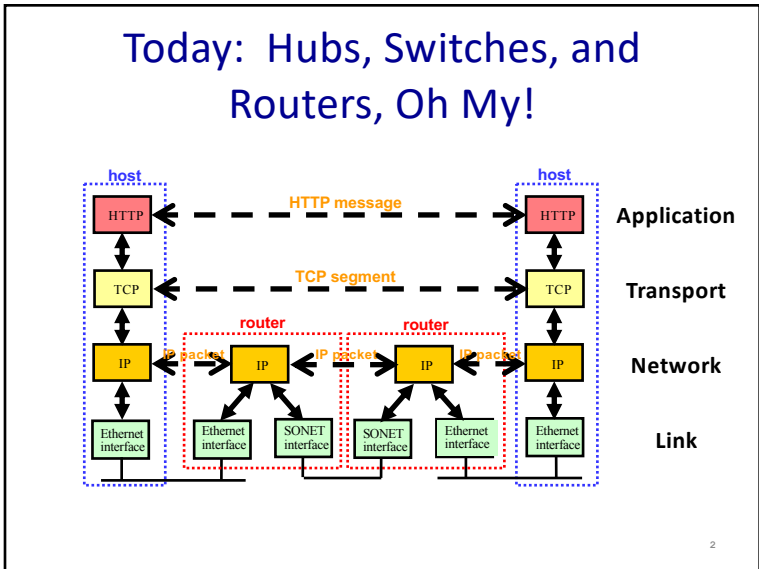


Hubs, Switches, and Routers, Oh My!

Mike Freedman
COS 461: Computer Networks

<http://www.cs.princeton.edu/courses/archive/spr20/cos461/>



OSI Model

Developed in late 1970s / early 1980s

	Data	Layer	
Host Layers	Data	Application Network Process to Application	Layer 7
	Data	Presentation Data representation and Encryption	
	Data	Session Interhost communication	
Media Layers	Segments	Transport End-to-End connections and Reliability	Layer 4
	Packets	Network Path Determination and IP (Logical addressing)	Layer 3
	Frames	Data Link MAC and LLC (Physical addressing)	Layer 2
	Bits	Physical Media, Signal and Binary Transmission	


- ## Terminology
- **Hubs and Repeaters**
 - Connect machines on same “layer 2” LAN
 - Broadcast: All frames are sent out all physical ports
 - **Switches and Bridges**
 - Connect machines on same “layer 2” LAN
 - Only send frames to selected physical port based on destination MAC address
 - **Routers**
 - Connect between LANs at “layer 3”, e.g., wide area
 - Only send packet to selected physical port based on destination IP address

“Layer 2” Hubs and Switches

5

Physical Layer: Repeaters


- **Distance limitation in local-area networks**
 - Electrical signal becomes weaker as it travels
 - Imposes a limit on the length of a LAN
- **Repeaters join LANs together**
 - Analog electronic device
 - Continuously monitors electrical signals
 - Transmits an amplified copy



6

Physical Layer: Hubs

- **Joins multiple input lines electrically**
 - Designed to hold multiple line cards
 - Do not necessarily amplify the signal
- **Very similar to repeaters**
 - Also operates at the physical layer



7


Limitations of Repeaters and Hubs

- **One large shared link**
 - Each bit is sent everywhere
 - So, aggregate throughput is limited
- **Cannot support multiple LAN technologies**
 - Does not buffer or interpret frames
 - Can't interconnect between different rates/formats
- **Limitations on maximum nodes and distances**
 - Shared medium imposes length limits
 - E.g., cannot go beyond 2500 meters on Ethernet

8

Link Layer: Bridges


- Connects two or more LANs at the link layer
 - Extracts destination address from the frame
 - Looks up the destination in a table
 - Forwards the frame to the appropriate segment
- Each segment can carry its own traffic



9

Link Layer: Switches


- Typically connects individual computers
 - A switch is essentially the same as a bridge
 - ... though typically used to connect hosts
- Supports concurrent communication
 - Host A can talk to C, while B talks to D



10

Bridges/Switches: Traffic Isolation


- Switch filters packets
 - Frame only forwarded to the necessary segments
 - Segments can support separate transmissions



11

Self Learning: Building the Table


- When a frame arrives
 - Inspect the *source* MAC address
 - Associate the address with the *incoming* interface
 - Store the mapping in the switch table
 - Use a timer to eventually forget the mapping



12

Self Learning: Handling Misses

- When frame arrives with unfamiliar destination
 - Forward the frame out all of the interfaces
 - ... except for the one where the frame arrived
 - Hopefully, this case won't happen very often!



Switches vs. Hubs

- Compared to hubs, Ethernet switches support
 - (Y) Larger geographic span
 - (M) Similar span
 - (C) Smaller span
- Compared to hubs, switches provides
 - (Y) Higher load on links
 - (M) Less privacy
 - (C) Heterogenous communication technologies

14

Looking closer...


15

Basic Router Architecture

- Each switch/router has a forwarding table
 - Maps destination address to outgoing interface
- Basic operation
 1. Receive packet
 2. Look at header to determine destination address
 3. Look in forwarding table to determine output interface
 4. Modify packet header (e.g., decr TTL, update chksum)
 5. Send packet to output interface

16

Basic Router Architecture

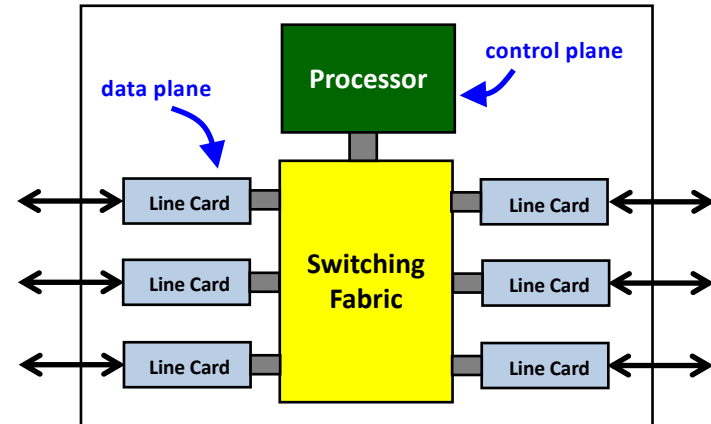


• Basic operation

1. Receive packet
2. Look at header to determine destination address
3. Look in forwarding table to determine output interface
4. Modify packet header (e.g., decr TTL, update chksum)
5. Send packet to output interface

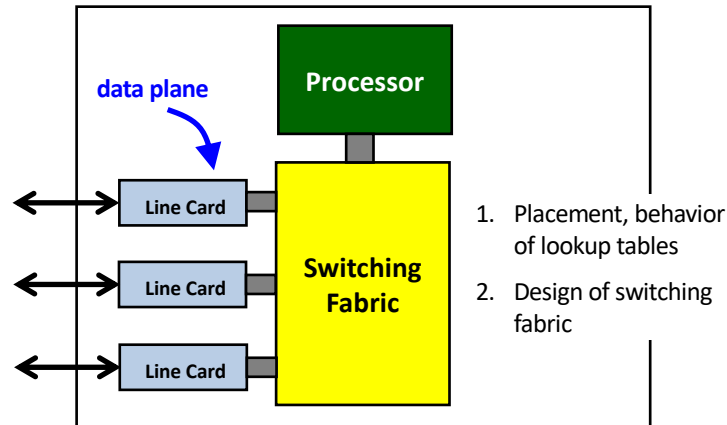
17

Router



18

Router



19

Lookup algorithm depends on protocol

Protocol	Mechanism	Techniques
Ethernet (48 bits) MPLS ATM	Exact Match	<ul style="list-style-type: none"> • Direct lookup • Associative lookup • Hashing • Binary tree
IPv4 (32 bits) IPv6 (128 bits)	Longest-Prefix Match	<ul style="list-style-type: none"> • Radix trie • Compressed trie • TCAM

20

Longest Prefix Match

- Each packet has destination IP address
- Router looks up table entry that matches address

68.211.6.120

Prefix	Output
68.208.0.0/12	1
68.211.0.0/17	1
68.211.128.0/19	2
68.211.160.0/19	2
68.211.192.0/18	1

21

Longest Prefix Match


- Each packet has destination IP address
- Router looks up table entry that matches address
- Benefit of CIDR allocation and LPR
 - Efficiency: Prefixes can be allocated at much finer granularity
 - Hierarchical aggregation: Upstream ISP can aggregate 2 contiguous prefixes from downstream ISPs to shorter prefix

22

How? Address lookup using trie

- Prefixes “spelled out” by following path from root
- To find the best prefix spell out address in trie

	Prefixes
A	111*
B	10*
C	1010*
D	10101

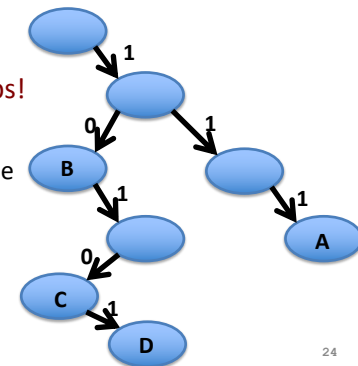


23

How? Address lookup using trie

- Prefixes “spelled out” by following path from root
- To find the best prefix spell out address in trie


- 1 per /vl = max 32 lookups!
- Too slow:
 - E.g., “Optical Carrier 48” line (2.5 Gbps) requires 160ns lookup ... or 4 memory accesses



24

K-ary trie (k=2)

	Prefixes
A	111*
B	10*
C	1010*
D	10101



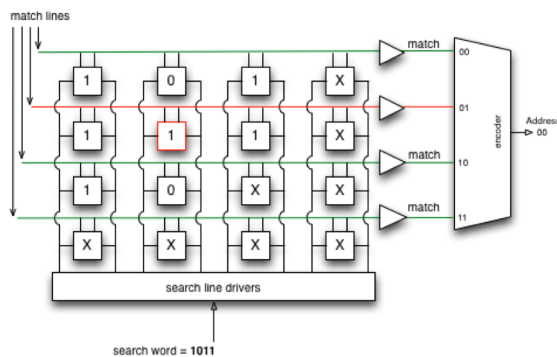
25

Alternatives to trie for LPM

- **Content-Address Memory (CAM)**
 - Input: tag (address)
 - Output: value (port)
 - Exact match, but $O(1)$ in hardware
- **Ternary CAM**
 - 0, 1, *
 - “value” memory cell and “mask” (care / don’t care) cell
- **LPM via TCAM**
 - In parallel, search all prefixes for all matches
 - Then choose longest match
 - Trick: choose first match, but already sorted by prefix length

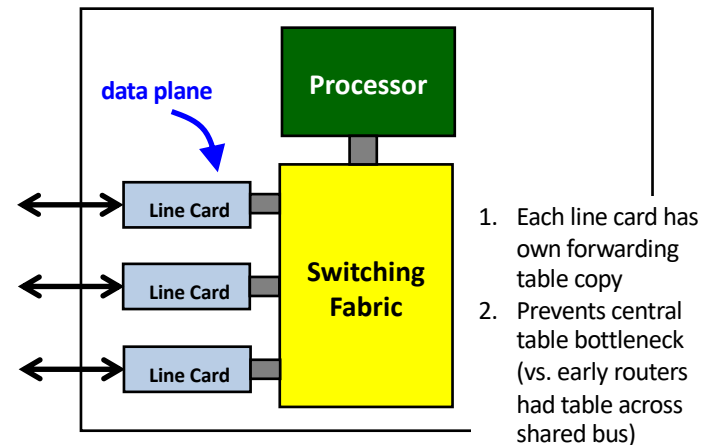
26

Alternatives to trie for LPM

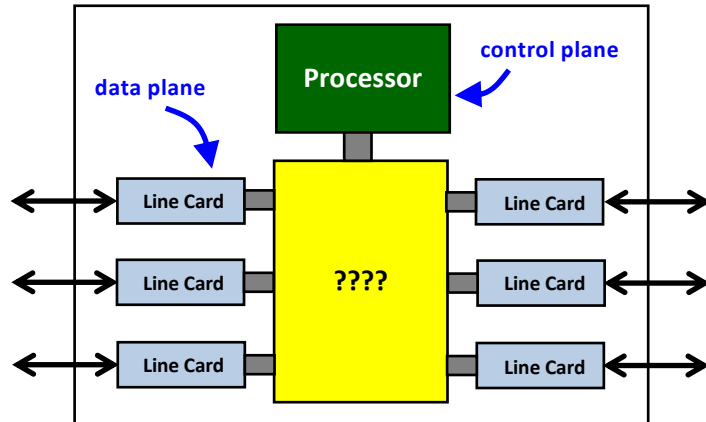


27

Decision: Forwarding table per line card



Decision: Crossbar switch




Decision: Crossbar switch

- Shared bus
 - Only one input can speak to one output at a time
- Crossbar switch / switched backplane
 - Input / output pairs that don't compete can send in same timeslot

30


Crossbar switching

- Every input port has connection to every output port
- In each timeslot, each input connected to zero or more outputs



Crossbar switching


- Every input port has connection to every output port
- In each timeslot, each input connected to zero or more outputs



- Good parallelism
- Needs scheduling

Everything gets complicated...


- **Problem: Head-of-line blocking**
 - The packet in front of queue blocks packets behind it from being processed
 - Say first packet at input 1 wants to go to output 5; second packet at 1 that wants 4 is still blocked



33


Everything gets complicated...

- **Solution: Virtual output queues**
 - One queue at input per output port
 - So avoids head-of-line blocking during crossbar scheduling



34

Data, Control, & Management Planes



	Data	Control	Management
Time-scale	Packet (ns)	Event (10 ms to sec)	Human (min to hours)
Tasks	Forwarding, buffering, filtering, scheduling	Routing, signaling	Analysis, configuration
Location	Line-card hardware	Router software	Humans or scripts

35

Cisco 8000 Series Routers



- Up to 648 400 GbE
- 260 Tbps backplane

36

Conclusions

- Physical devices sharing L2 & L3 networks have many common features
 - Forward table lookups
 - Queueing and backplane switching
 - Fast vs. slow paths
 - Switches and routers separate routing decisions (control plane) from forwarding actions (data plane)
- High speed necessitates innovation
 - Specialized hardware
 - Software algorithms

37