CS 43: Computer Networks

17: The Network Layer

November 6, 2018

SWARTHMORE COLLEGE

The Network Layer!

Application: the application (e.g., the Web, Email)

Transport: end-to-end connections, reliability

Network: routing

Link (data-link): framing, error detection

Physical: 1's and 0's/bits across a medium (copper, the air, fiber)

Network Layer

- Function: Route packets end-to-end on a network, through multiple hops

- Key challenge
  - How to route packets: Convergence
  - How to represent addresses: Scalability

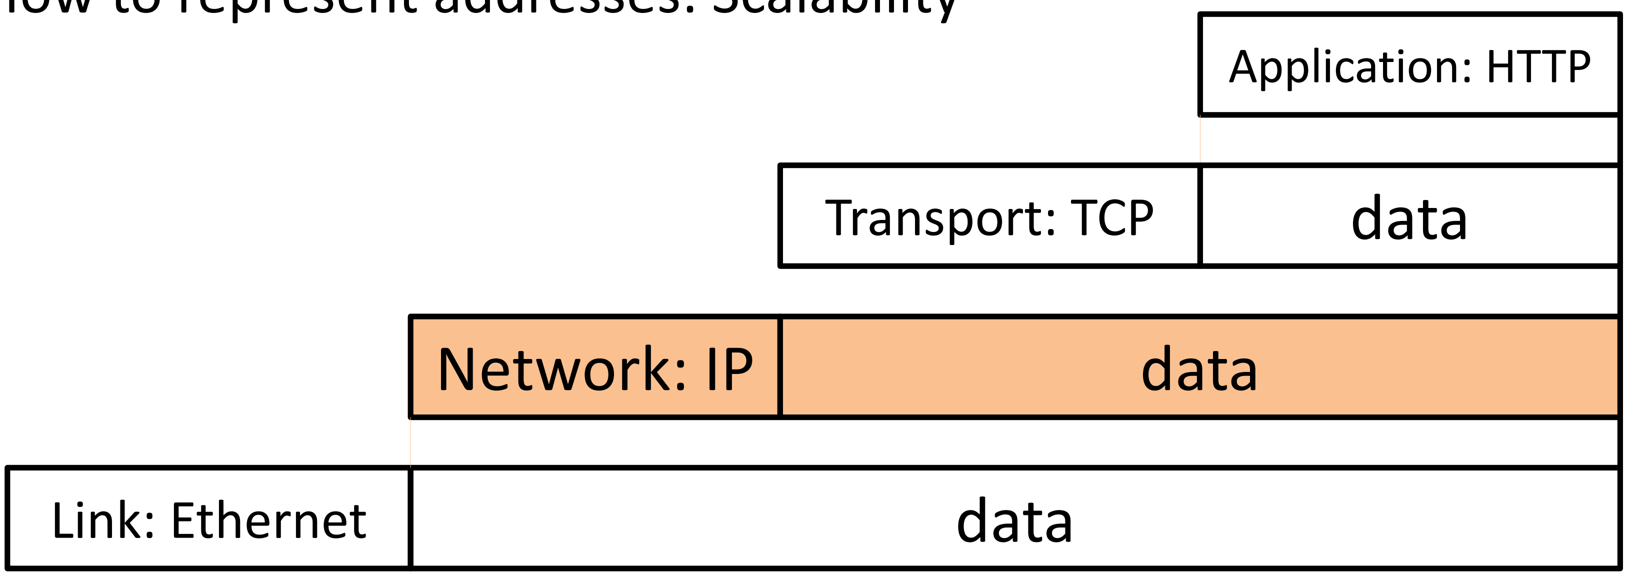

Network Layer

- DARPAnet Primary Goal: Connect Hosts
- "islands" of networks: SATNet, Packet Radio, Ethernet: how do we connect them?
- Routers forward packets using a common Internet Protocol
  - *Any* underlying data link protocol
  - *Any* higher layer transport protocol

Example of Internet Routing

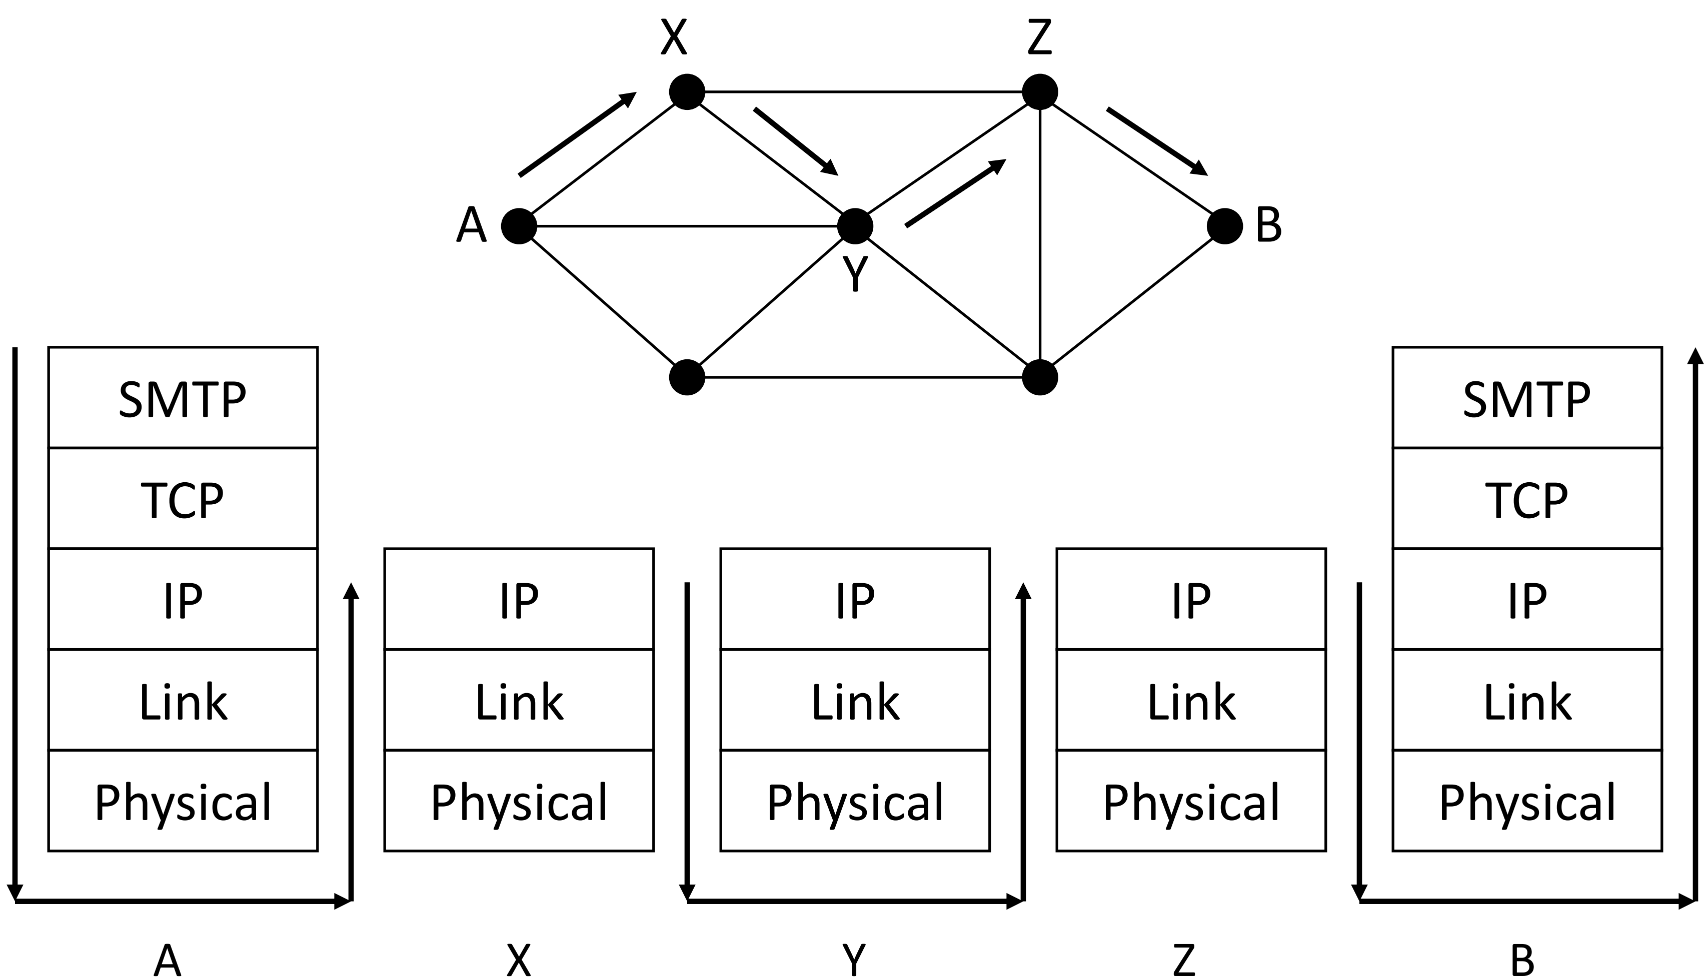

Network layer involved at every hop along the path.

Network Layer Functions

- Forwarding: move packets from router's input to appropriate router output
  ("data plane")

- Routing: determine route taken by packets from source to destination.
  ("control plane")

When should a router perform routing? Forwarding?

A. Do both when a packet arrives.

B. Route in advance, forward when a packet arrives.

C. Forward in advance, route when a packet arrives.

D. Do both in advance.

E. Some other combination

Network Layer Functions

- Forwarding: move packets from router's input to appropriate router output
  - Look up in a table
- Routing: determine route taken by packets from source to destination.
  - Populating the table

Interplay between routing and forwarding

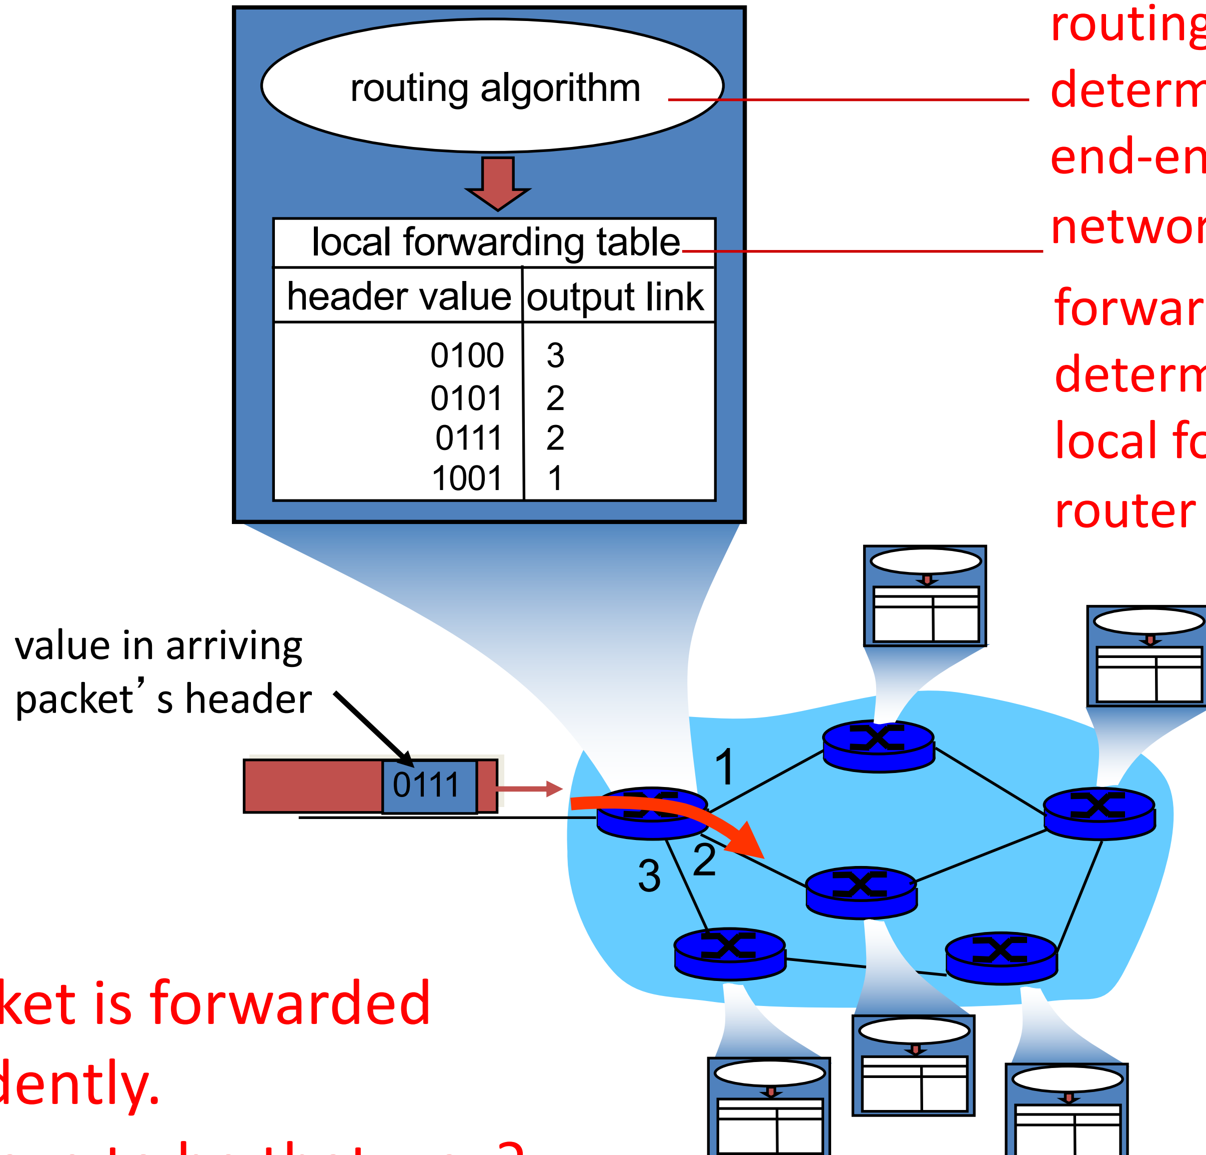

routing algorithm determines end-end-path through network

forwarding table determines local forwarding at this router

Each packet is forwarded independently.
Does it have to be that way?

Circuit Switching

- Reserve path in advance

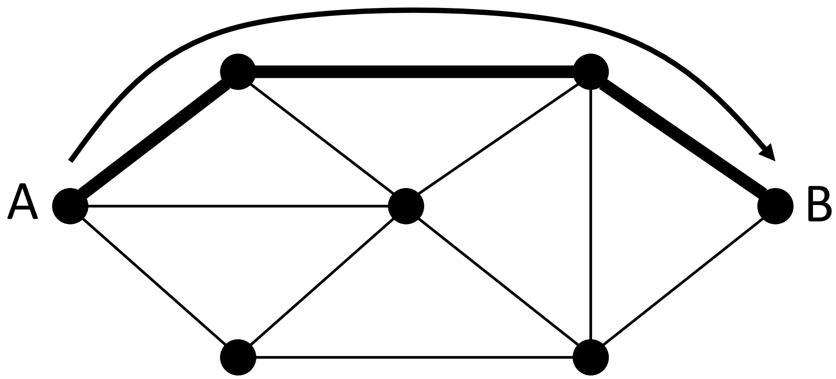

- (Old) telephone system

Why doesn't the Internet (typically) use circuits?

A. It's too slow to establish a connection.

B. It doesn't offer good enough performance.

C. It wastes resources.

D. It requires too many resources.

E. Some other reason.

Packet Switching

- Do we always need to reserve a link?

- <u>Statistical multiplexing</u>
  - Assign multiple conversations to a physical path
  - At any given time, one will have something to say

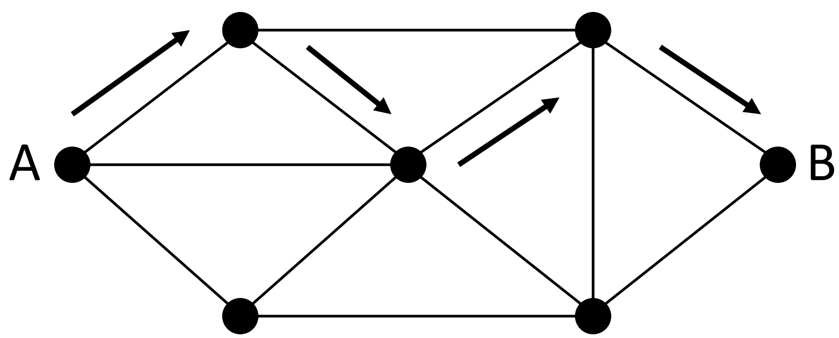

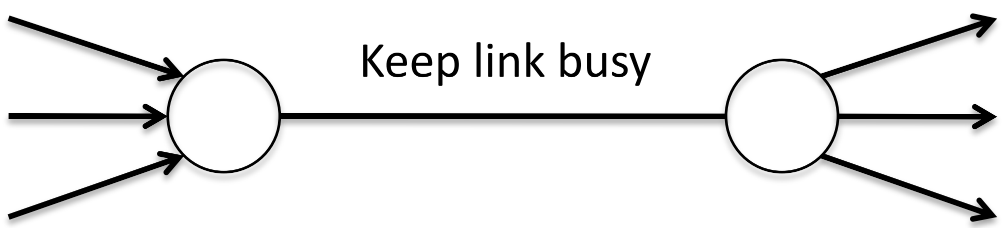

Packet Switching: Statistical Multiplexing

- Data traffic is bursty
  - Telnet, email, Web browsing, …
- Avoid wasting bandwidth
  - One host can send more when others are idle

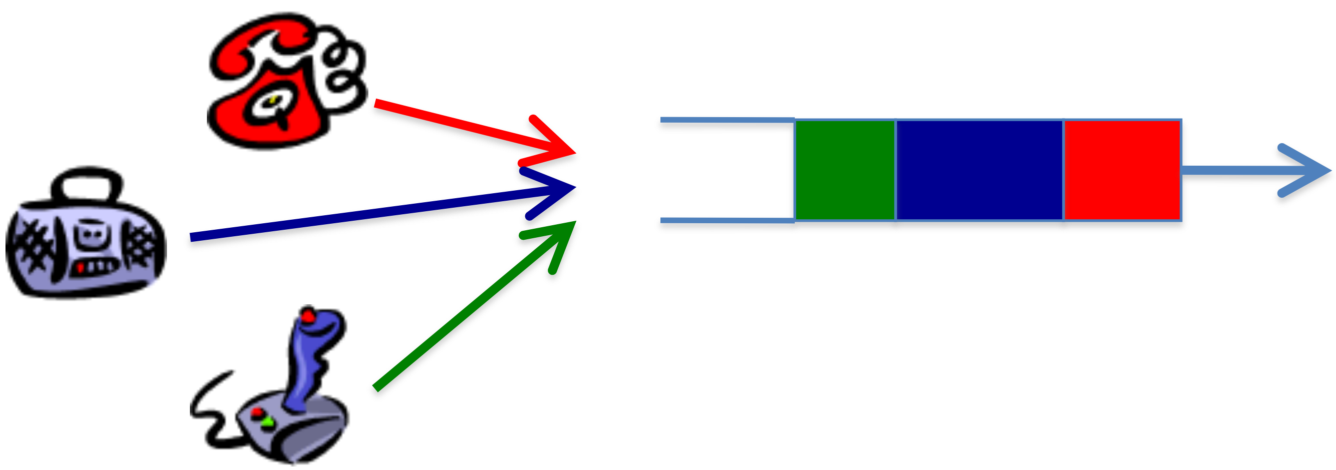

Which of the following is/are generally true of packet vs. circuit switching?

1. Packet switching has less variance in performance.
2. Circuit switching is less reliable.

A. Only 1 is true.
B. Only 2 is true.
C. Both 1 and 2 are true.
D. Neither 1 nor 2 are true.

Circuit-switching vs. Packet switching

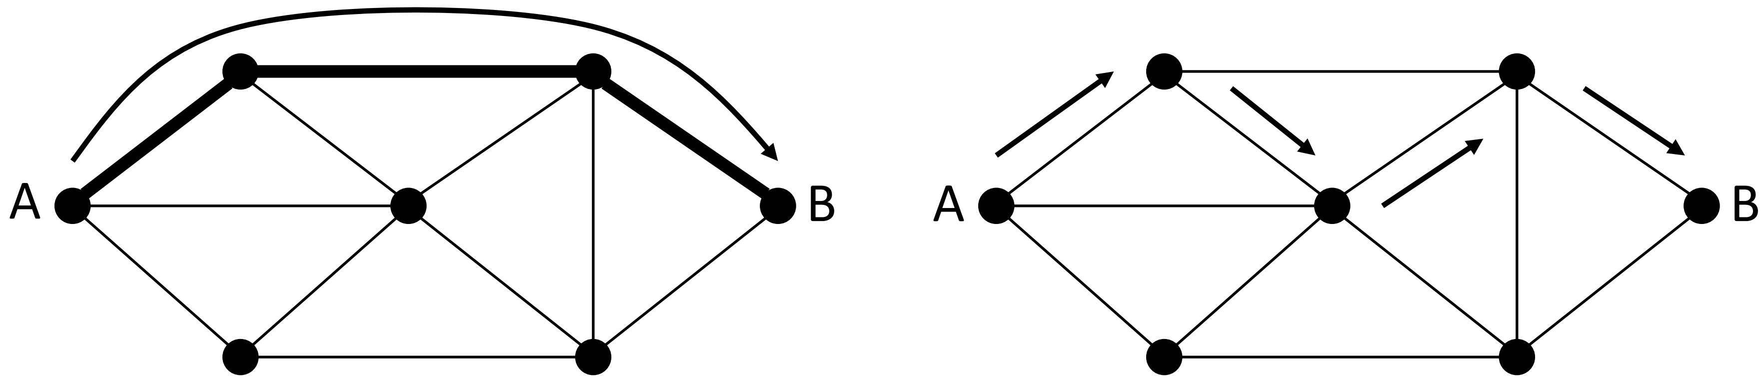

- Circuit switching: establish path, send data
  - Reserve resources, provide performance control
  - Example: telephone system
- Packet switching: forward packets hop by hop
  - Fair sharing despite bursts, statistical multiplexing
  - Example: postal system

Datagram vs. “Virtual Circuit”

- Datagram network provides network-layer connectionless service (packet switching)

- Virtual-circuit network provides network-layer connection service (like circuit switching)

Virtual circuits: Signaling Protocols

- Used to setup, maintain, teardown VC
- Used in ATM, frame-relay, X.25
- Less common in today's Internet

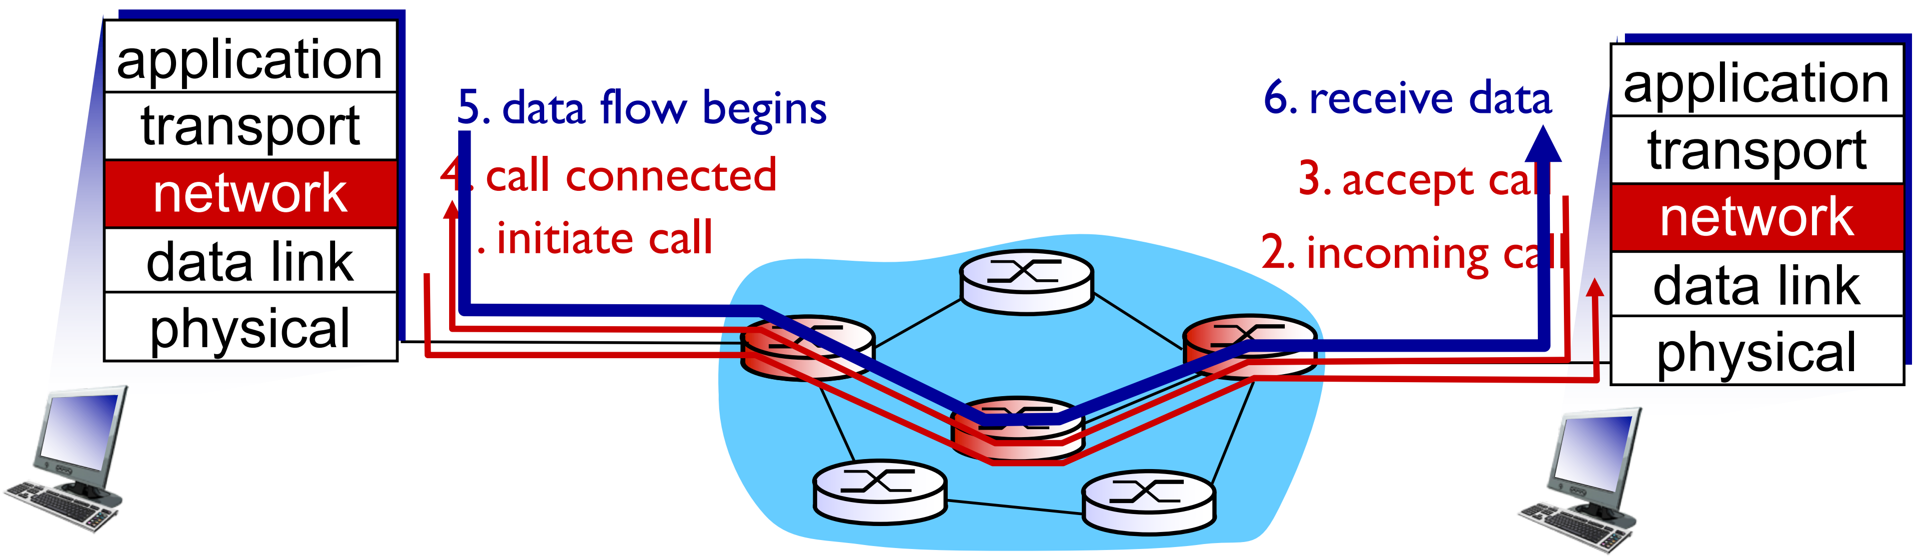

Datagram Networks

- No call setup at network layer
- Routers: no state about end-to-end connections
  - no network-level concept of "connection"
- Packets forwarded individually towards destination

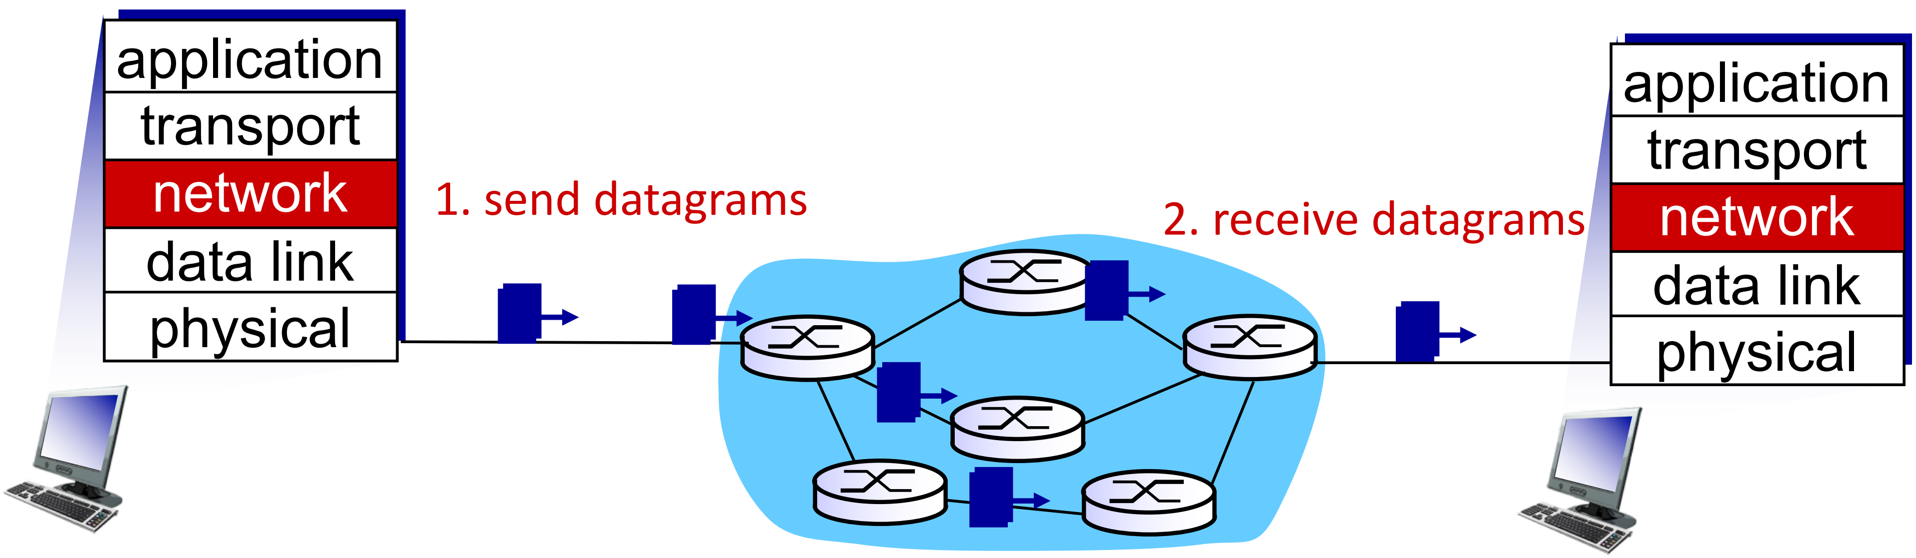

How should we populate a router's forwarding table?

A. A person should add entries to the table.

B. A program external to the router should add entries to the table.

C. Routers should communicate with each other to add entries to their tables.

D. Some other mechanism.

Routing

Traditional

- Routers run a routing protocol to exchange state.
- Use state to build up the forwarding table.

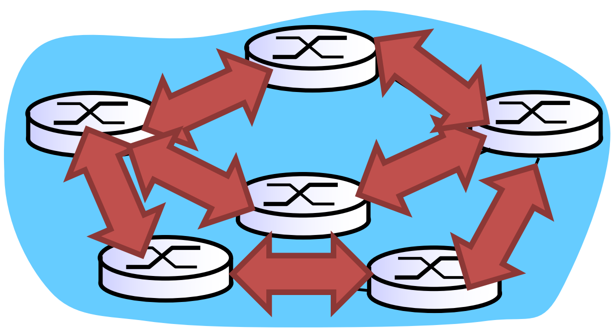

Assume this is the type of routing we're talking about unless we explicitly say otherwise!

Routing

Traditional

- Routers run a routing protocol to exchange state.
- Use state to build up the forwarding table.

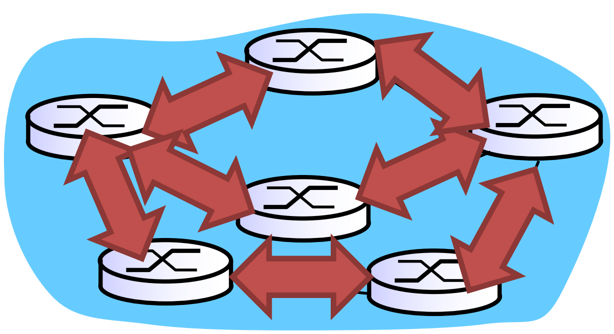

"Software-Defined"

- Routers are dumb, just do what they're told.
- Controller service explicitly tells each router what to do.
- Rare on the Internet, hot topic in data centers.

Datagram Forwarding

- Routers periodically exchange state.
- Use the state to build a forwarding table (FIB – Forwarding Information Base)

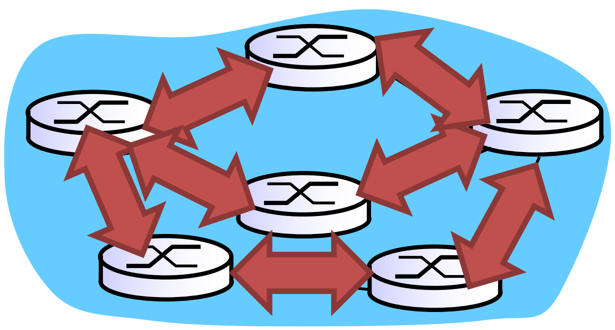

Datagram forwarding table

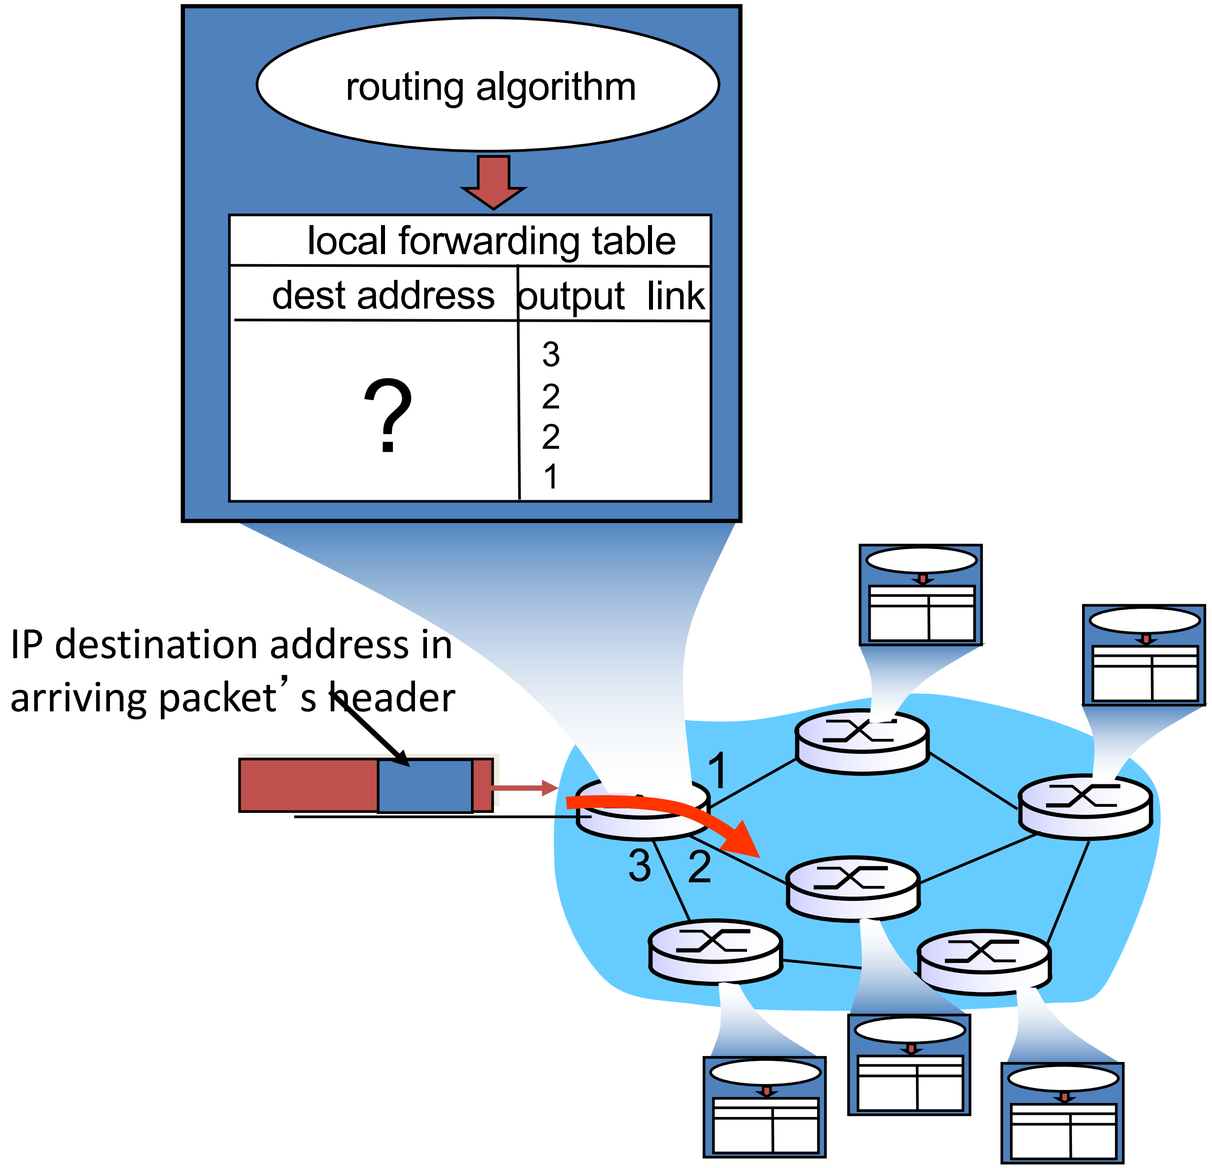

Routers exchange state (we'll save the what and when for later). They decide, for each destination, how to get there, and build a lookup structure for their forwarding table. What should they build?

A. A list – scan for the destination.

B. A hash table – look up the destination.

C. A tree – Follow branches that lead to the destination.

D. Some other software structure.

E. We can't do this in software, we need special hardware.

Aside: router architecture overview

- high-level view of generic router architecture:

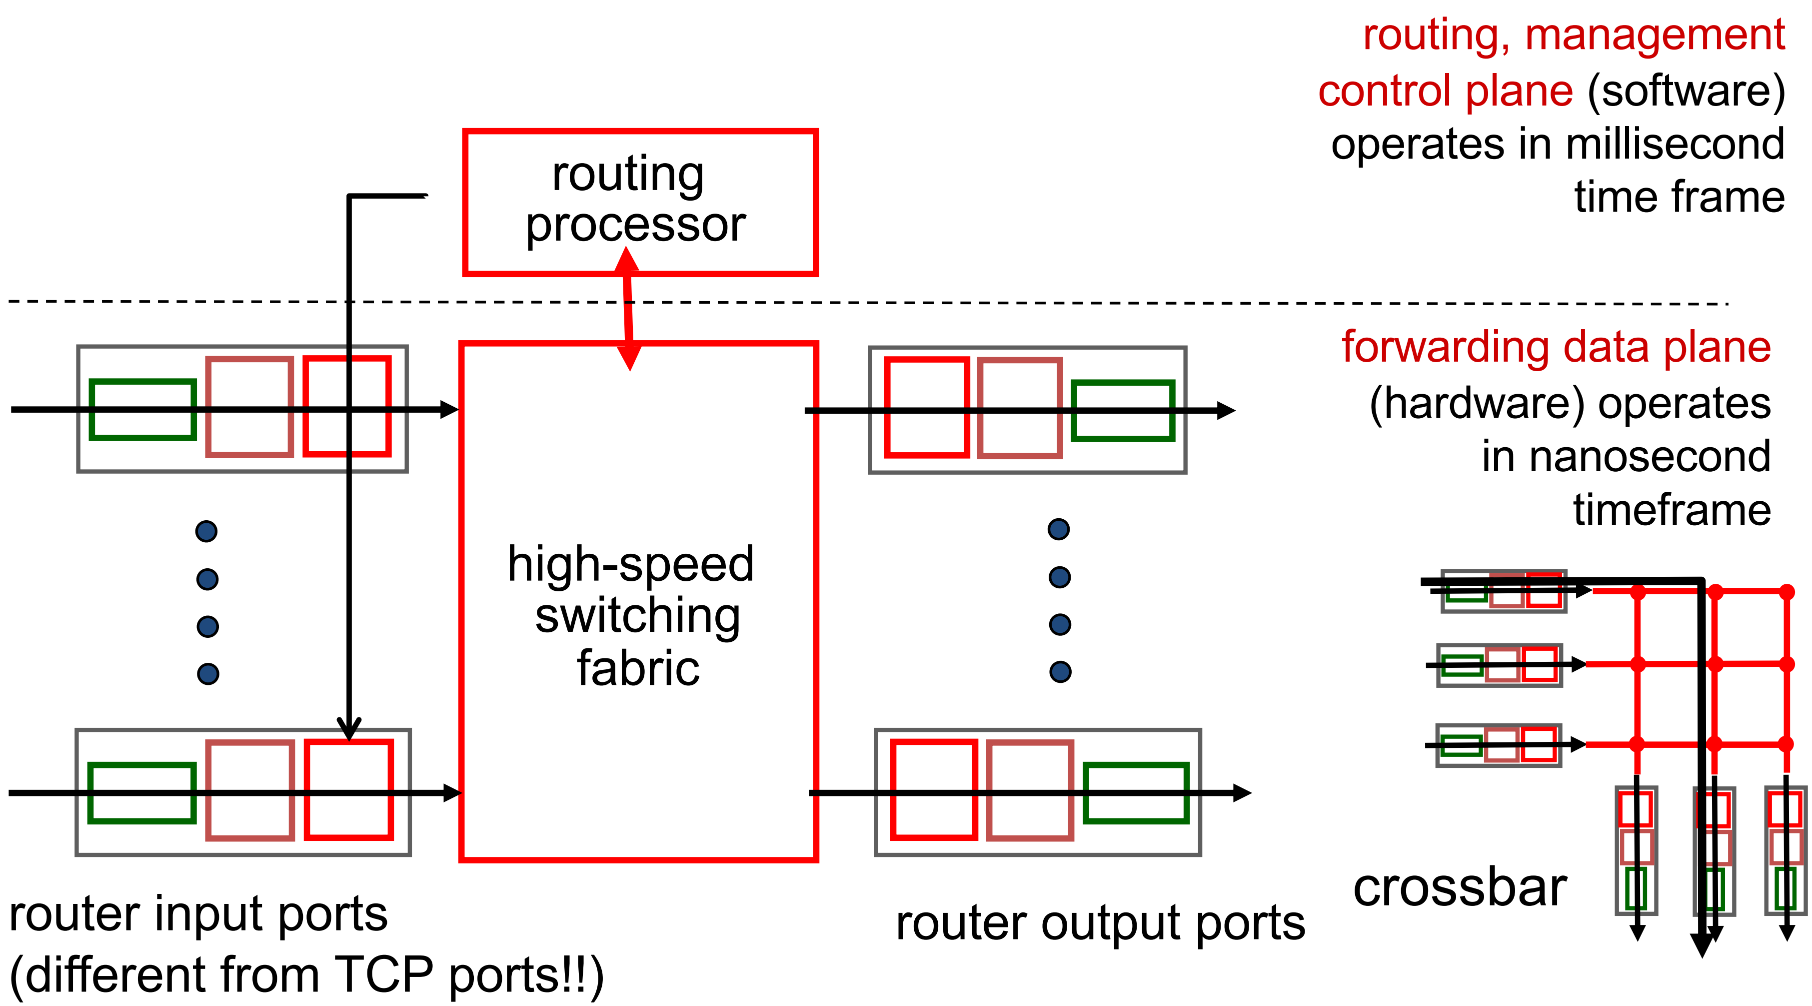

these are physical inputs/outputs to the router

Datagram forwarding table

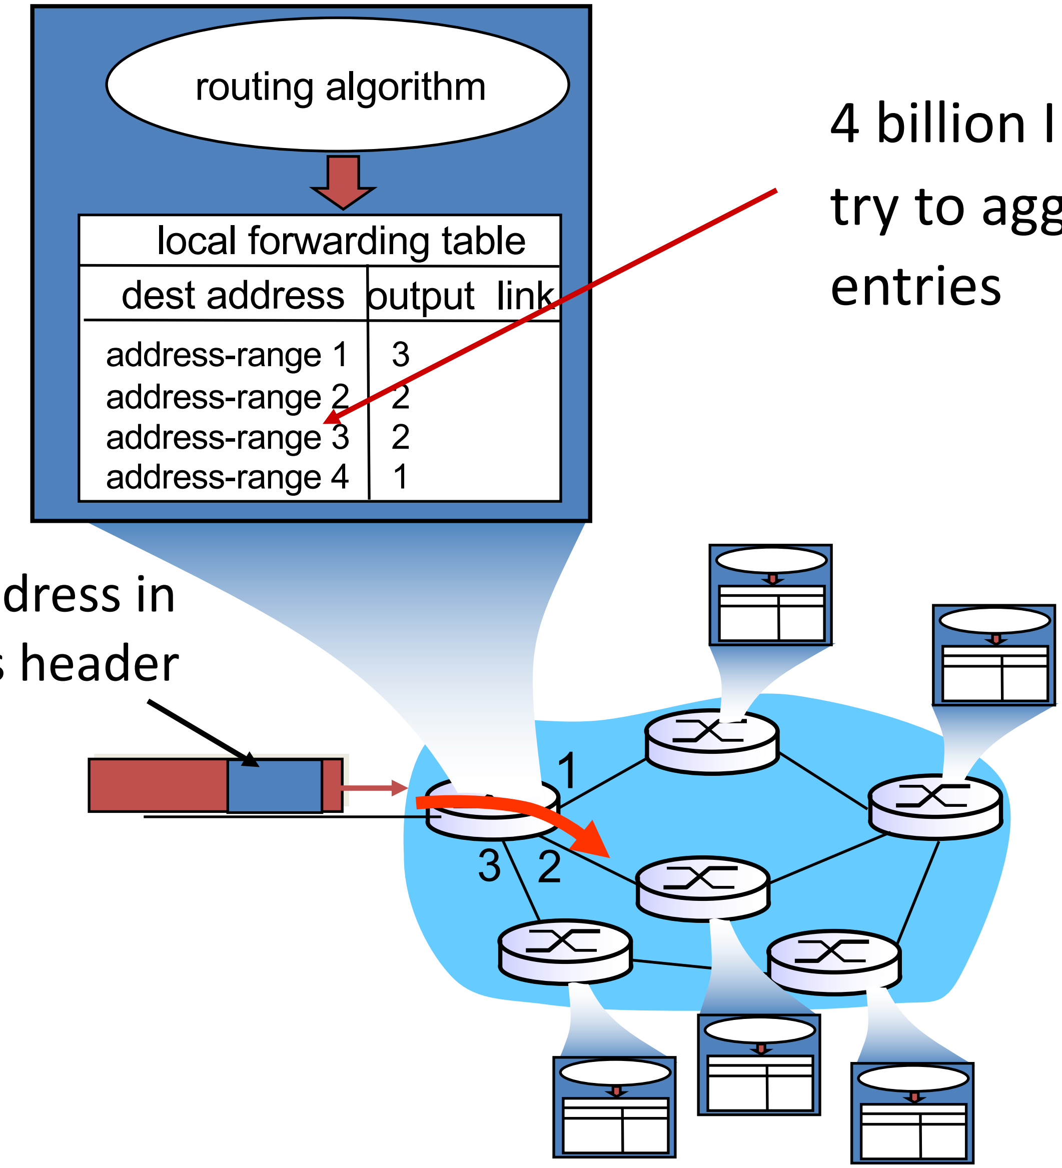

4 billion IP addresses, try to aggregate table entries

IP destination address in arriving packet's header

What's in a name?

- Host name: web.cs.swarthmore.edu
  - Domain: registrar for each top-level domain (e.g., .edu)
  - Host name: local administrator assigns to each host
- IP addresses: 130.58.68.164
  - Prefixes: ICANN, regional Internet registries, and ISPs
  - Hosts: static configuration, or dynamic using DHCP
- MAC addresses: D8:D3:85:94:5F:1E
  - OIDs: assigned to vendors by the IEEE
  - Adapters: assigned by the vendor from its block

What’s in a name? IP addresses

- Flat
  - e.g. each host is identified by a 48-bit MAC address
  - Router needs an entry for every host in the world

What's in a name? IP addresses

- Hierarchy
  - Addresses broken down into segments
  - Each segment has a different level of specificity
  - Usually tied to geographic location

IP Addressing

- IP: 32-bit addresses
  - Usually written in dotted notation, e.g. 192.168.21.76
  - Each number is a byte
  - Stored in Big Endian order (network byte order)

0 8 16 24 31

Decimal	192	168	21	76
Hex	C0	A8	15	4C
Binary	11000000	10101000	00010101	01001100

Datagram forwarding table

Destination Address Range	Link Interface
11001000 00010111 0001<u>0000 00000000</u> through 11001000 00010111 0001<u>0111 11111111</u>	0
11001000 00010111 00011000 <u>00000000</u> through 11001000 00010111 00011000 <u>11111111</u>	1
11001000 00010111 00011<u>001 00000000</u> through 11001000 00010111 00011<u>111 11111111</u>	2
Otherwise (default gateway)	3

Q: but what happens if ranges don't divide up so nicely?

Longest prefix matching

In a forwarding table entry, use the longest address prefix that matches destination address.

Destination Address Range	Link interface
11001000 00010111 00010*** *********	0
11001000 00010111 00011000 *********	1
11001000 00010111 00011*** *********	2
Otherwise (default gateway)	3

DA: 11001000 00010111 00011000 10101010

DA: 11001000 00010111 00010110 10100001

which interface?

Routing

<u>Traditional</u>

- Routers run a routing protocol to exchange state.

- Use state to build up the forwarding table.

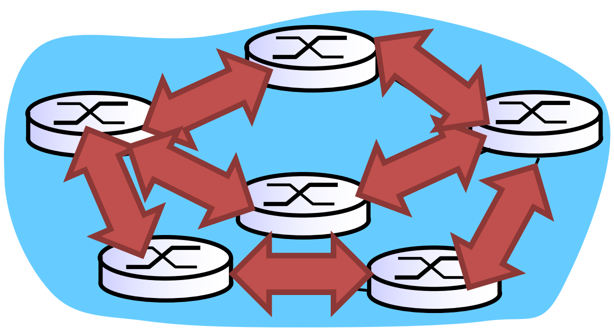

What services would we like a router to implement?

A. Basic connectivity: route packets to destination
B. Find policy-compliant paths (keep ISPs happy)
C. Traffic engineering
D. Impose limits on what can be accessed on the Internet vs. local ISP
E. All of the above

Nice things to have..

- Traffic engineering:
  - Want to avoid persistent overloads on links
  - Choose routes to spread traffic load across links
- Access Control:
  - Limit access to backend database machines.
  - Firewalls
- Network measurement

Number of published Internet Standards

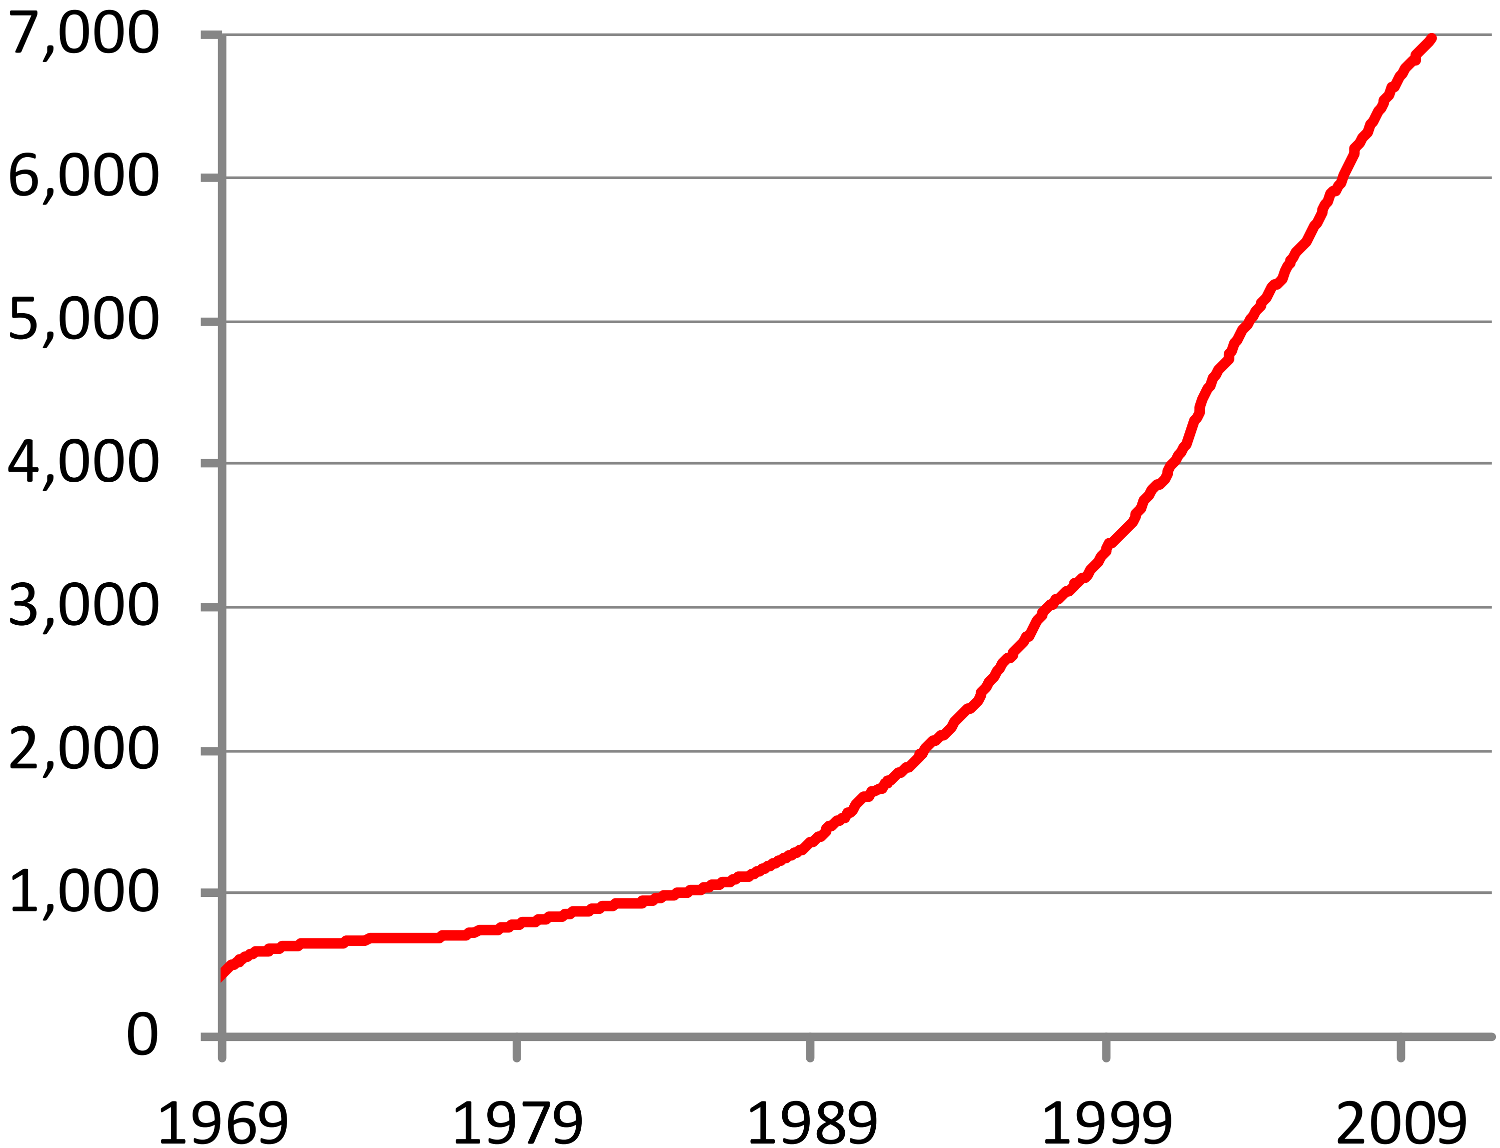

Graph from Nick McKeown

Routing

Traditional

- Routers run a routing protocol to exchange state.

- Use state to build up the forwarding table.

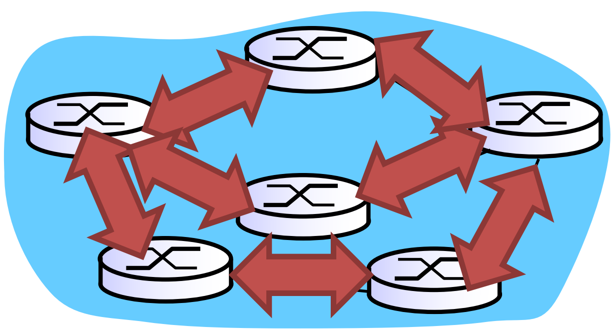

Software-Defined

- Routers are dumb, just do what they're told.

- Controller service explicitly tells each router what to do.

- Rare on the Internet, hot topic in data centers.

Software-Defined Networking (SDN)

Traditional Hardware

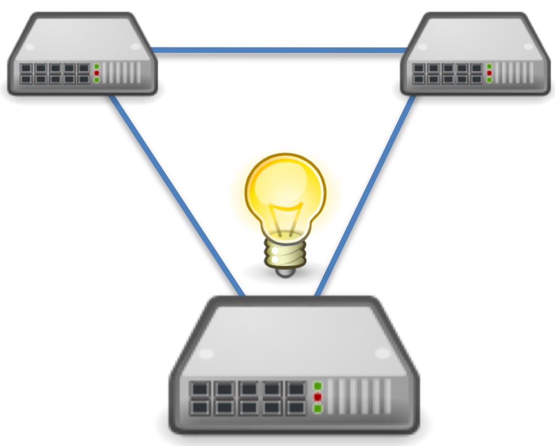

SDN Hardware

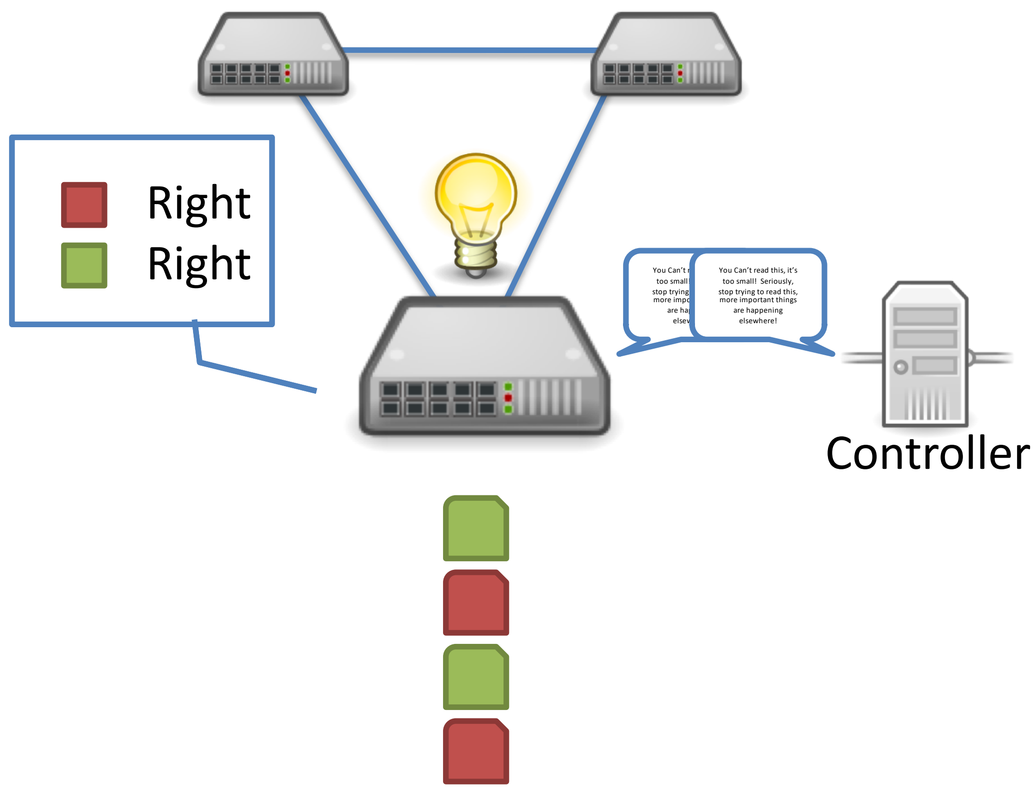

Summary

- Forwarding: moving packet from one interface to another (table lookup)
- Routing: Populating the table in advance

- On the Internet, best effort packet switching is the norm

- Hardware helps with quick forwarding using longest prefix matching