95-0973R1

Out-of-Rate RM Cell Issues and Effect of Trm, TOF, and TCR on Low Rate Sources

Raj Jain, Shiv Kalyanaraman, Sonia Fahmy, Fang Lu
Department of Computer and Information Sciences

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

- Transient response time
  - Time to rise from "Low Rate" to "High Rate"
- Effect of Trm
- Effect of TOF, ICR
- Effect of TCR
- Corrections to the behaviors and pseudocode

Trm

- Source Rule 3 (simply stated):
  Send one FRM cell after every Nrm cells or at least Trm ms and Mrm cells
- In all base vectors Trm = 100 ms
- At low rate, every 3rd cell is an FRM cell.
  At high rate, every 32nd cell is an FRM cell.
  Higher Trm ⇒ Less overhead
- Sources may get a low rate due to
  - Heavy VBR traffic,
  - Large # of ABR sources, or
  - Low bottleneck link speed (T1 links)

Effect of Trm

- Trm allows low rate sources to sense the network state more frequently than normal
- When the bandwidth becomes available, network may not be able to allocate the bandwidth at all until it sees an RM cell.
- Network may allocate the bandwidth unfairly if all active sources are not seen
- Lower Trm
  $\Rightarrow$ Lesser time between RM cells
  $\Rightarrow$ Faster transient response time
- Choice of Trm also depends upon link speed (OC-12)

Traffic Pattern: VBR + ABR

- Actual VBR cells are generated, queued, and share the link and switch resources
- VBR gets a preferential treatment
  ABR gets only left-overs

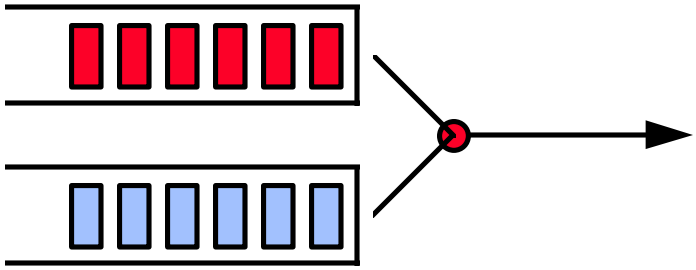

- Bi-directional Traffic

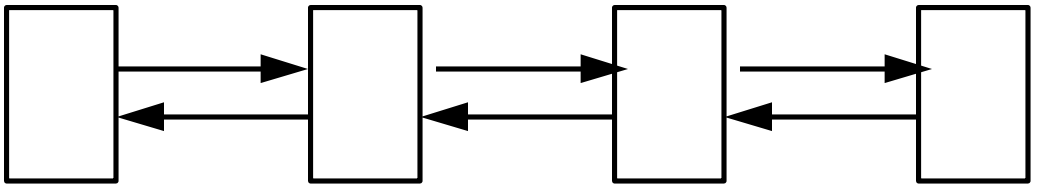

A Simple VBR Model

- On for x ms and off for y ms
- When on, VBR uses up C_{vbr} bandwidth
- In practice, x, y, C_{vbr} are random variables. We assumed constants.

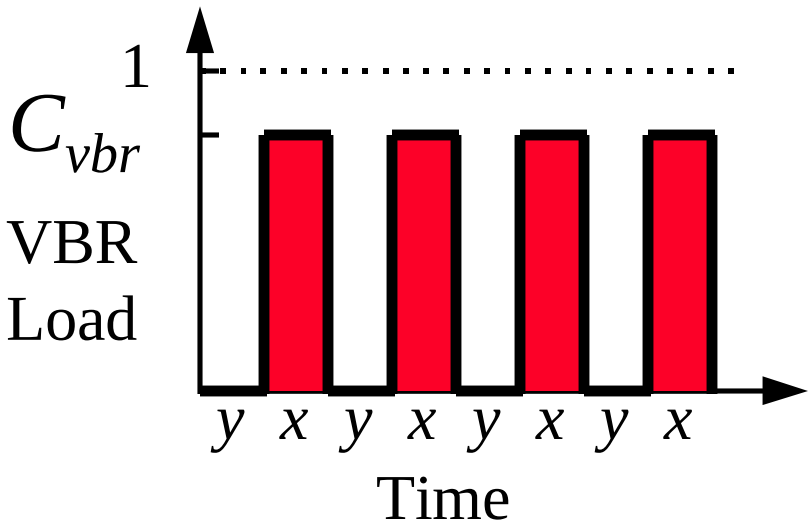

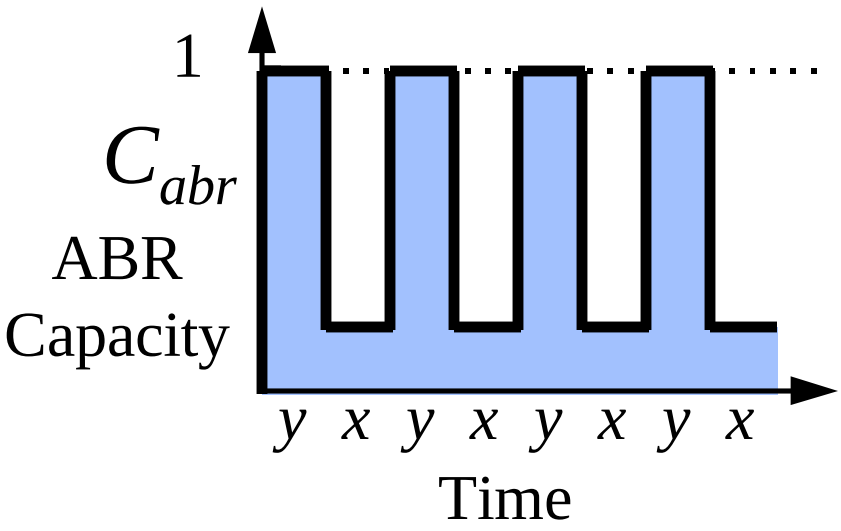

Simulation Parameters

- Source: Parameters selected for fast response
  Nrm = 32, RDF = 256 cells, TOF = 2, Xrm = 32,
  XDF = 1/16, TCR = 10 cps , ICR = PCR/20
  Trm = 1, 10, 100 ms
  AIRF = 1 $\Rightarrow$ Increases are not limited by AIR
  TDFF = 0 $\Rightarrow$ TOF decreases disabled
- Traffic: ABR: Infinite source, Bi-directional
  VBR: 20 ms off, 20 ms on, 89%, Bi-directional
  VBR starts at 2 ms $\Rightarrow$ On 2-22, 42-62, 82-102, ...
- Switch:
  Target Utilization = 90%
  Averaging interval = min{100 cells, 1 ms}

Two-Source Configuration

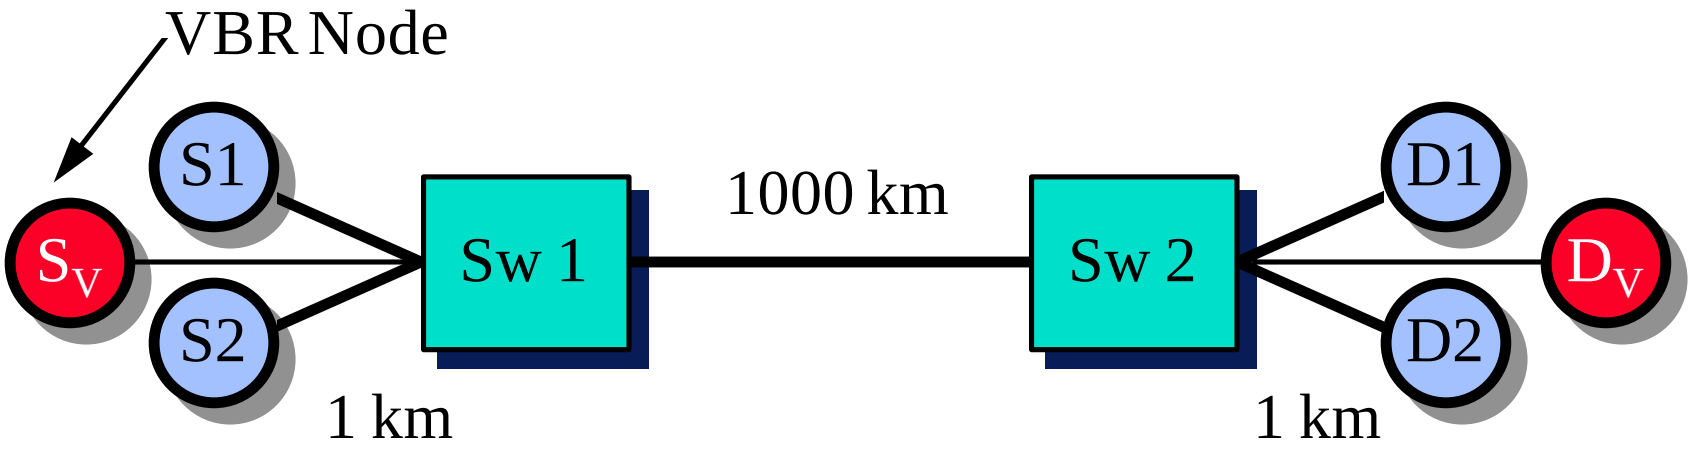

- All links 155.52 Mbps
- ABR sources go down to 0.8 Mbps when VBR comes on and go up to 70 Mbps when VBR goes away.
- Goal: Measure rise time for ABR sources

Simulation Results

- Available capacity may go unused for as long as 100 ms. (In our simulation, VBR comes back up every 20 ms and so unused time is 20 ms).
- Lower Trm $\Rightarrow$ More frequent RM cells
  $\Rightarrow$ Faster response

TOF

- Source Behavior 5a (Simply stated):
  If the time T since last FRM cell was sent is greater than TOF × Nrm×(1/ACR) then decrease by ACR × T × TDF down to ICR.
- TOF = 2 in all base vectors

Effect of TOF

- Effect 1: Rule triggered if the source rate is less than 1/TOF of ACR. Use it or loose it.
  $\Rightarrow$ ACR (and CCR) are close to source rate.
  (Some switch schemes are sensitive to this)
  $\Rightarrow$ Lower values of TOF are preferable.
- Effect 2: The rule is triggered on rate increase. Can't increase rate by more than a factor of TOF. Slows down ramp-up.
  $\Rightarrow$ Larger values of TOF provide faster transient response.

Simulation Parameters

All parameters same as that for Trm except:

- Source Parameters:
  - Trm = 100 ms
  - TOF = 2, 20, 100, 200
  - ICR = PCR, PCR/20, PCR/1000

Simulation Results

- With ICR = PCR, TOF has no effect.
  (TOF is effectively disabled for this ICR)
- With ICR = PCR/20, the sources push themselves back to ICR whenever network asks them to go up.
  Network and sources are at odds
  $\Rightarrow$ Oscillations
  - TOF too low
  - Formula gives large decreases
- Higher values of TOF help avoid these oscillations by triggering the decreases less often
- With ICR = PCR/1000, situation is worse.

Conclusion

- Higher values of TOF do provide better transient response for low ICR sources.
- The formula gives decreases that are too large

TCR

- Source Rule 11 (Simply stated):
  Out-of-rate FRM rate ≤ TCR
- TCR = 10 cps in all standard vectors

Effect of TCR

- Out-of-rate FRMs are not optional for sources (NICs). They are the only means to get out of ACR = 0 situation.
- Out-of-rate BRMs are not optional for destinations (NICs). They are the only means to control unidirectional ABR VCs.
- Use of out-of-rate RMs at non-zero ACR is optional and may improve transient response.

TCR Tradeoffs

- Higher TCR ⇒ More frequent feedback
  ⇒ More responsivity
- Lower TCR ⇒ Less out-of-rate cells
  ⇒ Less overhead

Simulation Parameters

- All parameters same as that for Trm except:
  - Source Parameters:
    Trm = 1 ms
    (To avoid confusion with TCR = 10 cps)
  - Switch Parameters:
    Averaging interval = min{30 cells, 1 ms}
  - Traffic Parameters:
    In one interval, we force VBR to use 90% of PCR
    $\Rightarrow$ Available bandwidth for ABR = 0
    $\Rightarrow$ Out-of-rate mechanism is triggered

Simulation Results

- The source, once stopped, is unable to use the bandwidth for 100 ms even when the bandwidth becomes available.
- This is not because there are no RM cells but because network feedback is ignored

Event Trace

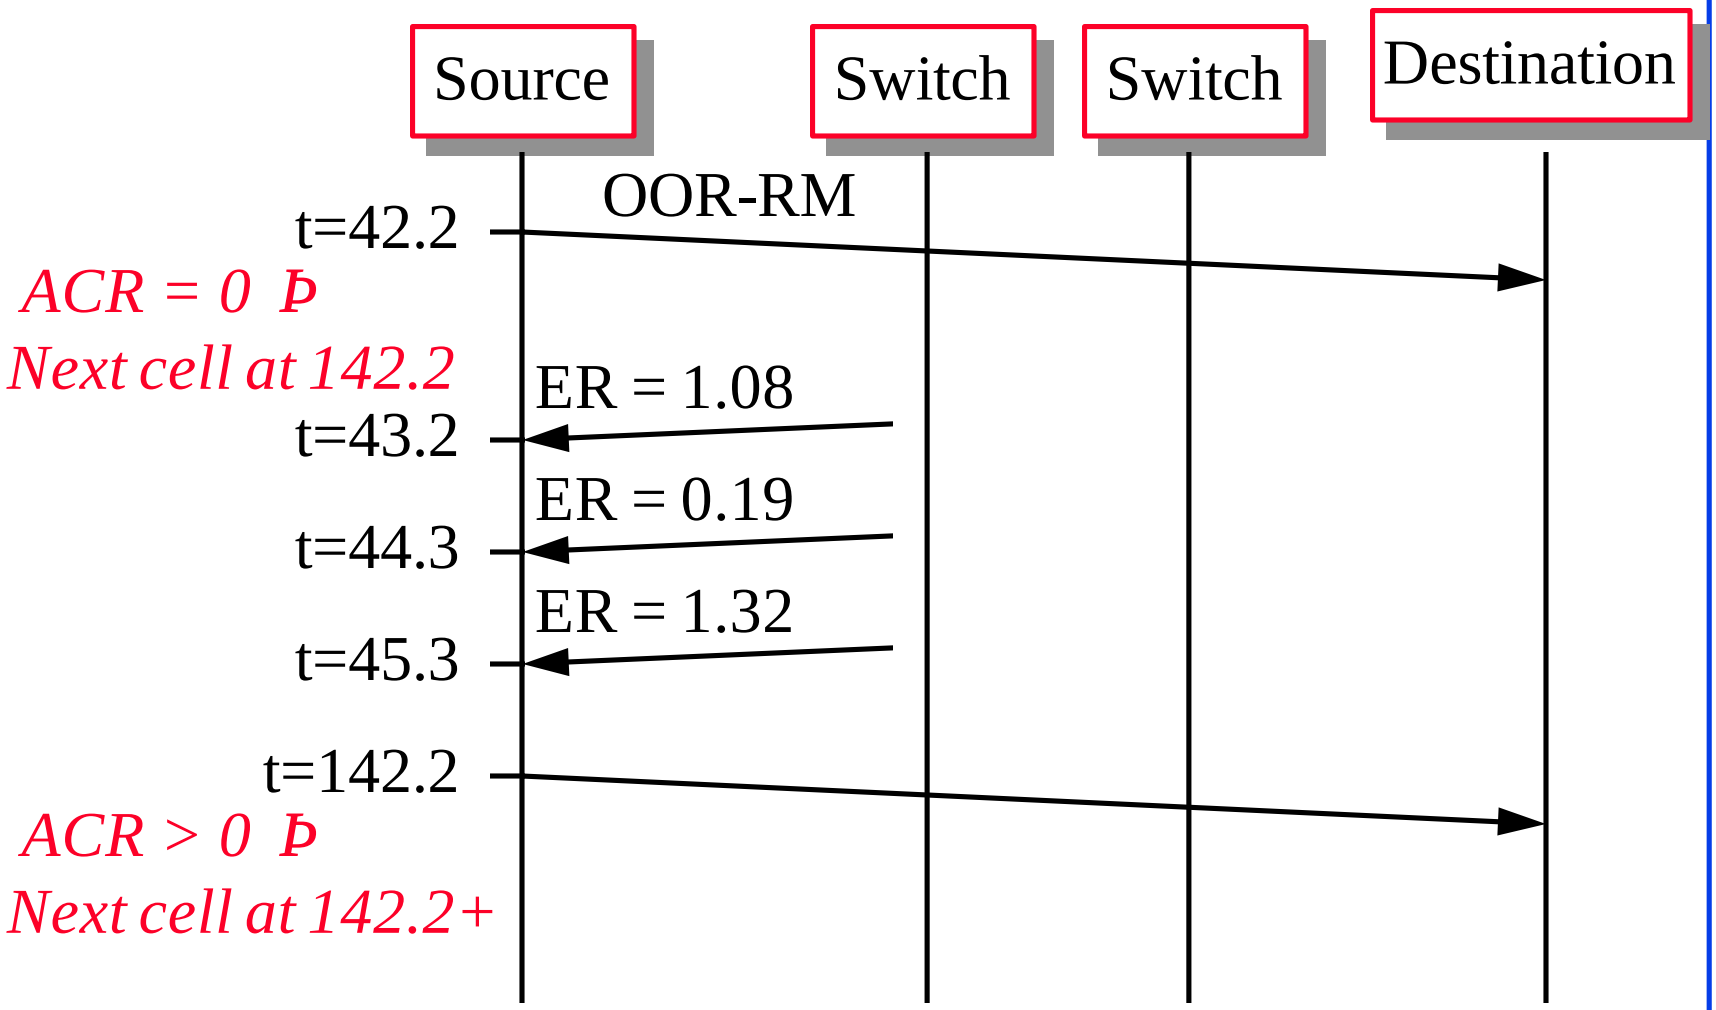
Source
Switch
Switch
Destination
OOR-RM
t=42.2
ACR = 0 Þ
Next cell at 142.2
ER = 1.08
t=43.2
ER = 0.19
t=44.3
ER = 1.32
t=45.3
t=142.2
ACR > 0 Þ
Next cell at 142.2+

Issue 1

- The text says nothing about how to schedule or reschedule the next cell
- The pseudocode chose to not reschedule cells on rate increases or decreases
- Four Possibilities: Reschedule if new rate will result in

Earlier Transmission	Later Transmission	
No	No	Pseudocode
No	Yes	Keep putting it off
Yes	*No*	*Recommended*
Yes	Yes	Keep putting it off

Recommendation 1

- Reschedule if the new ACR permits earlier scheduling (One late cell can make a big difference for the source)
- Do not reschedule if the new ACR will delay it further (One early cell can't hurt the network)

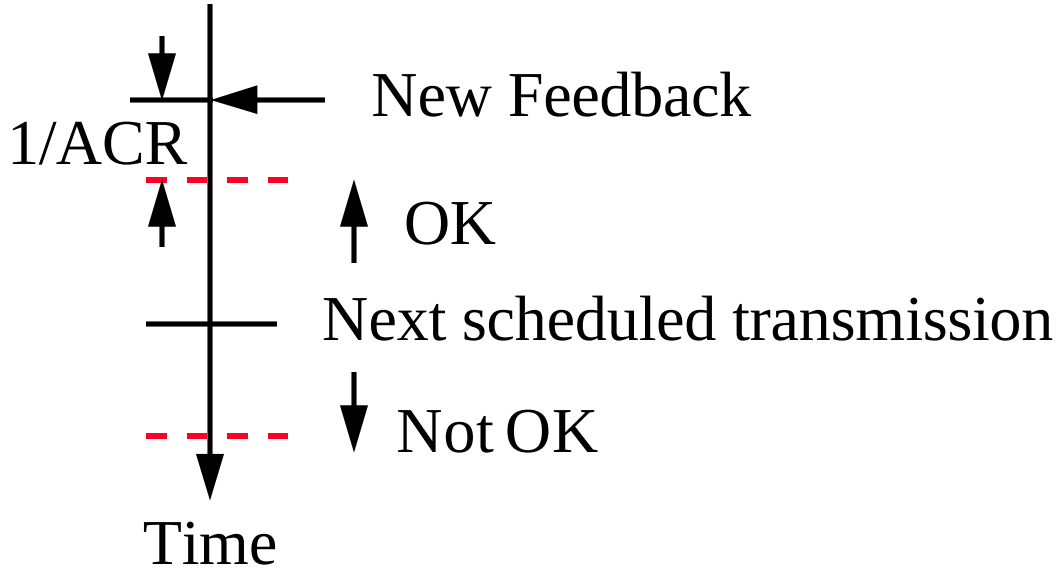

New Pseudocode

- Add the following to End-System Pseudocode - Receive (page 87 of ATMF95-0013R6):

  IF time_to_send > (now + 1/ACR)
  THEN time_to_send ← (now + 1/ACR)

Issue 2

- Source behavior requires interspersing FRM, BRMs and data even at low rate, the pseudocode does not implement it
- The pseudocode sends only OOR-RMs at ACR < TCR (no BRMs or data)

Issue 3

- Source behavior 11 permits OOR even if ACR > TCR
- Pseudocode does not implement it

Issue 4

- There are no guidelines on how to space out-of-rate RM cells. There are several possibilities:
  - Equally spaced 100 ms apart
  - 100 cells at 1ms then nothing for 9 s

  Are both choices valid?

Issue 5

- Is ACR = 0 legal?
- ATMF 95-0013R6 Section 5.10.3.1 (page 53) states that minimum ACR is 1 cps.
- The source or switch behavior say nothing about it
- Pseudocode does not impose this lower limit

Issue 6

- Source behavior 3a:
  The next in-rate cell shall be a forward RM-cell if and only if, since the last in-rate forward RM-cell was sent
  i) either at least Mrm in-rate cells have been sent
  or
  ii) at least Trm seconds have elapsed, or Nrm-1 in-rate cells have been sent.
- Pseudocode:
  If (Count >= Nrm) or
  ((count > Mrm) and (now >= last-RM + Trm)) ...

Recommendation

- Update source behavior:
  The next in-rate cell shall be a forward RM cell if and only if, since the last in-rate forward RM-cell was sent
  i) either at least Mrm in-rate cells have been sent and at least Trm seconds have elapsed, or
  ii) Nrm-1 in-rate cells have been sent.

Issue 7

- Do Nrm and Mrm include out-of-rate RM cells?
- Source behavior does not include them. Specifically asks for in-rate cells.
- Pseudocode includes out-of-rate cells in "count"

Summary

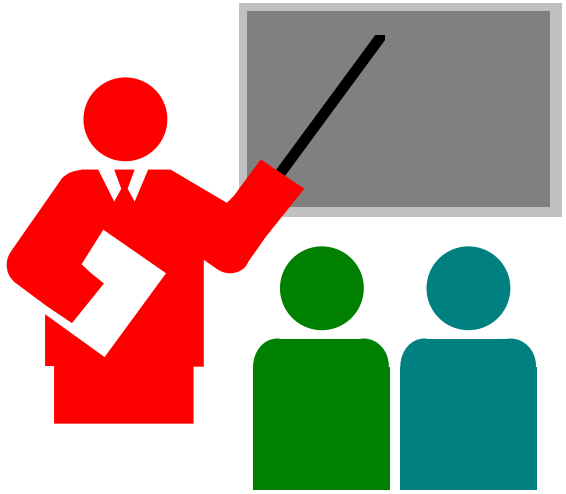

- Lower Trm gives better transient response
- TOF = 2 and low-ICR may cause oscillations
- OOR-RMs are not optional for NICs.
- Reschedule on rate increase.
- Numerous issues with low rate sources

Motion

- Add the following to the source behavior:

15. If ACR is increased according to Source Behavior #8, the source may use the new rate immediately even possibly rescheduling next scheduled transmission.

- Add the following to End-System Pseudocode - Receive (page 87 of ATMF95-0013R6):
  IF time_to_send > (now + 1/ACR)
  THEN time_to_send ← (now + 1/ACR)

Motion

- Update source behavior 3a:
  The next in-rate cell shall be a forward RM cell if and only if, since the last in-rate forward RM-cell was sent
  i) either at least Mrm in-rate cells have been sent and at least Trm seconds have elapsed, or
  ii) Nrm-1 in-rate cells have been sent.