95-1661R1

TBE vs Queue Sizes

Raj Jain, Sonia Fahmy, Shiv Kalyanaraman, Rohit Goyal, Fang Lu
The Ohio State University

Saragur M. Srinidhi
Sterling Software and NASA Lewis Research Center

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

Two Bursty Sources Configuration

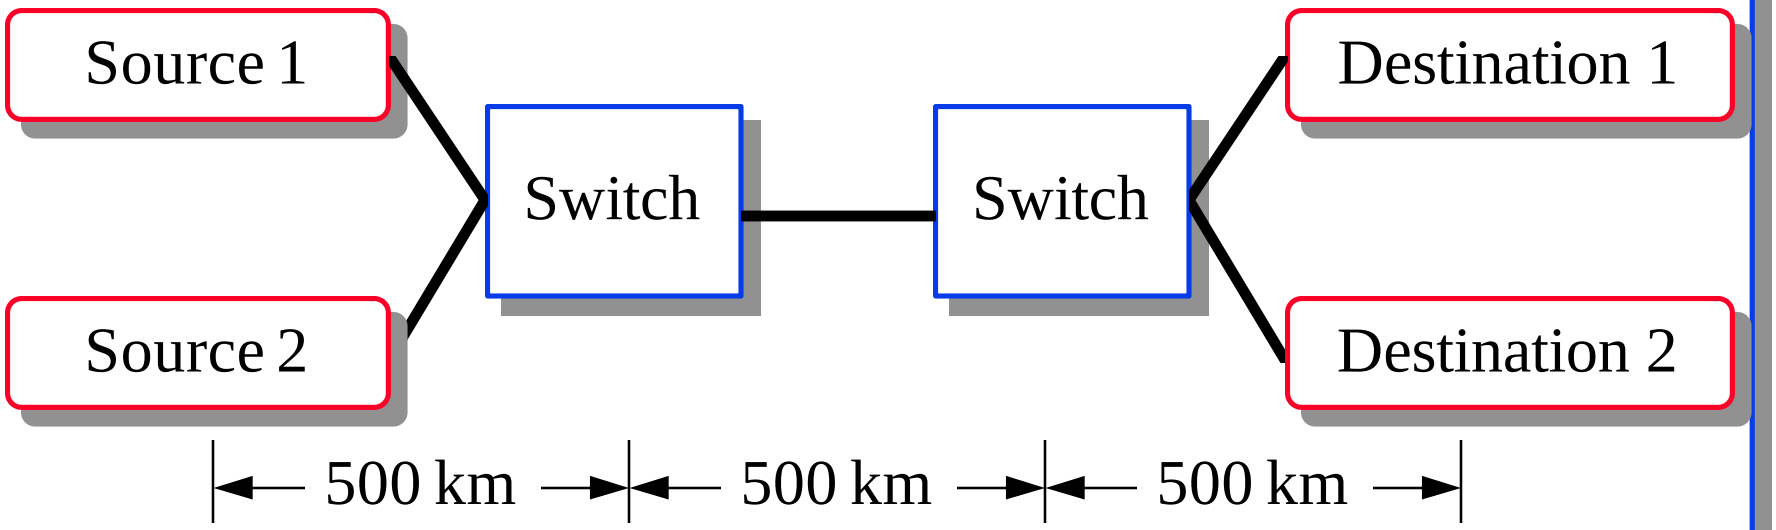

- All links 155 Mbps
- Both sources use closed-loop bursty request/response traffic
- All traffic bi-directional

Simulation Parameters

- Source: Parameters selected to maximize ACR
  TBE = {128, 512, 4096}
  CDF (XDF) = {0, 0.5}
  ICR = {155.2 Mbps, TBE/FRTT}
  CRM (Xrm) = Min{TBE/Nrm, PCR*FRTT/Nrm}
  TDF = 0, PNI = 1 $\Rightarrow$ Disable rule 5
  TOF = 2, PCR = 155.52 Mbps, MCR= 0, RIF (AIR) = 1,
  Nrm = 32, Mrm = 2, RDF = 1/512, Trm =100ms, TCR = 10 c/s
- Traffic: Bursty. Request size = 16 cells, Response size = 2000,
  Inter-cycle time = 10ms
- Switch: ERICA modified
  Target Utilization = 90%
  Averaging interval = min{30 cells, 200 μs}

Start-Up Transients

- Without rule 6, the queue length is close to the (# of sources -1)*burst size
- With rule 6, the queue length is close to TBE
  Actually, (# of sources -1)*Min{TBE, Burst Size}
- This applies if the network becomes empty after each cycle.
- Close-loop model $\Rightarrow$ Previous response has been received before new request is generated.

Two Source Configuration

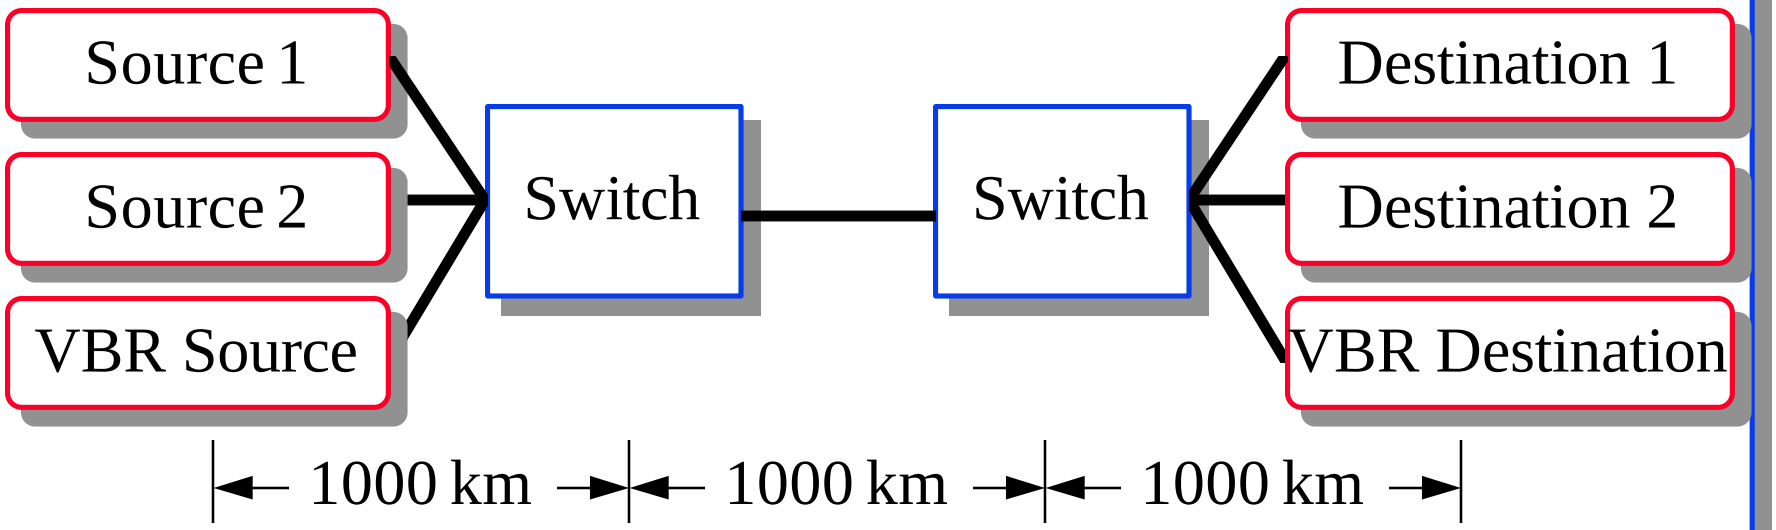

- All links 155 Mbps
- VBR background traffic, 20 ms on, 20ms off, start at t=2ms
- All traffic bi-directional, infinite ABR sources
- Parameters: Same as before. TBE = {128, 512, 1024}

Observations

- With TBE = 128 and two infinite sources (with VBR background) on a WAN, the ABR queue length can grow as large as 2500
- Rule 6 triggers only initially due to low ICR (high inter-RM cell time)
- Afterwards rule 6 triggers rarely and is overridden by the subsequent explicit feedbacks.
- Since the reverse flow is not stopped completely. The forward flow continues and keeps filling the queues.
- TBE does not significantly affect the maximum queue length.
- Whenever rate changes from low rate period to high rate, rule 6 triggers unnecessarily, e.g., after the low ICR

Analytical Explanation

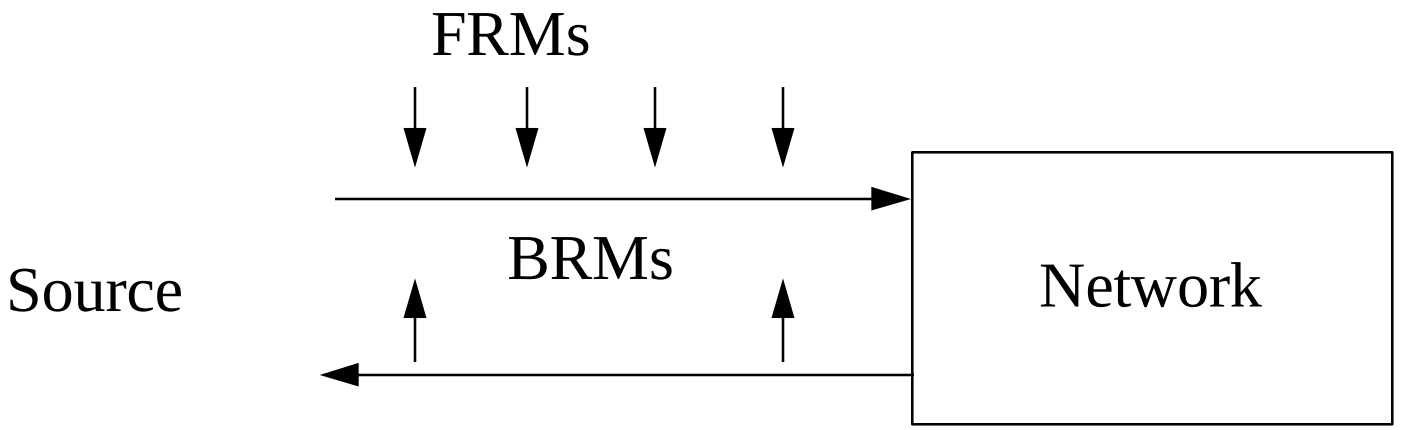

- Let R = FRM/sec
  r = BRM/sec
  R/r FRMs sent per BRM received
- Rule 6 triggered if and only if R/r ≥ Crm
- Maximum Q = fn(PCR, Target Utilization, VBR Amplitude)*Feedback delay
- Note maximum Q does not depend upon TBE

Finite Buffers

- All cells in excess of the buffers capacity are lost
- The effect of lost cells depends upon the frame size, higher layer protocol
- The exact number of cells lost varies slightly depending upon the RM cells lost, buffer management policy, bandwidth allocation policy (switch algorithm)
- FIFO queueing with cell loss
  Not all VCs have the same number of cells lost
  $\Rightarrow$ Unfairness

CLR with Finite Buffers

TBE	Buffer Size	CLR	
		W/O Rule 6	W Rule 6
128	256	35.63%	27.53%
	512	31.47%	22.41%
	1024	27.79%	18.81%
	2048	10.73%	4.78%
512	1024	27.15%	26.33%
1024	2048	9.61%	9.61%

Motion

- Add the following text on page 72 of R9
  "TBE limits the queue length only during initial startup and cannot be relied upon during the close-loop operation phase of a connection. During this latter phase, the contribution of a VC to the queue at a switch can be more than its TBE. The buffer usage at switch can be more than the sum of TBEs allocated to active VCs."

Motion

- Section 5.10.3.1 on page 43 of 95-0013R9
  Replace
  "Transient Buffer Exposure, TBE, is the negotiated number of cells that the network would like to limit the source to sending during idle startup period, before the first RM cell returns"
  by
  "Transient Buffer Exposure, TBE, is the negotiated number of cells that the network would like to limit the source to sending during initial startup period, before the first RM cell returns"

Summary

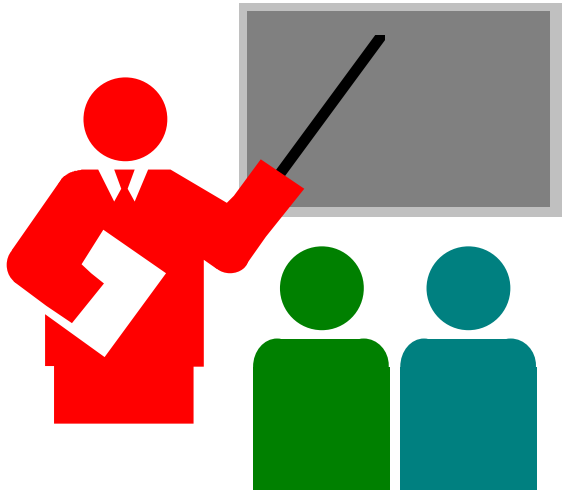

- Rule 6 limits cells mainly when the BRM flow is broken or stopped
- Low rate BRM flow ⇒ Queues can grow
- There is little relationship between TBE and queue length
- How does a switch allocate TBE?
  Not based on buffer availability.