Processing packets in packet switches

CS343
May 7th 2003

Nick McKeown
Professor of Electrical Engineering
and Computer Science, Stanford University

nickm@stanford.edu
www.stanford.edu/~nickm

1

Contents

2

The Network Layer View of the Internet

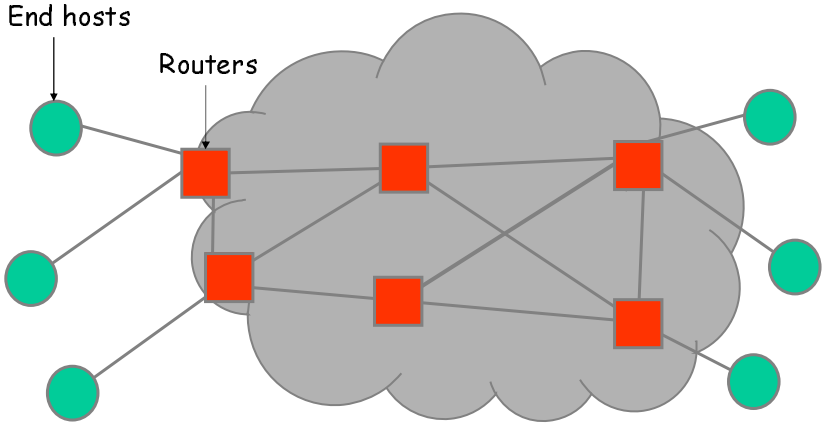

3

Hierarchical arrangement

A crude approximation

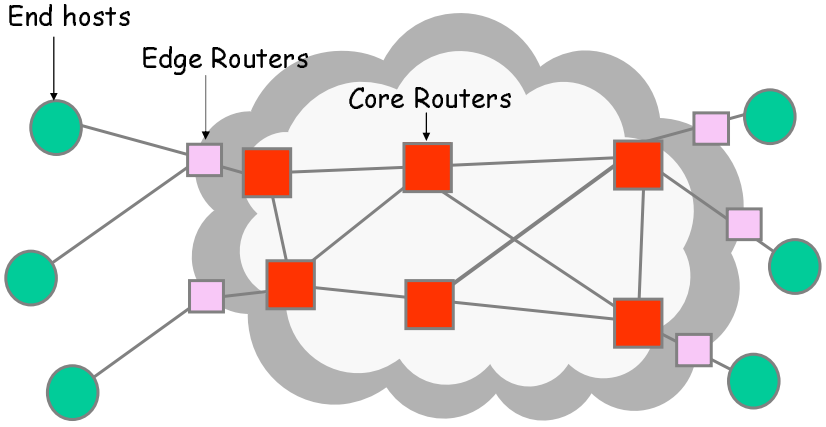

Core routers: Maximum capacity, minimum function.
Typically: 16 ports of 10Gb/s. Capacity 160Gb/s, 200Mpps. Price $1M.
Edge routers: Medium capacity, maximum flexibility and function.
Typically: 16 ports of 2.5Gb/s. Capacity 20-30 Gb/s, 10-20Mpps. Price $200k.

4

Hierarchical arrangement

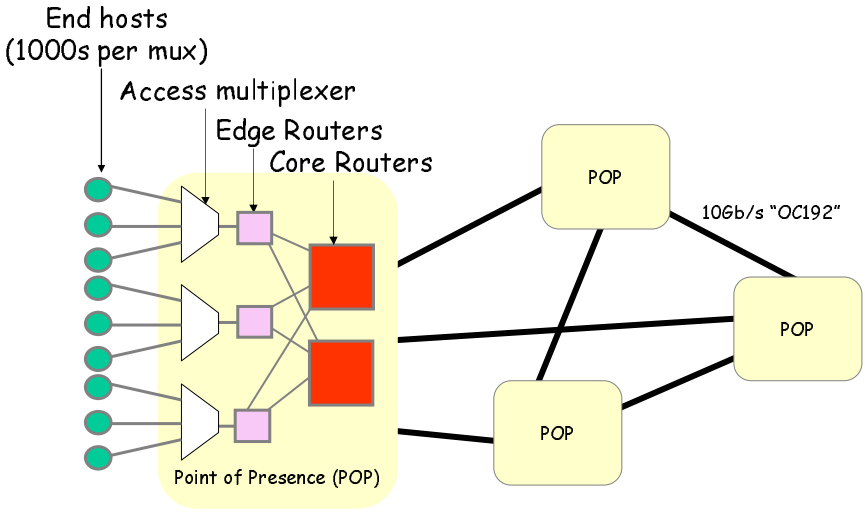

POP: Point of Presence. Richly interconnected by mesh of long-haul links.
Typically: 40 POPs per national network operator; 10-40 core routers per POP.

5

Autonomous Systems

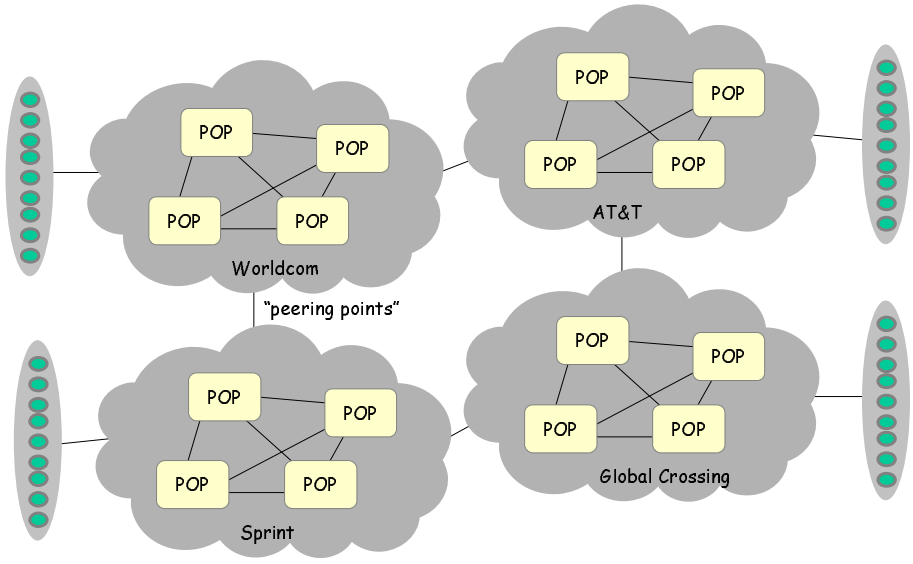

6

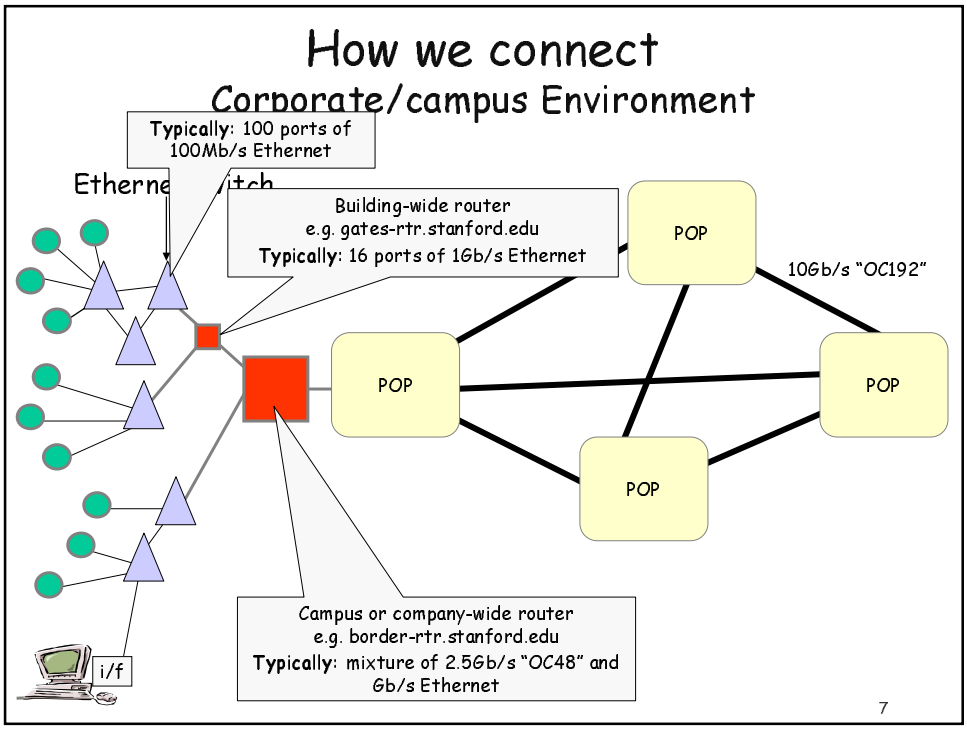
How we connect
Corporate/campus Environment
Typically: 100 ports of 100Mb/s Ethernet
Etherne
itch
Building-wide router
e.g. gates-rtr.stanford.edu
Typically: 16 ports of 1Gb/s Ethernet
POP
10Gb/s "OC192"
POP
POP
POP
Campus or company-wide router
e.g. border-rtr.stanford.edu
Typically: mixture of 2.5Gb/s "OC48" and Gb/s Ethernet
i/f
7

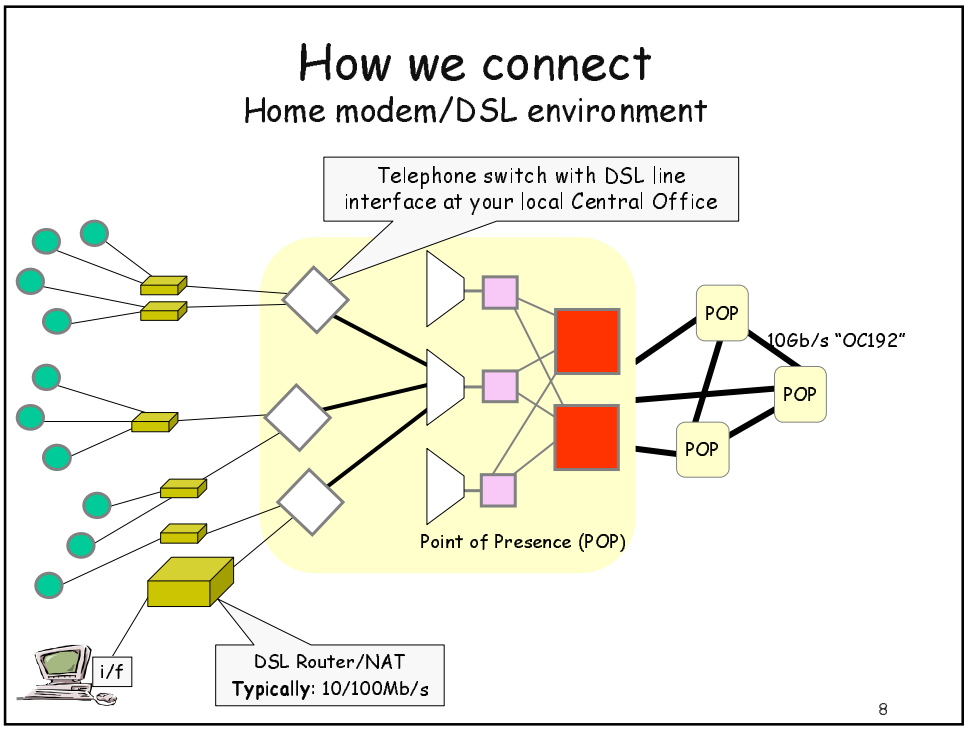
How we connect
Home modem/DSL environment
Telephone switch with DSL line interface at your local Central Office
POP
10Gb/s "OC192"
POP
POP
Point of Presence (POP)
DSL Router/NAT
Typically: 10/100Mb/s
i/f
8

Outline

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

9

What a High Performance Router Looks Like

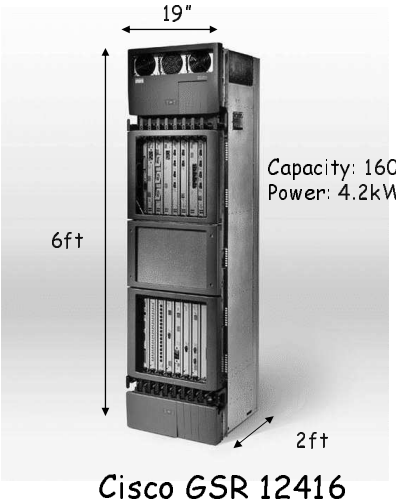

Cisco GSR 12416

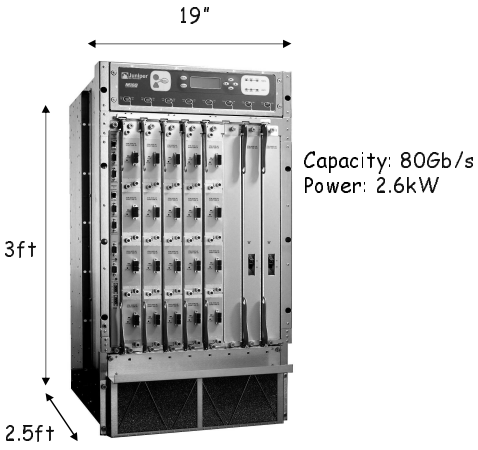

Juniper M160

10

Other packet switches

Cisco 7500 "edge" routers

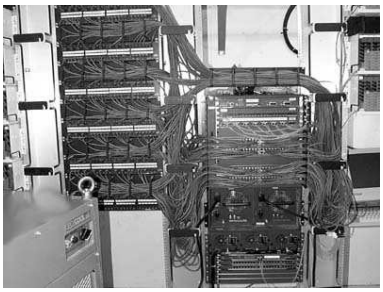

Wiring closet in Packard building

Lucent GX550 Core ATM switch

D-Link DSL router

11

Outline

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

12

The IP Datagram

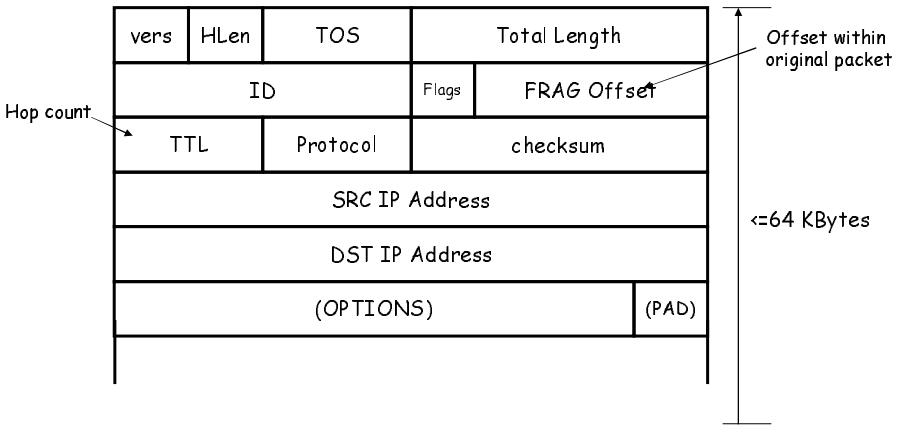

13

Forwarding in an IP Router

1. Lookup packet DA in forwarding table.
   - If known, forward to correct port.
   - If unknown, drop packet.
2. Decrement TTL, update header checksum.
3. Forward packet to outgoing interface.
4. Transmit packet onto link.

14

Ethernet Frame Format

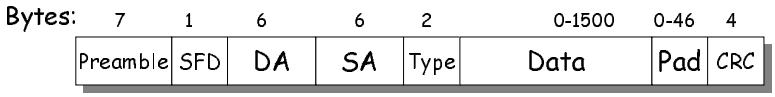

1. Preamble: trains clock-recovery circuits
2. Start of Frame Delimiter: indicates start of frame
3. Destination Address: 48-bit globally unique address assigned by manufacturer.
   - 1b: unicast/multicast
   - 1b: local/global address
4. Type: Indicates protocol of encapsulated data (e.g. IP = 0x0800)
5. Pad: Zeroes used to ensure minimum frame length
6. Cyclic Redundancy Check: check sequence to detect bit errors.

15

Encapsulation

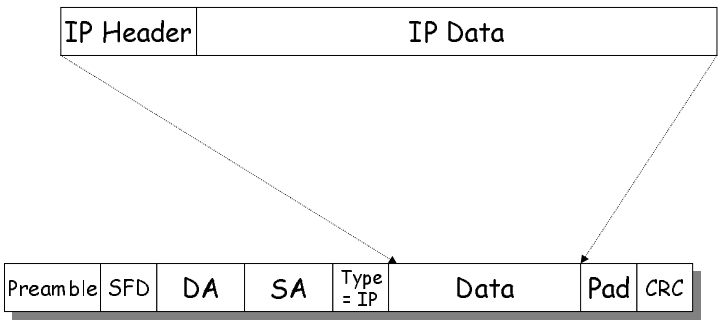

16

Outline

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

17

Generic Router Architecture

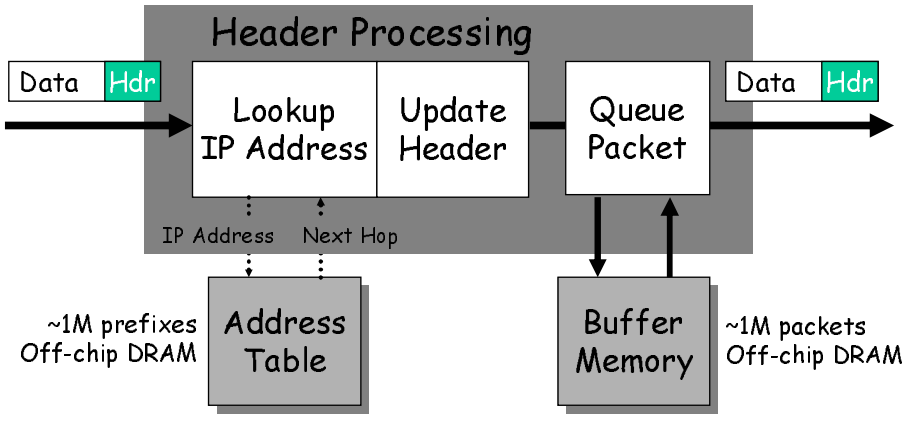

18

Generic Router Architecture

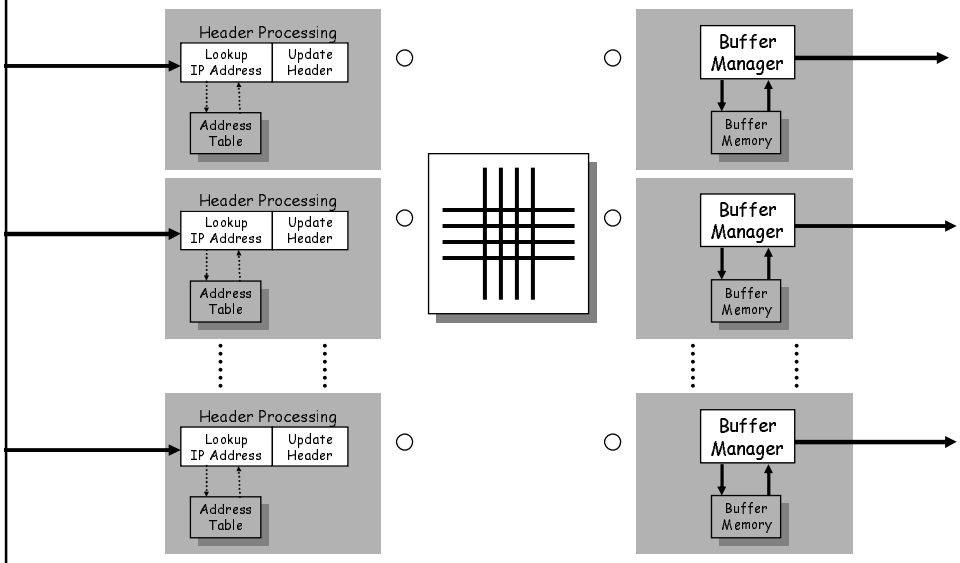

19

Contents

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

20

First Generation Routers

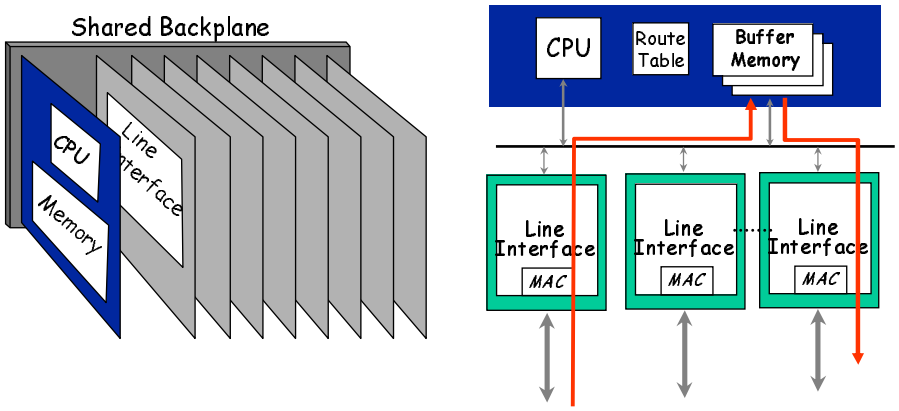

Typically <0.5Gb/s aggregate capacity

21

Second Generation Routers

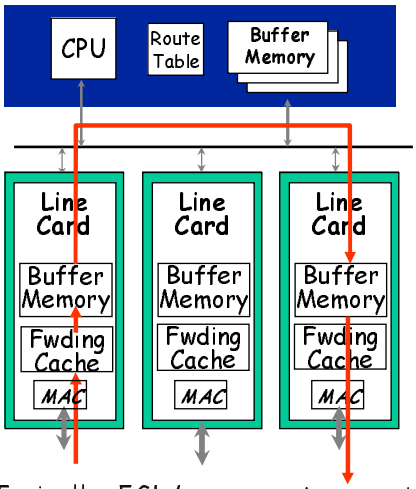

Typically <5Gb/s aggregate capacity

22

Third Generation Routers

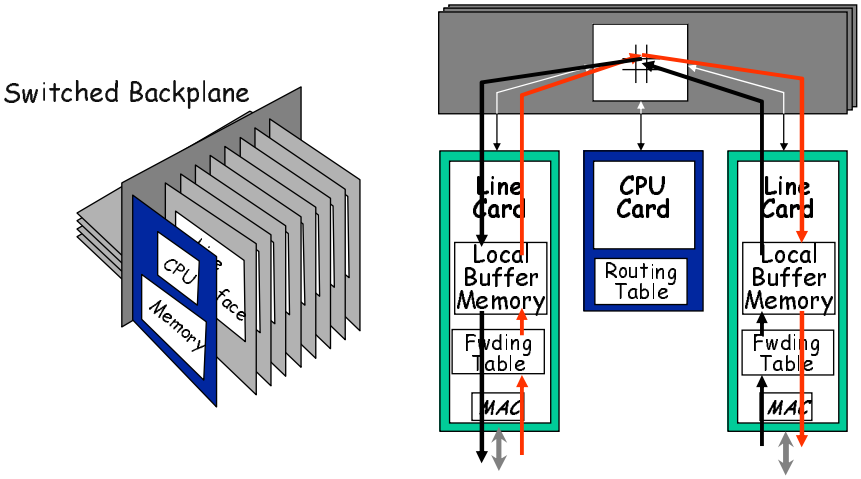

Typically <50Gb/s aggregate capacity

23

Fourth Generation Routers

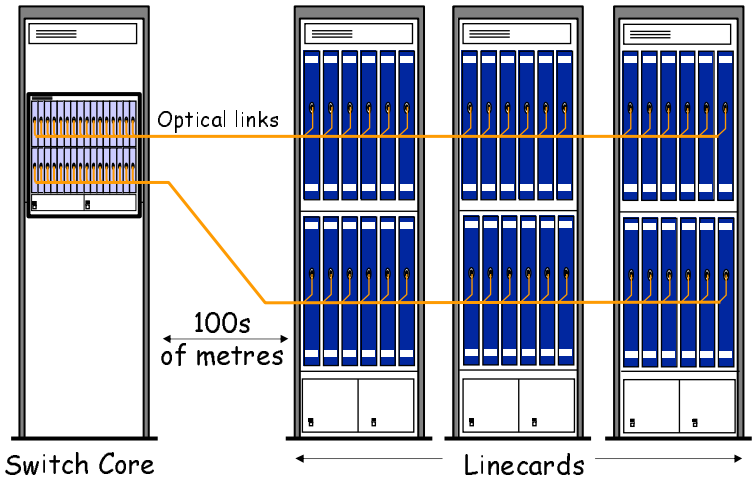

160Gb/s - 20Tb/s routers in development

Contents

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

25

Trends in Technology, Routers & Traffic

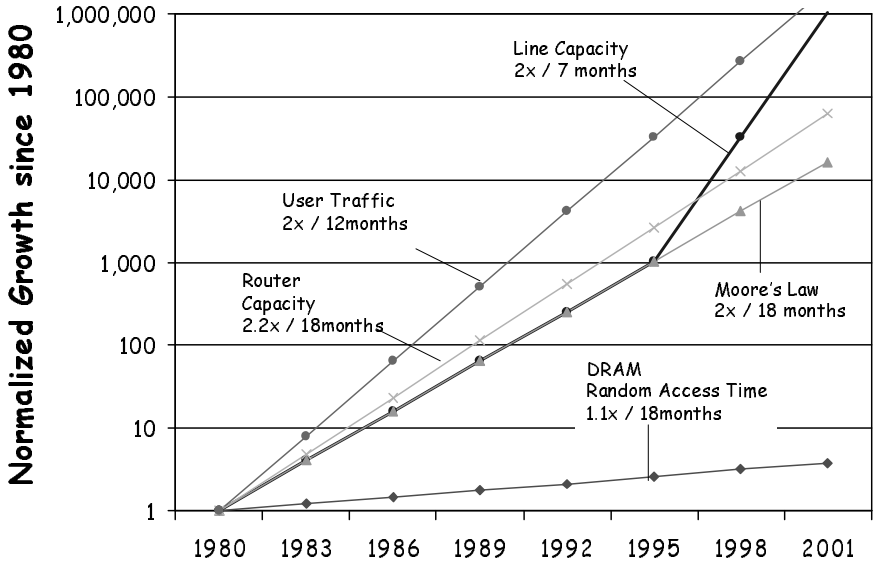

Trends and Consequences

1 CPU Instructions per minimum length packet

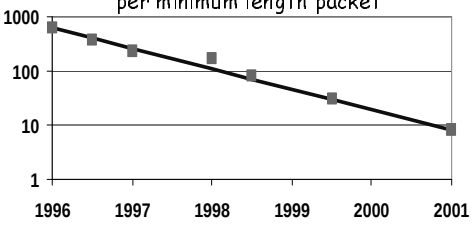

2 Disparity between traffic and router growth

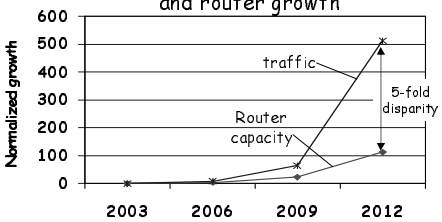

Consequences:

1. Packet processing is getting harder, and eventually network processors will be used *less* for high performance routers.
2. (Much) bigger routers will be developed.

Trends and Consequences (2)

3 Power consumption will Exceed POP limits

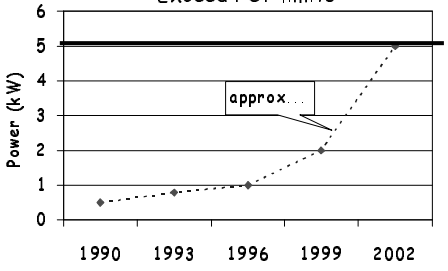

4 Disparity between line-rate and memory access time

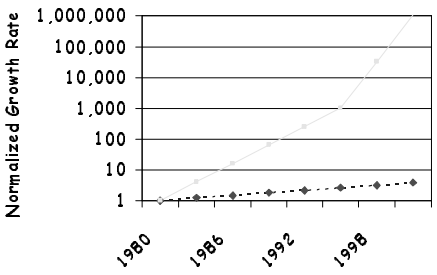

Consequences:

3. Multi-rack routers will spread power over multiple racks.
4. It will get harder to build packet buffers for linecards.

Contents

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

29

Technology Options

- General purpose processor
  - MIPS
  - PowerPC
  - Intel
- Network processor
  - Intel IXA and IXP processors
  - IBM Rainier
  - Control plane processors: SiByte (Broadcom), QED (PMC-Sierra).
- FPGA
- ASIC

30

Network Processors
Load-balancing

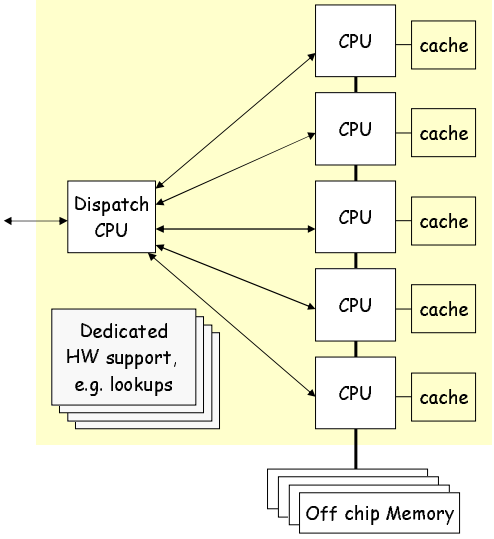

Incoming packets dispatched to:

1. Idle processor, or
2. Processor dedicated to packets in this flow (to prevent mis-sequencing).
3. Processor for processing needed by packet, e.g. security, transcoding, application-level processing.

31

Network Processors
Pipelining

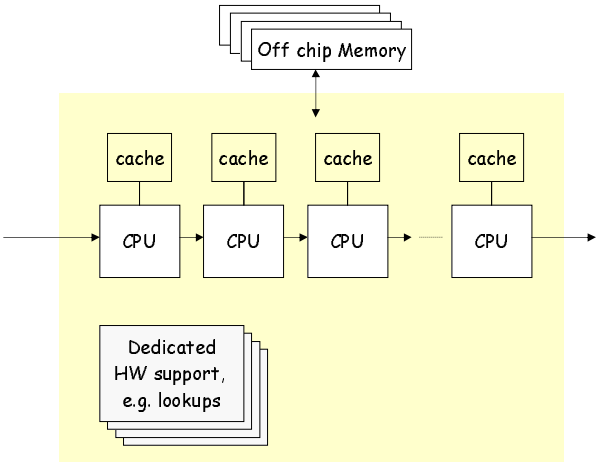

Processing broken down into (hopefully balanced) steps,
Each processor performs one step of processing.

32

Network Processors

Pros

- Flexibility: Protocols change, features are added.
- Reduced development time: In principle, should be quicker to develop software than design a custom chip.
- Reduces time-to-market, development costs, …

Cons

- Less efficient: slower than custom chip, more power.
- Usually designed using standard processors cores, not optimized for stream processing.
- Generally about 10x slower than general purpose CPU.
- Unusual development environments; hard to program.
- Often hard to partition functions over processors.

33

General Observations

- Up until about 1998,
  - Low-end packet switches used general purpose processors,
  - Mid-range packet switches used FPGAs for datapath, general purpose processors for control plane.
  - High-end packet switches used ASICs for datapath, general purpose processors for control plane.
- More recently,
  - 3rd party network processors now used in many low- and mid-range datapaths.
  - Home-grown network processors used in mid- and high-end.

34

Contents

1. What processing is done where?
2. What does a packet switch look like?
   - Examples of packet switches
   - What does a packet switch do?
   - Typical packet switch architecture
   - Evolution of high performance packet switch architecture
3. Trends and consequences
4. Technology options for processing packets
   - General purpose CPU
   - Network processors
   - FPGA
   - ASIC
5. My 2c

35

My 2c on network processors

- Is it clear that multiple small parallel processors are needed?
- When are 10 processors at speed 1 better than 1 processor at speed 10?
- Network processors make sense if:
  - Application is parallelizable into multiple threads/contexts.
  - Uniprocessor performance is limited by load-latency.
- If general purpose processors evolve anyway to:
  - Contain multiple processors per chip,
  - Support hardware multi-threading,
- ...then perhaps they are better suited because:
  - Greater development effort means faster general purpose processors,
  - Existing well-known development environments.

36

My 2c on network processors

The nail:

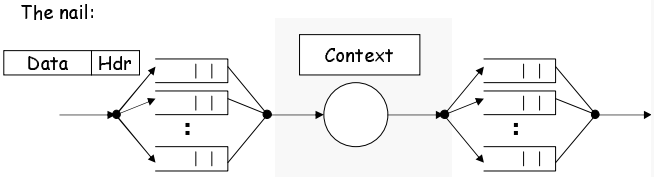

Characteristics:

1. Stream processing.
2. Multiple flows.
3. Most processing on header, not data.
4. Two sets of data: packets, context.
5. Packets have no temporal locality, and special spatial locality.
6. Context has temporal and spatial locality.

The hammer:

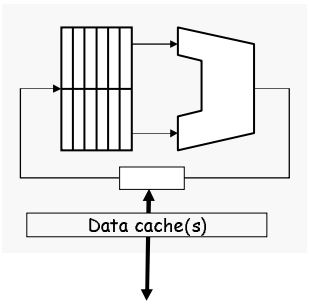

Characteristics:

1. Shared in/out bus.
2. Optimized for data with spatial and temporal locality.
3. Especially optimized for register accesses.

37

A network uniprocessor

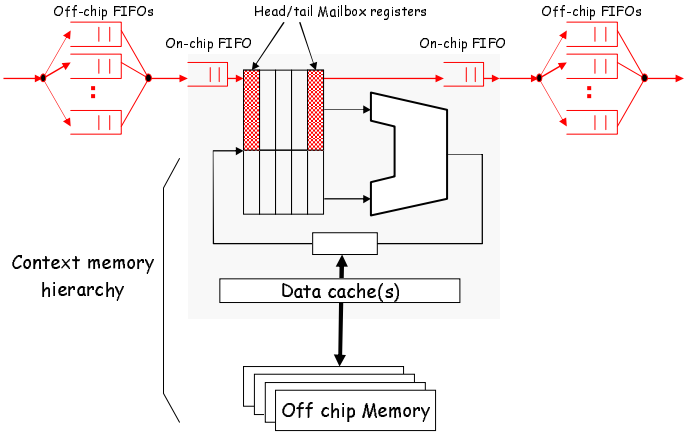

Add hardware support for multiple threads/contexts.

38