96-0518R1 TCP over UBR and Its Buffer Requirements

Raj Jain, Rohit Goyal, Shiv Kalyanaraman, Sonia Fahmy

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

- Performance for
  - WAN
  - LANs
  - With and Without EPD

n Source Configuration

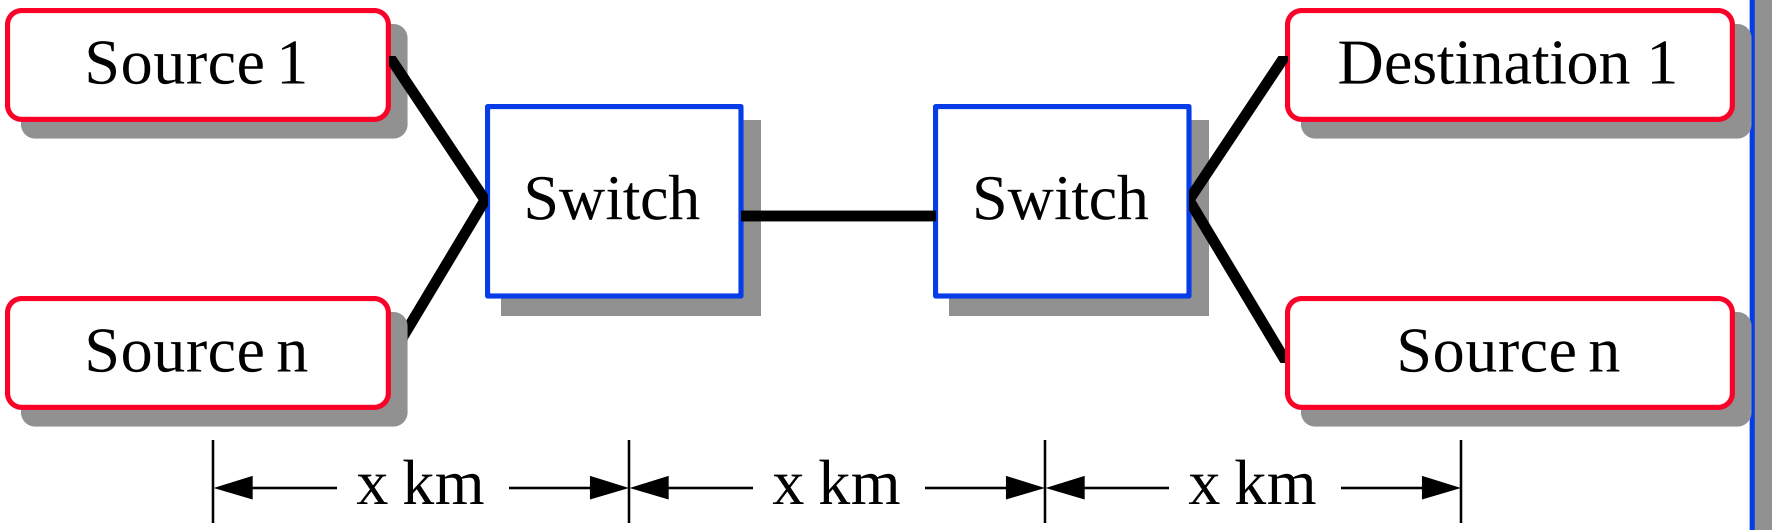

- All links 155 Mbps,
- WAN: 5 ms ⇒ RTT = 30 ms; LAN: 5 μs ⇒ RTT = 30 μs
- Unidirectional Infinite TCP sources.
  ⇒ TCP layer always has a packet to send if permitted by TCP window. Actual traffic on the network is bursty.
- No VBR

TCP/IP Parameters

- Maximum Segment Size = 512 bytes
- Timer granularity = 100 ms
- Early packet drop (EPD)
- No TCP processing time
- One-way delay = 15 ms = 291 kB
  Used window scaling option
- No delay ack timer
- Fast retransmit/recovery not completely experimented

Performance Metrics

- Efficiency = Sum of throughputs/Maximum possible throughput
  - Maximum Segment Size = 512 data
    = 512 data + 20 TCP + 20 IP + 8 LLC + 8 AAL5
    = 12 cells = 12*53 bytes = 636 bytes in ATM Layer
  - Maximum possible throughput = 512/636 = 80.5%
    = 125.2 Mbps on a 155.52 Mbps link
- Fairness = $\frac{(\Sigma x_i)^2}{n \Sigma x_i^2}$

  Where x_i = throughput of the ith TCP source

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 ≥ 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 ⇒ Lower throughput shows up as increased file transfer times = Lower capacity

Conclusion: UBR may be ok for: no VBR, Small number of sources, AND cheap implementation but not otherwise.

LAN Results

Buffer Size	Receiver Window	EPD	D1	D2	D3	D4	D5	Effici-ency	Fairness
1000	65535	N	1.8	1.9	1.8	1.8	18.1	20%	0.38
1000	32767	N	1.9	1.9	1.8	1.8	20.7	22%	0.36
500	65535	N	10.6	0.9	0.8	0.8	5.8	15%	0.48
500	32767	N	10.3	7.9	3.1	1.0	1.0	19%	0.60
1000	65535	Y	21.1	2.4	1.7	6.0	6.0	30%	0.52
1000	32767	Y	9.3	1.9	20.5	1.4	1.3	27%	0.46
500	65535	Y	3.1	0.8	8.1	8.1	10.7	25%	0.74
500	32767	Y	0.5	13.1	0.6	0.6	15.3	24%	0.44
10000	65535	N/A	25.0	25.0	25.0	25.0	25.0	100%	1.00
1000	8192	N/A	25.0	25.0	25.0	25.0	25.0	100%	1.00

- Low efficiency. High Unfairness.
  $\Rightarrow$ Do not use default (high) windows.
- EPD improves efficiency and fairness
- For full throughput: Need buffers = Σ receive windows

WAN 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

- Default windows are ideal for WANs.
  Fewer losses than LAN $\Rightarrow$ Better efficiency. Better fairness.
  $\Rightarrow$ EPD has less effect..
- For full throughput: Need buffers = Σ receive windows

Other Observations

- Setting EPD threshold as a fraction of buffer size is not useful. Better to set EPD Threshold = buffer size - n packets
- EPD improves fairness. But UBR+EPD is still unfair.

Summary

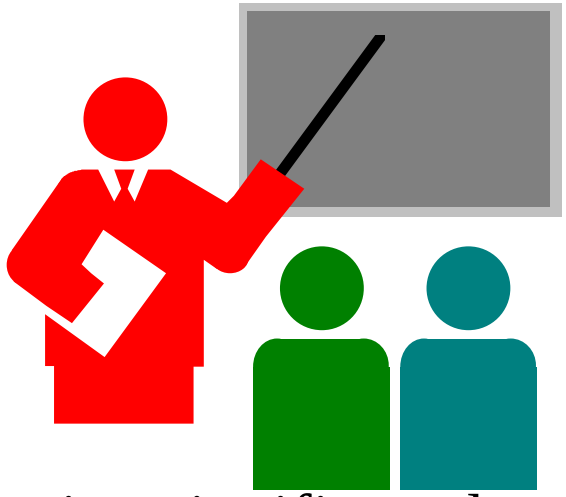

- Packet loss results in a significant degradation in TCP throughput. For best throughput, TCP needs no loss.
- For zero loss, need buffers = Σ receiver windows
- With enough buffers, ABR may guarantee zero loss for any number of TCP sources. With UBR there is no such guarantee.
- TCP + ABR is better than TCP + UBR.
  But, UBR may be OK for low-end products.

REFERENCES

- Raj Jain, S. Kalyanaraman, R. Goyal, S. Fahmy, F. Lu, and S. Srinidhi, "TBE and TCP/IP traffic," AF-TM 96-0177, February 1996.
- S. Kalyanaraman, Raj Jain, S. Fahmy, R. Goyal, F. Lu, and S. Srinidhi, "Performance of TCP/IP over ABR," Submitted to Globecom'96.
- Raj Jain, S. Kalyanaraman, R. Goyal, S. Fahmy, F. Lu, and S. Srinidhi, "Buffer requirements for TCP over ABR" AF-TM 96-0517, April 1996.

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