Quality of Service and Asynchronous Transfer Mode in IP Internetworks

Bruce A. Mah
`bmah@CS.Berkeley.EDU`
`http://http.CS.Berkeley.EDU/~bmah/`

The Tenet Group
Computer Science Division
University of California at Berkeley

Berkeley Multimedia and Graphics Seminar
27 November 1996

Motivation

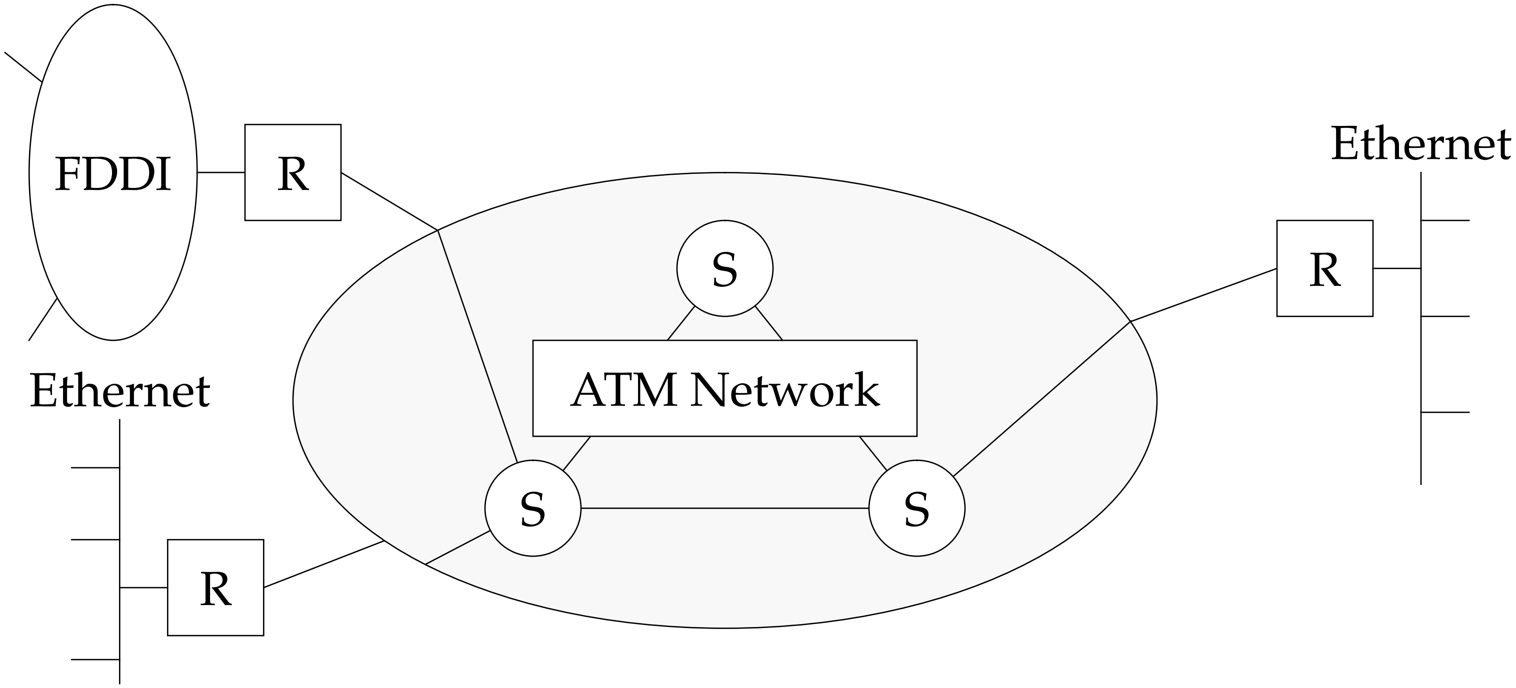

ATM growing in popularity, but Internet is ubiquitous and heterogeneous.

Want to use these types of networks together efficiently.

IP and ATM: Critical Differences

	Asynchronous Transfer Mode	Internet Protocol
Connections	Yes	No
Service Model	Quality of service support, can provide performance guarantees	Best-effort
Packets	Small, fixed-size cells	Variable-sized packets

Problems to be Solved

ATM QOS in an IP Internetwork

- How can IP applications benefit from ATM quality of service support?

Multiplexing

- What packets should share a virtual circuit?

Virtual Circuit Management

- When should virtual circuits be created and torn down?

Other Issues (not addressed here)

- Routing
- Address Resolution
- Multicast Support

Results

ATM QOS

QOS support can be helpful, if used carefully.

Multiplexing

Multiplexing eliminates virtual circuit setups, can help application performance.

Virtual Circuit Management

Caching idle virtual circuits can improve both network and application performance.

Outline

- ☞ Design Alternatives
- An Internet Simulated ATM Networking Environment
- Methodology
- Results and Analysis
    - ATM Quality of Service
    - Multiplexing
    - Virtual Circuit Management
- Conclusions

IP over ATM: A Software View

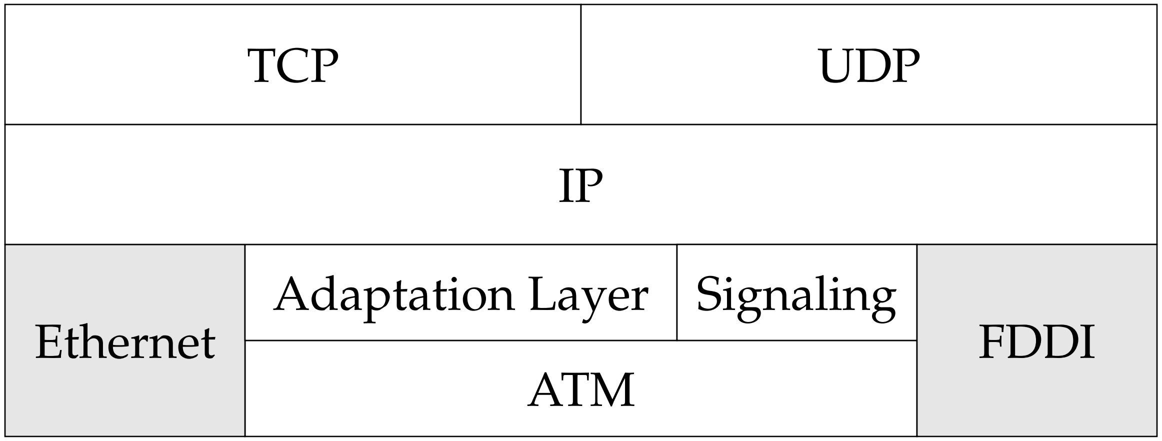

ATM stack treated by IP as a datalink layer.

Device driver puts IP packets in AAL frames, establishes virtual circuits as needed.

Policies and Design Alternatives

- ATM Quality of Service
  - Different service disciplines
  - Preference to different applications
- Multiplexing
  - ...per conversation
  - ...per application type per host pair
  - ...per router pair
- Virtual Circuit Management
  - Permanent virtual circuits
  - Switched virtual circuits
  - Switched virtual circuits with caching

IP over ATM Policy Space

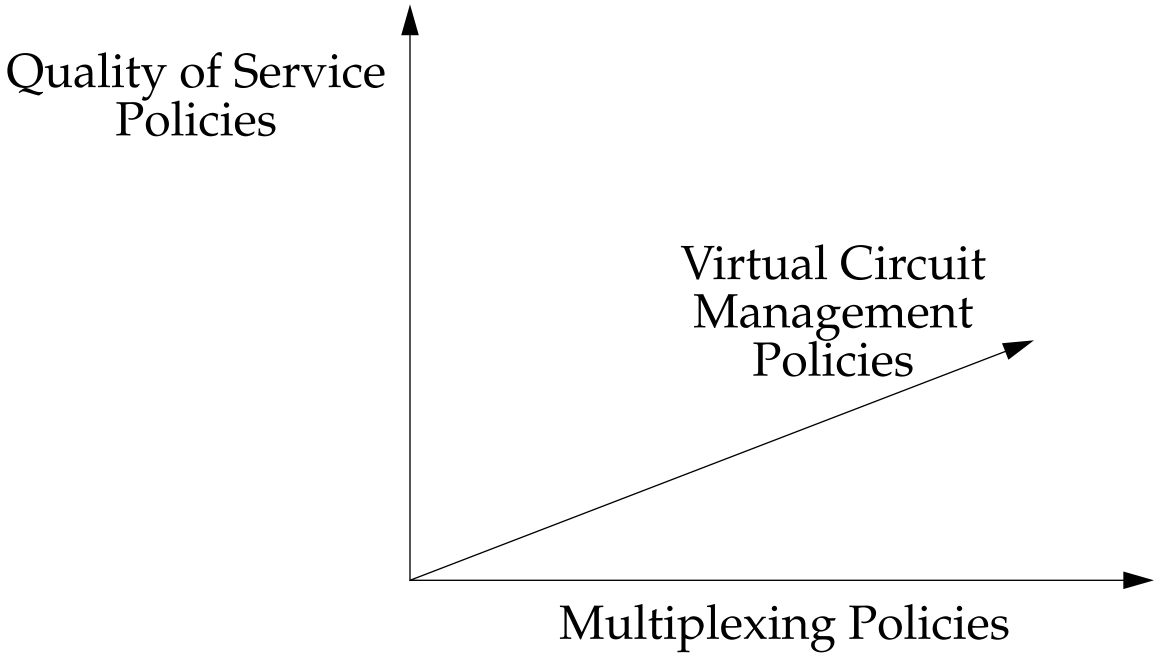

Policies to address problems: a three-dimensional space.

An Internet Simulated ATM Networking Environment (INSANE)

Functional Requirements

- ATM (cell transport, adaptation layer, signalling)
- TCP features (slowstart, congestion control, fast retransmit, etc.)
- Synthetic workload with application-specific traffic patterns

Logistical Constraints

- Easy to configure for large scenarios (1000+ hosts)
- Fast (hour-long simulations in reasonable time)
- Batch processing, off-line analysis

INSANE Protocol Stack

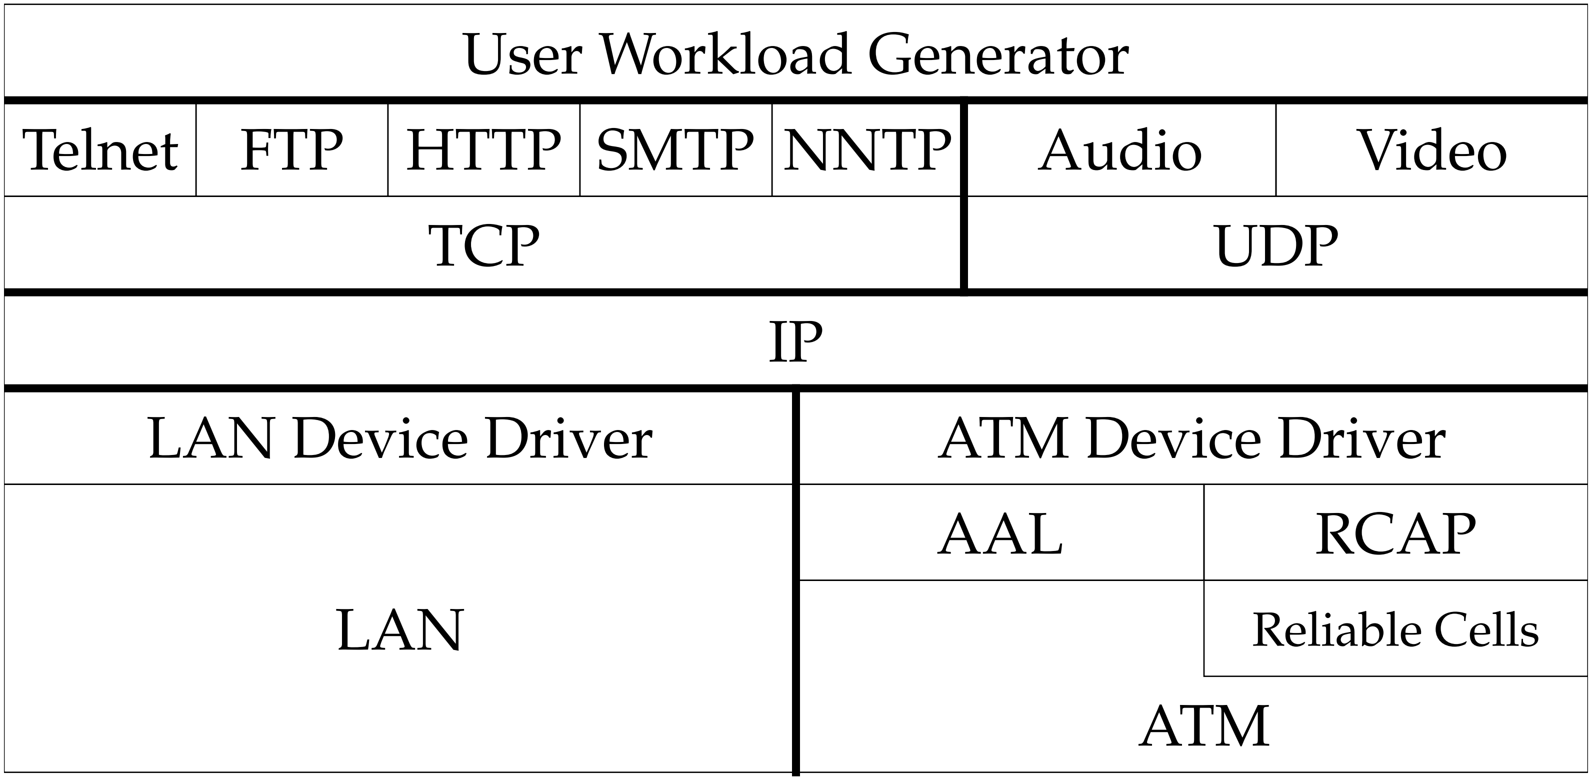

User Workload Generator

Telnet FTP HTTP SMTP NNTP Audio Video

TCP UDP

IP

LAN Device Driver ATM Device Driver

AAL RCAP

LAN Reliable Cells

ATM

Performance Notes

Hardware Platform	Running Time
Sun Ultra 1	3X
Sun Sparcstation 10	11X
100 MHz P5 (FreeBSD 2.1.0-RELEASE)	8X
DEC Alpha AXP 3000/400	8X

300 simulation runs (4000 seconds each)

1500 hours of CPU time on Sun Ultra 1 cluster

Estimated 10 GB raw data

Environment

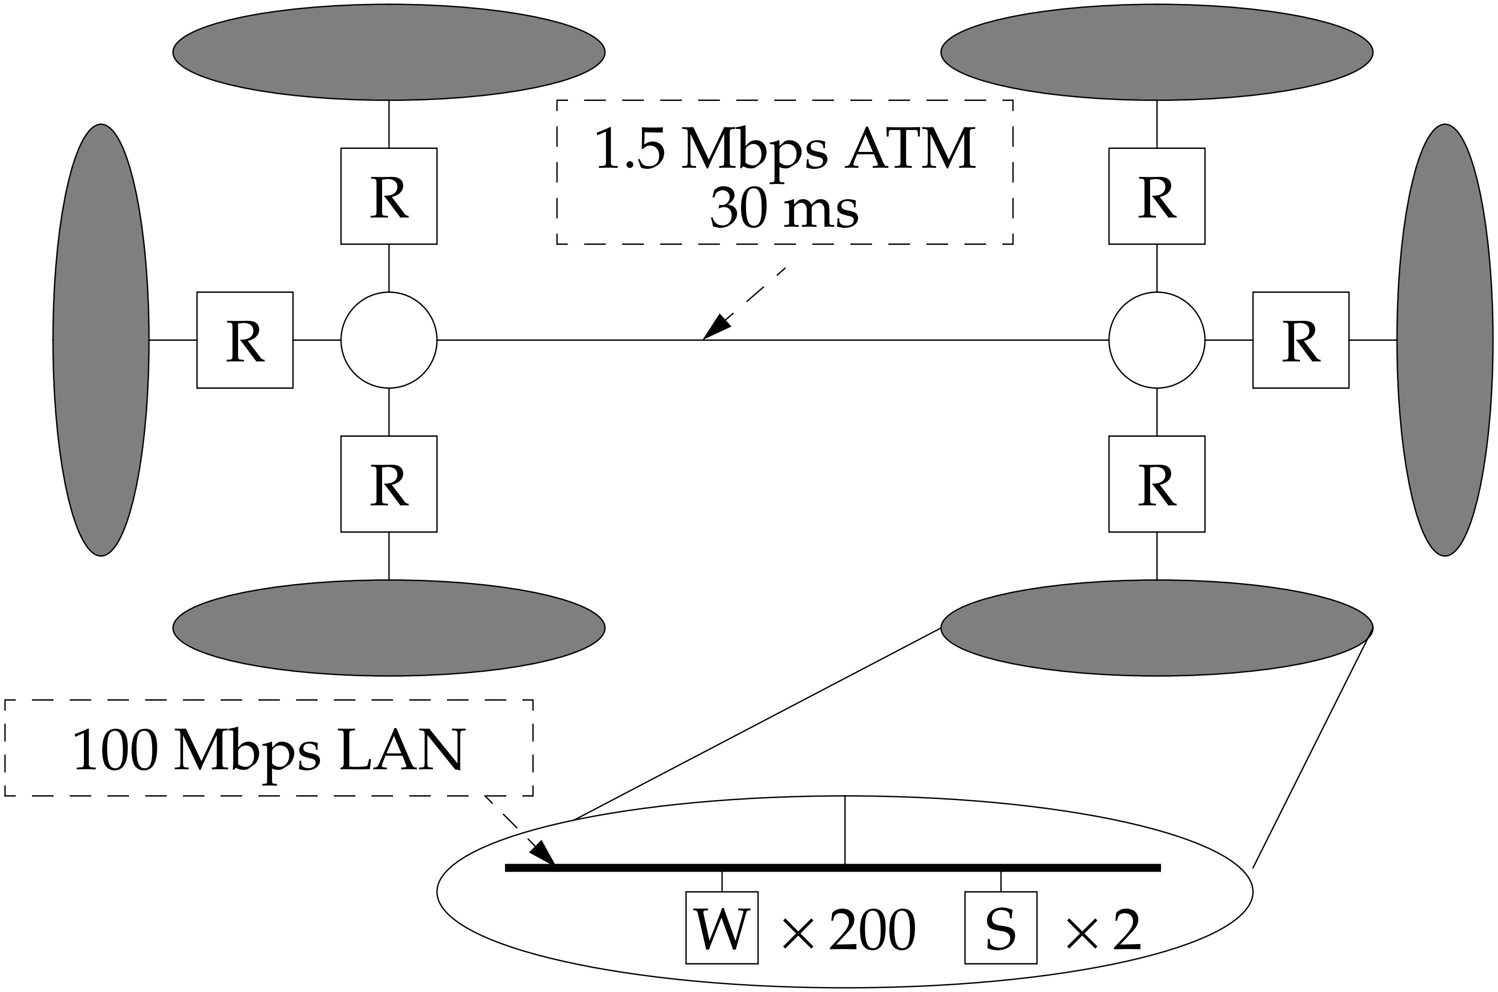

Workload

Application	Source	Interarrival Time (MM:SS, per site)
telnet	tcplib	0:10
FTP	tcplib	0:15
HTTP	empirical	0:05
Audio	tcplib	10:00
Video	empirical	10:00
SMTP	tcplib	0:04
NNTP	tcplib	3:45

Performance Metrics

Objective: Measure performance effects visible to applications and users.

telnet	Setup and round-trip time
FTP	File and session response times
HTTP	File, Web page transfer times
Audio	Loss rate, one-way end-to-end delay
Video	Loss rate

Outline

- Design Alternatives
- An Internet Simulated ATM Networking Environment
- Methodology
- ☞ Results and Analysis
  - ATM Quality of Service
  - Multiplexing
  - Virtual Circuit Management
- Conclusions

ATM Quality of Service

Schedulers

- Best Effort (First-Come-First-Served)
- Static Priority
- Rate-Controlled Static Priority
- Rate-Controlled Static Priority (Rate Jitter Control)

QOS parameters assigned based on application type

Results

- Static priority can give preference, but starvation a danger
- Rate control for bulk transfers yields inconclusive results
- Rate jitter control reduces losses in long TCP bulk transfers

High-Priority Telnet

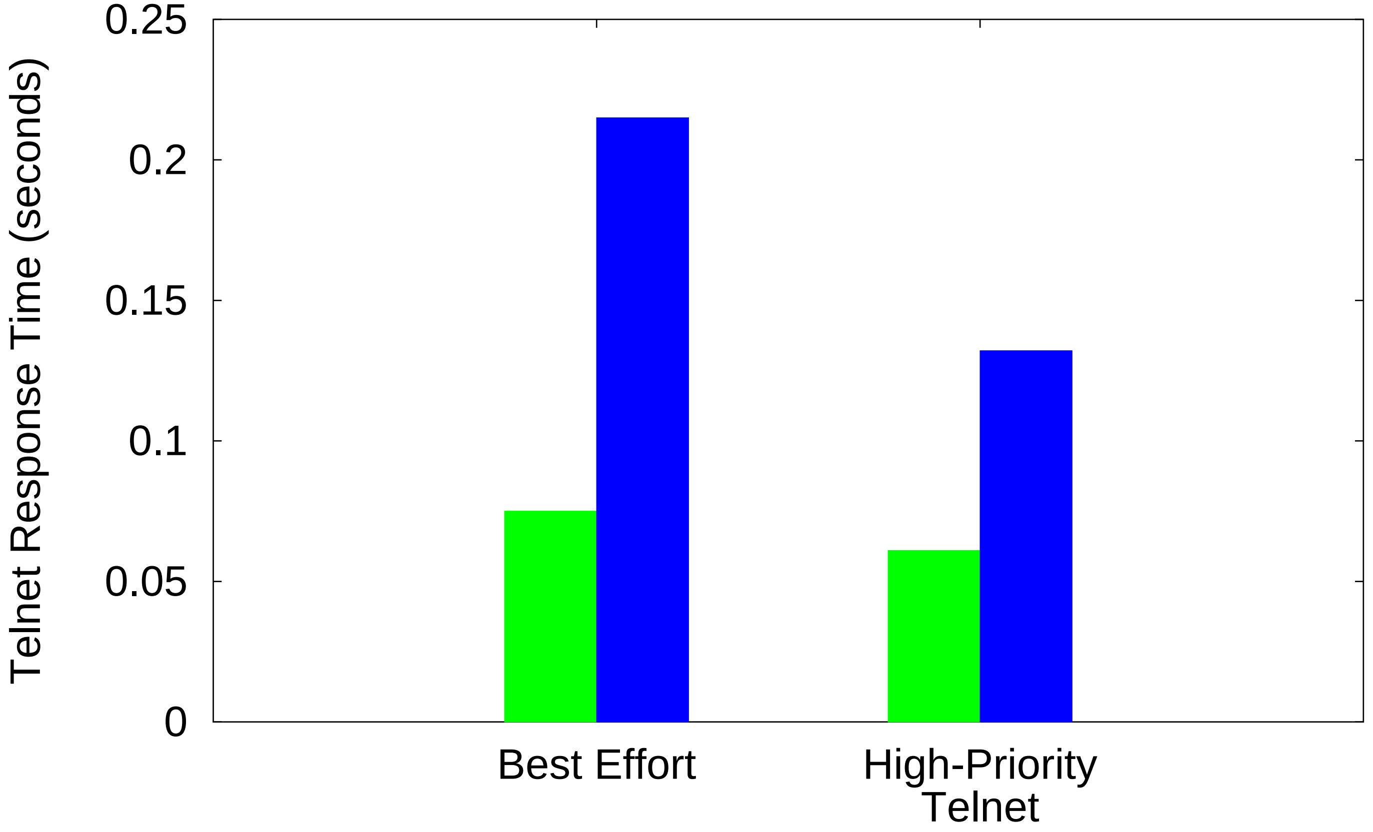

Interactive performance can be improved.

High-Priority FTP

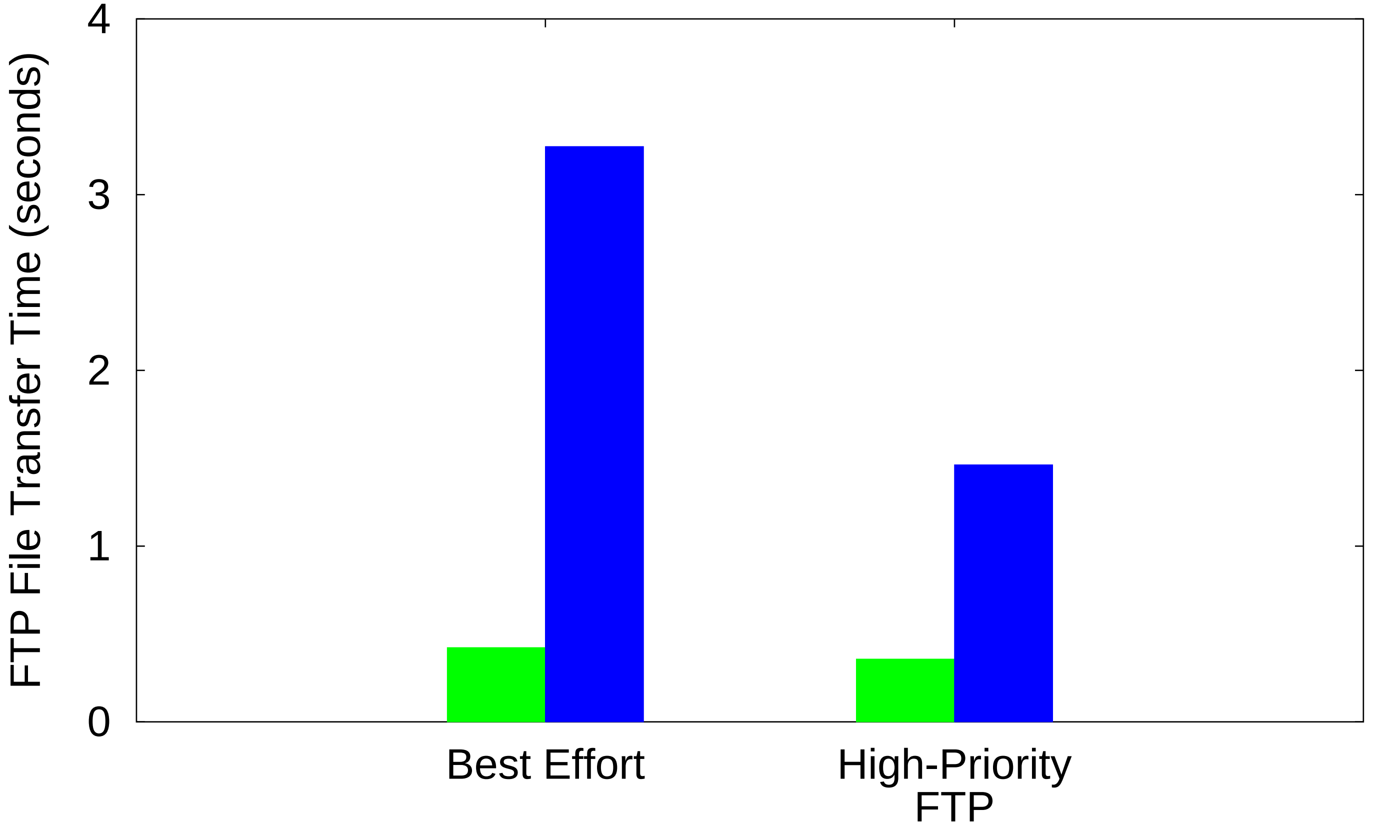

FTPs given high priority take less time.

High-Priority FTP

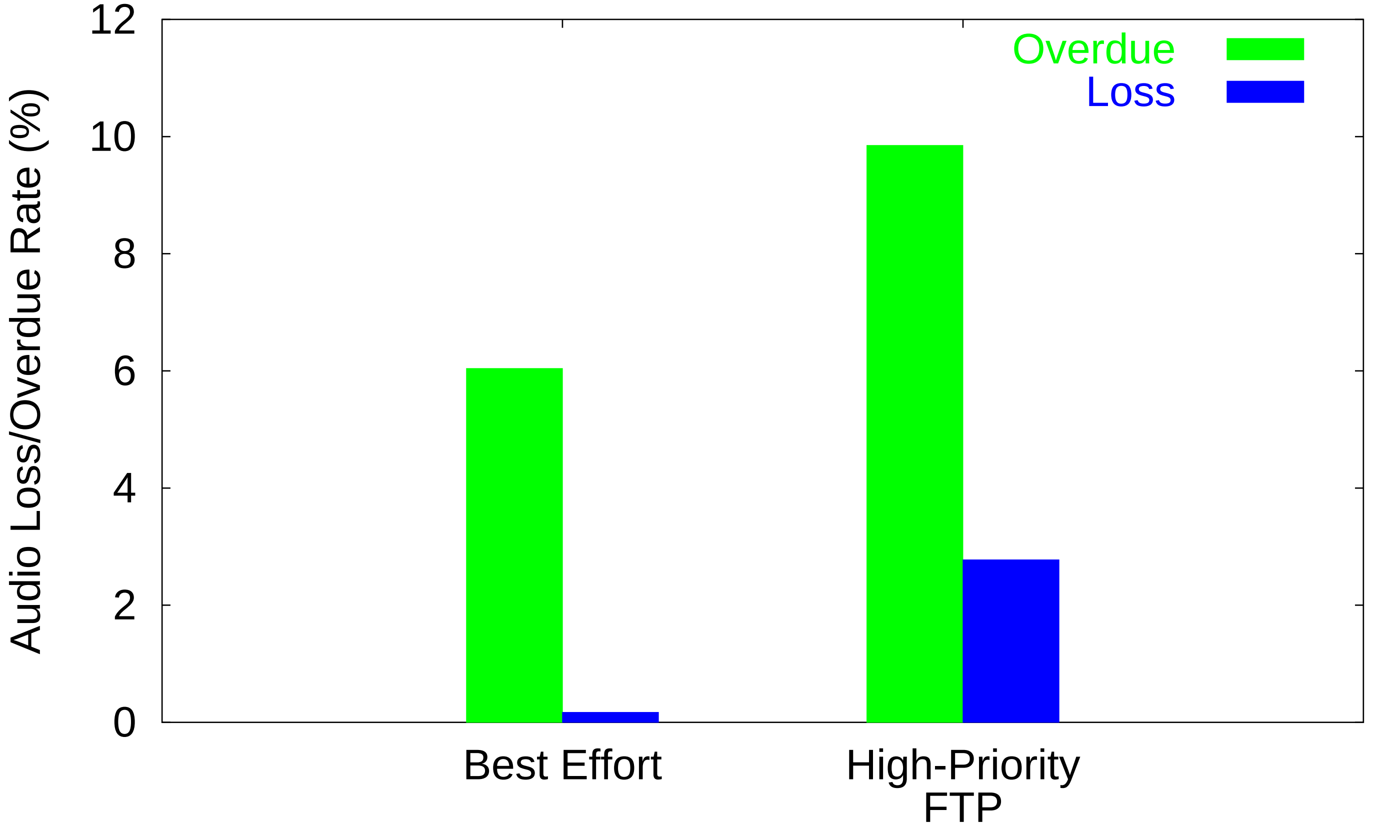

High-priority FTP can degrade others’ performance.

High-Priority Applications

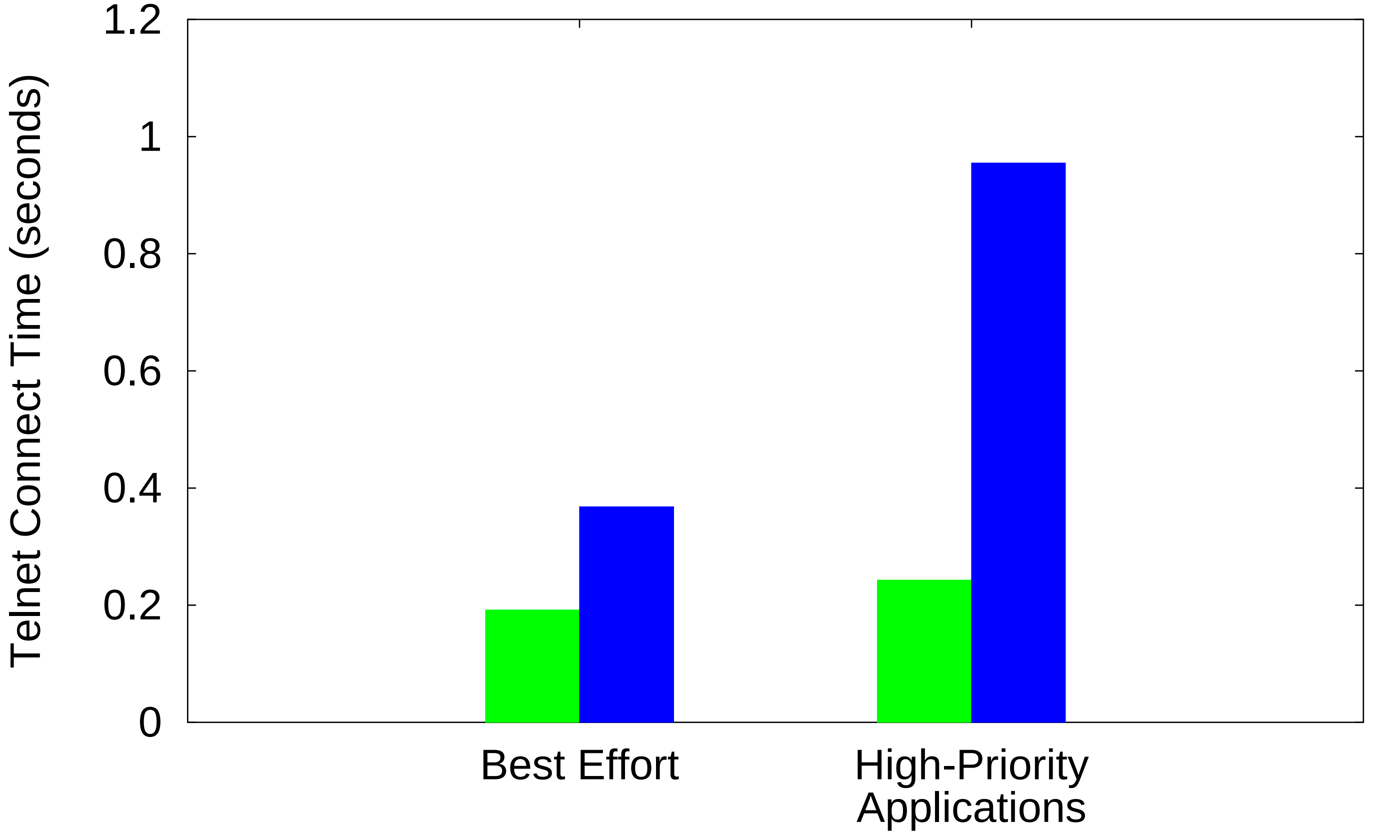

Signalling starved by higher priority traffic.

RCSP Rate-Limiting of Bulk Transfers

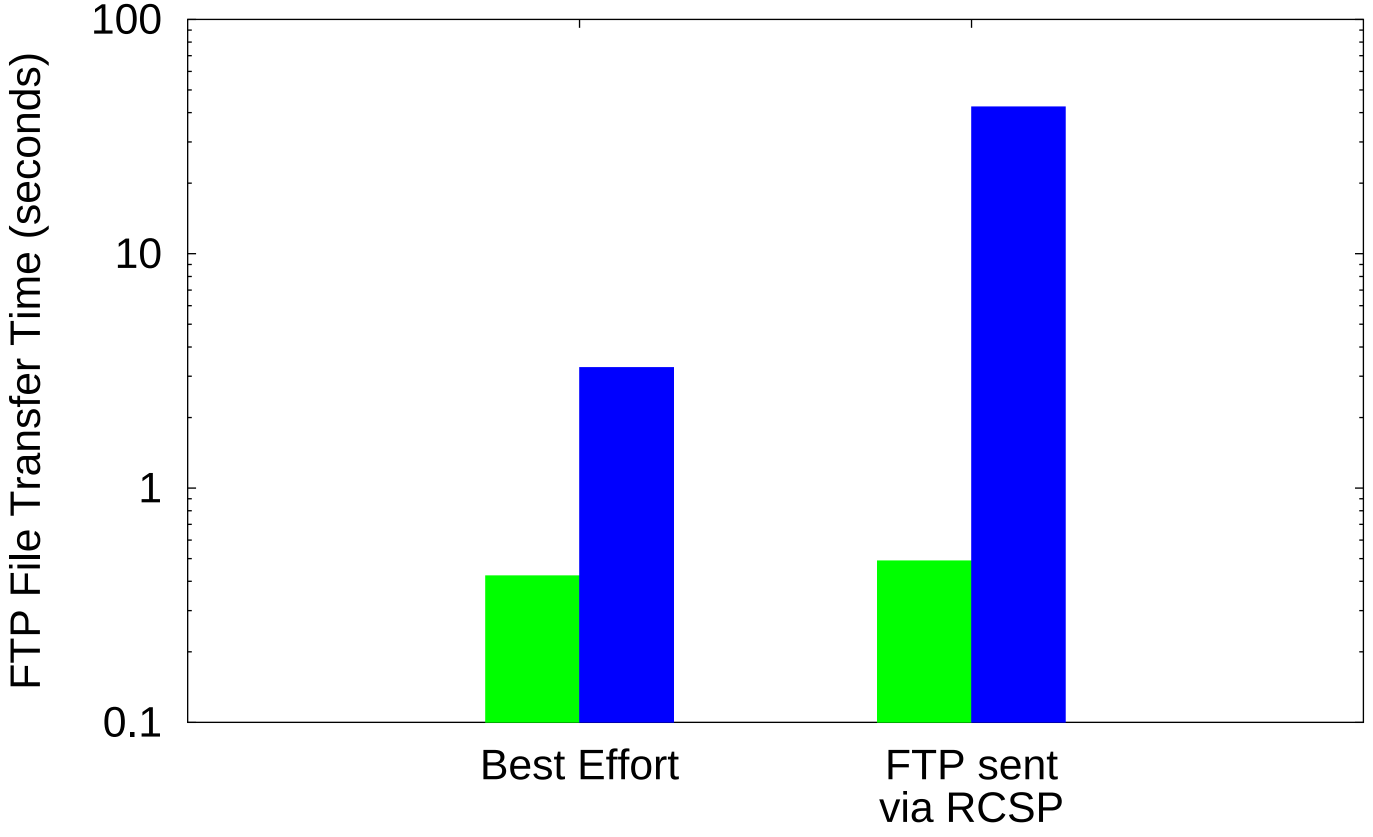

RCSP can be used to constrain bulk transfers.

RCSP Rate-Limiting of Bulk Transfers

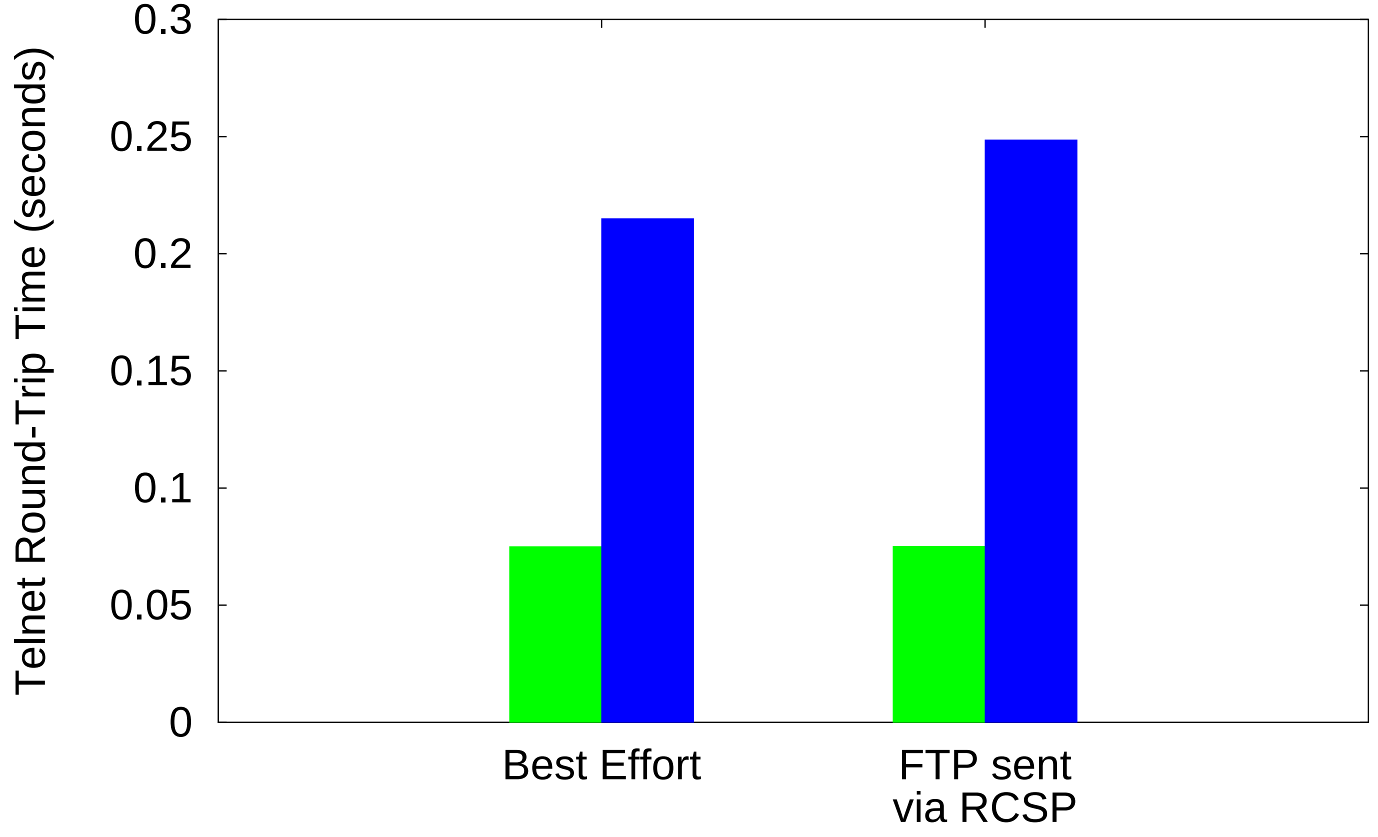

Rate-control effects on best-effort applications unclear.

Jitter Control Effects

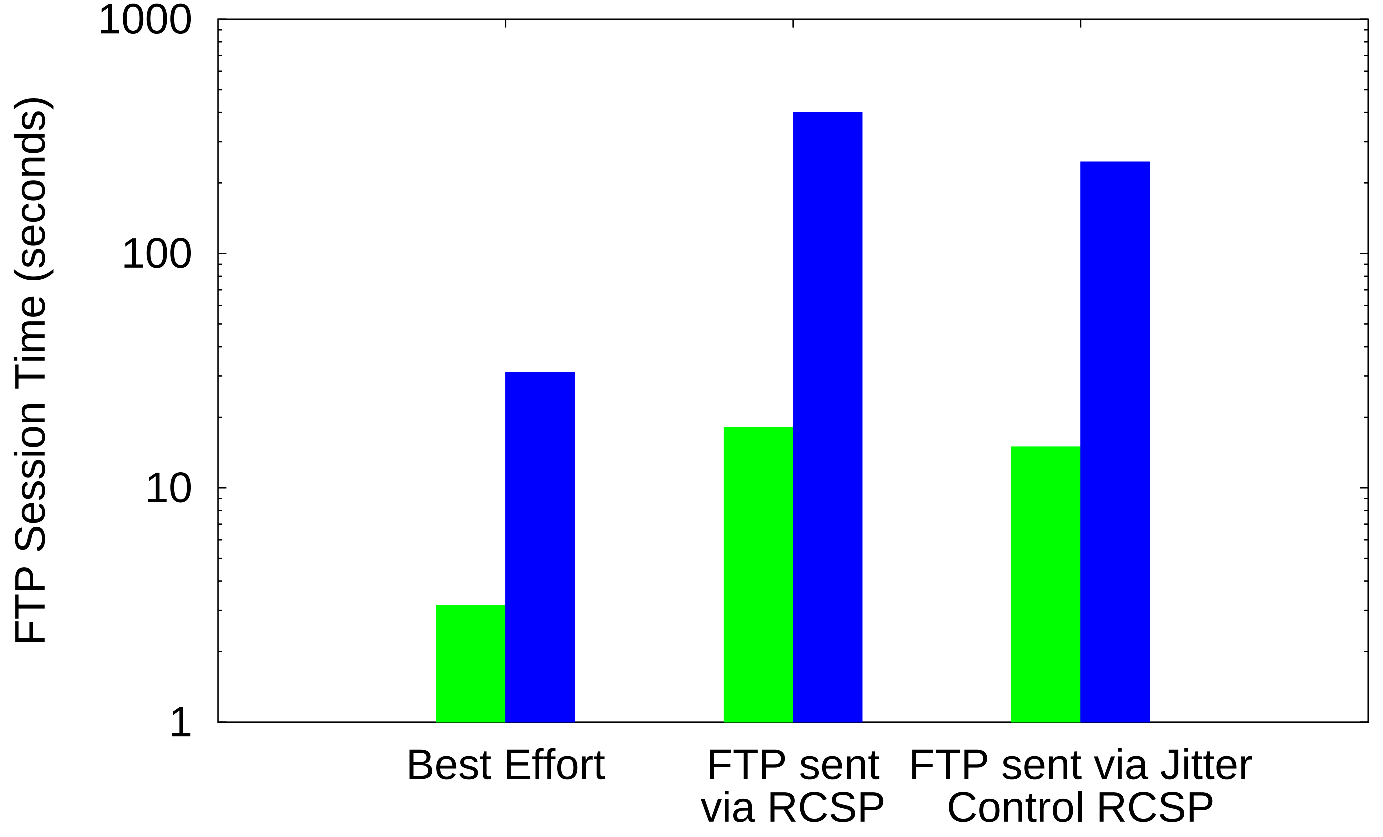

Bursts smoothed out by delay jitter control, fewer losses.

Outline

Design Alternatives

An Internet Simulated ATM Networking Environment

Evaluation

Results and Analysis

- ATM Quality of Service
- ☞ Multiplexing
- Virtual Circuit Management

Conclusions

Multiplexing Policies

Different levels of traffic aggregation

- Virtual circuit per conversation (e.g. TCP connection)
- Virtual circuit per application per end-host pair
- Virtual circuit per router pair

Results

- Multiplexing eliminates virtual circuit setups
- Interaction with policing degrades performance of long FTP, Web transfers
- Buffer contention causes losses with per-router pair multiplexing

Aggregation Improves Transfer Times

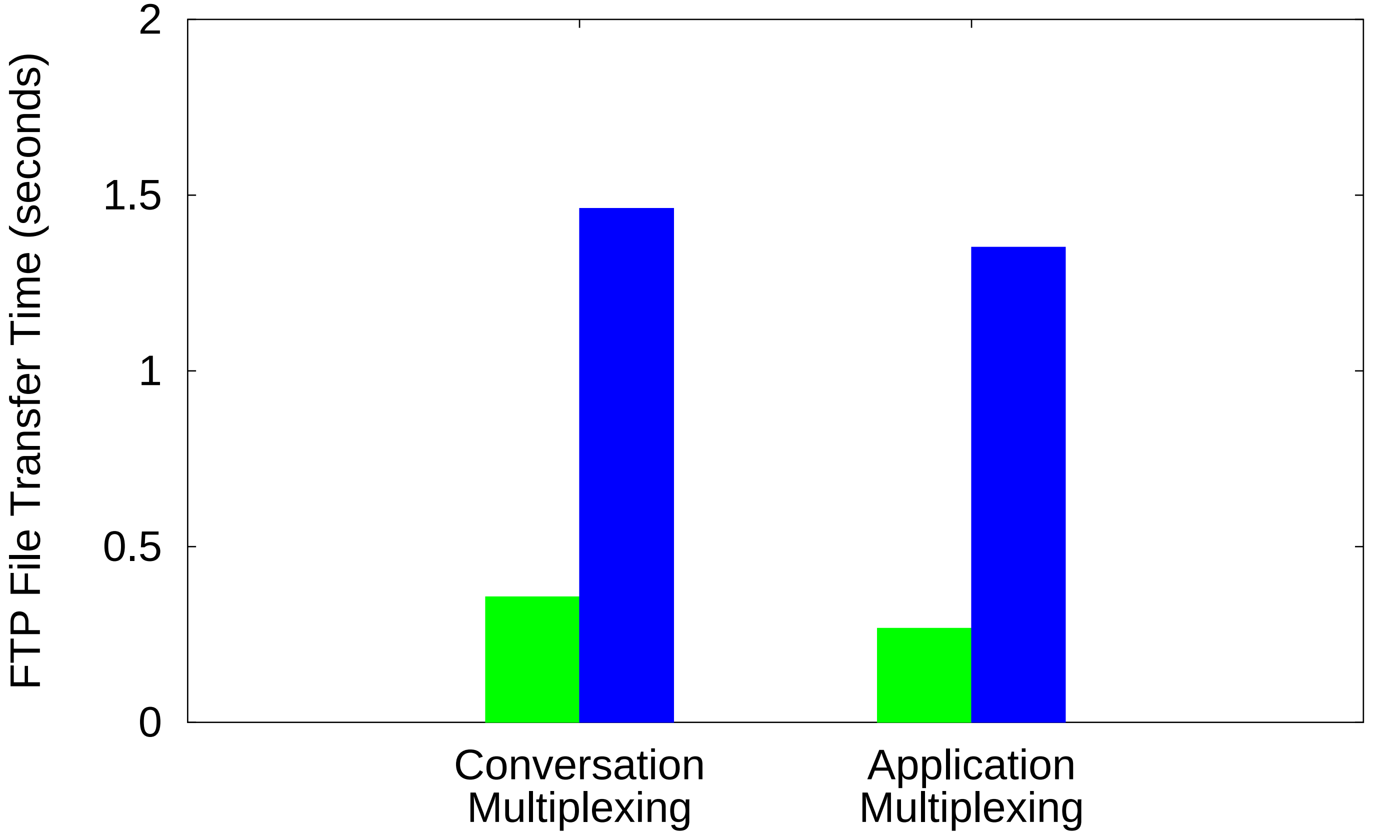

Reduced need for virtual circuit setups shortens file transfers.

Per-Application Interference

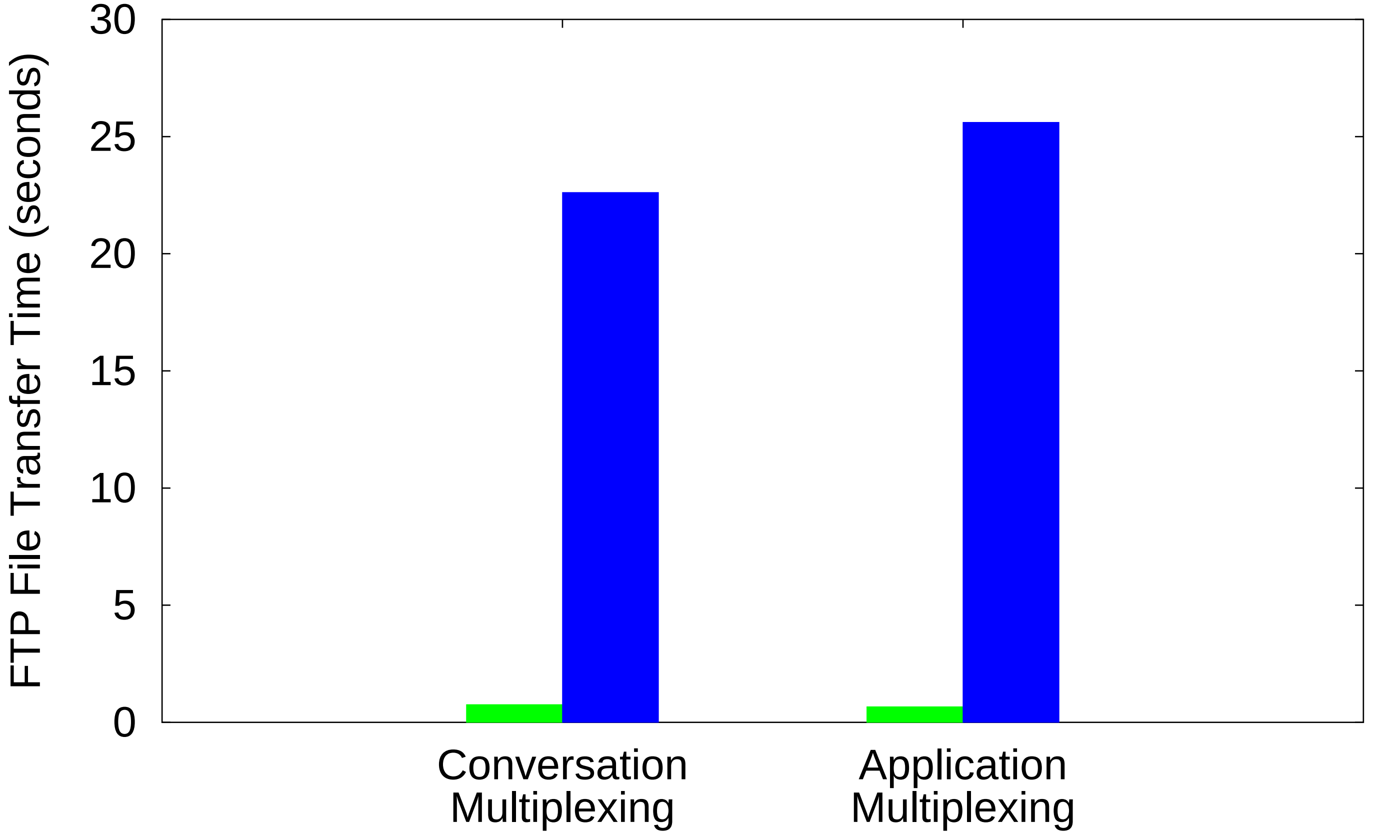

Previous gains disappeared or reversed themselves.

Per-Application Interference

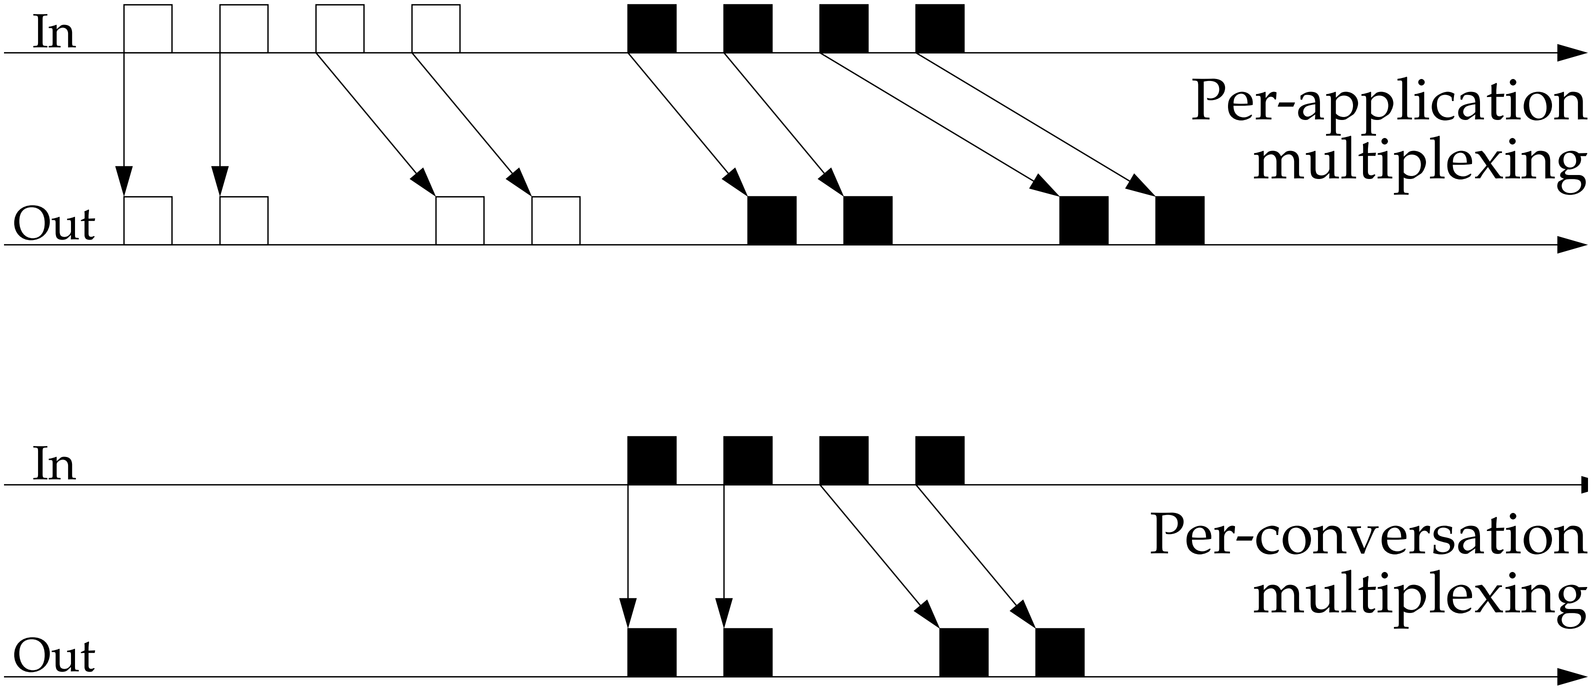

Rate control policing causes cells from one IP conversation to delay cells from another, even if not simultaneous.

Per-Router Multiplexing

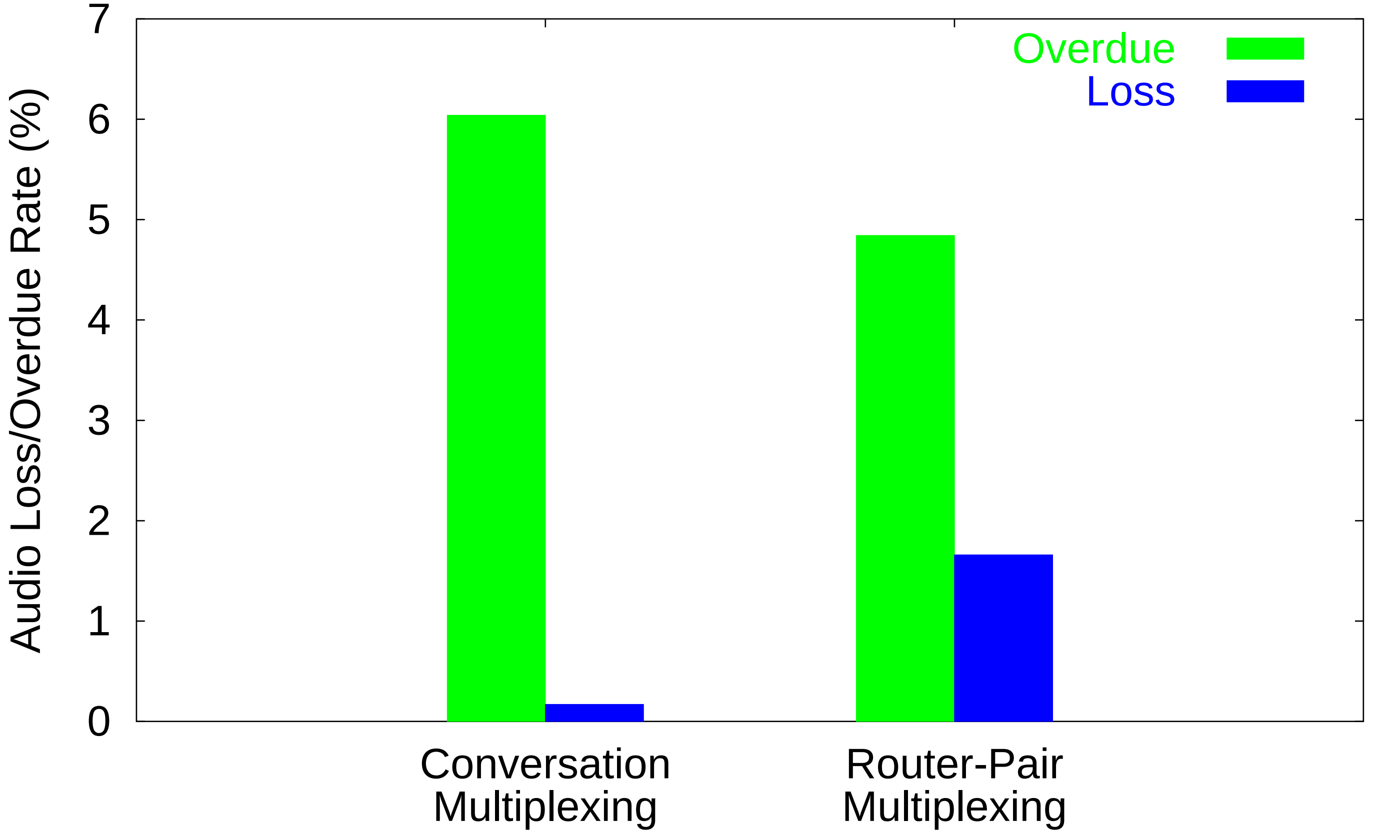

Loss rate multiplied by 3–8, due to buffer contention.

Outline

Design Alternatives

Evaluation

An Internet Simulated ATM Networking Environment

Results and Analysis

- ATM Quality of Service
- Multiplexing
- ☞ Virtual Circuit Management

Conclusions

Virtual Circuit Management Policies

Different policies for managing virtual circuits

- Permanent virtual circuits (set up at network start)
- Switched virtual circuits (set up on demand, torn down when idle 10 seconds)
- Switched virtual circuits (set up on demand, cache when idle 10 seconds)

Results

- Caching idle virtual circuits can eliminate about 95% of setups
- Significant speedups for applications doing repeated TCP connections

Virtual Circuit Cache

- Establish virtual circuit when needed
- After 10 seconds idle
  - Router “unbinds” virtual circuit from conversation(s)
  - Available for reuse
  - Match destination and QOS parameters
- After 300 seconds (5 minutes) idle
  - Tear down virtual circuit

Cache Statistics

Hit rates

For best-effort traffic	97–98%
For most other schemes	92–95%
Worst case	56–69%

Signalling Rates (setups per second, best-effort traffic)

	No caching	Caching
Per-application	3.0	0.06
Per-conversation	7.0	0.18

Per-router overhead: 20–40% more virtual circuits open.

Connection Caching and HTTP

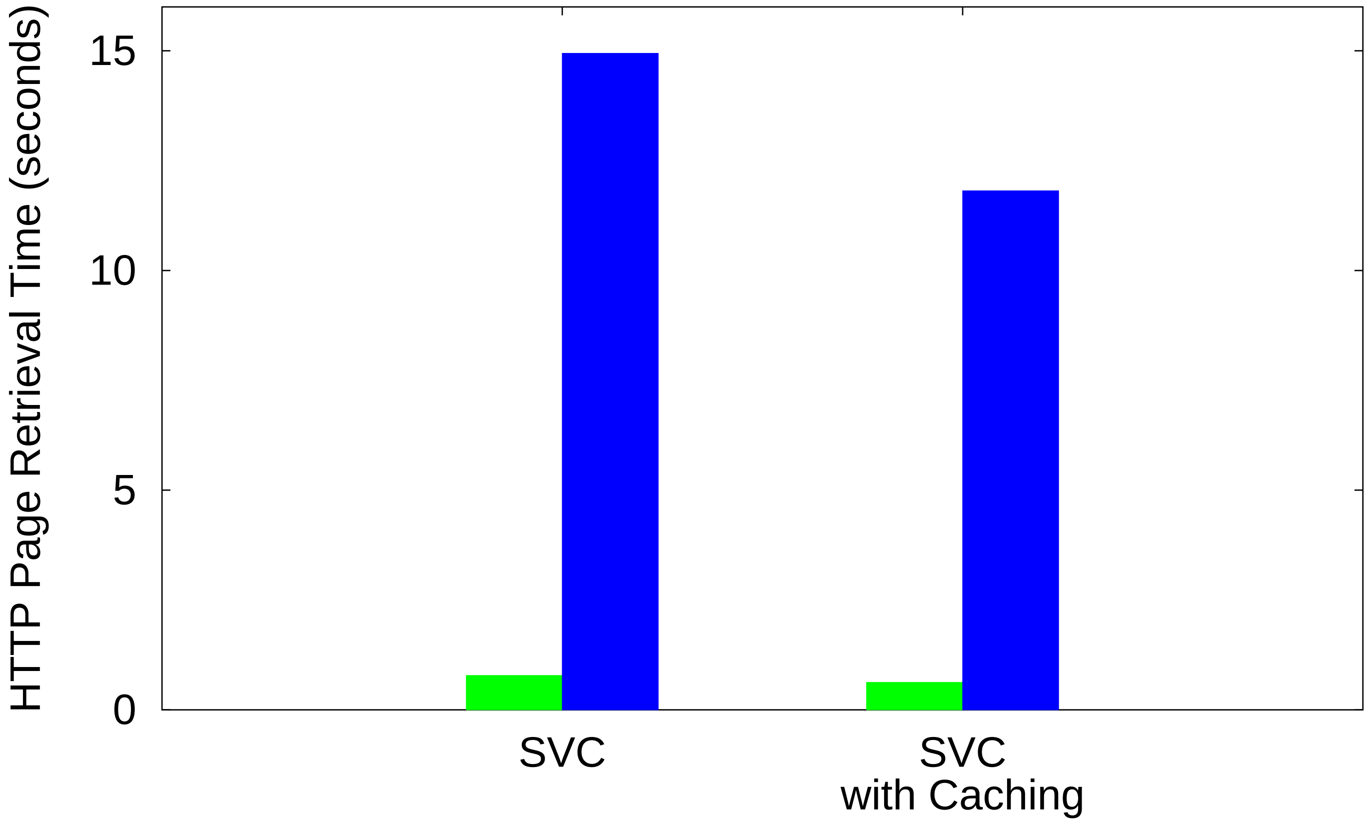

Cache has a beneficial impact on application performance.

Summary of Results

ATM Quality of Service

- Static priority effective for giving preference to any application, but starvation a danger
- Overall effects of rate control inconclusive
- Delay jitter control reduces losses for long bulk transfers

Multiplexing (traffic aggregation)

- Multiplexing shortens small transfers by eliminating VC setups
- Rate control effects cause long transfers to interfere with each other
- Contention for buffering increases loss

Virtual Circuit Management

- Caching unused connections eliminates vast majority of VC setups
- Improves application performance

My Related Work

- Network Traffic Measurement and Models
  - World Wide Web
  - Internet Video
  - IP Multicast
- The Tenet Real-Time Protocol Suite
  - First set of network protocols to give performance guarantees over a packet-switched internetwork

Conclusions

ATM quality of service can help Internet applications, if used carefully.

Multiplexing of traffic can help applications by reducing virtual circuit setups.

Virtual circuit caching improves application and network performance.

END OF TALK

(remaining slides used for answering questions)

Real-Time Guarantees

At connection setup time...

- “User” provides traffic specification, performance requirements
- Network performs admission control tests, returns “accept” or “deny”

During data transmission...

- Queues in network queue according to a particular service discipline
- Best-effort traffic fills “gaps” between cells for guaranteed connections

Rate-Controlled Static Priority

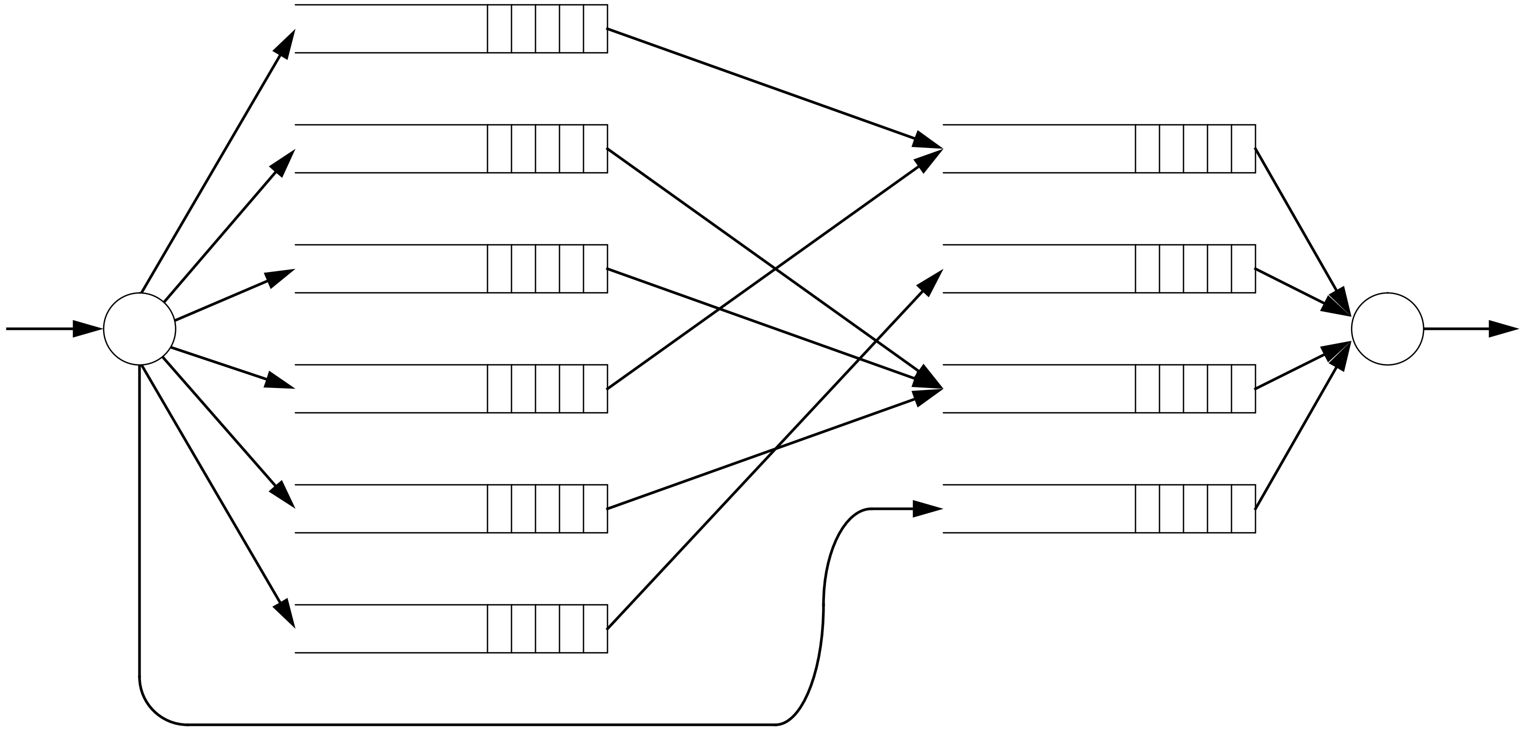

Input

Per-Connection
Rate Controllers

Priority Queues

Output

Static Priority Scheduling

Level	Type of Traffic
0	
1	Telnet
2	FTP control, Audio
3	FTP Data, HTTP, Video
Signalling	
Best-effort	

Work-Conserving RCSP

Application	Dir	Local Delay (ms)	Peak Rate (Kbps)	Average Rate (Kbps)	Avg. Interval (ms)
telnet	up	10	19.2	7.68	500
	down	10	38.4	19.2	1000
FTP (control)	up	20	3.84	0.77	5000
	down	20	3.84	0.77	5000
FTP (data)	up	80	3.84	3.84	2000
	down	80	19.2	19.2	2000
HTTP	up	40	7.68	3.84	10000
	down	40	96.0	48.0	1000
audio	any	20	96.0	76.8	100
video	any	50	148	110	2000

RCSP Rate-Limiting of Bulk Transfers

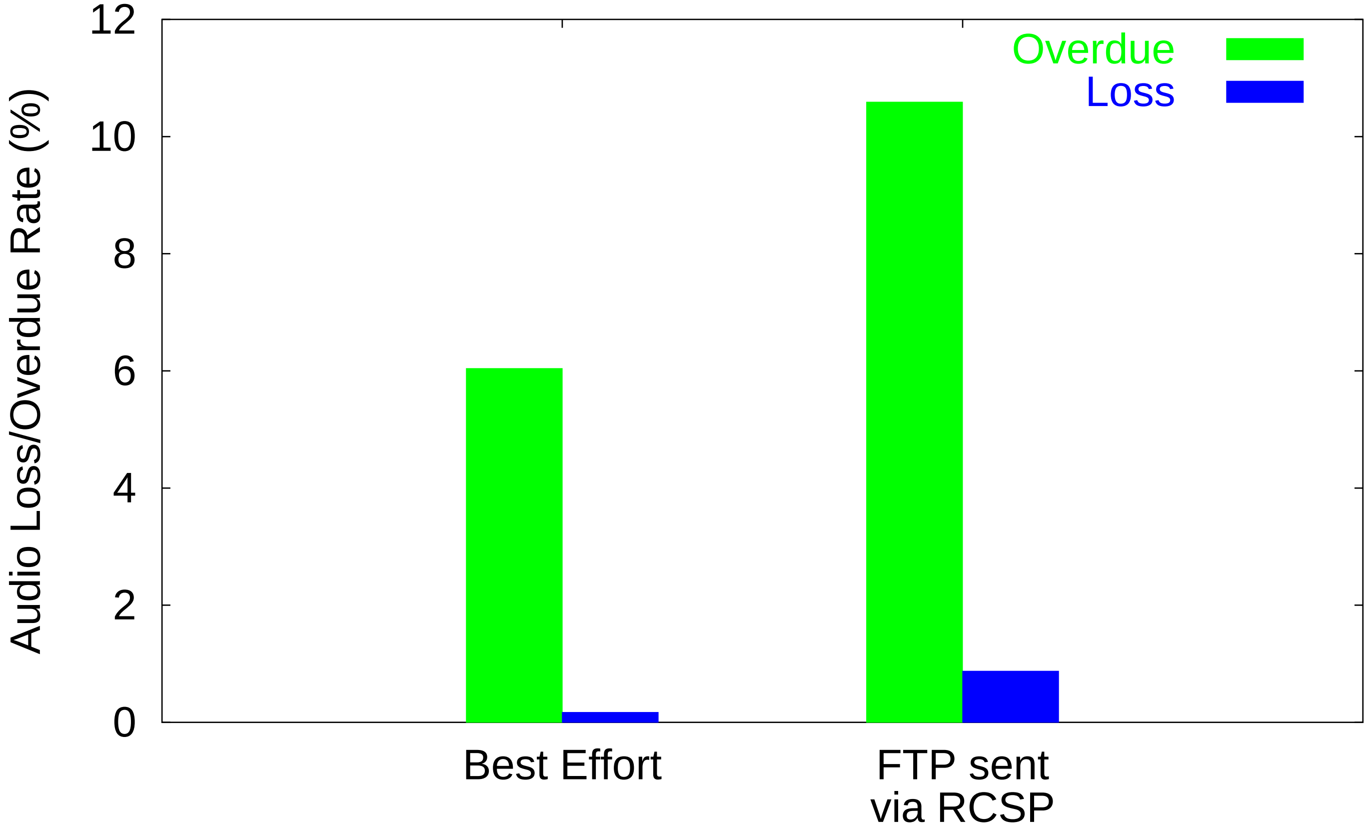

In most cases, no statistically significant effects.

ATM Switch Composite Object

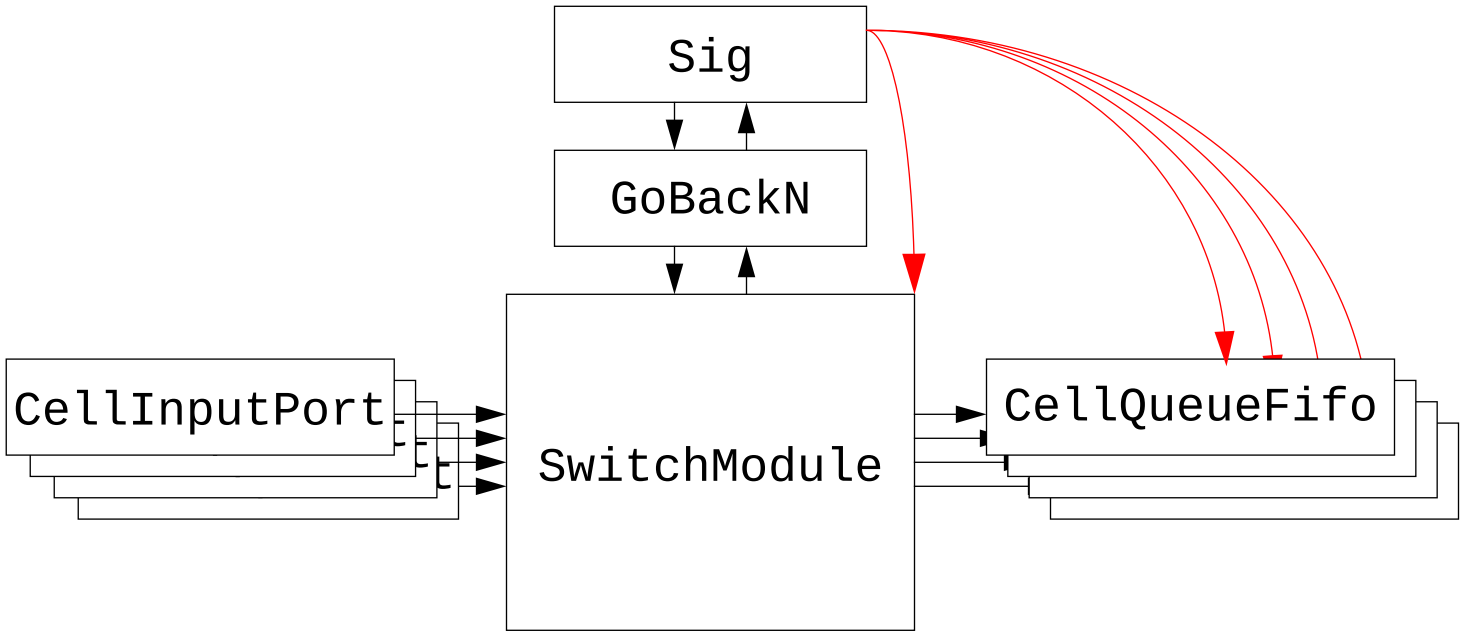

Composite objects created by invoking Tcl scripts

```
SwitchFifoN switch 8 1024 1000000
```

Interaction with Objects

Command line interface like debugger

Interact with objects and examine state

```
] router1.ip netstat -r
Destination      Mask             Gateway          Flags
128.32.150.0     255.255.255.0    128.32.150.254
128.32.131.0     255.255.255.0    128.32.131.254
127.0.0.1        255.0.0.0        127.0.0.1
] router1.ip netstat -i
Name             Address          Netmask          Ipkts Ierrs Opkts Oerrs
router1.lan1     128.32.150.254   255.255.255.0        0     0     1     0
router1.lan0     128.32.131.254   255.255.255.0        1     0     0     0
router1.lo0      127.0.0.1        255.0.0.0            0     0     0     0
```

‘State of the Art”

- XUNET II
    - Best-effort
    - Per-router-pair multiplexing
    - Permanent virtual circuits
- FORE Systems ATM LAN
    - Best-effort
    - Per-router-pair multiplexing
    - Permanent or switched virtual circuits (no caching)
- Ipsilon IP Switching
    - Best-effort (RSVP support in future)
    - Various multiplexing policies
    - Permanent or switched virtual circuits (not end-to-end, no caching)