Traffic Management in ATM Networks over Satellite Links

R. Goyal, R. Jain, M. Goyal, S. Fahmy, and B. Vandalore

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

TIA/CIS Meeting, July 7, 1999

Overview

- Why ATM?
- ATM service categories
- Improving TCP over UBR+ over Satellites
- Improving TCP over ABR over Satellites

This work was sponsored by
NASA Glenn Research Center, Cleveland.

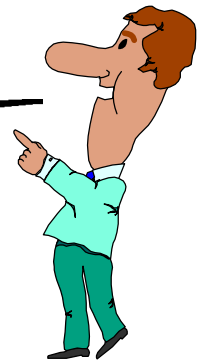

Why ATM?

- ATM vs IP: Key Distinctions
  - Traffic Management:
    Explicit Rate vs Loss based
  - Signaling: Coming to IP in the form of RSVP
  - PNNI: QoS based routing
  - Switching: Coming soon to IP
  - Cells: Fixed size or small size is not important

Old House vs New House

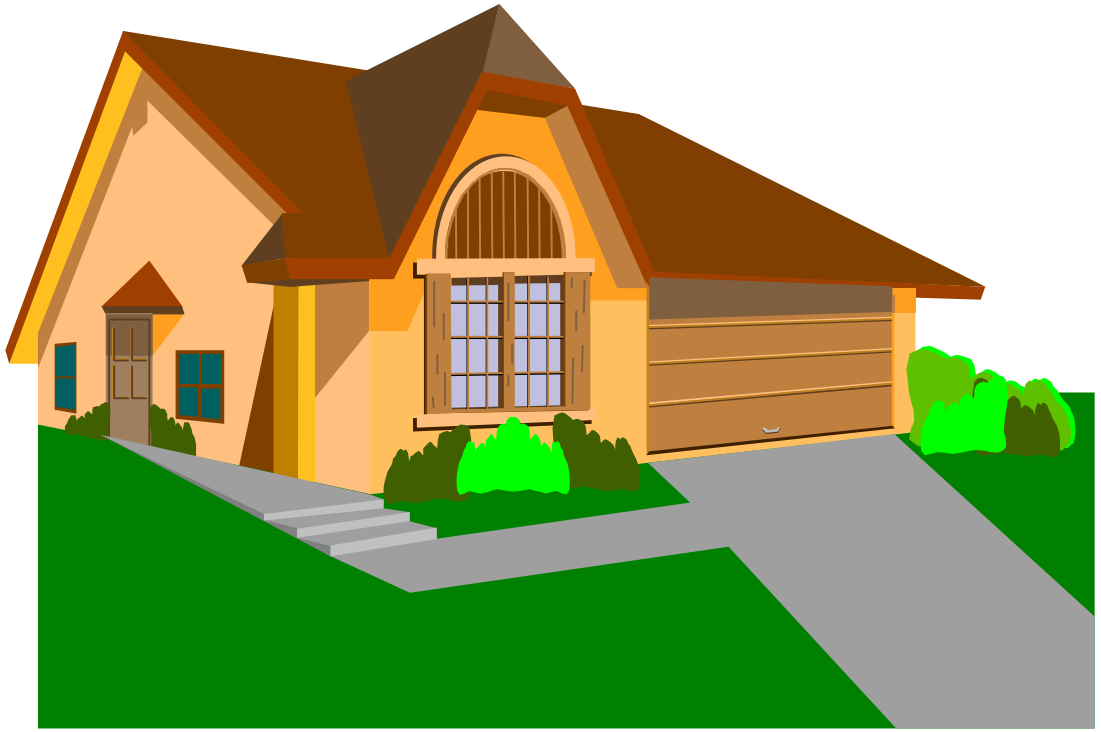

- New needs:

  Solution 1: Fix the old house (cheaper initially)

  Solution 2: Buy a new house (pays off over a long run)

Service Categories

Standby

Guaranteed

Joy Riders

Confirmed

Service Categories

- ABR (Available bit rate):
  Source follows network feedback.
  Max throughput with minimum loss.
- UBR (Unspecified bit rate):
  User sends whenever it wants. No feedback. No guarantee. Cells may be dropped during congestion.
- CBR (Constant bit rate): User declares required rate. Throughput, delay and delay variation guaranteed.
- VBR (Variable bit rate): Declare avg and max rate.
  - rt-VBR (Real-time): Conferencing.
    Max delay guaranteed.
  - nrt-VBR (non-real time): Stored video.

Buffer Management

- Accounting: Per-VC, Global
  Multiple or Single
- Threshold: Single or Multiple
- Four Types:
  - Single Accounting, Single threshold (SAST)
  - Single Accounting, Multiple threshold (SAMT)
  - Multiple Accounting, Single threshold (MAST)
  - Multiple Accounting, Multiple threshold (MAMT)

Buffer Mgmt Schemes

Group	Examples	Threshold	Drop Type	Tag Sensitive
SA ST	EPD, PPD	Static	Deterministic	No
	RED	Static	Probabilistic	No
MA ST	FRED Selective Drop,FBA	Dynamic	Probabilistic	No
	VQ+DEPD	Dynamic	Deterministic	No
MA MT	PME+ ERED	Static	Probabilistic	Yes
	DFBA	Dynamic	Probabilistic	Yes
SAMT	Priority Drop	Static	Deterministic	Yes

Slow Start Flow Control

- Window = Flow Control Avoids receiver overrun
- Need congestion control to avoid network overrun
- The sender maintains two windows:
  Credits from the receiver
  Congestion window from the network
  Congestion window is always less than the receiver window
- Starts with a congestion window (CWND) of 1 segment (one max segment size)
  $\Rightarrow$ Do not disturb existing connections too much.
- Increase CWND by 1 every time an ack is received

Slow Start (Cont)

- If segments lost, remember slow start threshold (SSThresh) to CWND/2
  Set CWND to 1
  Increment by 1 per ack until SSthresh
  Increment by 1/CWND per ack afterwards

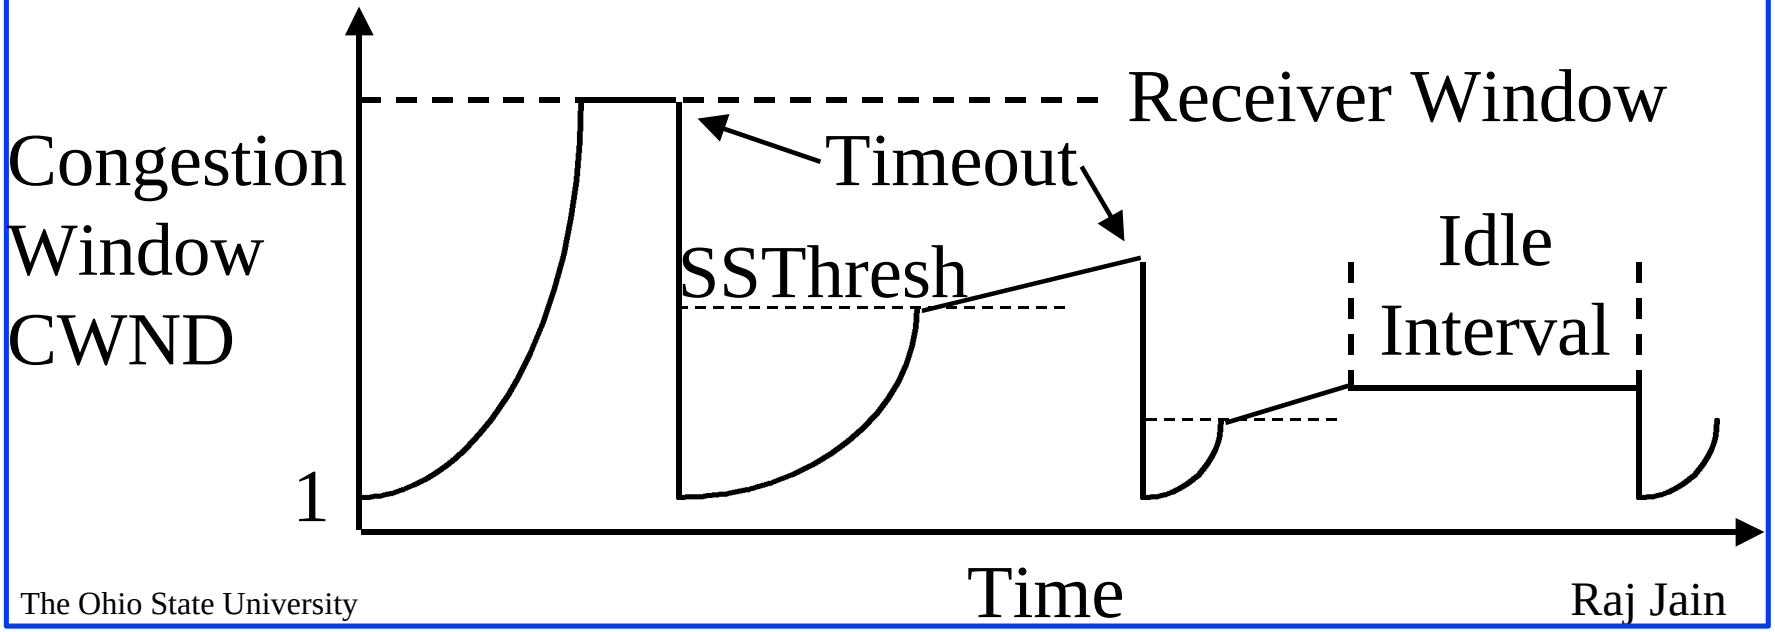

Slow Start (Cont)

- At the beginning, SSThresh = Receiver window
- After a long idle period (exceeding one round-trip time), reset the congestion window to one.
- Exponential growth phase is also known as "Slow start" phase
- The linear growth phase is known as "congestion avoidance phase"

Fast Retransmit and Recovery

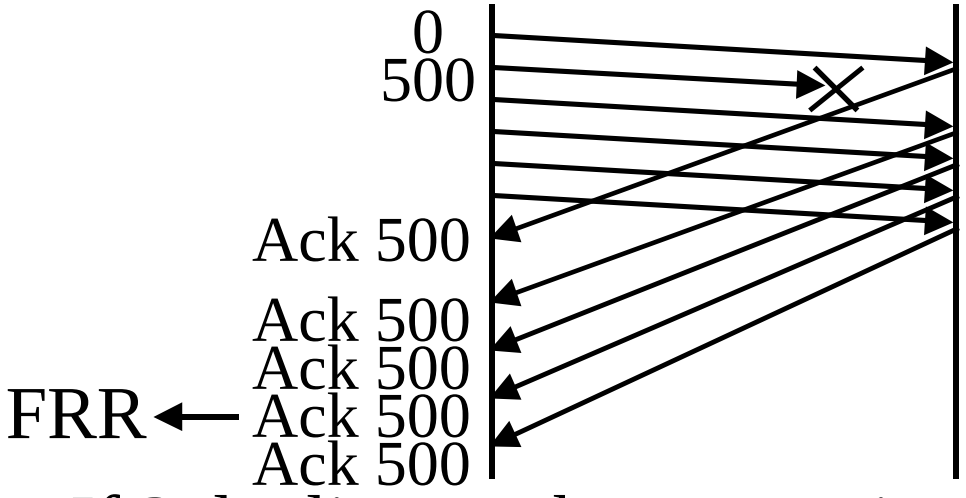

- If 3 duplicate acks are received for the same packet, assume that the next packet has been last. Retransmit it right away. Retransmit only one packet.
- Helps if a single packet is lost.
  Does not help if multiple packets lost.
- Ref: Stevens, Internet draft

Internet Protocols over ATM

- ATM Forum has designed ABR service for data
- UBR service provides no feedback or guarantees
- Internet Engineering Task Force (IETF) prefers UBR for TCP

ABR or UBR?

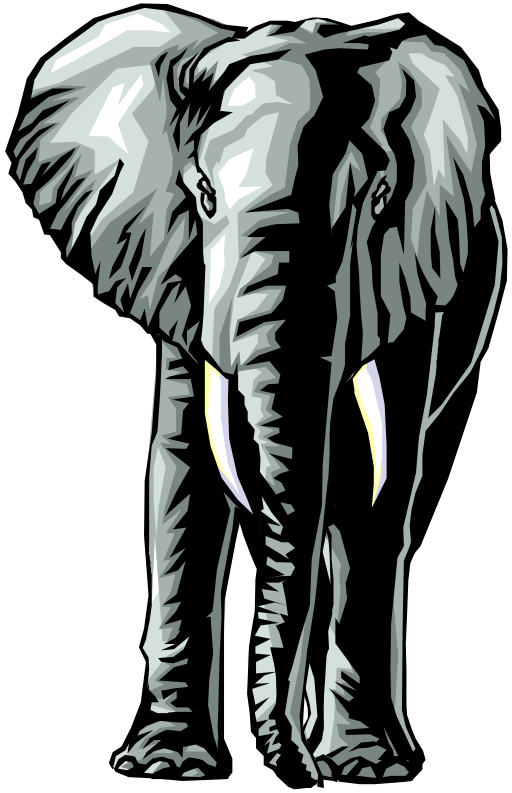

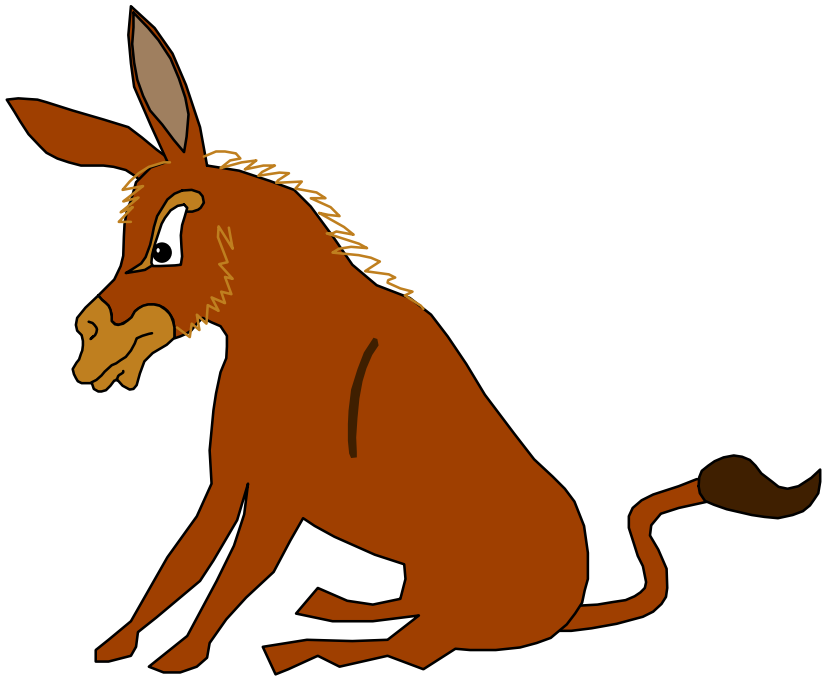

- Intelligent transport or not?

TCP over UBR+

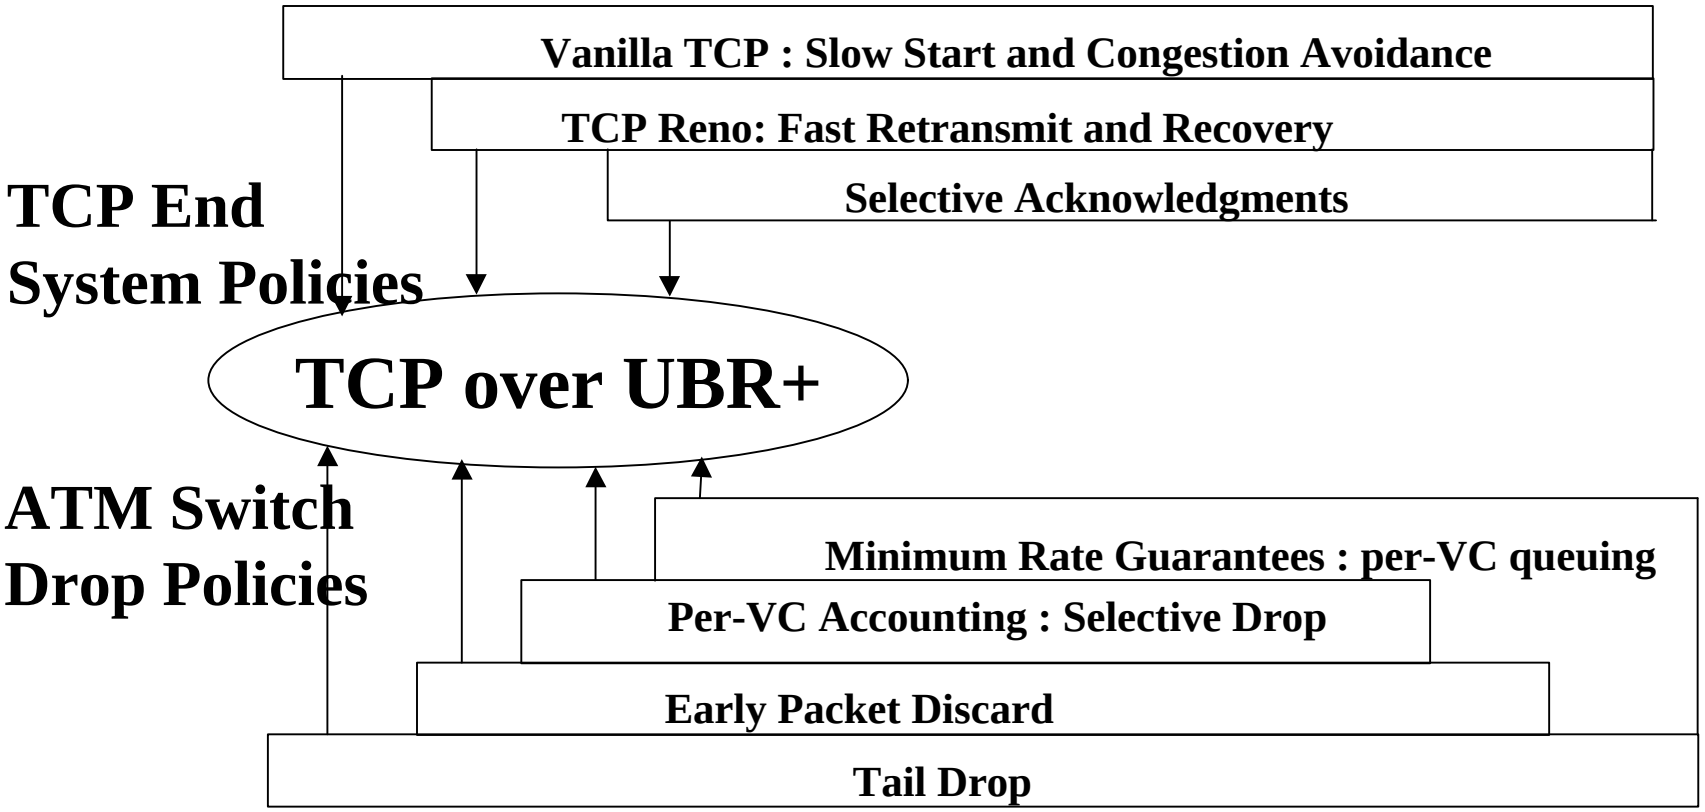
Vanilla TCP : Slow Start and Congestion Avoidance
TCP Reno: Fast Retransmit and Recovery
Selective Acknowledgments
TCP End System Policies
TCP over UBR+
ATM Switch Drop Policies
Minimum Rate Guarantees : per-VC queuing
Per-VC Accounting : Selective Drop
Early Packet Discard
Tail Drop

Policies

LEO, MEO, GEO
# of TCP Sources
Buffer Sizes

End-System Policies

Switch Policies		No FRR	FRR	New Reno	SACK + New Reno
No EPD					
	Plain EPD				
EPD	Selective Drop				
	Fair Buffer Allocation				

Policies: Results

- In LANs, switch improvements (PPD, EPD, SD, FBA) have more impact than end-system improvements (Slow start, FRR, New Reno, SACK). Different variations of increase/decrease have little impact due to small window sizes.
- In satellite networks, end-system improvements have more impact than switch-based improvements
- FRR hurts in satellite networks.
- Fairness depends upon the switch drop policies and not on end-system policies

Policies (Continued)

- In Satellite networks:
  - SACK helps significantly
  - Switch-based improvements have relatively less impact than end-system improvements
  - Fairness is not affected by SACK
- In LANs:
  - Previously retransmitted holes may have to be retransmitted on a timeout
    $\Rightarrow$ SACK can hurt under extreme congestion.

Buffer Requirements: Results

- Very small buffer sizes result in low efficiency
- Moderate buffer sizes (less than 1 RTT)
  - Efficiency increases with increase in buffer size
  - Efficiency asymptotically approaches 100%
- 0.5*RTT buffers provide sufficiently high efficiency (98% or higher) for SACK TCP over UBR even for a large number of TCP sources

Guaranteed Frame Rate (GFR)

- UBR with minimum cell rate (MCR)
  $\Rightarrow$ UBR+
- Frame based service
  - Complete frames are accepted or discarded in the switch
  - Traffic shaping is frame based.
    All cells of the frame have CLP =0 or CLP =1
  - All frames below MCR are given CLP =0 service.
    All frames above MCR are given best effort (CLP =1) service.

Guaranteed Rate

- Guaranteed Rate (GR): Reserve a small fraction of bandwidth for UBR class.

GR	GFR
per-class reservation	per-VC reservation
per-class scheduling	per-VC accounting/scheduling
No new signaling Can be done now	Need new signaling In TM4+

Guaranteed Rate: Results

- Guaranteed rate is helpful in WANs.
- For WANs, the effect of reserving 10% bandwidth for UBR is more than that obtained by EPD, SD, or FBA
- For LANs, guaranteed rate is not so helpful. Drop policies are more important.
- For Satellites, end-system policies seem more important.

4 Ways to Improve UBR over Satellites

1. Implement "Selective Acknowledgement" in end-systems
2. Disable "Fast retransmit and recovery" in end-systems
3. Reserve a small fraction of bandwidth for UBR in the switches
4. Fix slow start implementations in end-systems to avoid errors due to integer arithmetic

The Explicit Rate ABR

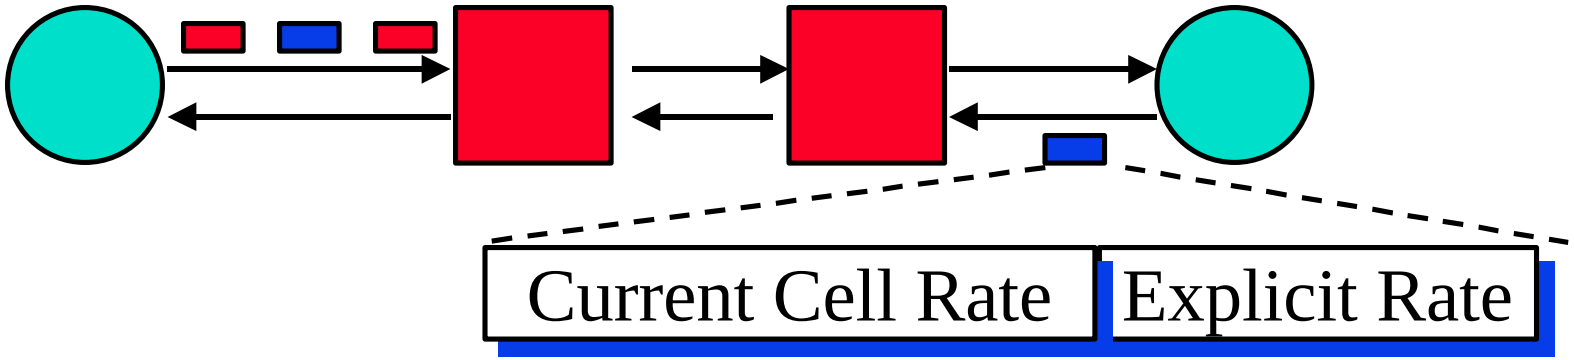

- Sources send one RM cell every n cells
- The RM cells contain "Explicit rate"
- Destination returns the RM cell to the source
- The switches adjust the rate down
- Source adjusts to the specified rate

Rule 5

- If you haven't used the link for ADTF (ACR Decrease Time Factor) seconds, reset your rate to ICR (Initial Cell Rate)
- Default ADTF = 0.5 seconds
- Allowed Range = 0.01 to 10.23 s
- On satellites, it takes a long time to ramp up
  $\Rightarrow$ Avoid triggering rule 5
  $\Rightarrow$ Use larger values of ADTF

Source Rule 6 Study

- CRM limits the number of cells lost if the link is broken
- Source Rule (6):
  If you do not receive feedback from the network after CRM×Nrm cells, reduce your ACR:

  $$ACR = \max\{MCR, ACR - ACR \times CDF\}$$

- The CRM is a 8-bit parameter ⇒ Max 256×32 cells ⇒ Very low throughput on satellite links
- For very high-speed satellite links, CRM must be 32-bit ⇒ CRM was changed to 24-bit

4 Ways to Improve ABR over Satellite

1. Increase the limit on the number of outstanding cells before decreasing ⇒ Large TBE

   The size of was increased from 8 bit to 24 bit to accommodate satellite paths.
2. Use larger increase factor
   ⇒ RIF=1 ⇒ Fast transient Response
3. Implement backward congestion notification (BECN)
4. Implement Virtual Source/Virtual Destination

VS/VD

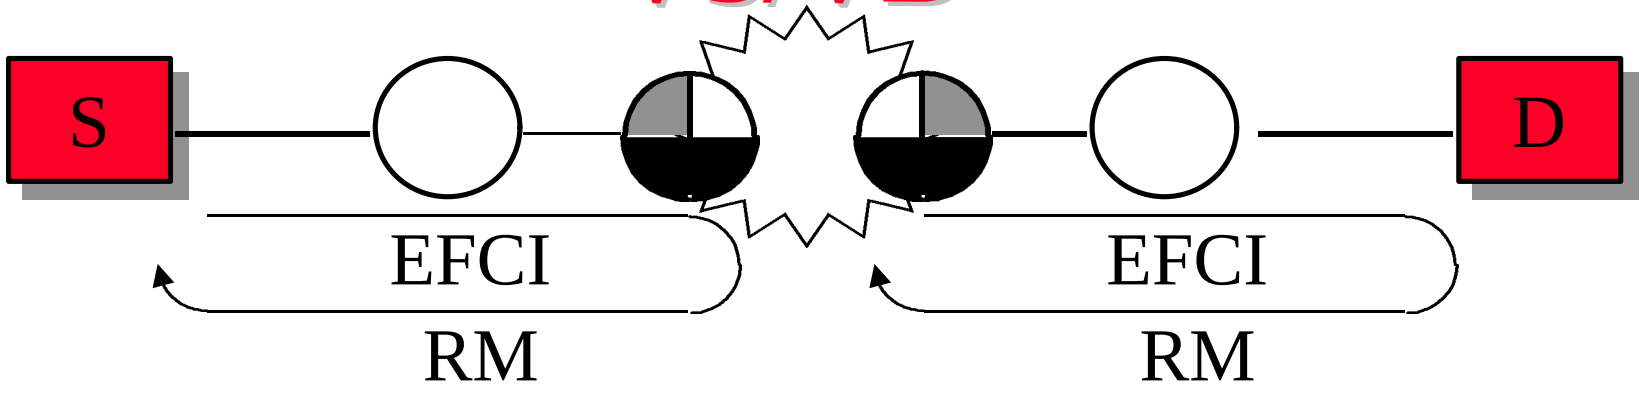

- Virtual source/virtual destinations (VS/VD) follow all notification/control rules
- Can be hop-by-hop

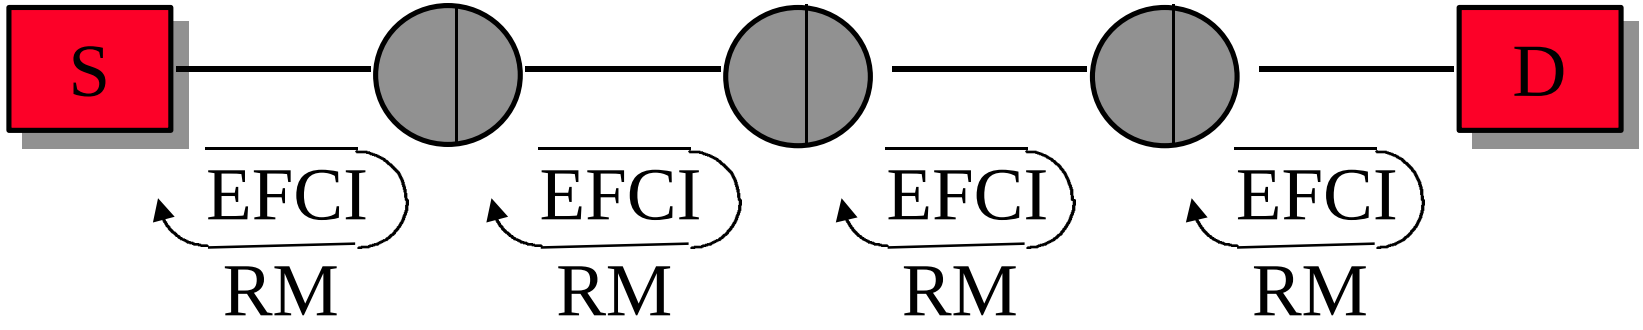

- Virtual dest/sources maintain per-VC queues.

VS/VD: Results

- Without VS/VD:

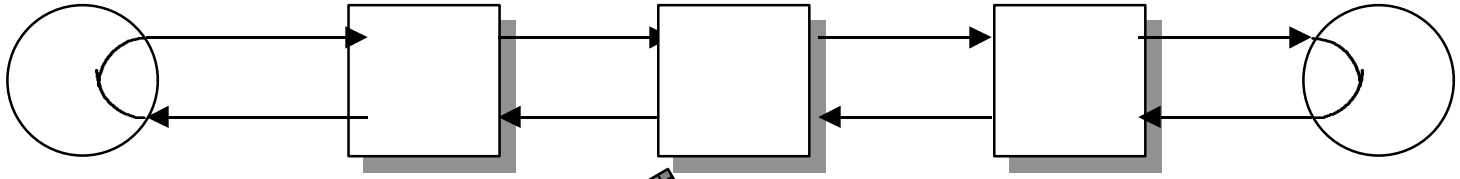

- With VSVD:

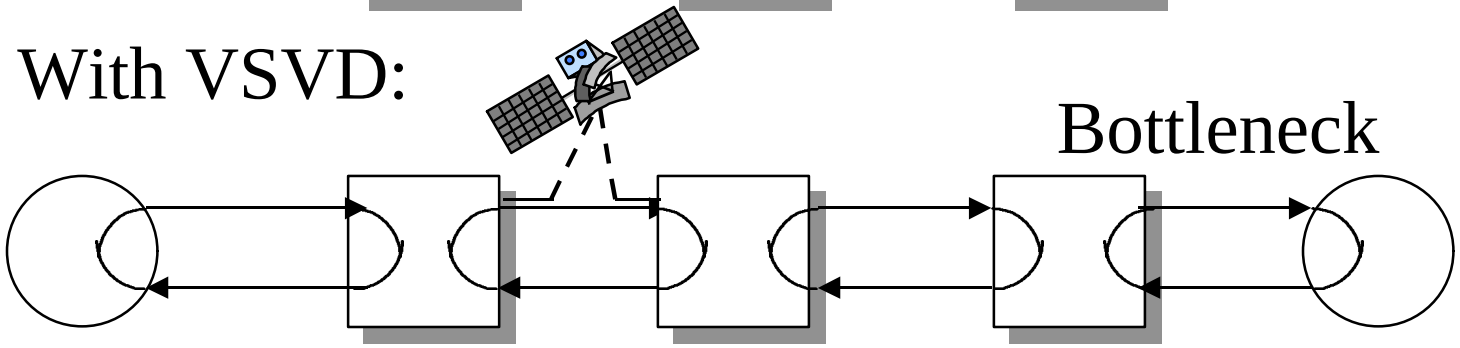

- With VSVD, the buffering is proportional to the delay-bandwidth of the previous loop
  ⇒ Good for satellite networks

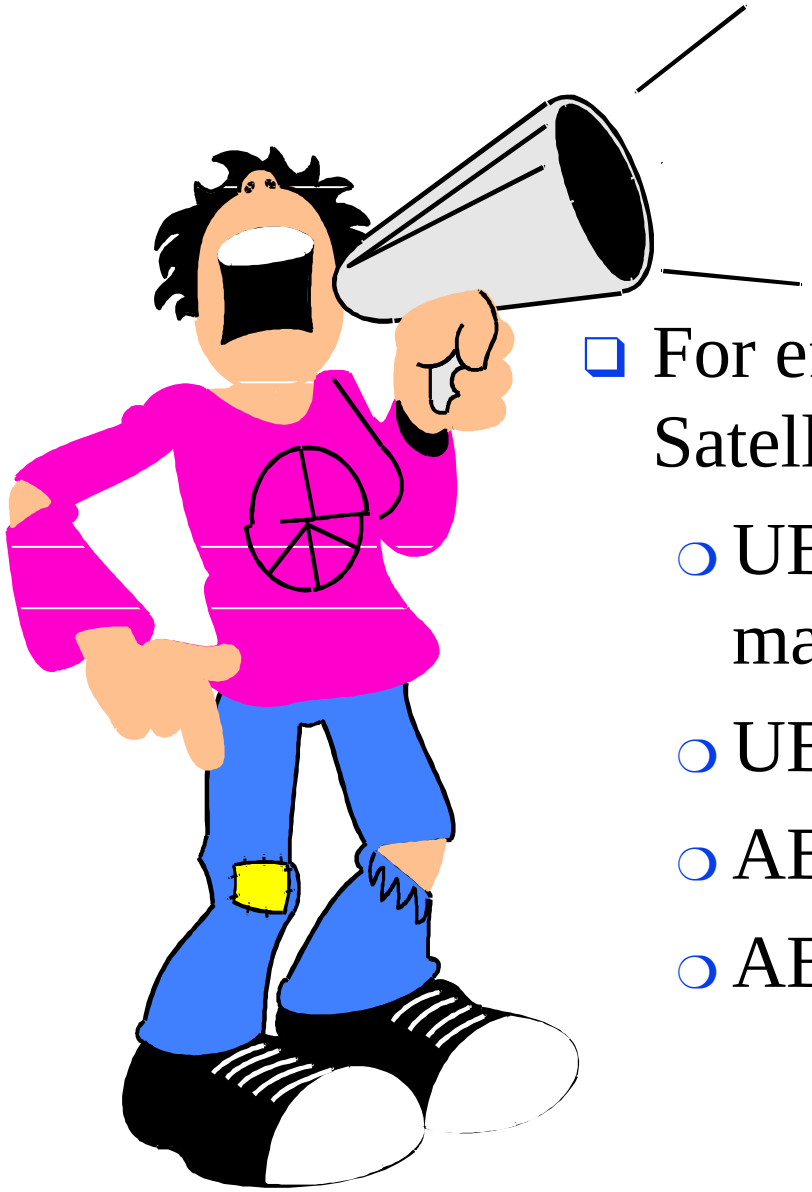

Summary

- For efficient support of TCP over Satellite Links:
  - UBR with intelligent buffer management
  - UBR with guaranteed rate
  - ABR with proper parameters
  - ABR with VS/VD

Summary (Cont)

- Implement VS/VD, BECN, RIF=1, TBE=Large to improve ABR over satellites
- Implement SACK, Disable FRR, reserve bandwidth for UBR, and correct TCP implementations to improve UBR over satellites.

Thank You!