95-1343R1

Straw-Vote Comments on TM4.0 R8

Raj Jain, Shiv Kalyanaraman, Sonia Fahmy, 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/

- XRM Range
- ICR Initialization
- XRM Initialization
- Rule 5
- Minor Comments

XRM Range: History

On Wednesday, August 9, 1995: We presented results of two months of analysis.

- Motion was passed with many to 2 votes:
  XRM is an integer whose size is implementation dependent.
- Motion that was specifically not passed:
  *XRM is an integer.**
  **Note: 8 bits are sufficient 24 bits are required for long-delay path.*

On Thursday, August 10, 1995: Motion was reversed without any analysis.

Six Reasons to Remove the Note

- The first format does not recommend any particular size. Eight-bit or 24-bit both are allowed. The second format tries to recommend some particular implementation.
- XRM is an internal parameter and unless we decide to signal it, there is no reason to recommend or enforce any particular value that implementors can choose themselves. Leave it to implementors. They are not dumb.
- Eight-bit XRM is inconsistent:
  XRM = CIF/Nrm
  Max XRM = (Max CIF)/(Min Nrm) = $2^{24}/2 = 2^{23}$

- Eight-bit XRM is limiting the distance and bandwidth:
  $XRM = PCR \times RTT/Nrm$
  Max $XRM = (\text{Max } PCR) \times (\text{Max } RTT)/(\text{Min } Nrm)$
  Eight-bit XRM is not sufficient even on 155 Mbps over North-American continent, NRM of 32 and no load:
  XRM = 155 Mbps × 30 ms one way × 2 /32
  = 366 c/ms × 60 ms /32
  = 682
- Users may want to use ATM networks (and your products) for longer distances, higher speeds, different Nrms, and higher loads

Motion

- On page 51 of ATMF 95-0013R8, Xrm row, column 3, delete:

  Range: 0 to 255 (0 to 1677215 - See Note 2)

- Re-place the following text :

  XRM is an integer. Its size is implementation dependent.

- On page 52, Delete:

Note 2: For Xrm in large delay-bandwidth product situations a 24-bit implementation may be preferable. Otherwise an 8-bit implementation is generally sufficient.

ICR Initialization

- ICR is the most influential parameter = Influential Cell Rate
- It controls the life-time performance for bursty sources
- It controls the performance when available bandwidth varies due to VBR
- ICR = Min{PCR, a*CIF/RTT}
- RTT is a random number.
  95-percentiles and monthly averages are either too complex or meaningless
- On LAN's: RTT low $\Rightarrow$ ICR = PCR
  Network has practically no direct control over ICR
- Your VC's performance depends upon the time it was initialized $\Rightarrow$ Start all you VC's at 4AM.

Three Source Configuration

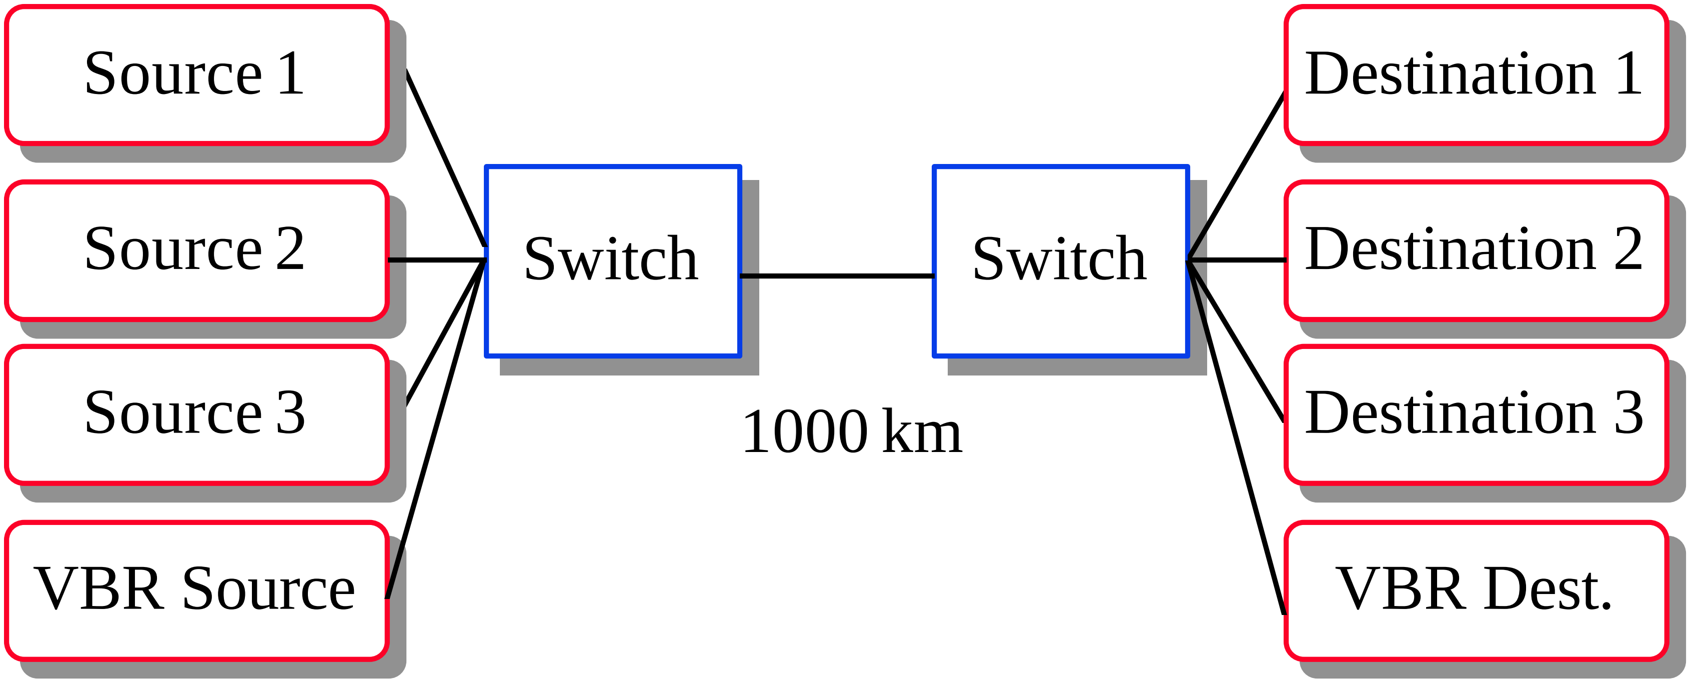

- All links 155 Mbps
- Three ABR sources started at different RTTs.

Simulation Parameters

- Source: Parameters selected to maximize ACR
  Nrm = 32
  AIRF = 1 $\Rightarrow$ AIR = PCR/Nrm $\Rightarrow$ ACR is not limited by AIR
  RDF= 512 cells
  {TDFF, PNI} = {1/8, 0} or {0, 1} $\Rightarrow$ Rule 5 on or off
  CIF = 512
  RTT = Propagation delay multipliers of 1, 10 or 100
  XDF = 1/2
- Traffic: Bi-directional
- Switch:
  Target Utilization = 90%
  Averaging interval = min{30 cells, 200 μs}

Motion

- Remove:
  ICR = Min{PCR, a*CIF/RTT}
- ICR is directly negotiated between the source and the network.

XRM Initialization

$XRM = Min\{CIF/Nrm, PCR * RTT/Nrm\}$

- RTT is a random number
- Birth time conditions affect the life-time performance. In this case, it is better to be born during bad times. Birth at good time ⇒ Low RTT ⇒ 2nd term smaller ⇒ XRM Low ⇒ Easily triggered during congestion.
- On LANs, RTT is low.
  At 155 Mbps, 3 km LAN, Nrm = 32
  XRM = 0.366 cells/μs * 15μs /32=0.17
  ⇒ XRM = 0 or 1
  XRM triggered every 32 cells under heavy load

- The formula is based on the assumption that
  Optimal CIF = PCR * RTT
  While the formula is correct, the RTT in the CIF formula should not include queueing delay. It should be pure propagation delay.

Propagation Vs Queueing Delay

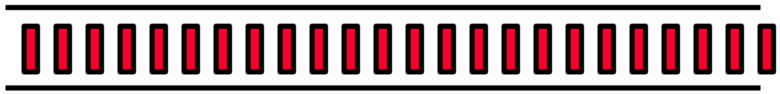

- To optimally fill the pipe, number of cells in flight is proportional to the pipe size (propagation delay).

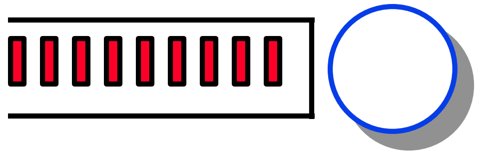

- To optimally fill the queue, number of cells is exactly one. Every extra cell is adding to the CTD without increasing the throughput
- Longer queues $\neq$ Put more of your cells in the queue. If done, repeatedly, the system will become unstable.
- **Conclusion**: The RTT in the formula should be propagation delay only.

Motion

- The RTT in XRM formula should be propagation delay only.

Minor Comments

- $XRM = Min\{CIF/Nrm, PCR * RTT/Nrm\}$
  The right hand side is a floating point number. Should it be rounded up or down?

Rule 5

- Designed to avoid "ACR retention"
  Also known as "Use it or loose it" policy
  Source rate much below ACR $\Rightarrow$ Decrease ACR
- Source Rate = Average rate over the last Nrm cells
  = Nrm/T, where T = time since the last RM sent
- Undesirable effects:
  - Unnecessarily triggered on network directed ACR increases
  - Current formula and parameter values $\Rightarrow$ ACR $\leftarrow$ ICR
  - Gaps between bursts $\Rightarrow$ Rule 5 easily triggered
  - The decrease is proportional to T $\propto$ 1/source rate

- May not trigger for sources with ACR retention (high ICR) and may trigger for sources without ACR retention problem (normal increase) :
  - The rule is ineffective for VCs born during good times or high CIF.
    ICR = Min{PCR, a*CIF/RTT)
    Low RTT at birth time $\Rightarrow$ High ICR = PCR
    These VC's are automatically exempt from Rule 5.
    Similarly VC's allowed large CIF are exempt from Rule 5.
  - The rule may be repeatedly triggered unnecessarily:

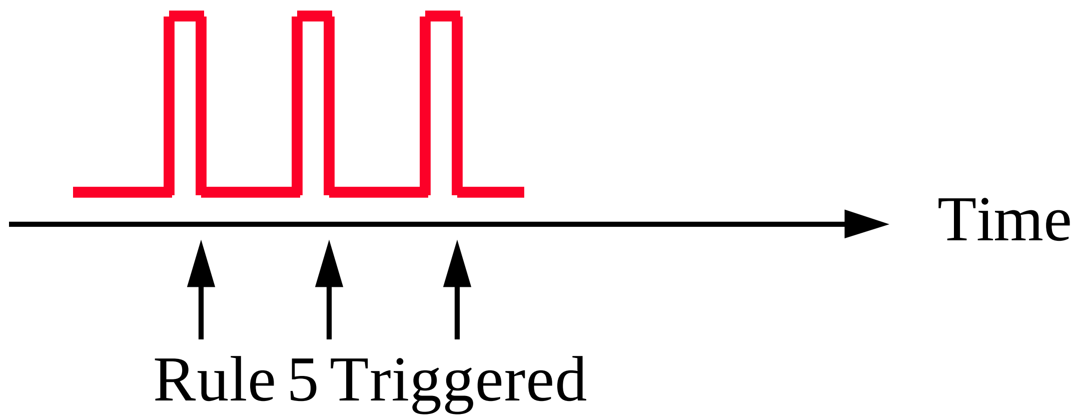

Summary

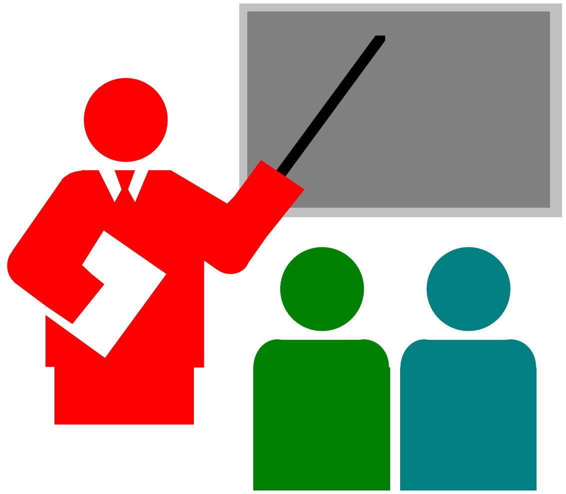

- XRM = 256 is insufficient even for North-America
- ICR should be directly negotiated
- Round trip delay with queueing is highly random
  Network performance is unpredictable
  Remove round-trip from ICR and XRM formulae.
- Rule 5 should be fixed.