

Quality of Service and Asynchronous Transfer Mode in IP Internetworks

Bruce A. Mah
bmah@CS.Berkeley.EDU

The Tenet Group
University of California at Berkeley



Qualifying Examination
23 November 1994

Outline

Introduction

Previous and Related Work

Research Issues

Research Plan

Summary

Introduction

Asynchronous Transfer Mode (ATM)

A new networking technology for integrated services networks

Internet Protocol (IP)


Network layer protocol of the Internet

Communication over heterogeneous internetworks

IP over ATM

The use of an ATM network to transmit IP datagrams

Asynchronous Transfer Mode (ATM)



Point-to-point links between switches in a switching fabric

Connection-oriented network

Potential for quality-of-service guarantees

No specific QOS model implied or assumed

Possible to specify general delay, bandwidth, and loss requirements

Some number of service classes

The Need for Internetworking

ATM is not (will not be?) ubiquitous and dominant

Time to deploy a new technology

Legacy networks (e.g. Ethernet)

New, non-ATM networks

An internetworking solution is still required

The Internet Protocol (IP)

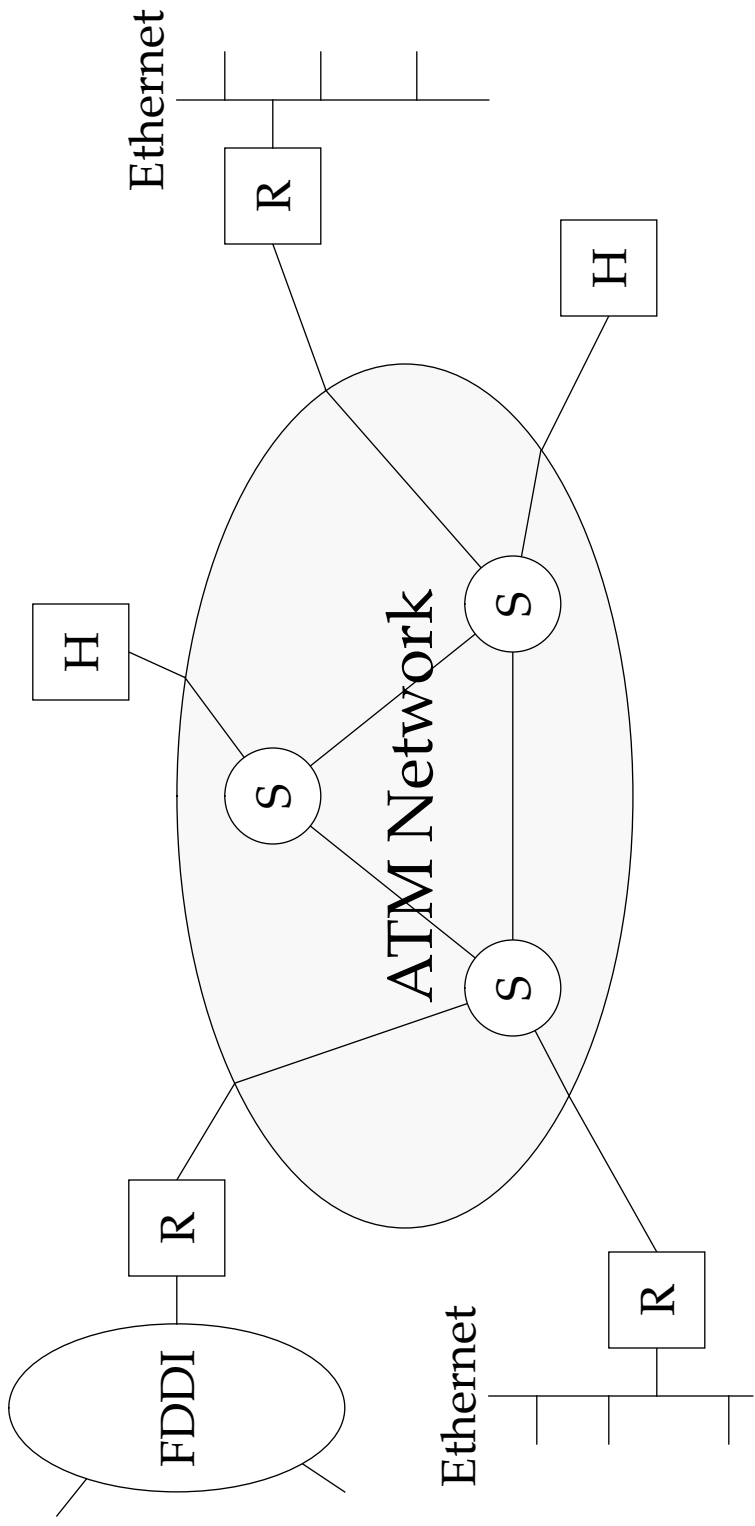
TCP	UDP
IP	
Ethernet	FDDI
	AAL
	ATM

Supports heterogeneous internetworks

Datagram-oriented network-layer protocol

No performance guarantees or quality of service support

Environment



Conventional data (e.g. ftp, telnet) as well as multimedia traffic (e.g. nv, vat)

IP over ATM

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

Idea: Use quality of service features of ATM to improve IP performance

Different types of virtual circuits to meet needs of different applications

telnet \Rightarrow low latency bound

Video playback \Rightarrow bandwidth guarantee

Related Work

Models for IP over ATM

Conventional Model—RFC 1577 (1994)

Subnet Model—R. Cole (1994)

Connectionless Servers—D. Box, et al. (1993), D. Omundsen, et al. (1994)

Traffic Measurement and Analysis

A. Schmidt and R. Campbell (1993)

Traffic Multiplexing

R. Cáceres (1992)

K. Claffy (1994)

Virtual Circuit Management

C. Lund, et al. (1994)

S. Keshav, et al. (1994)

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

Examples

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

Version	Hdr Ln	Prec	TOS	Total Length	
ID		Flags		Fragment Offset	
TTL	Protocol		Header Checksum		
Source Address					
Destination Address					
Source Port			Destination Port		
TCP Sequence Number					
TCP Acknowledgment Number					
Hdr Ln	Rsrvd	Flags		Window Size	
TCP Checksum			Urgent Pointer		

Determination of Requirements

Goal is *not* end-to-end performance guarantees, but improved performance across ATM subnet

Well-known applications

For applications whose needs and traffic characteristics are known

Identified by fields in upper-layer protocol headers

Example: telnet

Monitoring traffic

Adaptive algorithm for determining bandwidth requirements

Long-lived conversations only

Example: Video transmission with user adjustments

Determination of Requirements


Explicit signalling

- Disadvantage: needs support by application
- In-band (e.g. IP TOS/Precedence, IP options)
- Out-of-band with a signalling protocol (e.g. RSVP, RSCAP, 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—Most commercial ATM LANs


Virtual circuit per host pair

Virtual circuit per application per host pair—R. Cáceres (1992)

Virtual circuit per IP conversation

Combinations

Virtual Circuit Per Router Pair



All traffic between a pair of r 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

Resource reservation 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 \Rightarrow start

Timeout \Rightarrow 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

Research Plan

Approach

Evaluate IP over ATM with QOS policies in an actual ATM network

Completed Work

Packet classifier for XUNET II

MBONE Measurements

Future Work

IP over QOS-guaranteed virtual circuits on XUNET

Evaluation of policies

Research Plan Approach

Evaluate policies in an actual wide-area ATM network

- Quality of service mapping

- Multiplexing

- Virtual circuit management


Implementation and Evaluation: XUNET II

- A wide-area ATM backbone network

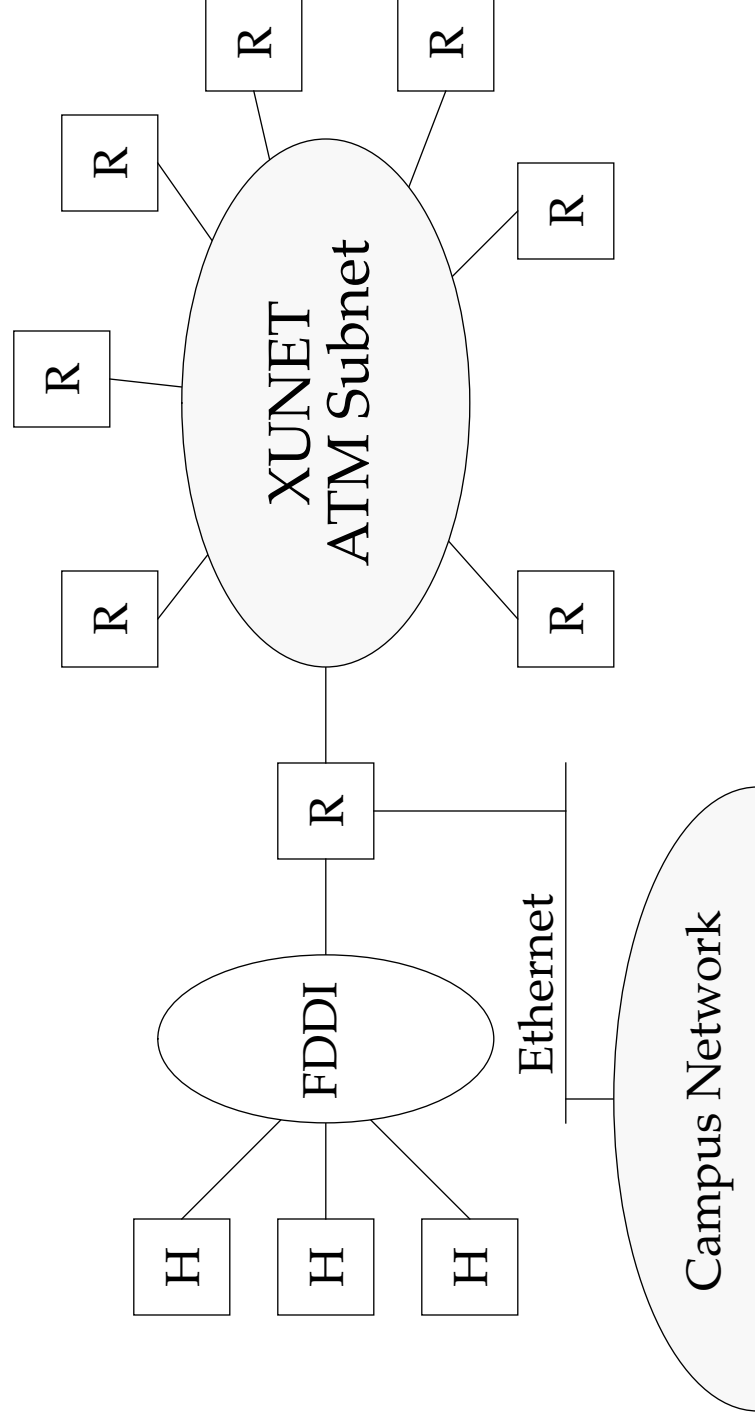
- ATM/DS3 connecting FDDI LANs

- IP over ATM service in place now (PVCs only)

XUNET II Geographic Topology



XUNET II Logical Topology (IP Layer)



Completed Work-Packet Classifier

XUNET IP driver modified to discriminate among 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

Static timeout (15 seconds for initial testing)

Completed Work—MBONE Measurements

Traditional data applications (ftp, telnet) well-studied

Little published about audio and video on Internet


Traces taken of MBONE conversations

Video—nv

Audio—vat

Image distribution—imm

Video Bitrate (nv 128 Kbps)



Peak bitrate (one-second samples) 270 Kbps

Future Work

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

Future Work

Experimental setup

- Transmit IP traffic from applications

- Vary intensity and type of cross traffic

- Measure performance as seen by applications

Evaluate performance for IP conversations with respect to:

- End-to-end latency

- Throughput

- Setup latency

Evaluate overall network performance with respect to:

- Utilization

- Efficiency and overhead

- Best-effort performance

Contributions

Policies and subpolicies for the transmission of IP datagrams over QOS-guaranteed ATM virtual circuits

QOS mapping

Multiplexing

Virtual circuit management

Criteria and a methodology for evaluating IP over ATM policies

A working implementation of both a set of IP over ATM policies and any necessary underlying framework

An evaluation of policies on a wide-area ATM network

Summary

ATM will be incorporated into IP internetworks

Issue: Can the use of QOS guarantees within an ATM subnet be beneficial to IP applications?

Use of virtual circuits with QOS parameters to carry IP data

Multiplexing

Virtual circuit use and management

Evaluate policies on XUNET II, an experimental wide-area ATM network

Performance of individual conversations

Performance of network as a whole