Traffic Management of Internet Protocols over ATM

Raj Jain

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

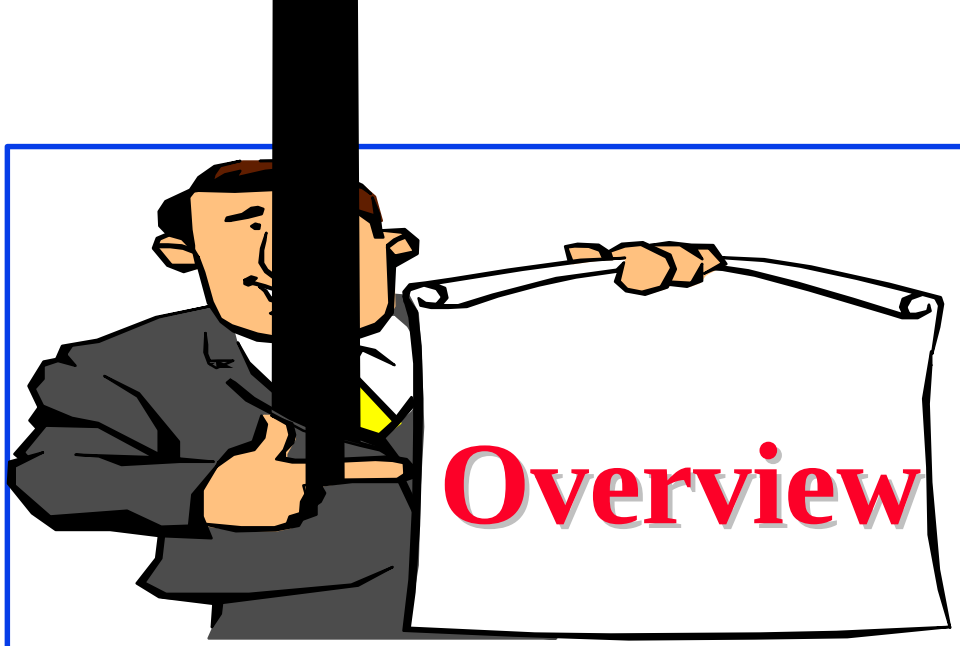

Overview

- Why ATM?
- ATM Service Categories: ABR and UBR
- Binary and Explicit Feedback
- ABR Vs UBR
- TCP/IP over ABR
- TCP/IP over UBR

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

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.

Binary Rate Scheme

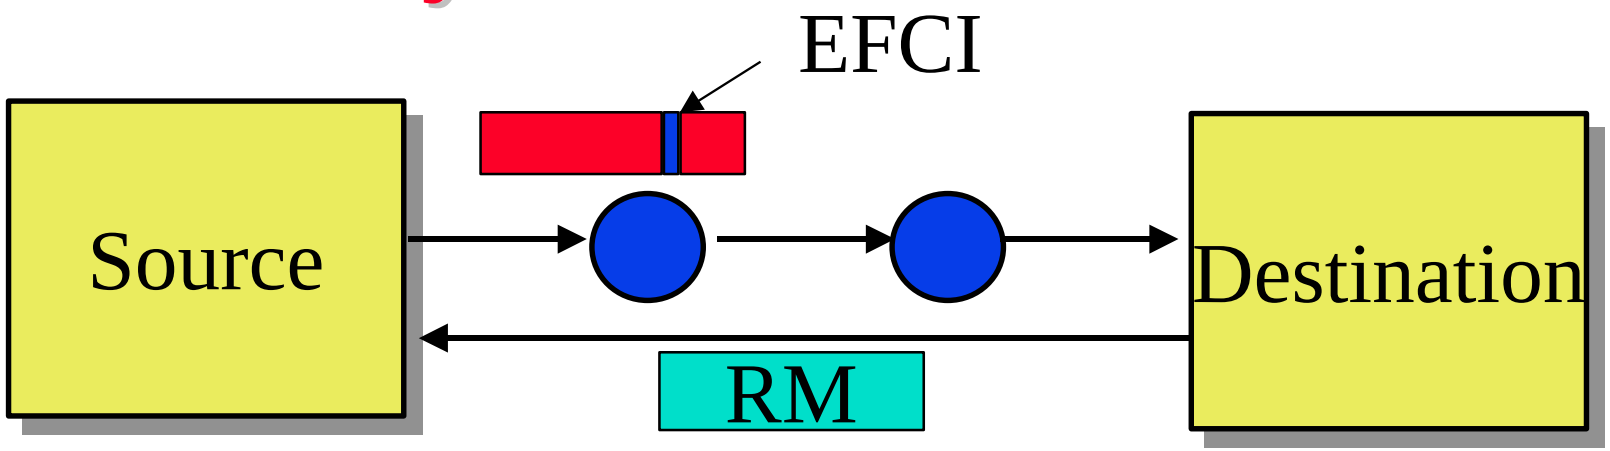

- DECbit scheme in many standards since 1986.
- Forward explicit congestion notification (FECN) in Frame relay
- Explicit forward congestion indicator (EFCI) set to 0 at source. Congested switches set EFCI to 1
- Every nth cell, destination sends an resource management (RM) cell to the source

The Explicit Rate ABR

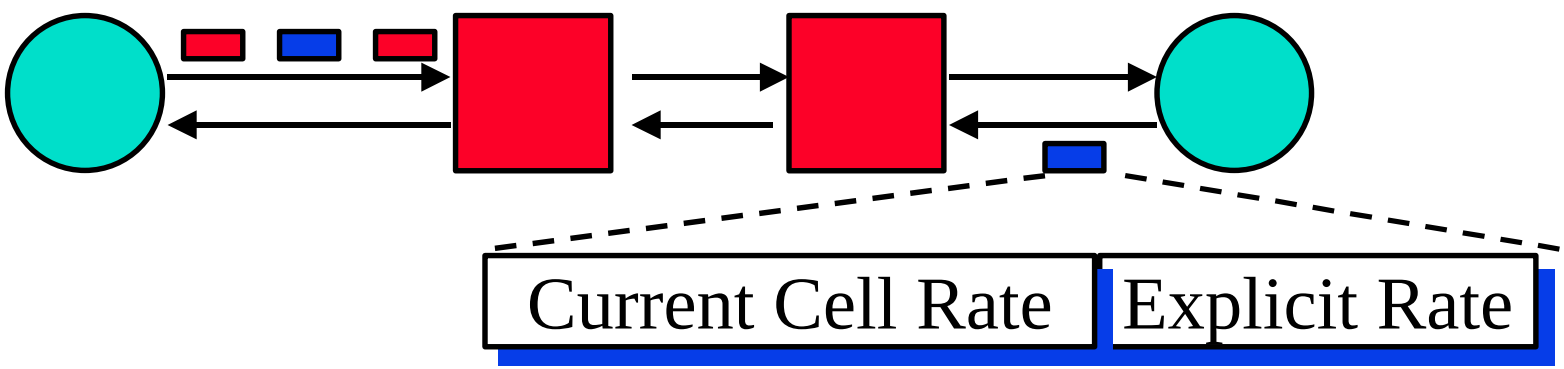

- Proposed in July 1994
- Sources send one RM cell every n cells
- The RM cells contain "Explicit rate"
- Destination returns the RM cell to the source
- The switches adjust the rate down
- Source adjusts to the specified rate

Why Explicit Rate Indication?

- Longer-distance networks
  $\Rightarrow$ Can't afford too many round-trips
  $\Rightarrow$ More information is better
- Rate-based control
  $\Rightarrow$ Queue length = ΔRate $\times$ ΔTime
  $\Rightarrow$ Time is more critical than with windows

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

Observations About ABR

- ABR performance depends upon the switch algorithm. Assuming *ERICA*. (Ref: http://www.cis.ohio-state.edu/~jain/)
- No cell loss for *TCP* if switch has buffers ≈ 4 × RTT.
- No loss for any number of TCP sources w 4 × RTT buffers.
- No loss even with VBR background. W/o VBR, 3×RTT buffers will do.
- Under many circumstances, 1× RTT buffers may do.
- Required buffers depend upon RTT, feedback delay, switch parameters, and characteristics of VBR.

Observations about UBR

- No loss for TCP if Buffers
  = Σ TCP receiver window
- Required buffering depends upon number of sources.
- Receiver window ≥ RTT for full throughput
- Unfairness in many cases.
- Fairness can be improved by proper buffer allocation, selective drop policies, and scheduling.
- No starvation ⇒ Lower throughput shows up as increased file transfer times = Lower capacity

Conclusion: UBR may be OK for: LAN, w/o VBR, Small number of sources, AND cheap implementation

ABR vs UBR

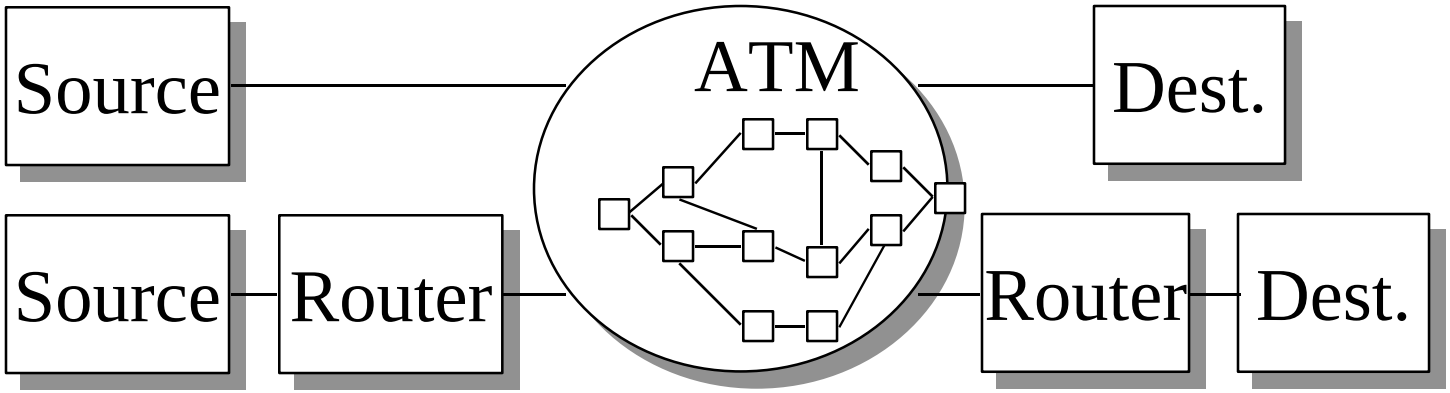

ABR	UBR
Queue in the source	Queue in the network
Pushes congestion to edges	No backpressure
Good if end-to-end ATM	Same end-to-end or backbone
Fair	Generally unfair
Good for the provider	Simple for user

Improving Performance of TCP over UBR

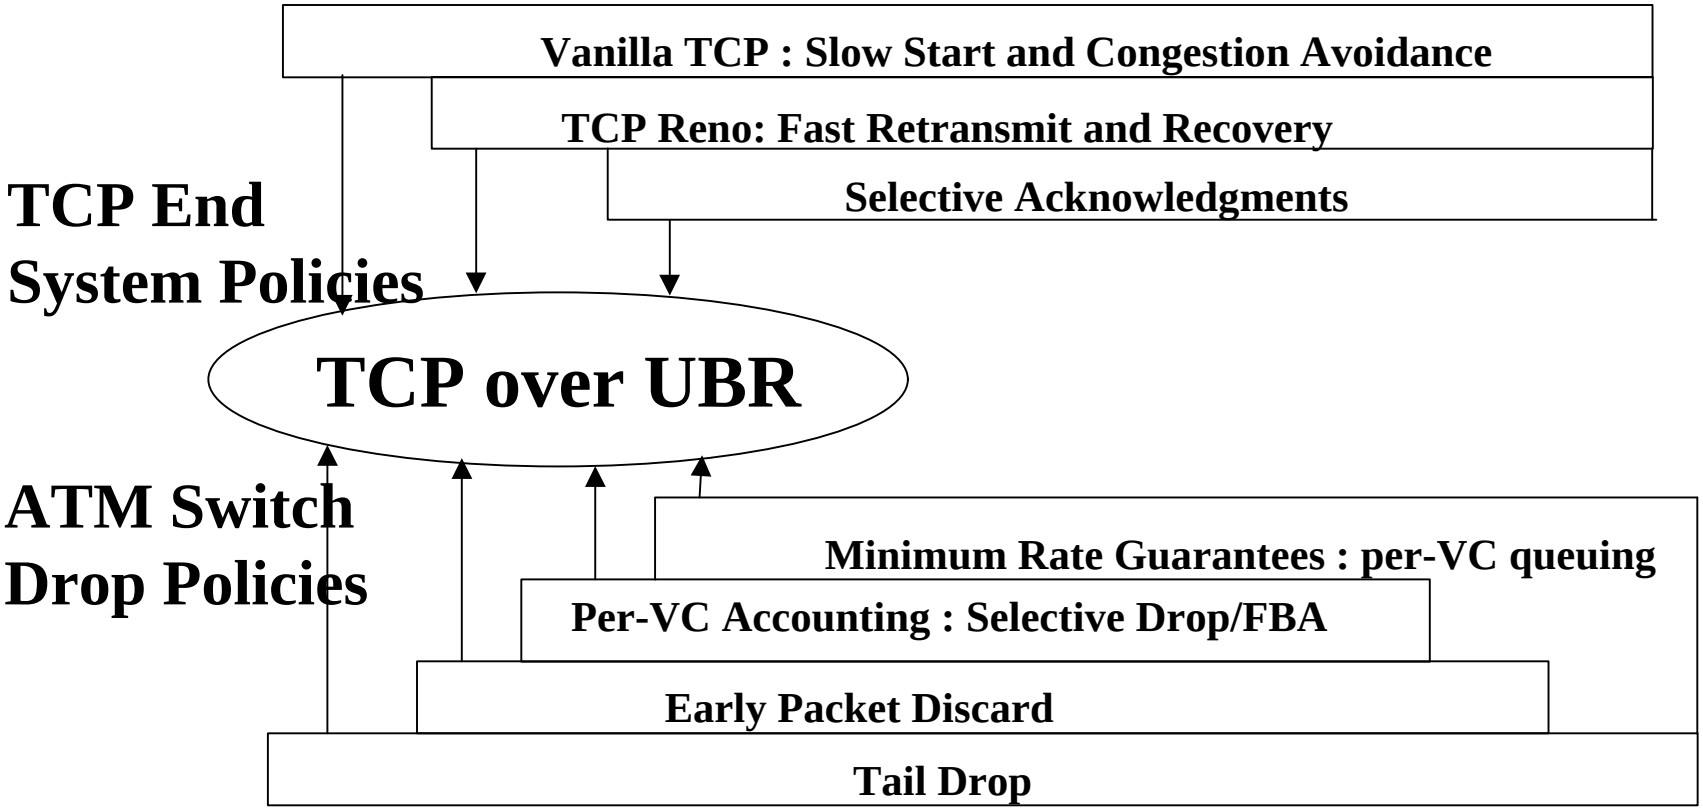

Policies

Switch Policies		End-System 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 large bandwidth-delay networks, end-system improvements have more impact than switch-based improvements
- FRR hurts in large bandwidth-delay networks.

Policies (Continued)

- Fairness depends upon the switch drop policies and not on end-system policies
- In large bandwidth-delay 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.

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 Service

- 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	Need new signaling
Can be done now	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.

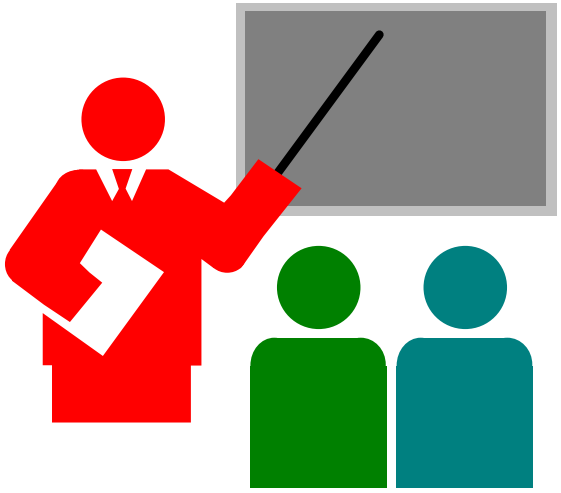

Summary

- Traffic management distinguishes ATM from its competition
- Binary feedback too slow for rate control.
  ER switches better for high bandwidth-delay paths.
- ABR pushes congestion to edges.
  UBR+ may be OK for LANs but not for large bandwidth-delay paths.

Our Contributions and Papers

- All our contributions and papers are available on-line at http://www.cis.ohio-state.edu/~jain/
- See Recent Hot Papers for tutorials.

Thank You!

ATM Research at OSU

- Traffic Management:
  - ERICA+ Switch Algorithm
  - Internet Protocols over ATM
  - Multi-class Scheduling
- Voice/Video over ATM
- Performance Testing
- ATM Test bed: OCARnet

Multi-class Scheduling

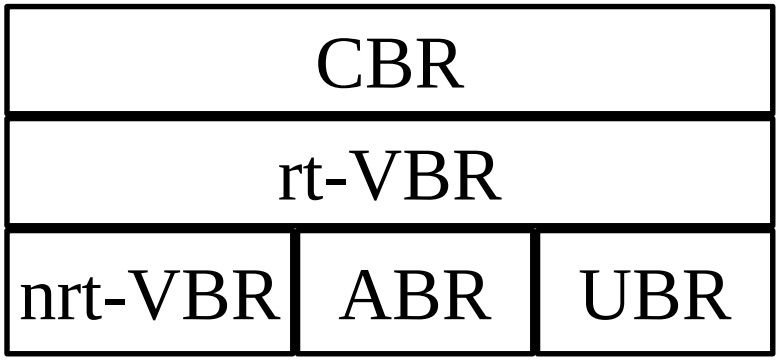

- Ensures *no-starvation* for all classes even under overload.
- Each class has an *allocation* = Guaranteed under overload
- Some classes need minimum delay ⇒ have *priority*.
- Some classes are greedy.
  Left-over capacity is *fairly* allocated.

Voice/Video over ATM

- Speech suppression
  ⇒ Unused bandwidth can be used by data
  Cannot be used by voice.
- Hierarchical compression of Video
  Different users can see different bandwidth video
- Multipoint ABR
- Real-time ABR

Real-Time ABR

- Compressed video is VBR.
  VBR is subject to connection denial.
- Compression parameters can be adjusted dynamically
- In situations, where reduced service is preferable over connection denial, such as in tactical environments, Video over ABR is preferable over no Video.
- ABR divides the available bandwidth fairly among contending connections
- By proper control, ABR can be designed to reduce delay $\Rightarrow$ Real-time ABR

OSU National ATM Benchmarking Lab

- Started a new effort at ATM Forum in October 1995
- Defining a new standard for frame based performance metrics and measurement methodologies
- We have a measurement lab with the latest ATM testing equipment. Funded by NSF and State of Ohio.
- The benchmark scripts can be run by any manufacturer/user in our lab or theirs.
- Modeled after Harvard benchmarking lab for routers

OCARnet

- Ohio Computing and Communications ATM Research Network
- Nine-Institution consortium lead by OSU
  - Ohio State University
  - Ohio Super Computer Center
  - OARnet
  - Cleveland State University
  - Kent State University
  - University of Dayton
  - University of Cincinnati
  - Wright State University
  - University of Toledo

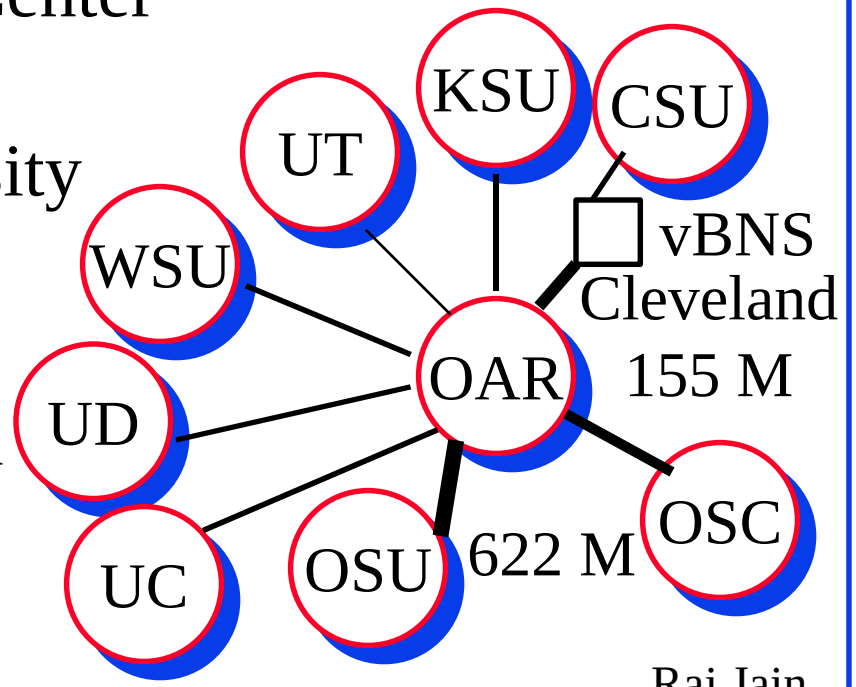

OCARnet

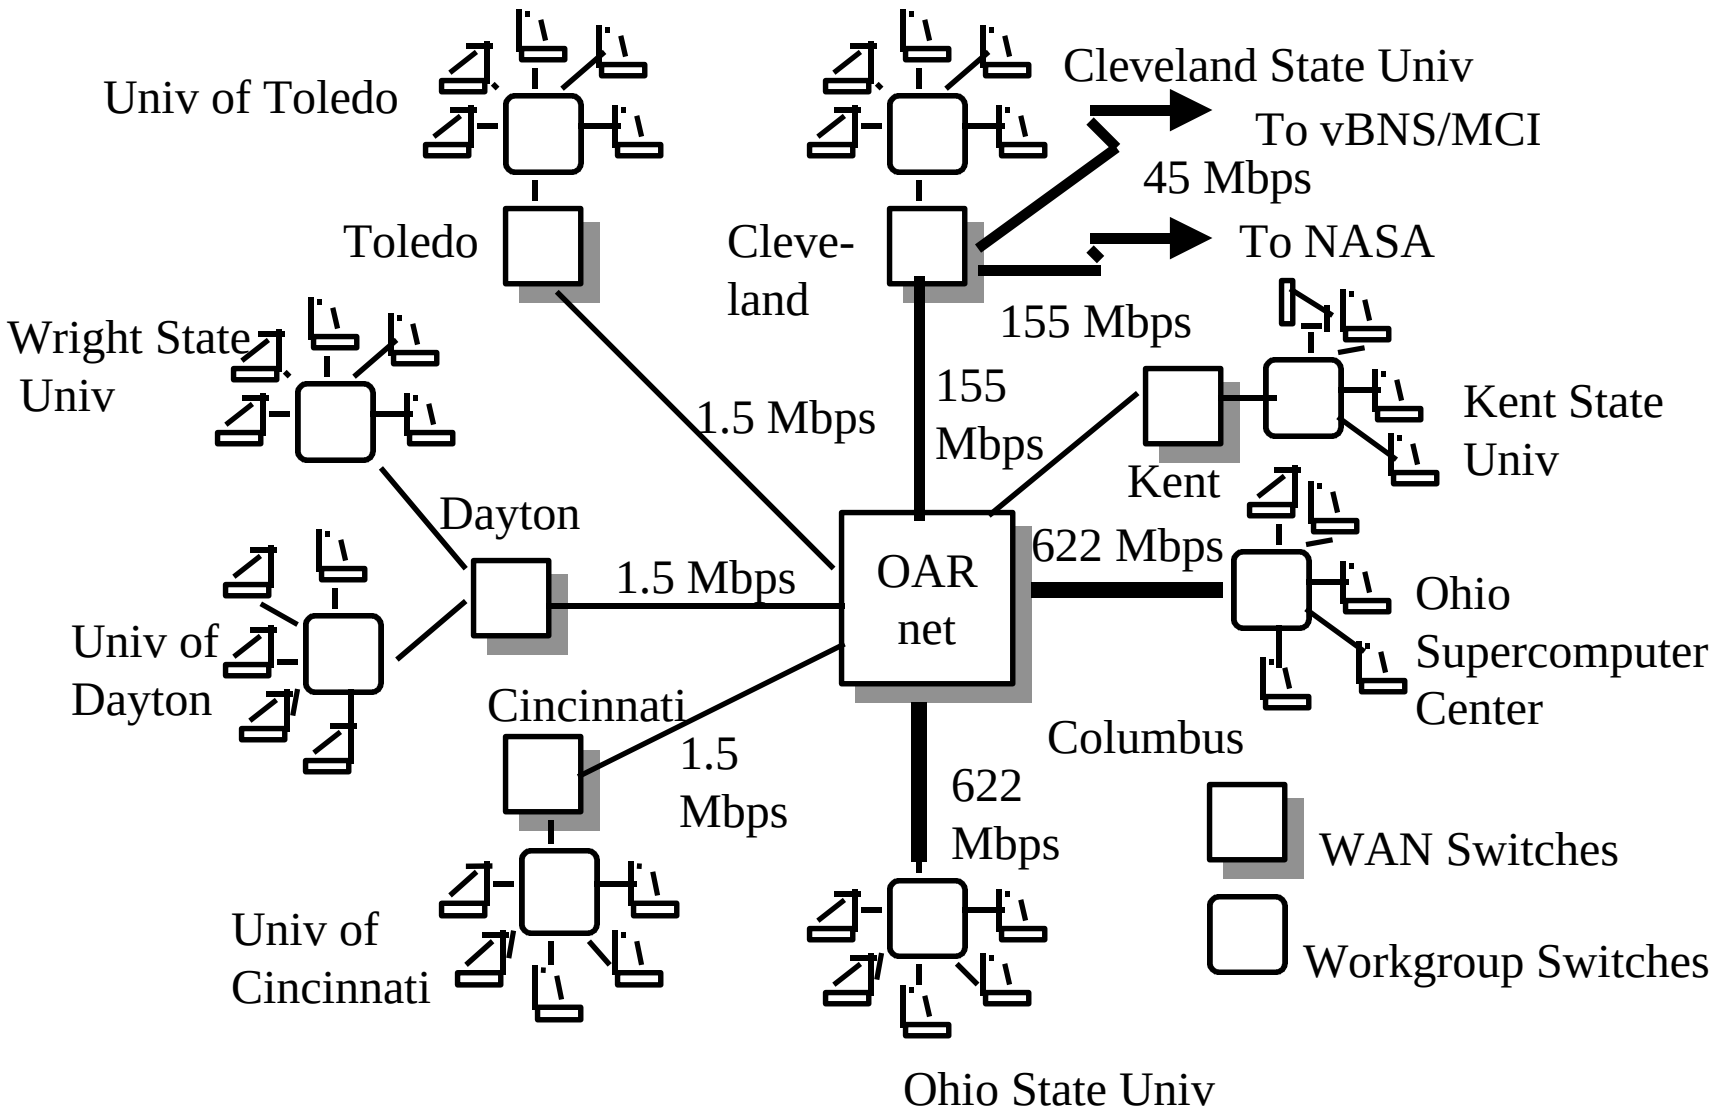
Univ of Toledo
Toledo
Cleveland State Univ
To vBNS/MCI
45 Mbps
To NASA
Cleve-
land
155 Mbps
Wright State
Univ
1.5 Mbps
155
Mbps
Kent State
Univ
Kent
Dayton
622 Mbps
OAR
net
1.5 Mbps
Univ of
Dayton
Ohio
Supercomputer
Center
Cincinnati
Columbus
1.5
Mbps
622
Mbps
WAN Switches
Workgroup Switches
Univ of
Cincinnati
Ohio State Univ