

source: computer-networks-webdesign.com




CSCI 6760 - Computer Networks Spring 2017

Instructor: Prof. Roberto Perdisci
perdisci@cs.uga.edu

These slides are adapted from the textbook slides by J.F. Kurose and K.W. Ross

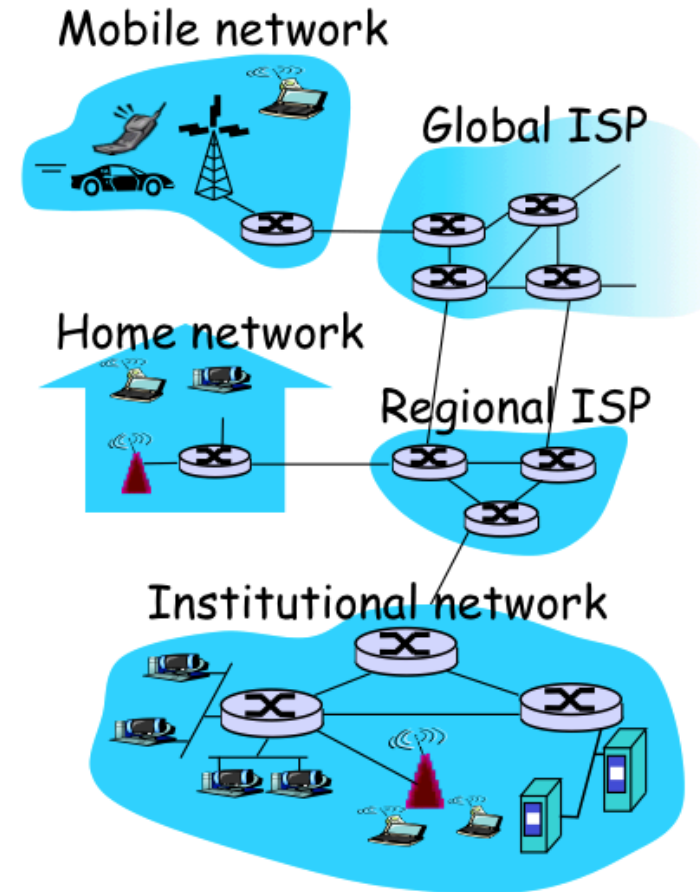
Introduction

- ▶ What is the Internet?
 - ▶ World-scale “network of networks”
 - ▶ Each network is essentially independent
 - ▶ No central authority (Registrars have some saying...)
 - ▶ Hundreds of millions of devices
 - ▶ Likely billions, considering mobile devs
 - ▶ Infrastructure that provides communication services to apps
- ▶ Host nodes (*hosts* for short)
 - ▶ Called *end systems*
 - ▶ run apps
 - ▶ Used to be computers, now include TVs, smart-phones, washing machines...
- ▶ Routers
 - ▶ Forward network packets
 - ▶ Make it possible to connect one network to another



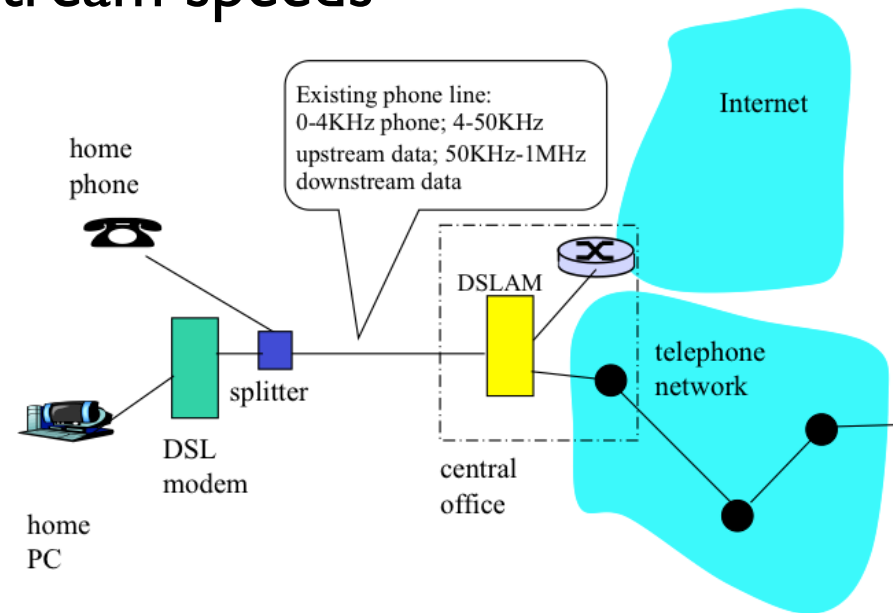
Introduction

- ▶ **How do hosts connect to the net?**
 - ▶ ISP = Internet Service Provider
 - ▶ Global vs. Regional ISPs
 - ▶ (e.g., AT&T, Comcast, Verizon, etc...)
- ▶ **Types of connections**
 - ▶ Dial-up (not common anymore)
 - ▶ DSL
 - ▶ Cable
 - ▶ Fiber
 - ▶ Wireless (3G, IMAX, WiFi)
 - ▶ Direct Ethernet access






Digital Subscriber Line (DSL)

- ▶ Uses existing PSTN infrastructure
- ▶ Dedicated physical line to telephone central office
- ▶ Asymmetric upstream/downstream speeds
 - ▶ 125kbps / 1.5Mbps
 - ▶ 256kbps / 3Mbps
 - ▶ ...
- ▶ Speed in bits per second (bps)
 - ▶ Typically limited by physical constraints
 - ▶ Rate-limited on purpose based on costs
 - ▶ Depending on contract




Cable

- ▶ Leverages cable TV infrastructure
- ▶ Asymmetric upstream/downstream speeds
 - ▶ 1Mbps-100Mbps
 - ▶ Upstream/Downstream speeds often differ (asymmetric)



Direct Ethernet-based Access

- ▶ Typical of companies, universities, etc.
- ▶ 10Mbps to 10Gbps
- ▶ End systems typically connect to a switch
- ▶ Access to Internet provided through institutional router
 - ▶ EITS provides access to UGA hosts




Wireless Access

- ▶ Mobile devices connect to an access point
- ▶ Access point connects to router

- ▶ Wireless LAN
 - ▶ 802.11b/g (up to 54Mbps)


- ▶ Wide-area Access

- ▶ Cellular system
 - ▶ GPRS, 3G, 4G
- ▶ Satellite



The Network Edge

- ▶ Communication models
 - ▶ Client / Server
 - ▶ Peer-to-Peer (P2P)
- ▶ Client / Server
 - ▶ Client host requests service
 - ▶ Server host provides service
 - ▶ E.g., Browser = Web Client
- ▶ P2P (often used for file sharing)
 - ▶ Minimal or no use of centralized servers
 - ▶ E.g., Skype, BitTorrent, Emule, ...



The Network Core

- ▶ Set of interconnected routers
- ▶ Forward data from one network to another


- ▶ Data transfer approaches:

1. Circuit Switching

- ▶ Communication resources between end hosts are *reserved*

2. Packet Switching

- ▶ *Shared* resources
- ▶ *Best effort* delivery



Circuit Switching

- ▶ **Dedicated communication resources**
 - ▶ Resources are reserved for the entire duration of the communication
 - ▶ E.g., phone call through PSTN uses circuit switching
- ▶ **Network resources (bandwidth) are “sliced”**
 - ▶ Circuit uses one or more slices
 - ▶ Access to resources using FDM or TDM
- ▶ **Performance**
 - ▶ Circuit setup time required
 - ▶ Guaranteed performance
 - ▶ No sharing
 - ▶ Resource idle if not used: **potential waste!**




Circuit Switching


Example:

FDM

4 users



TDM



Packet Switching


- ▶ End-to-end (or host-to-host) communications split into data chunks or ***packets***
- ▶ Each packet uses full link bandwidth
- ▶ Network users share resources
 - ▶ Resources used *as needed* (no reservation)
 - ▶ Aggregate demand may exceed available resources
 - ▶ Congestion may occur
 - ▶ wait for resources to become available
 - ▶ if too much congestion, packets may be lost
- ▶ Packets move one hop at a time
 - ▶ Store and forward
 - ▶ Nodes wait to receive entire packet before forwarding it



Packet Switching

▶ Statistical Multiplexing

- ▶ Packets arrive with no fixed timing pattern
- ▶ Bandwidth shared *on demand*
- ▶ Different from FDM/TDM, for which resource are guaranteed for entire “call time”

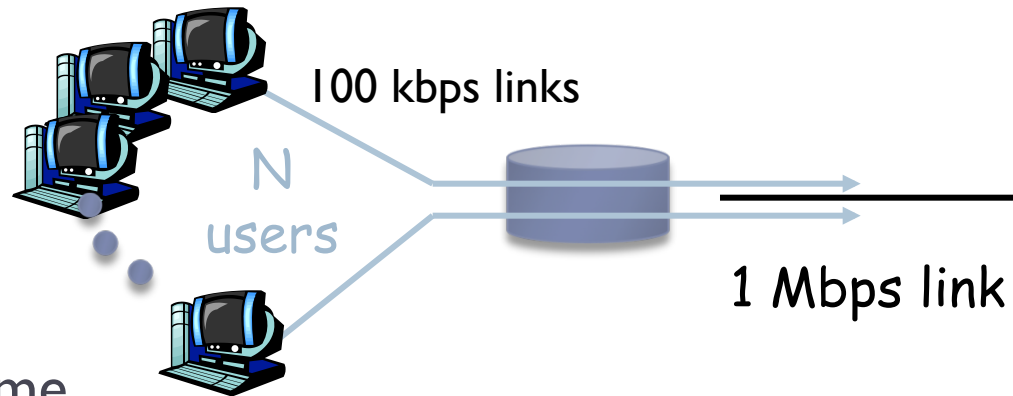


Packet Switching vs. Circuit Switching

- ▶ Packet switching allows more users to use the network

- ▶ **Example**

- ▶ 1 Mbps shared link
- ▶ N users
- ▶ Each user active 10% of time
- ▶ Users send 100 kbps each when active



- ▶ **Circuit switching**

- ▶ Allows only $N = 10$ users



- ▶ **Packet switching**

- ▶ How many users can share the same 1 Mbps link?



Packet Switching vs. Circuit Switching

- ▶ Packet switching allows more users to use the network




Packet Switching vs. Circuit Switching

- ▶ Packet switching allows more users to use the network

- ▶ **Example**

- ▶ 1 Mbps shared link
- ▶ N users
- ▶ Each user active 10% of time
- ▶ Users send 100 kbps each when active



- ▶ **Circuit switching**

- ▶ Allows only $N = 10$ users


- ▶ **Packet switching**

- ▶ Assuming $N = 35$, probability that more than 10 users are active at any given time is ~ 0.0004
 - ▶ Why?
- ▶ Therefore, more than 10 users are allowed to use the network




Packet Switching vs. Circuit Switching

- ▶ Packet switching does not waste bandwidth
- ▶ Example
 - ▶ Only 1 active user
 - ▶ User needs to send 1 MB of data
 - ▶ With TDM can only send 100 kbps = 80 sec
 - ▶ With packet switching can use entire bandwidth = 8 sec



Packet delays


- ▶ *Store-and-Forward*: the entire packet must arrive and stored, before a router can forward it to the next node



$$d_{\text{node}} = d_{\text{proc}} + d_{\text{queue}} + d_{\text{trans}} + d_{\text{prop}}$$

Packet delays

- ▶ d_{proc} : processing time
 - ▶ check for bit errors
 - ▶ lookup next hop link
- ▶ d_{queue} : queuing delay
 - ▶ time waiting at the output link packet queue
 - ▶ depends on link congestion
- ▶ d_{trans} : transmission delay
 - ▶ How long to copy packet on the link?
 - ▶ L: packet length (bits)
 - ▶ R: link bandwidth (bps)
 - ▶ $d_{\text{trans}} = L/R$




- ▶ d_{prop} : link propagation
 - ▶ How long for each bit to arrive to destination?
 - ▶ d: physical length of link
 - ▶ s: propagation speed (depends on type of link material)
 - ▶ $d_{\text{prop}} = d/s$
- ▶ $d_{\text{trans}} \neq d_{\text{prop}}$

$$\text{Bandwidth-Delay Product} = R * d_{\text{prop}}$$

Packet delays: Example

- ▶ NYC to London (5,500km) on Optical Fiber
- ▶ propagation speed $\sim 200,000\text{km/s}$
 - ▶ $d_{\text{prop}} = 5,500/200,000 = 27.5\text{ms}$
- ▶ Assume 15Mbps link bandwidth
- ▶ 1,500-byte packet
 - ▶ $d_{\text{trans}} = 8 * 1500 / 15E6 = 0.8\text{ms}$



- ▶ Assume also d_{queue} and d_{proc} are negligible
$$d_{\text{node}} = d_{\text{trans}} + d_{\text{prop}} = 28.3\text{ms}$$



Queuing delay

- ▶ R : link bandwidth (bps)
- ▶ L : packet length (bits)
- ▶ a : avg packet arrival rate

▶ La/R : **Traffic Intensity**

- ▶ $La/R \ll 1$ causes small avg delay
- ▶ As La/R increases towards 1 delay goes up
- ▶ $La/R > 1$ means more traffic arrives than can be handled by the link
 - ▶ Infinite delay == packet loss!



$La/R \sim 0$




$La/R \rightarrow 1$



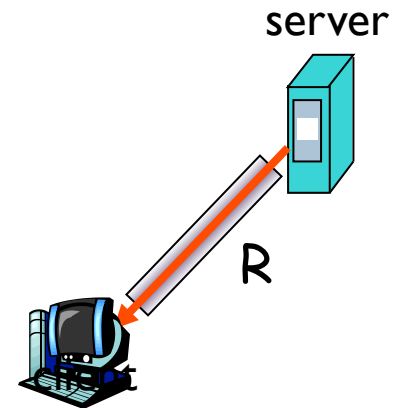
Packet Loss

- ▶ A and B are sharing the Internet connection
- ▶ Traffic Intensity $\lambda_a/R > 1$
- ▶ Router's buffer gets full
- ▶ B send packet, but router's buffer is full
- ▶ The packet will be discarded




End-to-End Throughput

- ▶ **Effective rate (bps) at which data is transferred between client and server**
 - ▶ Instantaneous throughput
 - ▶ bps that client receives at any given instant of time
 - ▶ Average throughput
 - ▶ overall throughput for a data transfer process
- ▶ **Example: file transfer**
 - ▶ F = file size, t = time taken to receive the entire file
 - ▶ Avg throughput = F/t
 - ▶ Inst. throughput may vary significantly from a given time instant to another
 - ▶ The higher the avg throughput, the better
- ▶ **Example2:VoIP**
 - ▶ High quality calls requires a constant minimum instant throughput and low delays between packets



End-to-End Throughput

- ▶ Effective rate (bps) at which data is transferred between client and server
- ▶ Assume that
 - ▶ $R_s = 2\text{Mbps}$, $R_c = 1\text{Mbps}$
 - ▶ $R = 5\text{Mbps}$ (equally shared)
 - ▶ $N = \#$ of clients and servers
 - ▶ $T = ???$
 - ▶ What is the effective throughput?




Example: $N=10$ connections share same link

- 10 simultaneous file downloads!

End-to-End Throughput

- ▶ Effective rate (bps) at which data is transferred between client and server
- ▶ Assume that
 - ▶ $R_s = 2\text{Mbps}$, $R_c = 1\text{Mbps}$
 - ▶ $R = 5\text{Mbps}$ (equally shared)
 - ▶ $N = \#$ of clients and servers
 - ▶ $T = \min(R_c, R_s, R/N)$




- Example: $N=10$ connections share same link
- 10 simultaneous file downloads!

The Internet is a network of networks

▶ Organized in a hierarchy

- ▶ Tier-1 ISPs (Level3, AT&T, etc...) and large content providers (Akamai, Google, etc.) are on top
- ▶ They peer (i.e., exchange traffic) directly or at IXPs
- ▶ IXP = Internet eXchange Point (check IXPs list on Wikipedia)

Tier-1 ISPs form the **Internet backbone**




Characteristics of Tier-1 ISPs

- directly connect to other Tier-1
- connect to lots of Tier-2
- international coverage



Internet Exchange Points

- ▶ Provide a facility where ISPs can “peer”



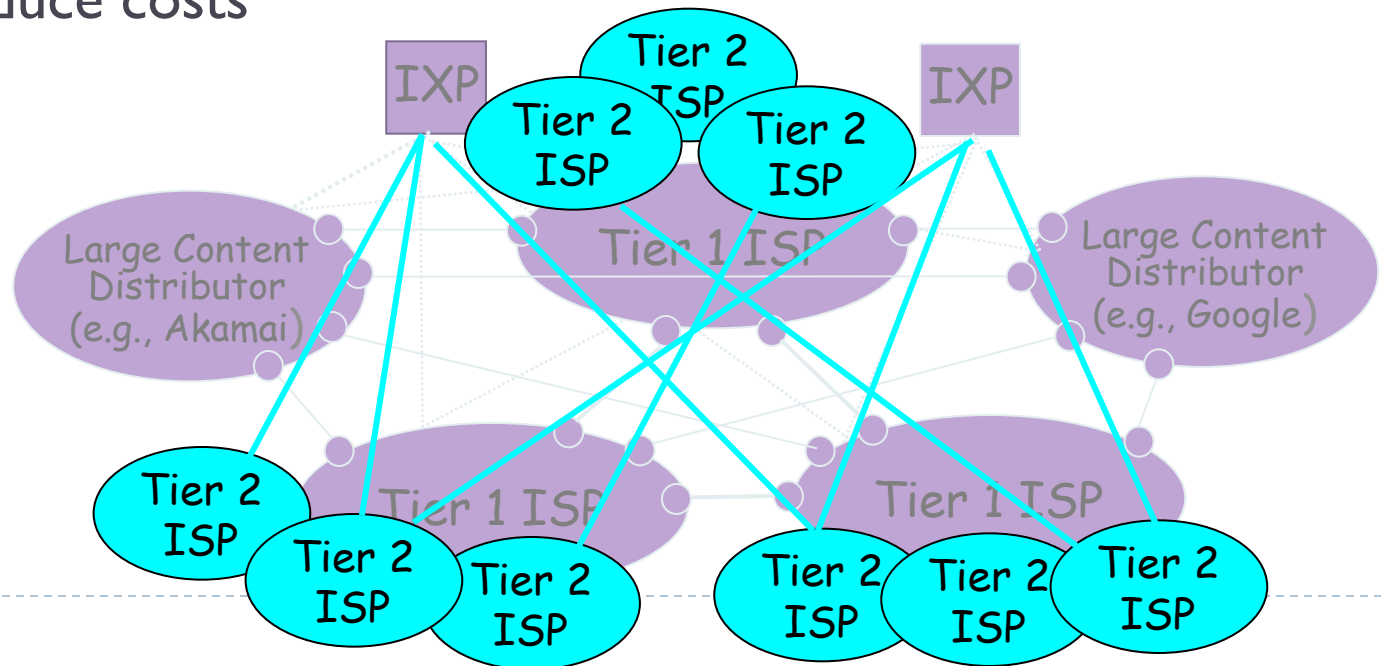
- 1 Alcatel-Lucent 7210 SAS-M
- 2 ADVA FSP3000R7 for Remote-Locations
- 3 Alcatel-Lucent 7950XRS20 Core-Node
- 4 Alcatel-Lucent 7950XRS40 Edge-Node
- 5 Alcatel-Lucent 7210 SAS-M
- 6 ADVA FSP3000R7 for Interconnect-Connections
- 7 Alcatel-Lucent 7950XRS20 Edge-Node



The Internet is a network of networks

▶ Tier-2 ISPs


- ▶ Smaller, often regional/national ISPs
- ▶ Pay to connect to one or a few Tier-1 ISPs
- ▶ Tier-1 ISPs have many Tier-2 ISP customers
- ▶ Tier-2 ISPs sometimes *peer* directly or at IXPs to bypass Tier-1 and reduce costs



The Internet is a network of networks

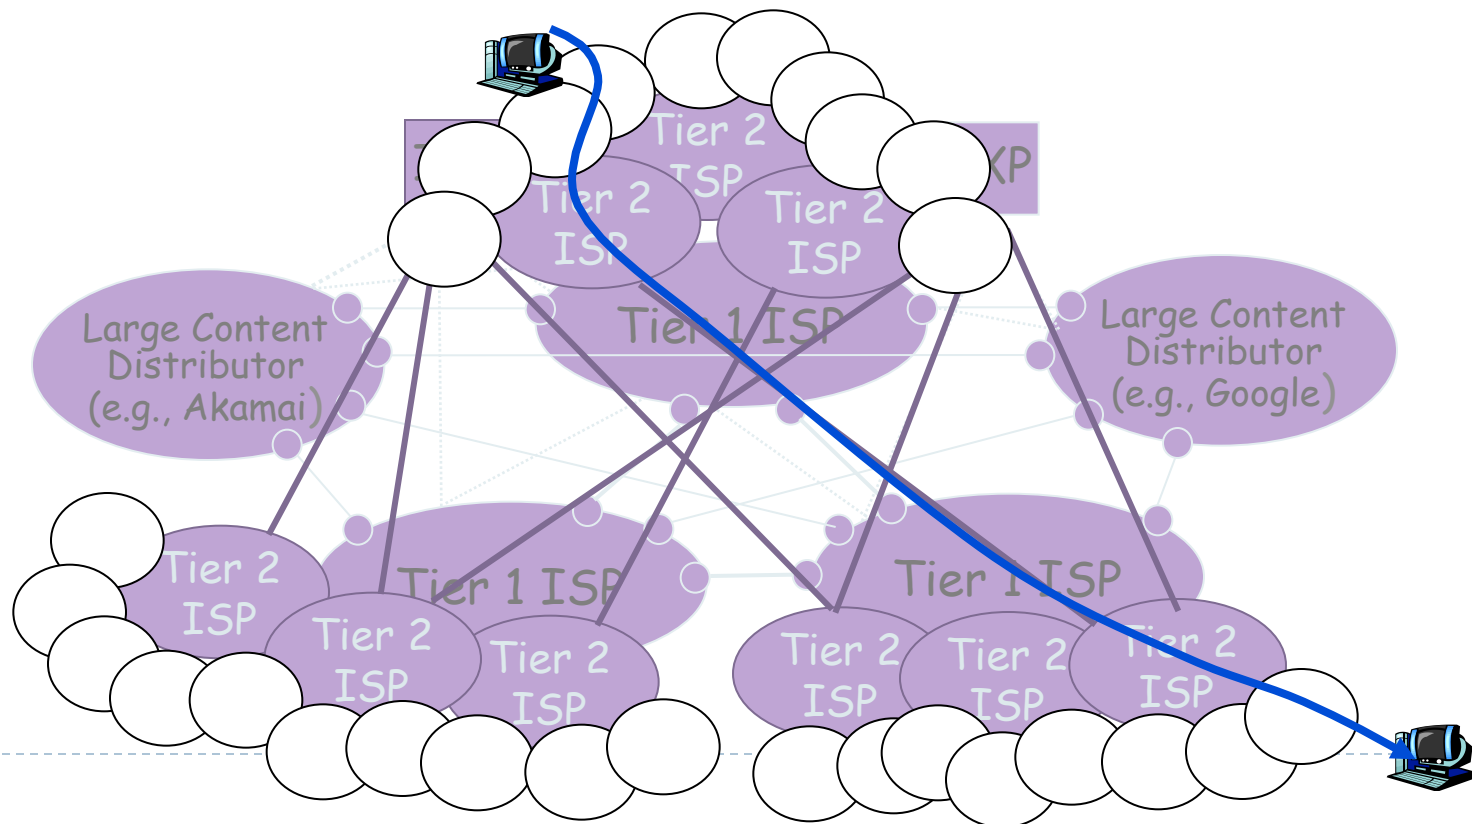
- ▶ Tier-3 ISPs are local ISPs
- ▶ Pay Tier-1 or Tier-2 ISPs to send/receive data
- ▶ Last hop, closest to end hosts

Some Tier-1 ISPs also offer lower-Tier type services (e.g., AT&T is also a local access ISP)



The Internet is a network of networks

- ▶ Packets from one end host to another usually traverse several networks at different levels



How do packets get to destination?*

traceroute to www.italia.gov.it (94.86.40.47), 30 hops max, 40 byte packets

```
1 128.192.76.129 (128.192.76.129) 0.525 ms 0.638 ms 0.747 ms
...
4 eboydf.net.uga.edu (128.192.166.69) 1.637 ms 2.250 ms 2.376 ms
5 h70-33-127-97.paws.uga.edu (70.33.127.97) 1.860 ms 1.494 ms 2.556 ms
6 spnetx.net.uga.edu (128.192.166.1) 3.473 ms 2.992 ms 3.204 ms
7 131.144.206.45 (131.144.206.45) 120.032 ms 119.842 ms 3.581 ms
8 xe-3-1-921.r00.atInga05.us.bb.gin.ntt.net (204.2.241.33) 3.573 ms 4.159 ms 4.117 ms
9 ae-0.r20.atInga05.us.bb.gin.ntt.net (129.250.3.176) 4.683 ms 4.104 ms 4.078 ms
10 p64-0-1-0.r21.dllstx09.us.bb.gin.ntt.net (129.250.5.26) 31.783 ms 31.931 ms 31.931 ms
11 ae-2.r08.dllstx09.us.bb.gin.ntt.net (129.250.3.81) 31.277 ms 31.917 ms 30.932 ms
...
14 te8-1.ashburn1.ash.seabone.net (89.221.40.3) 147.764 ms 147.870 ms 149.757 ms
15 te0-1-0-7.newyork50.new.seabone.net (195.22.206.3) 153.118 ms 156.356 ms 156.260 ms
16 pos0-10-0-0.milano50.mil.seabone.net (195.22.216.215) 214.033 ms 214.035 ms 214.105 ms
17 ibs-resid.milano50.mil.seabone.net (93.186.128.162) 145.318 ms 145.968 ms 142.934 ms
18 ***
...
21 80.21.5.86 (80.21.5.86) 162.830 ms 162.743 ms 165.573 ms
22 host106-35-static.58-88-b.business.telecomitalia.it (88.58.35.106) 151.528 ms 154.627 ms 154.625 ms
23 ***
24 host47-40-static.86-94-b.business.telecomitalia.it (94.86.40.47) 166.256 ms 161.895 ms 162.122 ms
```




More on Traceroute....

- ▶ **Demo Time!**
 - ▶ ping
 - ▶ traceroute
 - ▶ throughput test



Internet Protocol Stack

- ▶ **Application:** supports network applications
 - ▶ Example: HTTP, FTP, SMTP, ...
- ▶ **Transport:** process-to-process data transfer
 - ▶ Example: TCP, UDP
- ▶ **Network:** routing of datagrams from source host to destination host
 - ▶ IP
- ▶ **Link:** data transfer between neighbor nodes
 - ▶ Ethernet, 802.11x (WiFi), PPP
- ▶ **Physical:** bits on the wire



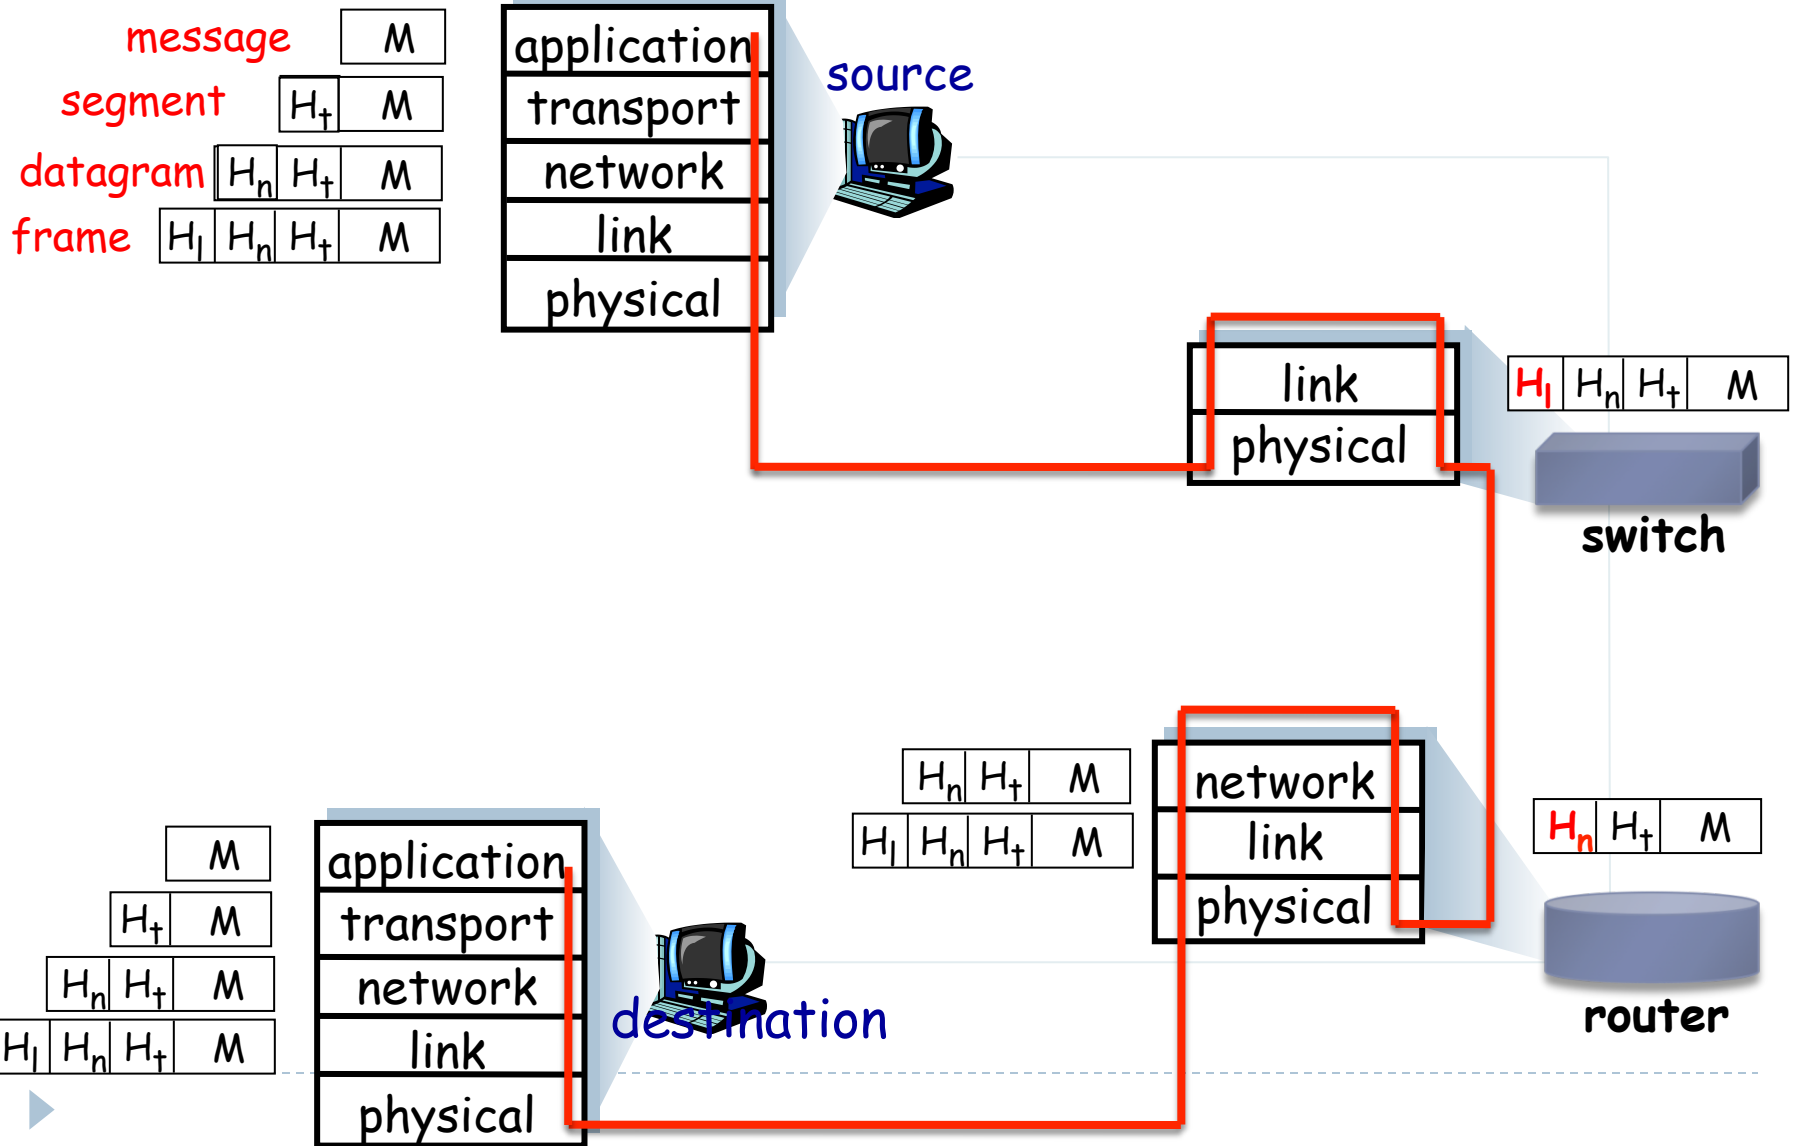
ISO/OSI reference model

- ▶ **Presentation:** allows application level protocol to correctly interpret/send data
 - ▶ Convert data format according to a specific encryption/compression algorithm
 - ▶ Machine-specific encoding
- ▶ **Session:**
 - ▶ synchronization, checkpoint, recovery of data exchange
- ▶ These two layers are missing from the Internet Stack
 - ▶ When needed, must be implemented at the application level

Application
Presentation
Session
Transport
Network
Link
Physical

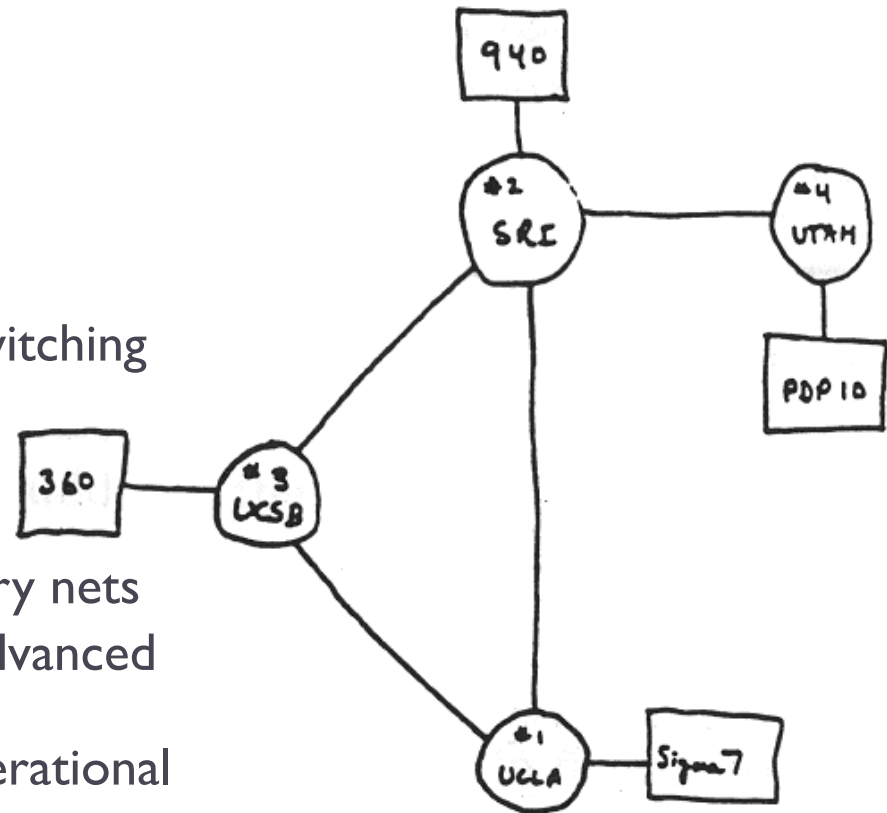


Encapsulation



Internet History

- ▶ 1960s – early 1970s
 - ▶ Development of early packet-switching principles
 - ▶ 1961 : queuing theory
 - ▶ 1964 : packet switching in military nets
 - ▶ 1967 : ARPANet conceived by Advanced Research Project Agency
 - ▶ 1969 : first ARPANet node is operational
 - ▶ 1972 : ARPANet public demonstration
 - ▶ NCP (Network Control Protocol) is the first host-to-host protocol
 - ▶ First e-mail application
 - ▶ ARPANet has 15 nodes
 - ▶ 1976: Ethernet developed at Xerox (competing with Token Ring at IBM)
 - ▶ 1979 : 200 ARPANet nodes



THE ARPA NETWORK

Internet History

- ▶ **1980s : new protocols developed, new networks**
 - ▶ 1982: SMTP protocol defined
 - ▶ 1983: deployment of TCP/IP
 - ▶ 1983: DNS defined for name-to-IP mapping
 - ▶ 1985: FTP protocol
 - ▶ 1988: TCP congestion

- ▶ **Several new national networks**
 - ▶ Cernet, BITnet, NSFnet, Minitel
 - ▶ 100,000 nodes interconnected



Internet History

- ▶ 1990s – 2000's: the Web, new apps, commercialization
- ▶ Early 1990s
 - ▶ Hypertext
 - ▶ HTML, HTTP
 - ▶ Mosaic (Netscape)
- ▶ Late 1990s – 2000's
 - ▶ Commercialization of the Web
 - ▶ P2P applications
 - ▶ Instant messaging
 - ▶ Internet backbone at Gbps
 - ▶ Network Security becomes super important
 - ▶ Evolved from few trusted nodes to millions of untrusted ones



Internet History

- ▶ ~750 Million hosts
- ▶ Real-time apps: VoIP (Skype), Video Streaming (PPLive)
- ▶ Web 2.0
 - ▶ Youtube, Online games, ...
 - ▶ Social Networks: Facebook, Twitter, MySpace, Linked-in, ...
- ▶ The Cloud: Gmail, Amazon
- ▶ Wireless Internet is becoming pervasive
 - ▶ 3G, 4G, WiMAX, ...



Internetworking design principles

- ▶ **Minimalism and autonomy**
 - ▶ No internal changes needed to interconnect networks
- ▶ **Simplicity**
 - ▶ Best effort model
- ▶ **Stateless routers**
- ▶ **Decentralized control**

- ▶ **Principles set by Vint Cerf and Bob Kahn (TCP/IP)**
 - ▶ Define today's Internet architecture

- ▶ **Simplicity vs. Security ?**




Who “controls” the Internet?

- ▶ Internet = large set of interconnected networks
- ▶ No central management
 - ▶ Each network is operated and managed independently
- ▶ However, a number of things need to be **coordinated**
 - ▶ Assignment of IP addresses, AS numbers
 - ▶ Registration of domain names



Internet Assigned Numbers Authority



- ▶ *IANA is responsible for coordinating some of the key elements that keep the Internet running smoothly*
 - ▶ IP addresses, AS numbers
 - ▶ Domain Names
 - Roots, .int, .arpa, IDNs
 - ▶ Protocol Assignments in collaboration with IETF
- ▶ *IANA is supervised by Internet Corporation for Assigned Names and Numbers (ICANN)*
- ▶ *Internet Engineering Task Force: The mission of the IETF is to make the Internet work better by producing documents that influence the way people design, use, and manage the Internet*

Network Security

- ❖ **field of network security:**
 - how bad guys can attack computer networks
 - how we can defend networks against attacks
 - how to design architectures that are immune to attacks
- ❖ **Internet not originally designed with (much) security in mind**
 - *original vision:* “a group of mutually trusting users attached to a transparent network” 😊
 - Internet protocol designers playing “catch-up”
 - security considerations in all layers!

Fundamental Security Components

▶ Confidentiality

- ▶ Secrecy of information (usually achieved using crypto)

▶ Integrity

- ▶ Trustworthiness of data
 - ▶ Prevention: deny unauthorized changes
 - ▶ Detection: identify if unauthorized changes happened



▶ Availability

- ▶ Ability to access data/resources

▶ Authentication

- ▶ Verification of someone's identity

▶ Authorization

- ▶ Check if user has permission to perform a certain action



Bad guys: compromise hosts via Internet

- ❖ malware can get in host from a **virus, worm, or trojan horse**.
- ❖ **spyware malware** can record keystrokes, web sites visited, upload info to collection site.
- ❖ infected host can be enrolled in **botnet**, used for spam and DDoS attacks.
- ❖ malware often **self-replicating**: from one infected host, seeks entry into other hosts




Bad guys: compromise hosts via Internet

Denial of Service (DoS): attackers make resources (server, bandwidth) unavailable to legitimate traffic by overwhelming resource with bogus traffic

1. select target
2. break into hosts around the network (see botnet)
3. send packets to target from compromised hosts


Example:
Recent Events connected to Wikileaks



The bad guys can sniff packets

Packet sniffing:

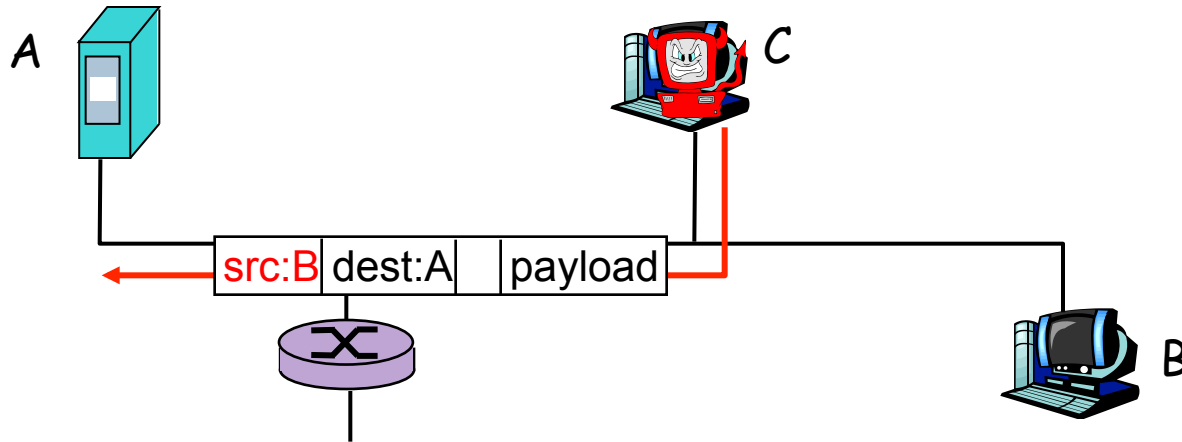
- ▶ broadcast media (shared Ethernet, wireless)
- ▶ promiscuous network interface reads/records all packets (e.g., including passwords!) passing by



- ❖ Wireshark software used for end-of-chapter labs is a (free) packet-sniffer

The bad guys can use false source addresses


IP spoofing: send packet with false source address



The bad guys can record and playback

record-and-playback: sniff sensitive info (e.g., password), and use later

- ▶ password holder is that user from system point of view



... lots more on security (throughout, Chapter 8)