ECE 697J – Computer Networks Overview

10th September, 2002

University *of* Massachusetts Amherst

Last Class' Survey

- Suggestions:
  - Discuss alternatives to Active Networks
    - Added 1 or 2 papers: Overlay Networks, IPv6
- Concerns:
  - Not enough background in networking / architecture
    - Looks as if people have sufficient background
    - Make sure to contribute in your area
    - Try to catch up in other areas as we go
- Secure shell:
  - telnet sends password in cleartext
  - ssh encrypts everything (including password)

Last Class' Survey

- Background knowledge:

	I never heard of this.	I need review on this.	I'm comfortable that I understand this.
Computer network (i.e., routers, LANs)			10
TCP		1	9
IP		1	9
Processor architecture		7	3
Programmable logic	2	6	2

Class Overview

- Overview
  - Internet structure
  - Internet protocols: TCP/IP
  - Routers
  - Applications and Active Networks
- Paper assignments
- Reading and discussion

Internet Structure

- End-systems
  - Applications talk to other applications on other end-system
  - Clients, servers, peers
  - Increasing number of hosts: computers, wireless phones / PDAs, sensors
- Routers
  - Forward packets
  - Implement additional functions:
    - Quality of service
    - Firewalls
    - Load-balancing
  - Belong to different organizations

Internet Structure

- Backbone providers (NBP)
  - National / international
  - “core”
  - High-bandwidth, long-distance
  - UUNet/Worldcom, Sprint, AT&T
  - NBPs interconnect at public Network Access Points (NAPs) or private peering points
- Regional Internet Service Provider (ISP)
- Local ISPs
  - Local dial-up / broadband access
  - Uplink to regional ISP / backbone

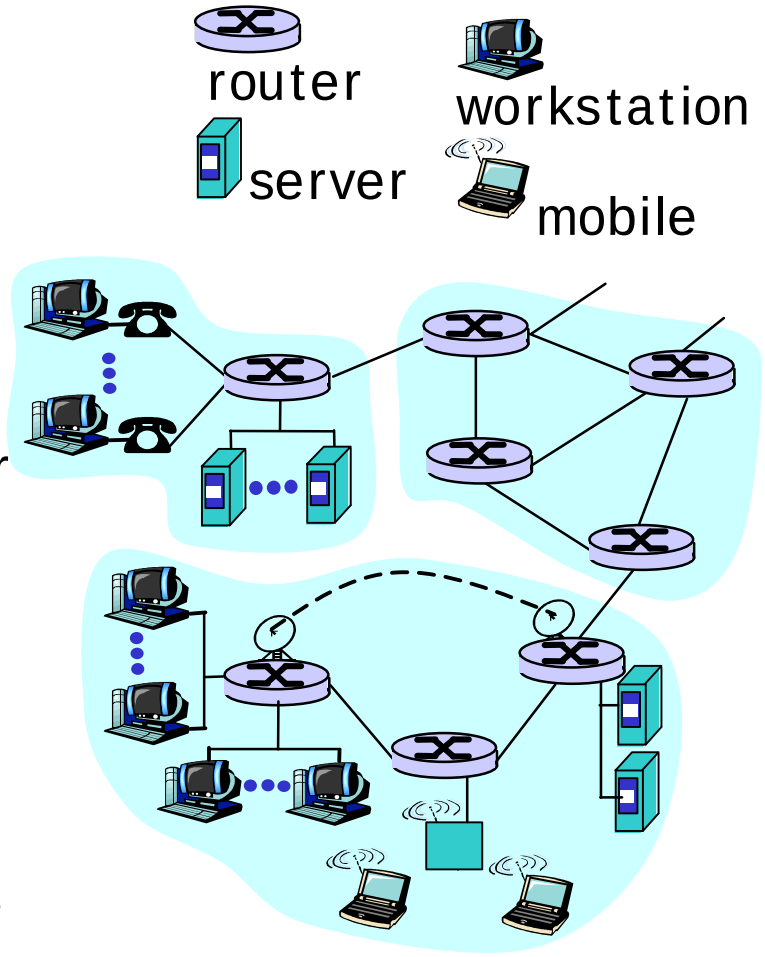

Sprint Backbone

OC-3 =
150Mb/s

OC-12 =
600Mb/s

OC-48 =
2.4Gb/s

(end 2001)

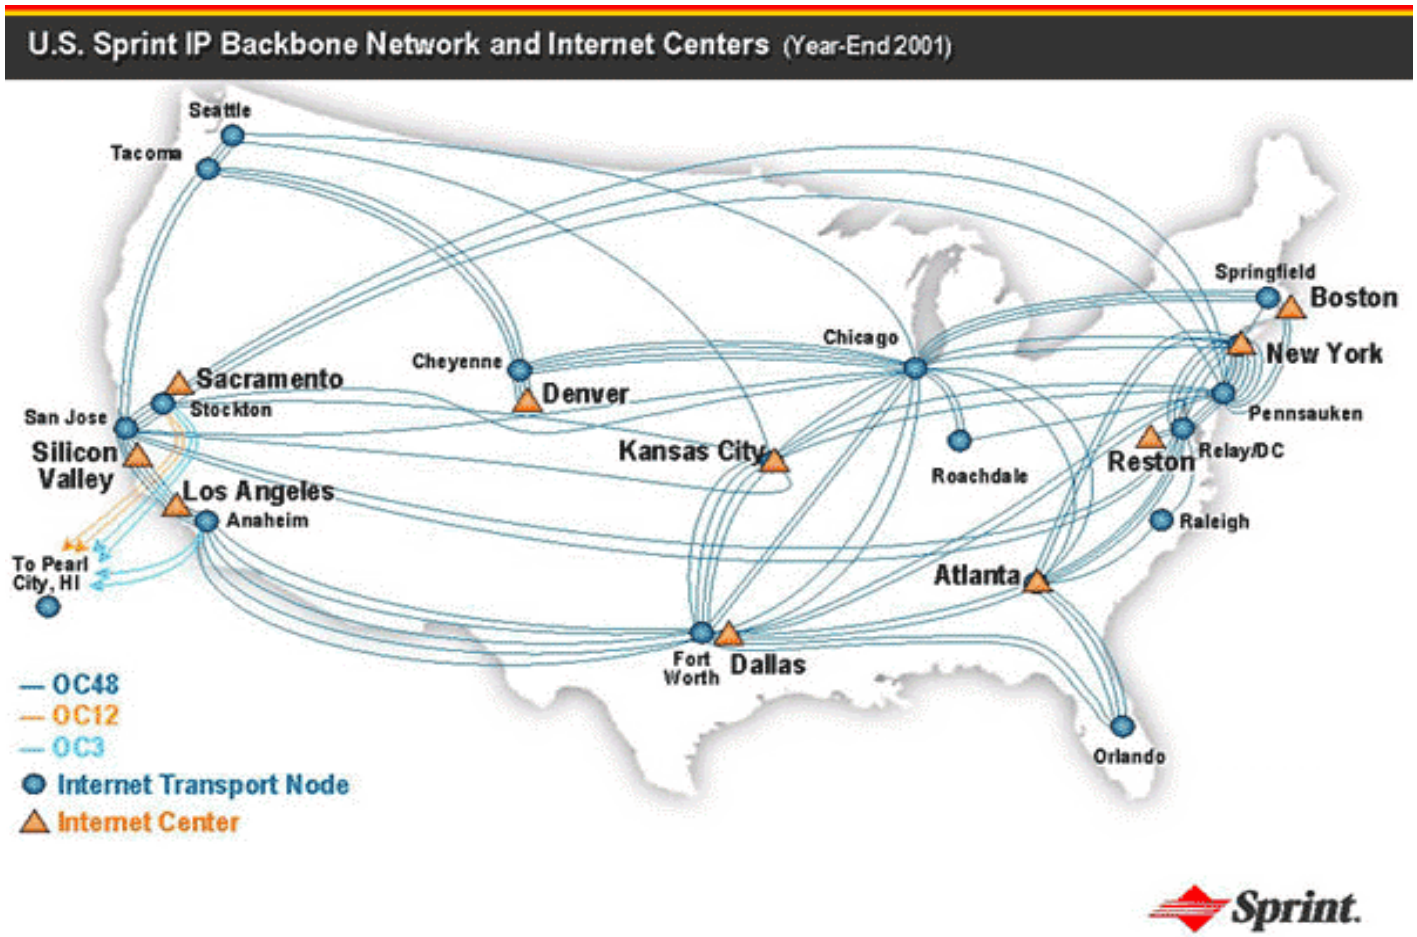

Backbone Growth

OC-3 = 150Mb/s

OC-12 = 600Mb/s

OC-48 = 2.4Gb/s

(end 2000)

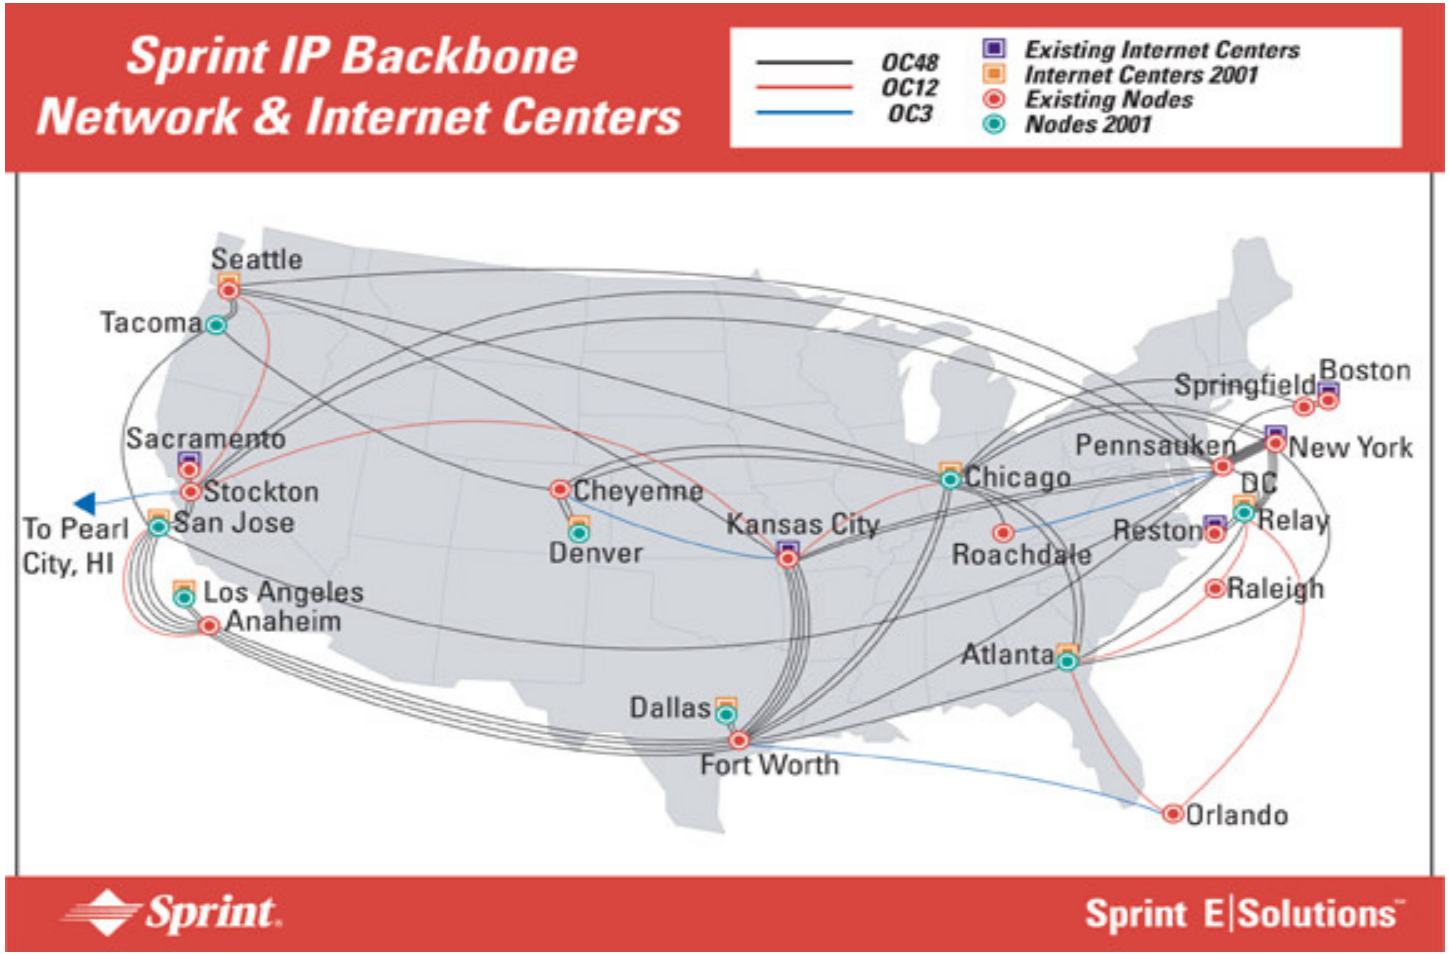

University *of* Massachusetts Amherst

International Backbone

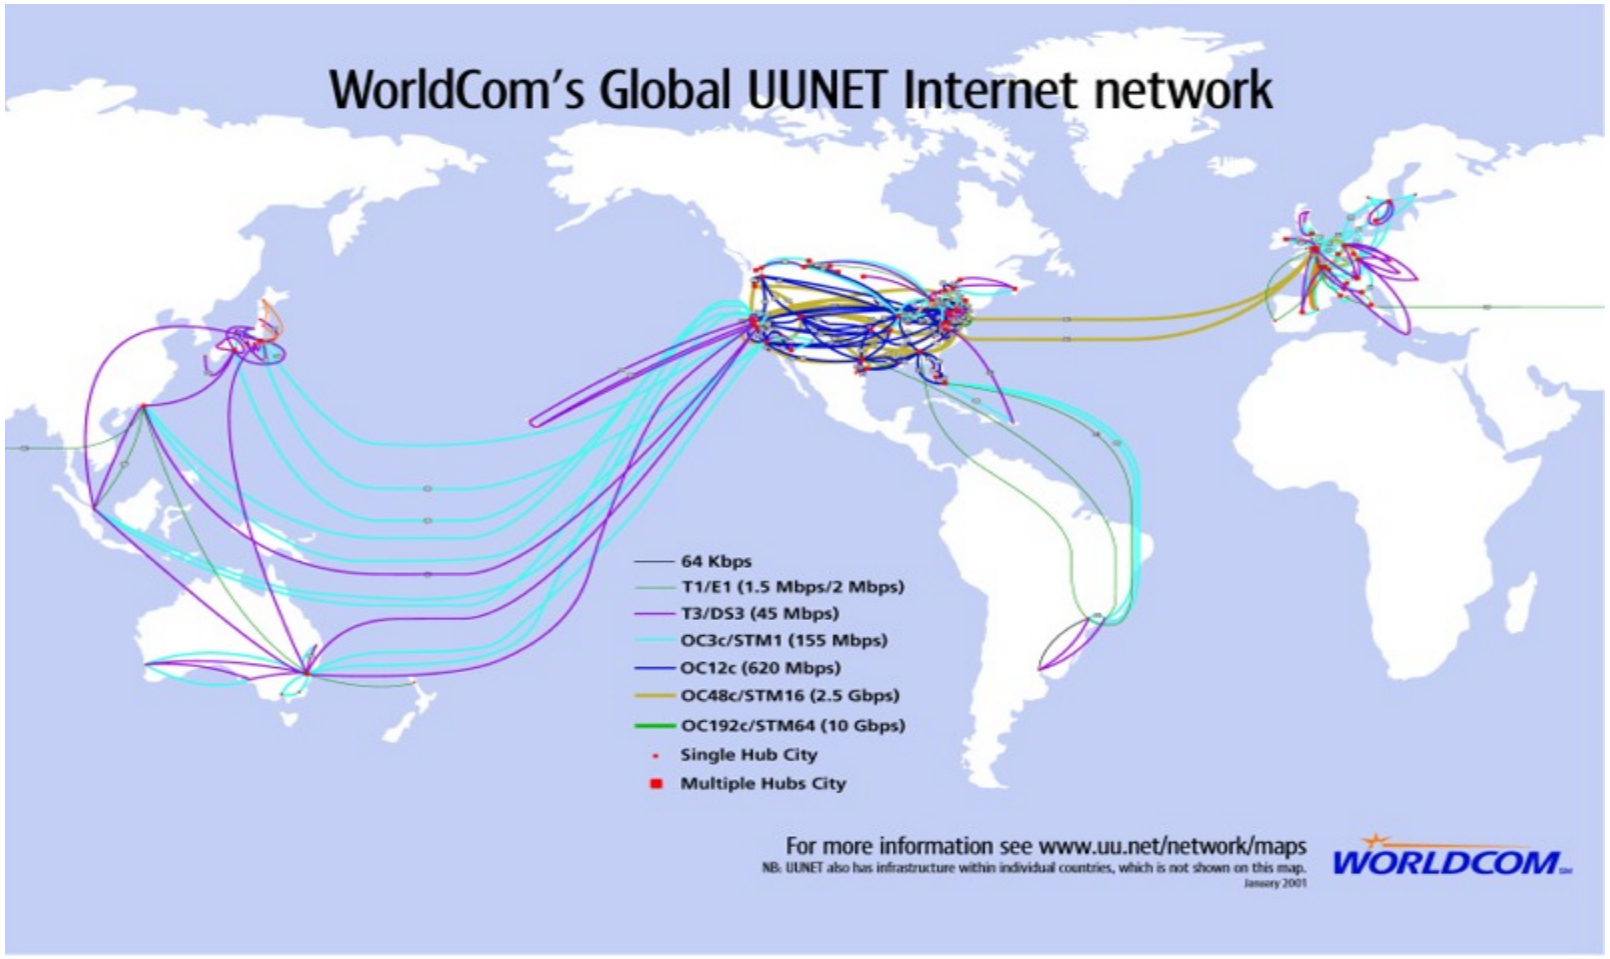

Internet Protocols

- Protocol stack:
  - Application Layer
    - Network applications (FTP, HTTP)
  - Transport Layer
    - Host-to-host data transfer (TCP, UDP)
  - Network Layer
    - Routing of datagrams from source to destination (IP)
  - Link Layer
    - Transfer of data between neighboring network elements (Ethernet, PPP)
  - Physical
    - Actual transmission of bits on the medium

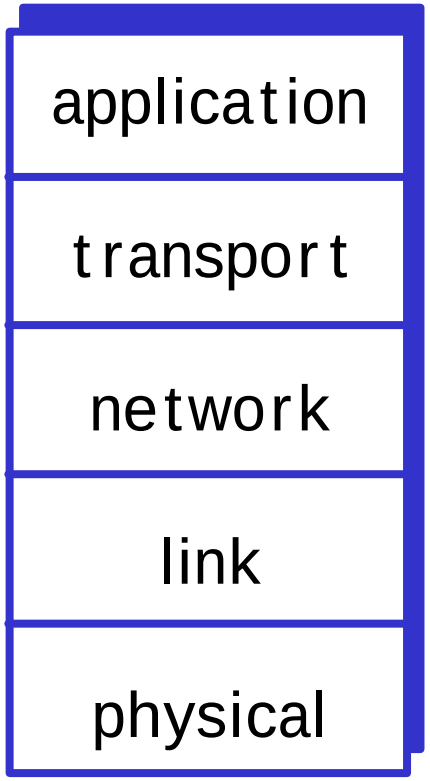

End-to-End Communication

- Protocol stack is traversed partially
  - Depends on functionality of network element
- Bridges
  - Layer 2
- Routers
  - Layer 3
- Gateways
  - Layer 4
- Application Layer Gateways
  - Layer 7

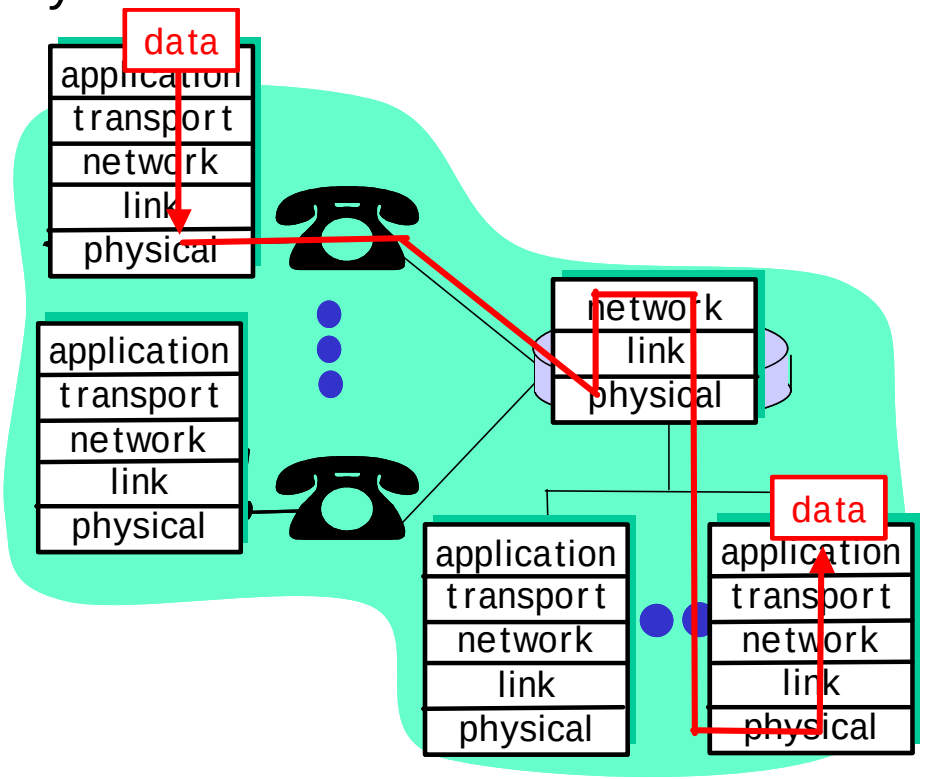

IP Protocol

- What does a router need to do to implement IP?

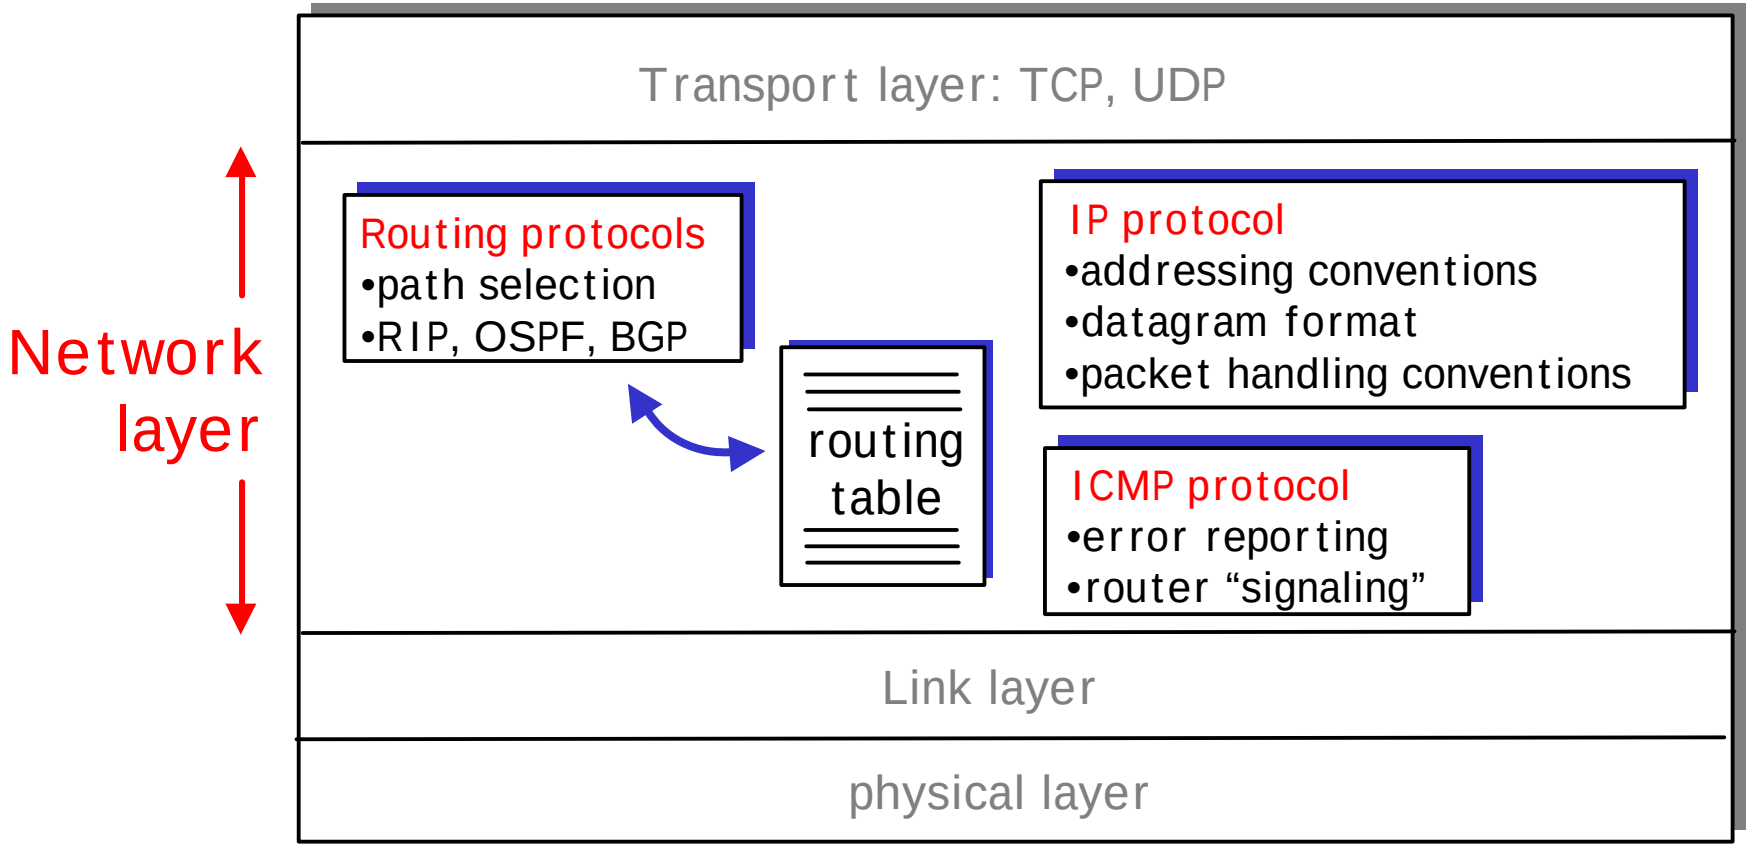

IP Processing

- Data path:
  - IP header checksum computation
  - IP destination address lookup
  - Decrement TTL field
  - Adjust header checksum
  - Forward packet to output port
  - Queue packet
  - Send packet on outgoing link
- Control path:
  - Routing message processing
  - ICMP processing

TCP Protocol

- Demultiplexes to individual application
- Reliable stream communication on top of unreliable IP
- TCP handles:
  - Reliability
  - Reordering
  - Packet loss
  - Duplicate packets
  - Flow control
  - Congestion control
- TCP level processing
  - QoS
  - Firewalling

TCP Segment Header

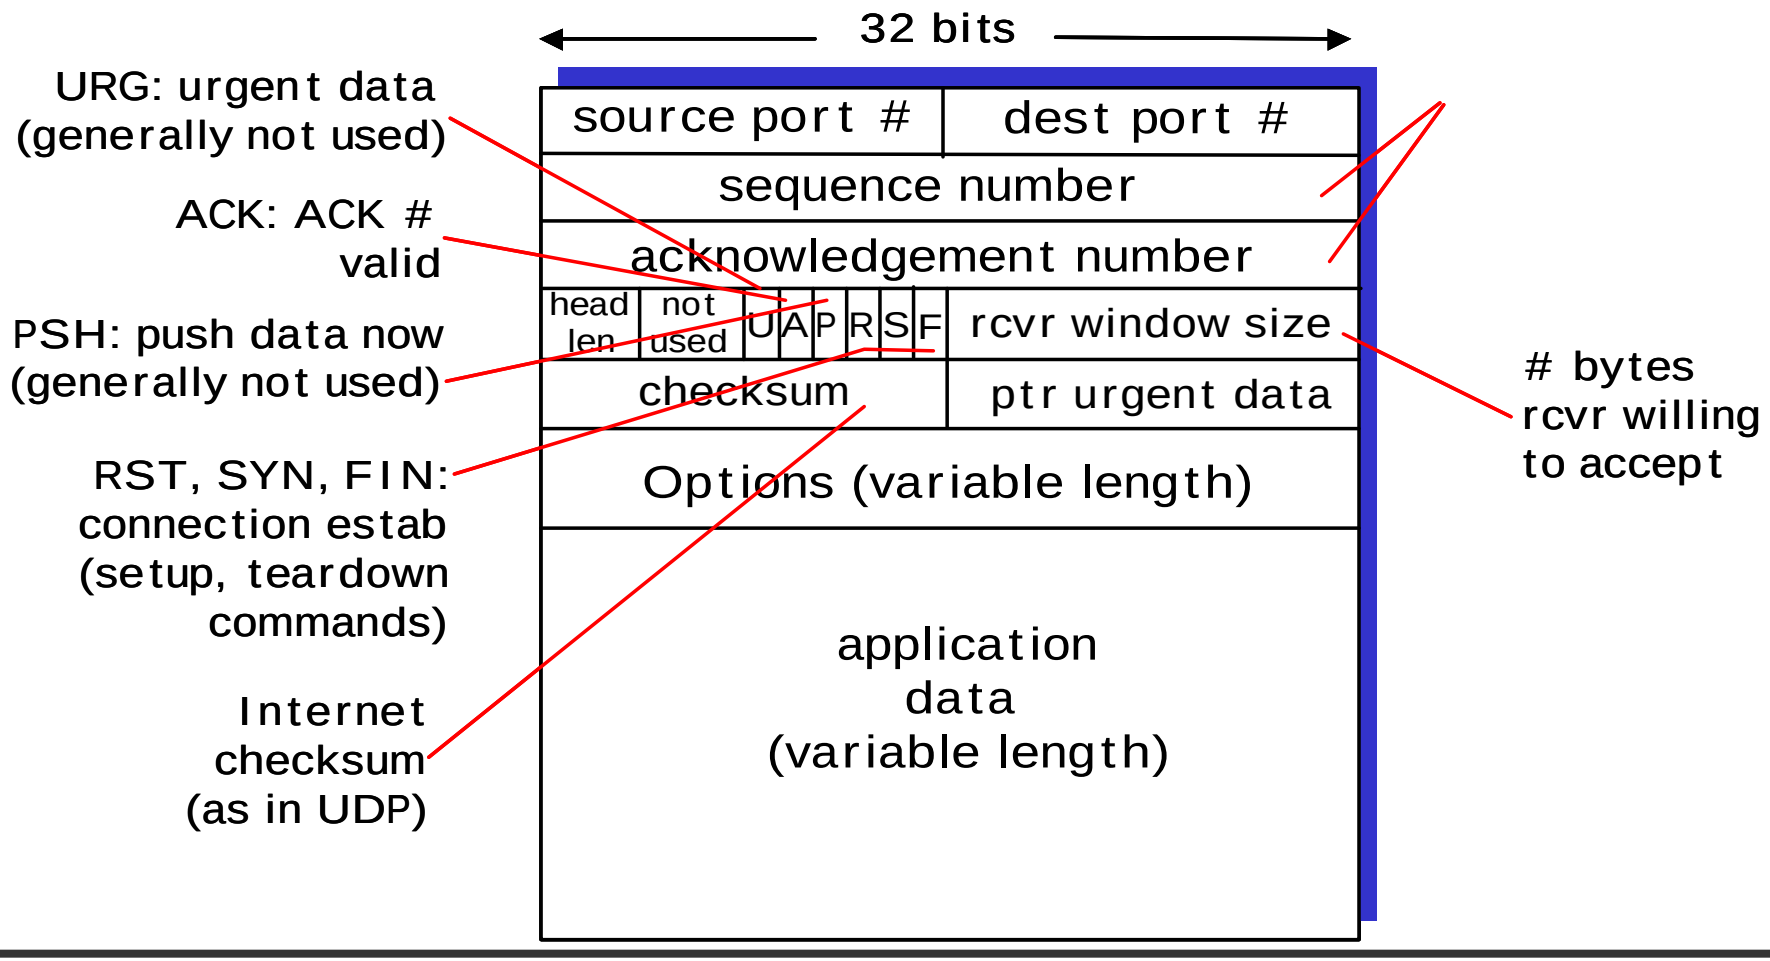

TCP Example

- Sequence #
  - Next expected byte
- ACK #
  - Last received byte
- Connection establishment
  - 3-way handshake

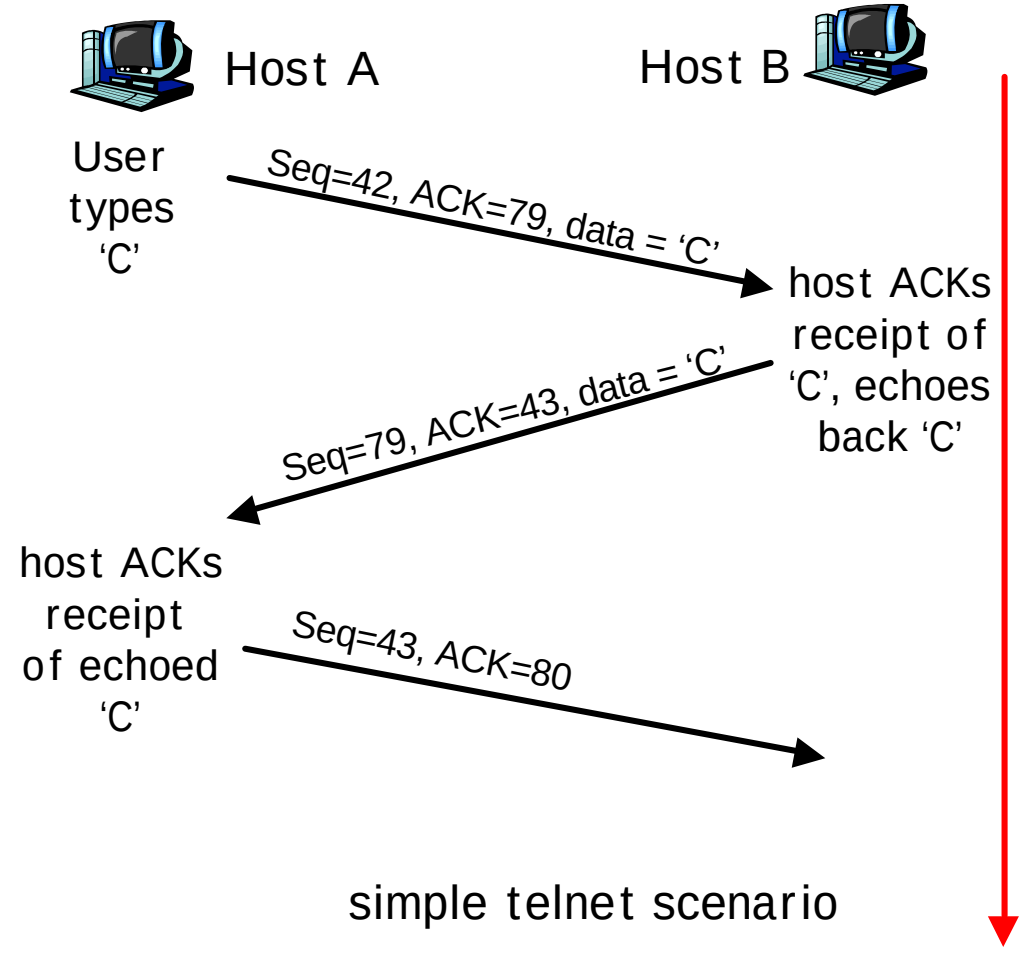

TCP Congestion Control

- TCP slow-start increases sending rate
- If “too much” is sent, packets are lost
- Sender reduces rate on packet loss
- Routers can intentionally reduce rate
  - Drop packet (Random Early Discard (RED))
- Packet loss due to other reasons causes problems
  - Wireless links (packet corruption, not congestion)
- Overall TCP achieves fairness among competing flows

Routers

- Simplest case
  - Workstation with 2+ network interfaces
  - Operating system performs routing (IP processing)
  - Easy to program and extend
  - Functions can be added by programming OS
    - Good for experimental purposes
- Problems
  - Very limited throughput
    - PCI bus bandwidth
    - Software efficiency
  - Not scalable
- Backbone routers need much more performance

Router Architecture

- Switching Fabric connects input and output ports
- Input side
  - Lookup
  - IP processing
  - Queuing*
- Output side
  - Queuing*
- Switching Fabric
  - High-speed
  - Non-blocking

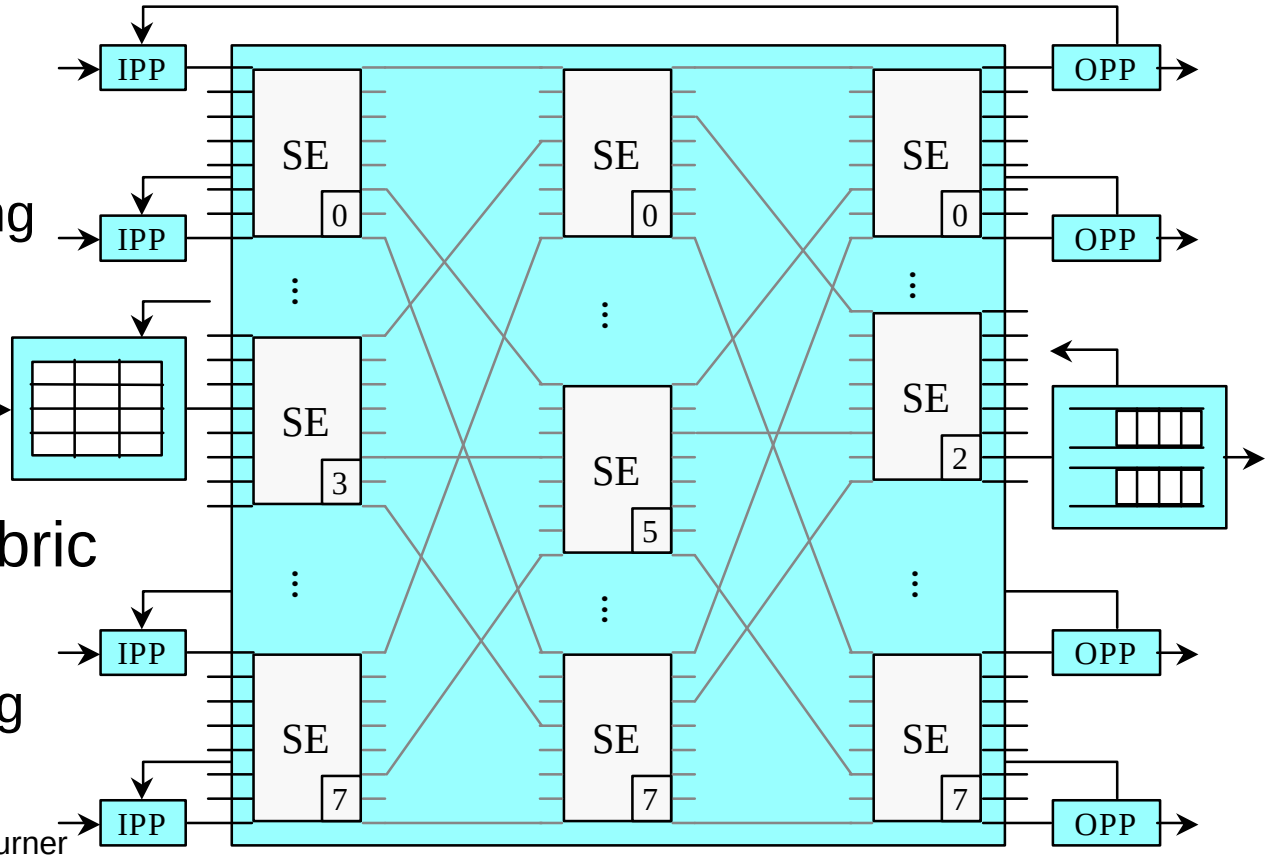

Figure by Jonathan S. Turner

Router Architecture

- Data path
  - Forwarding engine on each port
  - Distributed processing of packets
  - Scalable to large number of ports
- Control path
  - Control processor
  - Processes OSPF and/or BGP data
  - Maintains routing tables
  - Distributes forwarding tables to ports
- Challenges
  - Route lookup speed
  - Switch fabric throughput and cost
  - Queue memory bandwidth and size (TCP oscillation)

Why Active Networks?

- Internet designed for IP forwarding
  - IP is “common ground” – clearly defined
  - Difficult to add functionality
- Worldwide and commercial use requires additions
  - Security, firewalls
  - Quality of service
  - Network address translation (NAT), IPv6
- Typically: hardware solution for each problem
  - Firewall, multicast routers, QoS router, etc.
  - Costly replacement / additions
- Better: programmable router
  - Software determines functionality
  - Scalable design
- More on Thursday

Paper assignments

- Two preferences

Presentations

- Typical components
  - Introduction / motivation
  - Talk overview
  - Background / related work
  - Problem statement
  - Main contribution
  - Discussion
  - Extensions / future work
  - Summary
- Depends on paper / individual style / preferences

Discussion

- Homework from last time:
  - Questions to ask after presentation…

Reading Assignment

- Dates are on web page
- Please read paper
  - Won’t be announced as homework every time
- “Reading companion” sheet
  - Helps focus on typical questions
  - Optional
  - If you use it, bring it to class