ABR or UBR?

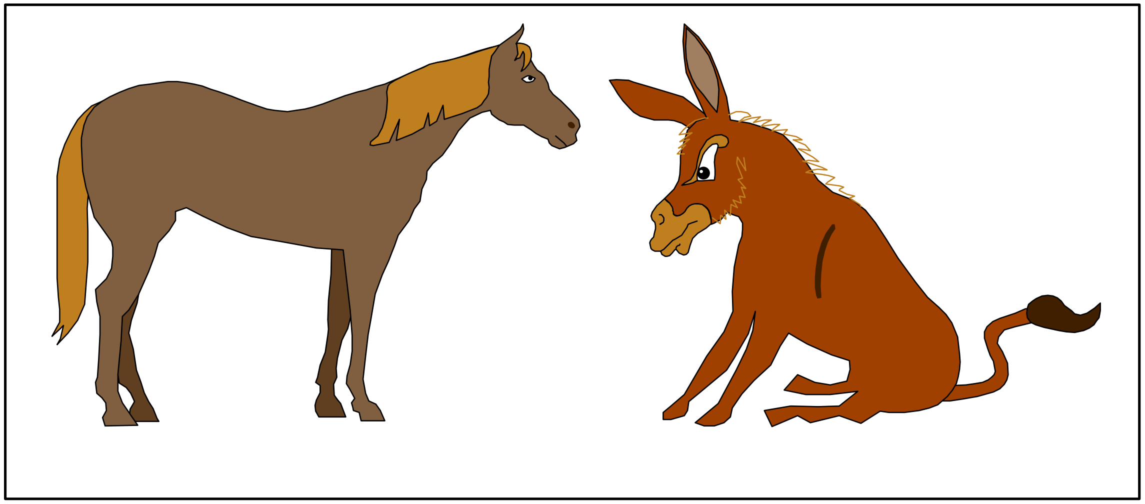

Raj Jain

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

- Seven facts about TCP
- Three facts about ATM
- Seven observations about ABR
- Seven observations about UBR

Our Quest

- TCP has window-based congestion control.
- ABR provides rate-based control,
  while UBR provides no control.
- Is TCP + ABR better than TCP + UBR?

Seven Facts 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 packet loss can reduce throughput considerably.
- Slow start limits the packet loss but loses considerable time.
  With TCP, you may not lose too many packets but you loose time.
- Bursty losses cause more throughput degradation than isolated losses.
- Fast retransmit/recovery helps in isolated losses but not in bursty losses.
- Timer granularity is the key parameter in determining time lost.

Three Facts about ATM

These apply to ABR as well as UBR:

- Cell loss rate (CLR) gives no indication of throughput loss.
  1% cell loss can cause 50% throughput loss.
  10% cell loss may result in only 10% throughput loss.
- Dropping all cells of a packet is better than dropping randomly (EPD).
- Never drop the EOM cell of a packet.
  It results in two packet losses.

Seven Observations About ABR

- ABR performance depends heavily upon the switch algorithm.
  Following statements are based on our *modified ERICA* switch algorithm.
  (For ERICA, see http://www.cis.ohio-state.edu/~jain/)
- Other key parameters: Round-trip Time, Number of sources, feedback delay from bottleneck.
- 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.
- Drop policies improve throughput but are not critical.

Seven Observations about UBR

- Switch queues may be as high as the sum of TCP windows
  No cell loss for TCP if Buffers = Σ TCP receiver window
- Required buffering depends upon the number of sources.
- TCP receiver window $\geq$ RTT for full throughput with 1 source.
- Unfairness in many cases.
- Fairness can be improved by proper buffer allocation, drop policies, and scheduling.
- Drop policies are more critical (than ABR) for good throughput
- 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 but not otherwise.

Summary

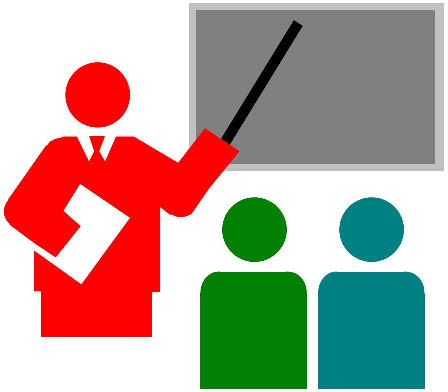

- Packet loss results in a significant degradation in TCP throughput. For best throughput, TCP needs no loss.
- ABR performance depends upon switch algorithm.
- With enough buffers, ABR may guarantee zero loss for any number of TCP sources. With UBR there is no such guarantee.
- TCP + ABR is better than TCP + UBR.
  But, UBR may be OK for low-end LANs.
- How much improvement with UBR+? Coming soon...

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
- Bursty losses are less damaging than scattered losses