The Use of ATM Quality of Service to Support IP Internetworks

Bruce A. Mah
bmah@CS.Berkeley.EDU

The Tenet Group
University of California at Berkeley
and
International Computer Science Institute

XUNET Student Meeting, Chicago, IL
20 February 1995

Outline

- Introduction
- Issues
  - Using ATM QOS for IP conversations
  - Multiplexing
  - Virtual circuit management
- Current and Ongoing Work
- Summary

Environment: A Heterogeneous Internetwork

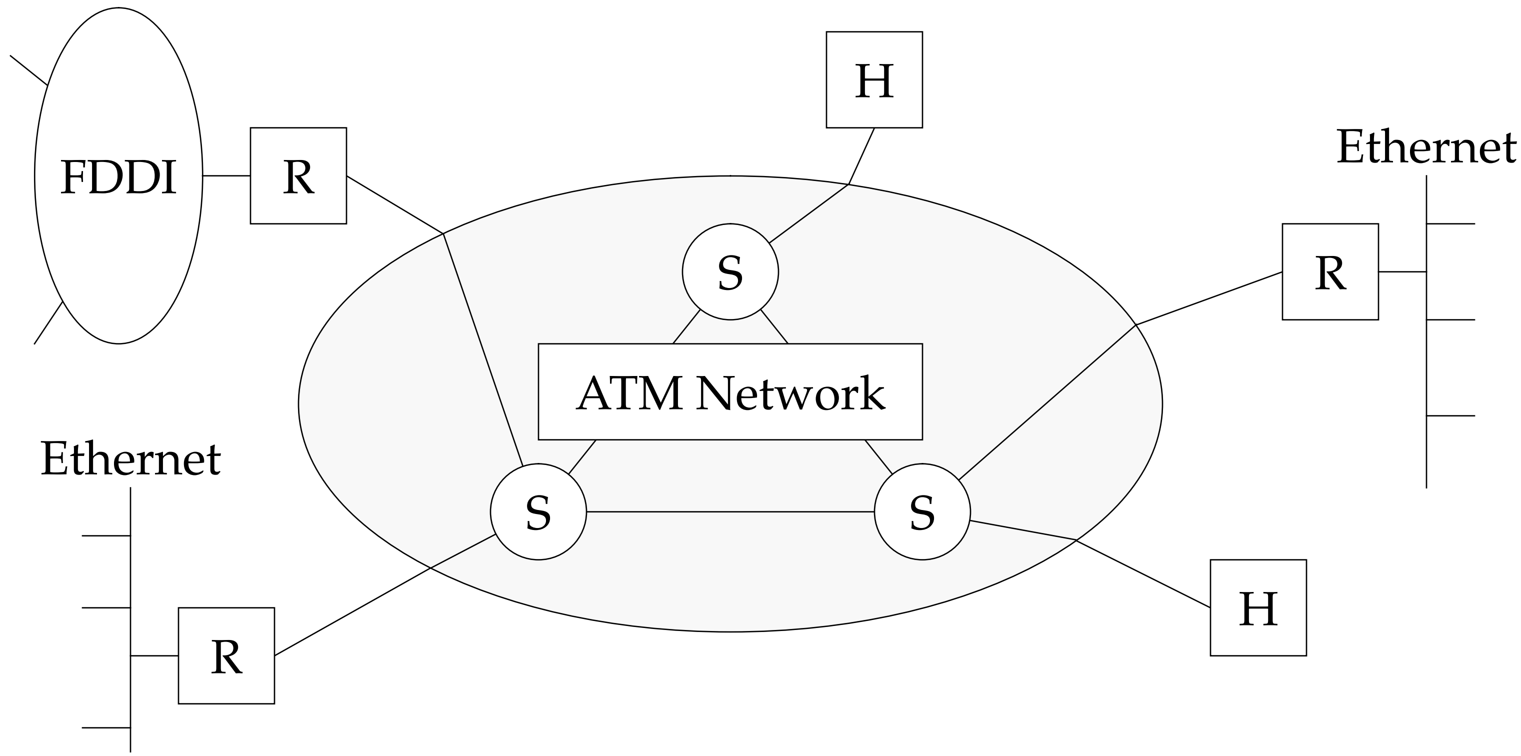

Mixed data and multimedia traffic (e.g. `telnet`, `ftp`, `vic`, `vat`)

IP over ATM

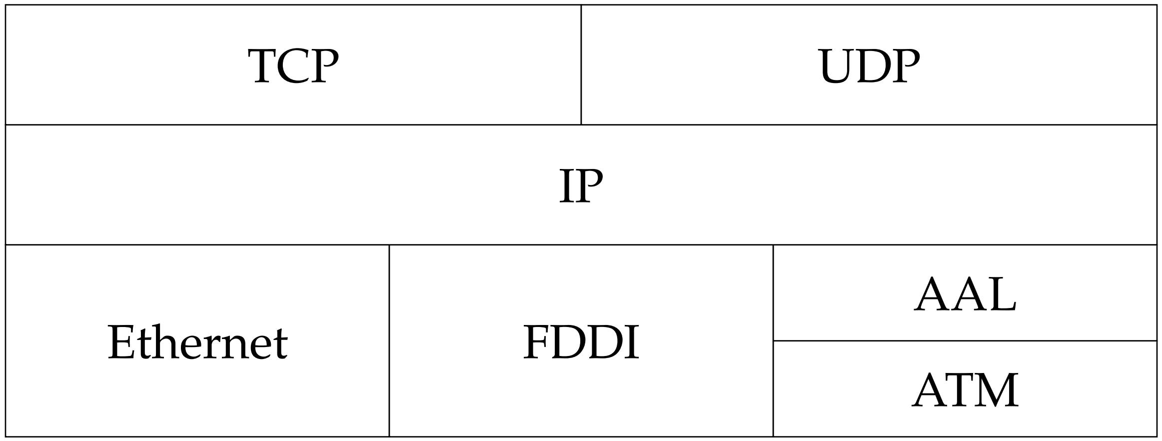

Question: How can IP applications get the best performance over an ATM subnet?

Idea: Use quality of service features of ATM

Different types of virtual circuits for different applications

- `telnet` ⇒ low latency bound
- Video playback ⇒ bandwidth guarantee

Research Issues

The Use of ATM QOS for IP Conversations

- How to map a stream of IP datagrams onto a virtual circuit?
- What performance requirements do IP conversations have?

Multiplexing

- How and when should multiple conversations share a virtual circuit?

Virtual Circuit Management

- When should virtual circuits be created?
- When should virtual circuits be torn down?

Mapping IP Datagrams onto a Virtual Circuit

- “IP Conversation”
  - A stream of related IP datagrams between common endpoints with some definable set of QOS parameters
  - Hosts and routers place packets of IP conversations onto ATM VCs
- For example:
  - All datagrams for a given TCP connection
  - All datagrams between a pair of UDP ports
  - All datagrams on an IP multicast tunnel
  - All ICMP datagrams between a host pair

Mapping IP Datagrams onto a Virtual Circuit

IP (and TCP, UDP, ICMP) headers to determine conversations

<table>
<tr><td>Version</td><td>Hdr Ln</td><td>Prec</td><td>TOS</td><td colspan="4">Total Length</td></tr>
<tr><td colspan="4">ID</td><td>Flags</td><td colspan="3">Fragment Offset</td></tr>
<tr><td colspan="2">TTL</td><td colspan="2">Protocol</td><td colspan="4">Header Checksum</td></tr>
<tr><td colspan="8">Source Address</td></tr>
<tr><td colspan="8">Destination Address</td></tr>
<tr><td colspan="4">Source Port</td><td colspan="4">Destination Port</td></tr>
<tr><td colspan="8">TCP Sequence Number</td></tr>
<tr><td colspan="8">TCP Acknowledgment Number</td></tr>
<tr><td>Hdr Ln</td><td>Rsrvd</td><td colspan="2">Flags</td><td colspan="4">Window Size</td></tr>
<tr><td colspan="4">TCP Checksum</td><td colspan="4">Urgent Pointer</td></tr>
</table>

Determination of Requirements

Well-known applications

- For applications whose needs and traffic characteristics are known
- Example: `telnet`

Monitoring traffic

- Adaptive algorithm for determining bandwidth requirements
- Example: Video transmission with user adjustments

Explicit signalling

- In-band (e.g. IP TOS/Precedence, IP options)
- Out-of-band with a signalling protocol (e.g. RCAP, RSVP)

Default requirements

- For traffic not covered by other means
- “Best effort”?

Datagram to Virtual Circuit and QOS Mapping

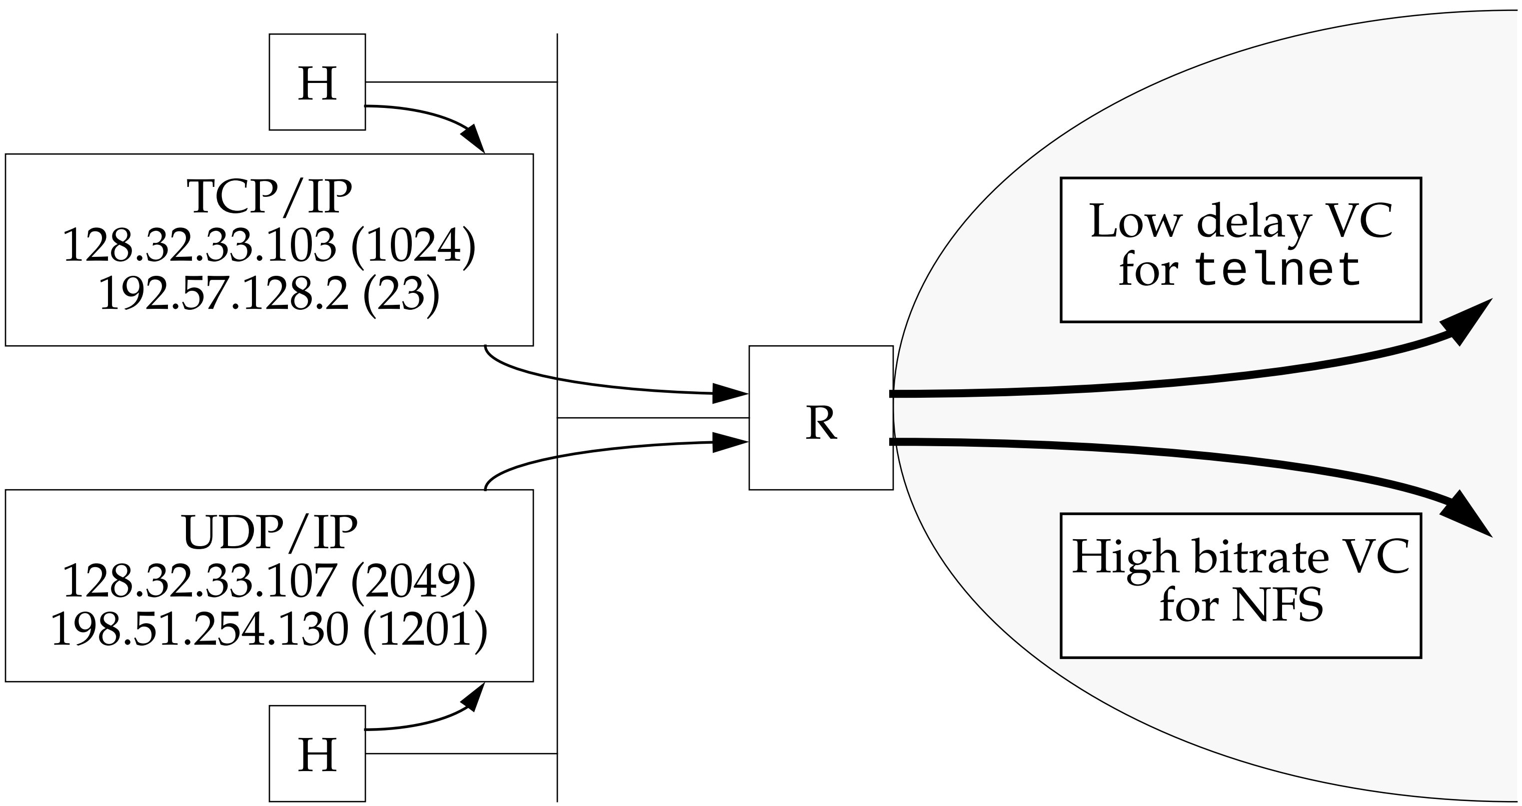

Multiplexing with QOS Considerations

Tradeoff

- Protection of individual IP conversations
- Increased utilization of reserved resources due to statistical multiplexing

Levels of Multiplexing

- Virtual circuit per router pair
- Virtual circuit per host pair
- Virtual circuit per application per host pair
- Virtual circuit per IP conversation
- Combinations

Virtual Circuit Per Router Pair

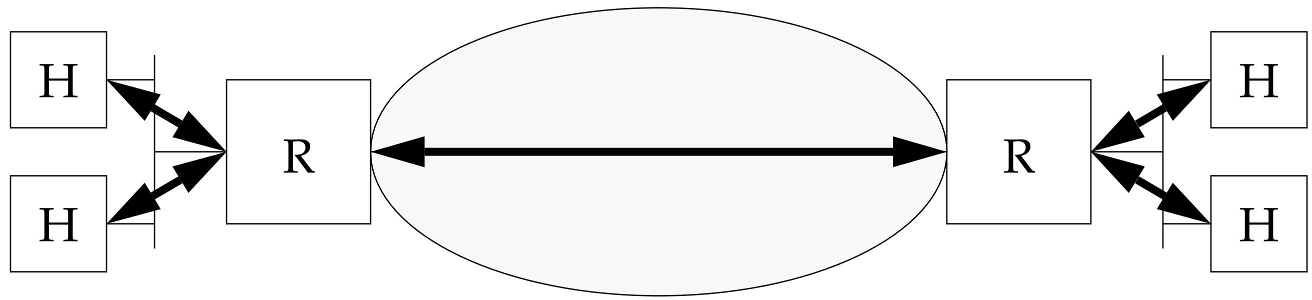

All traffic between a pair of routers routed over same virtual circuit

Statistical multiplexing of conversations over virtual circuit

No protection among conversations sharing a router pair

Uses:

- Low-bitrate or bursty traffic (ICMP)
- Background best-effort traffic (electronic mail)

Virtual Circuit Per Conversation

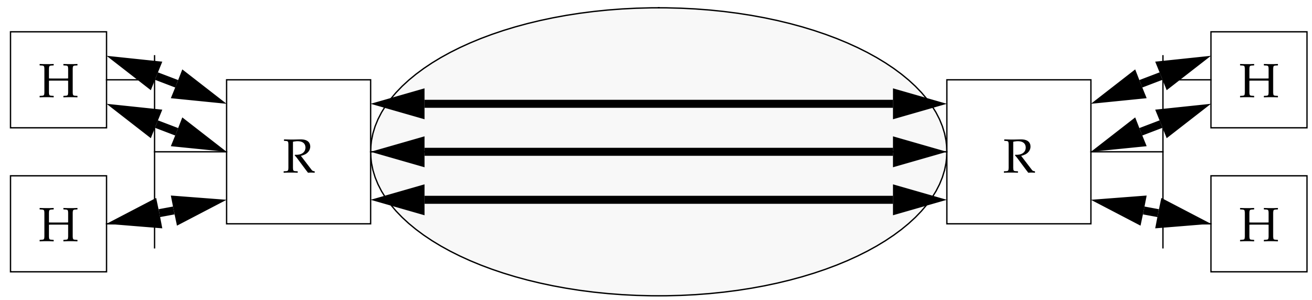

Each IP conversation seen by a router uses a separate virtual circuit

IP conversations protected from each other over ATM subnet

No statistical multiplexing gain within virtual circuit

Uses:

- Real-time video
- Interactive file transfer

Virtual Circuit Management with QOS Considerations

- Paradigm shift: ATM connections vs. IP datagrams
  - When to set up and tear down connections for datagrams?
- Tradeoff
  - Tie up resources for QOS-guaranteed virtual circuits
  - Connection setup latency
- Alternatives
  - Permanent Virtual Circuits (PVCs)
  - Switched Virtual Circuits (SVCs)
  - Switched Virtual Circuits with connection caching
  - Combinations

Permanent Virtual Circuits

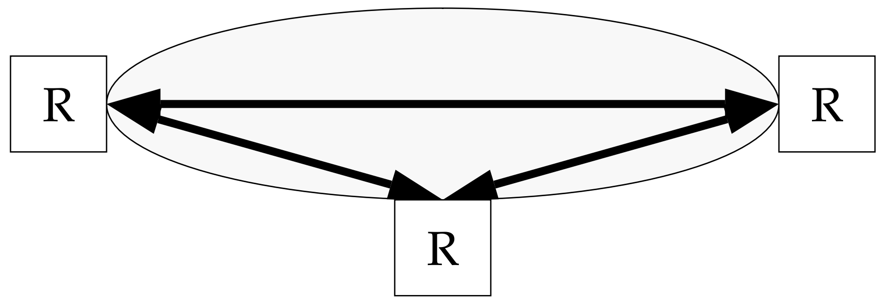

No connection setup latency

Resources always reserved

Scalability a problem: $O(n^2)$ connections

Switched Virtual Circuits

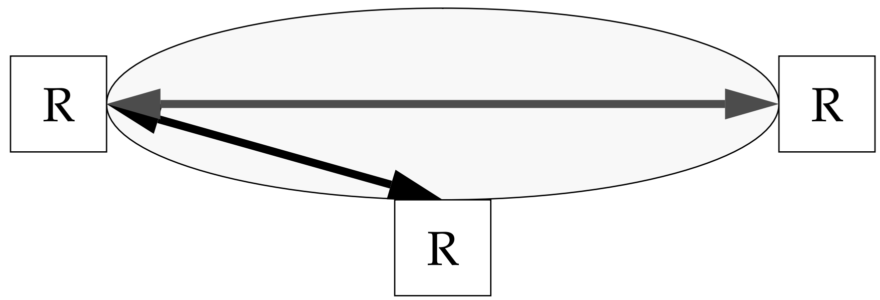

Determine start and end of IP conversations

- First packet ⇒ start
- Timeout ⇒ end

Connection setup latency incurred

- For first packet of conversation
- For other packets if SVC closed too early

Enhancement: connection caching for other IP conversations

Implementation and Evaluation

- Research issues imply IP-over-ATM policies
  - Quality of service mapping
  - Multiplexing
  - Virtual circuit management
- How do these perform in an actual ATM network?
- XUNET II
  - A wide-area ATM backbone network
  - ATM/DS3 connecting FDDI LANs
  - IP over ATM service in place now (PVCs only)

Completed Work

- XUNET IP driver modified to discriminate IP conversations
  - All TCP connections map to unique virtual circuits
  - All UDP streams map to unique virtual circuits
  - Aggregation of conversations possible
  - No performance degradation observed
- XUNET IP over ATM service enhanced to use SVCs
  - Virtual circuits created on demand
  - Fixed idle timeout
- MBONE Measurements
  - Video—`nv`
  - Audio—`vat`
  - Image distribution—`imm`

Next Steps

- Obtain QOS-guaranteed virtual circuits over XUNET
  - Scheduling disciplines—H. Saran, et al. (1993)
  - Signalling support—S. Keshav (1994)
- Formulate IP over ATM policies
  - Mapping from IP conversation(s) to QOS-guaranteed virtual circuit
  - Multiplexing policy
  - Virtual circuit setup and teardown policy

Evaluation

- Performance for IP conversations
  - End-to-end latency
  - Throughput
  - Setup latency
- Overall network performance
  - Utilization
  - Efficiency and over head
  - Best-effort performance

Summary

Idea: Use QOS guarantees for an IP-over-ATM service

- Use of virtual circuits with QOS parameters to carry IP data
- Multiplexing
- Virtual circuit use and management

Plan: Evaluate policies on XUNET II

- Performance of individual conversations
- Performance of network as a whole