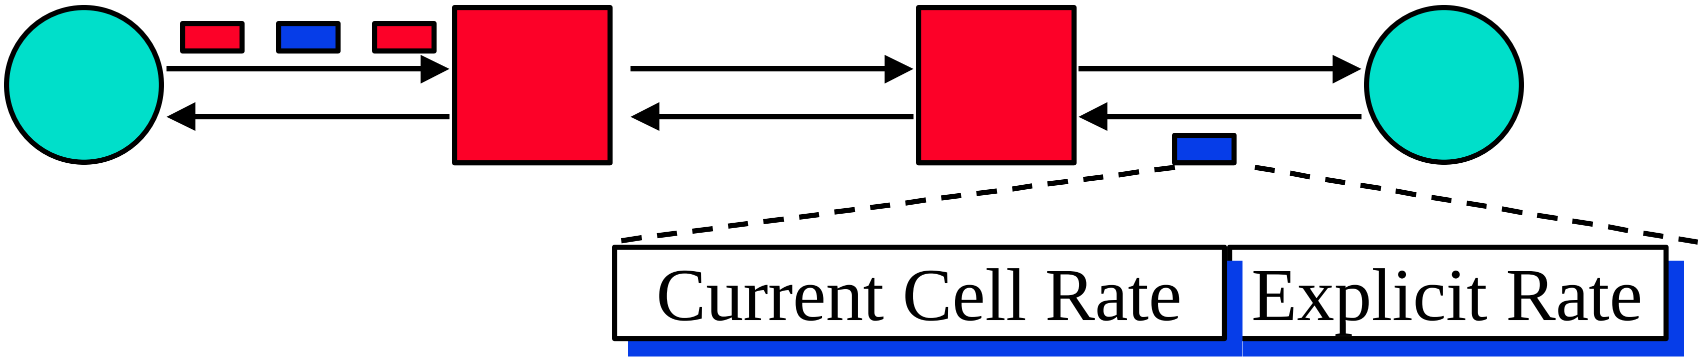

- Sources send one RM cell every n cells
- The RM cells contain "Explicit rate"
- Destination returns the RM cell to the source
- The switches adjust the rate down
- Source adjusts to the specified rate
- Interoperates with all switch algorithms.

UBR

- No specifications on switch or source behavior
- The sources send at peak rate.
- Switches drop cells if buffers full.
- Switch behavior similar to current routers.
- Intelligent protocols can use loss as implicit congestion indication and reduced load
- TCP is one such intelligent protocol
- UBR+:
  - Early packet discard (EPD)
  - EPD + Selective discard (Fair buffer allocation)

Observations about TCP

- TCP successfully avoids congestion collapse.
- TCP can automatically fill any available capacity.
- TCP performs best when there is NO packet loss. Even a single loss can reduce throughput considerably.
- Slow start limits the packet loss but loses time. You may not lose too many packets but you loose time.
- Fast retransmit/recovery helps in isolated losses but not in bursty losses.
- Bursty losses cause more degradation
- Timer granularity is the key in determining time lost.

n Source + VBR Configuration

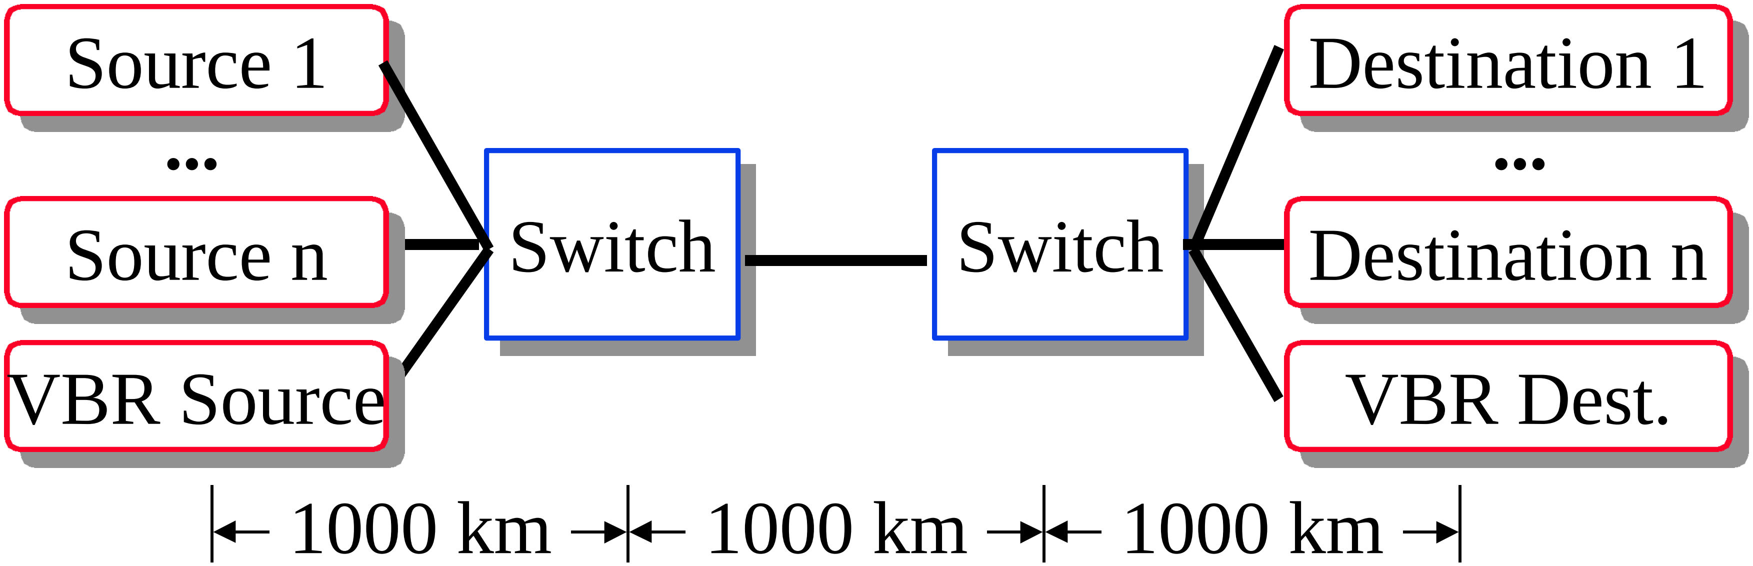

- All links 155 Mbps
- If VBR background , 100 ms on (80%), 100 ms off, start at t = 2ms
- All traffic unidirectional, Large file transfer.

Simulation Results: Summary

# srcs	TBE	Buffer Size	T1	T2	T3	T4	T5	Throughput	% of Max	CLR.
2	128	256	3.1	3.1				6.2	10.6	1.2
2	128	1024	10.5	4.1				14.6	24.9	2.0
2	512	1024	5.7	5.9				11.6	19.8	2.7
2	512	2048	8.0	8.0				16.0	27.4	1.0
5	128	640	1.5	1.4	3.0	1.6	1.6	9.1	15.6	4.8
5	128	1280	2.7	2.4	2.6	2.5	2.6	12.8	21.8	1.0
5	512	2560	4.0	4.0	4.0	3.9	4.1	19.9	34.1	0.3
5	512	5720	11.7	11.8	11.6	11.8	11.6	58.4	100.0	0.0

- CLR has high variance
- CLR does not reflect performance. Higher CLR does not necessarily mean lower throughput
- CLR and throughput are one order of magnitude apart

Observations About ABR

- ABR performance depends upon the switch algorithm. Following statements are based on our *ERICA* algorithm. (For ERICA, see http://www.cis.ohio-state.edu/~jain/)
- No cell loss for *TCP* if switch has Buffers = 4 × RTT.
- No loss for any number of TCP sources w 4 × RTT buffers.
- No loss even with VBR.
  W/o VBR, 3×RTT buffers will do.
- Under many circumstances, 1× RTT buffers may do.
- Required buffers depend upon RTT, feedback delay, switch parameters, and characteristics of VBR.

UBR Results

Buffer Size	Receiver Window	EPD	D1	D2	D3	D4	D5	Effici- ency	Fair- ness
12000	600000	N	16.9	17.9	17.9	19.2	17.4	71%	1.00
12000	1800000	N	16.9	17.9	17.9	19.2	17.4	74%	1.00
36000	600000	N	21.3	21.3	21.3	21.3	21.2	85%	1.00
36000	1800000	N	27.2	28.1	11.0	12.1	27.9	85%	0.88
12000	600000	Y	31.8	15.9	15.3	15.8	15.4	75%	0.89
12000	1800000	Y	31.8	15.9	15.3	15.8	15.4	75%	0.89
36000	600000	Y	21.1	21.1	21.7	21.2	20.8	85%	1.00
36000	1800000	Y	13.3	31.9	14.5	14.5	31.7	85%	0.86
12000	120000	N/A	24.0	24.1	24.0	24.1	24.0	96%	1.00
36000	360000	N/A	23.9	24.2	23.9	24.2	23.9	96%	1.00

- For full throughput: Need buffers = Σ receive windows
- EPD improves throughput but not fairness.

Observations about UBR

- No loss for TCP if Buffers = Σ TCP receiver window
- Required buffering depends upon number of sources.
- Receiver window $\geq$ RTT for full throughput
- Unfairness in many cases.
- Fairness can be improved by proper buffer allocation, selective drop policies, and scheduling.
- No starvation $\Rightarrow$ Lower throughput shows up as increased file transfer times = Lower capacity

Conclusion: UBR may be OK for: LAN, w/o VBR, Small number of sources, AND cheap implementation

ABR vs UBR

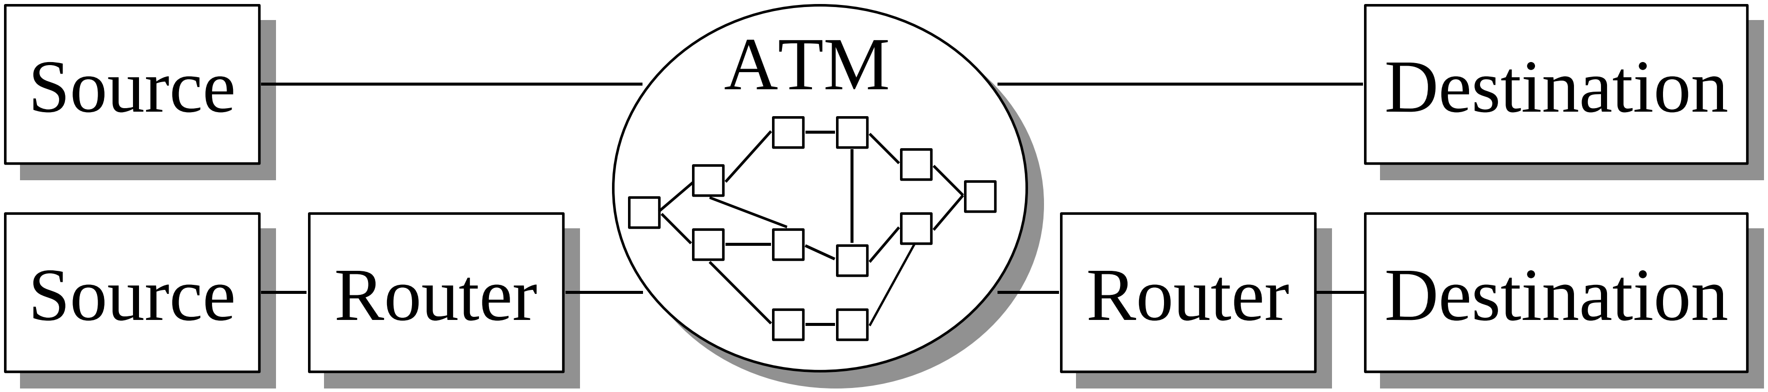

ABR	UBR
Queue in the source	Queue in the network
Pushes congestion to edges	No backpressure
Good if end-to-end ATM	Same end-to-end or backbone
Fair	Generally unfair
Good for the provider	Simple for user

Summary

- Packet loss results in a significant degradation in TCP throughput. For best throughput, TCP needs no loss.
- With enough buffers, ABR may guarantee zero loss for any number of TCP sources.
- Performance of ABR depends on the switch algorithm
- For zero loss, UBR need buffers = Σ receiver windows

Our Papers/Contributions

All our past ATM forum contributions, papers and presentations can be obtained on-line at http://www.cis.ohio-state.edu/~jain/

- S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal and S. Kim, "Performance and Buffering Requirements of Internet Protocols over ATM ABR and UBR Services," Submitted to IEEE Communications Magazine, September 1, 1996.
- S. Kalyanaraman, R. Jain, R. Goyal, S. Fahmy and S. Kim, ``Performance of TCP/IP Using ATM ABR and UBR Services over Satellite Networks," submitted to IEEE Communication Society Workshop on Computer-Aided Modeling, Analysis and Design of Communication Links and Networks, McLean, VA, October 20, 1996.

- S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal and S. Kim, "Buffer Requirements For TCP/IP Over ABR," Proc. IEEE ATM'96 Workshop, San Francisco, August 23-24, 1996.
- S. Kalyanaraman, R. Jain, S. Fahmy, R. Goyal, F. Lu and S. Srinidhi, ``Performance of TCP/IP over ABR," To appear Globecom'96, London, November 1996.
- R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal , "Buffer Requirements for TCP over ABR," ATM Forum/96-0517, April 1996.
- R. Jain, R. Goyal, S. Kalyanaraman, S. Fahmy, "Performance of TCP over UBR and buffer requirements," ATM Forum/96-0518, April 1996.
- R. Jain, S. Kalyanaraman, S. Fahmy, R. Goyal , "TBE and TCP/IP traffic," ATM Forum/96-0177, February 1996.
- ATM Forum/96-1172: ERICA Switch Algorithm: A Complete Description (August 1996)

TCP over ATM: 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.
- 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,"
- 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/