

AN OVERVIEW OF HARDWARE ISSUES FOR IP AND ATM

"Name one thing you could achieve with ATM that you couldn't with IP!"



Nick McKeown

Assistant Professor of Electrical Engineering
and Computer Science
Stanford University

nickm@ee.stanford.edu
<http://ee.stanford.edu/~nickm>

Outline

Part I


- The need
- Trends in ATM switches
- Trends in IP routers
- Merging of the two

Part II

- Some key technologies
- Summary

2/25: Merging IP and ATM


What's the Problem?



3/25: Merging IP and ATM

The demand


The San Jose NAP



Source: <http://www.mfsdatanet.com/MAE/west.stats.html>

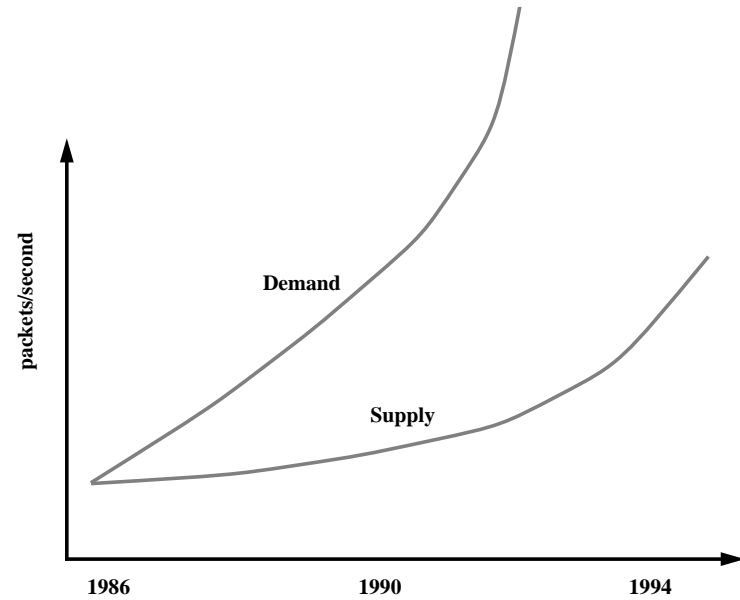
4/25: Merging IP and ATM

The supply




5/25: Merging IP and ATM

Why we need faster routers



6/25: Merging IP and ATM

The race is on...




Ascend (Netstar), Ipsilon, Toshiba, BBN, [Cisco, Bay, Juniper, Torrent] ?.....

7/25: Merging IP and ATM

Trends in ATM switches


Generic ATM Switch:



8/25: Merging IP and ATM


Trends in IP routers

Generic IP Router:




9/25: Merging IP and ATM

Merging of IP & ATM



10/25: Merging IP and ATM


Merging of IP & ATM



11/25: Merging IP and ATM


Trends in IP routers

Trend 1: Move CPU off forwarding path



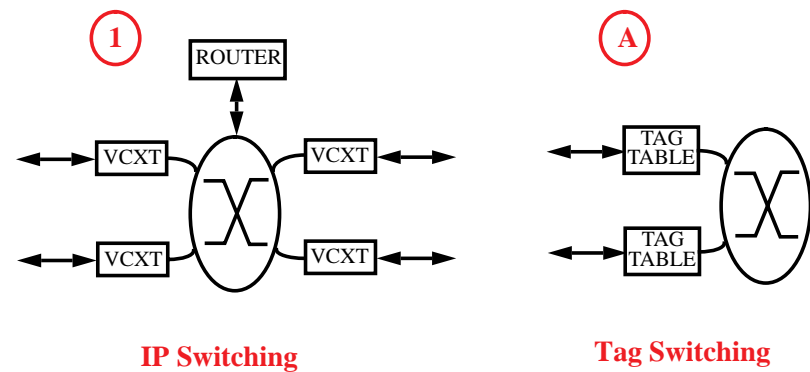
12/25: Merging IP and ATM

Trends in IP routers
Trend 2: AVOID SHARED BUS




13/25: Merging IP and ATM

Consequence



14/25: Merging IP and ATM

The fundamental hardware difference between IP & ATM



BUT:

ATM VCI \cong 16 bits AND we get to choose!

15/25: Merging IP and ATM

Claim

If we could do FAST longest prefix matches
Then we wouldn't be here!

16/25: Merging IP and ATM

Merging of IP & ATM

Why?

If you could get the whole IP forwarding table in fast memory (and update it invisibly!) then who needs ATM?

Removing IP forwarding engine from the datapath is one thing: but still need Quality Processor on the critical path.

Large number of individual flows (ISMP Flow Type 1)

=> aggregation onto coarser src-dest flows (Type 2)

=> need reassembly/frame-mode switches

17/25: Merging IP and ATM


Some key hardware technologies

- Memory bandwidth
- Serial link technology

- Special-purpose memories
- CPU vs special-purpose processors

18/25: Merging IP and ATM


Memory Bandwidth



19/25: Merging IP and ATM

Trends in ATM switches

How fast can this architecture go?




How fast can a 16 port switch operate with this architecture?

cell per port

20/25: Merging IP and ATM

Serial link technology

“wires are becoming a scarce resource”



Standard interfaces are a bottleneck:

- 100 Mbit/sec per pin is tough.
- 100 signal pins for 10 Gbit/sec.
- Large chip packages and board connectors drive up costs.

Serial interfaces are efficient:

- 1 Gbit/sec per pin.
- 10 signal pins for 10 Gbit/sec.

21/25: Merging IP and ATM

Serial link technology

Problems with High Speed Links


- **Precise timing needed**
 - Hard to determine which bit is which.
 - Distributed clocks have skew (1 ns/bit time at 1 Gb/s).
 - Solution: recover timing from data.
- **High speed signals**
 - Need to transmit and receive 1 Gb/s.
 - Solutions: low swing signals, good terminations on transmission lines.

Much ongoing R&D in this area.


22/25: Merging IP and ATM

Special purpose memories

1 FIFOs




2 Programmable FIFOs




23/25: Merging IP and ATM

CPU vs. special-purpose processor

1



2



24/25: Merging IP and ATM

Summary

- **Trends in switching**
 - Single-stage shared memory and busses are running out of steam.
 - Combined I/O Queueing and/or multistage switches are require.
- **Trends in routing**
 - Removal of CPUs from forwarding path.
 - Quality Processor still on the main forwarding path.
- **Key Technologies**
 - Memory Bandwidth
 - Serial link technology
 - Special-purpose memory
 - Special-purpose processors