

Course Review: Chapter 1

- What's a protocol?
- Circuit Switching and Packet Switching
- Network Taxonomy
- Four Sources of Packet Delay
- Internet Protocol Stack

Introduction 1-1

What is a Protocol?

*protocols define format,
order of msgs sent and
received among network
entities, and actions
taken on msg
transmission, receipt*

Introduction 1-2

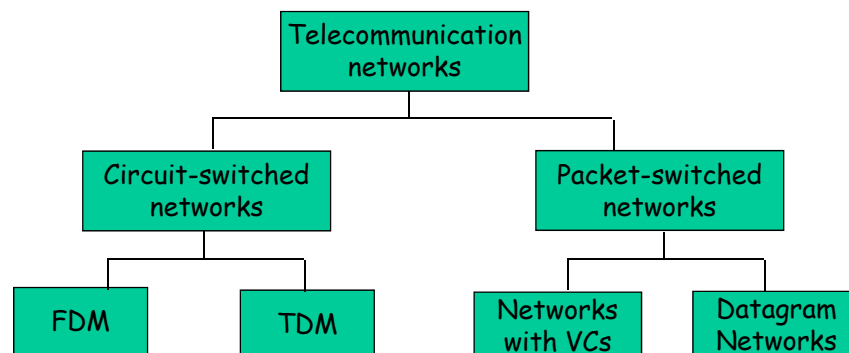
Packet switching versus circuit switching

Is packet switching a "slam dunk winner?"

- ❑ Great for bursty data
 - resource sharing
 - simpler, no call setup
- ❑ Excessive congestion: packet delay and loss
 - protocols needed for reliable data transfer, congestion control
- ❑ Q: How to provide circuit-like behavior?
 - bandwidth guarantees needed for audio/video apps
 - still an unsolved problem (chapter 6)

Introduction 1-3

Network Taxonomy



- Datagram network is not either connection-oriented or connectionless.
- Internet provides both connection-oriented (TCP) and connectionless services (UDP) to apps.

Introduction 1-4


Four sources of packet delay

1. nodal processing:

- check bit errors
- determine output link

2. queueing

- time waiting at output link for transmission
- depends on congestion level of router



Introduction 1-5

Delay in packet-switched networks


3. Transmission delay:

- R = link bandwidth (bps)
- L = packet length (bits)
- time to send bits into link = L/R

4. Propagation delay:

- d = length of physical link
- s = propagation speed in medium ($\sim 2 \times 10^8$ m/sec)
- propagation delay = d/s

Note: s and R are very different quantities!



Introduction 1-6

Nodal delay


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

- ❑ d_{proc} = processing delay
 - typically a few microsecs or less
- ❑ d_{queue} = queuing delay
 - depends on congestion
- ❑ d_{trans} = transmission delay
 - $= L/R$, significant for low-speed links
- ❑ d_{prop} = propagation delay
 - a few microsecs to hundreds of msecs

Introduction 1-7

Internet protocol stack

- ❑ **application:** supporting network applications
 - FTP, SMTP, STTP
- ❑ **transport:** host-host data transfer
 - TCP, UDP
- ❑ **network:** routing of datagrams from source to destination
 - IP, routing protocols
- ❑ **link:** data transfer between neighboring network elements
 - PPP, Ethernet
- ❑ **physical:** bits "on the wire"



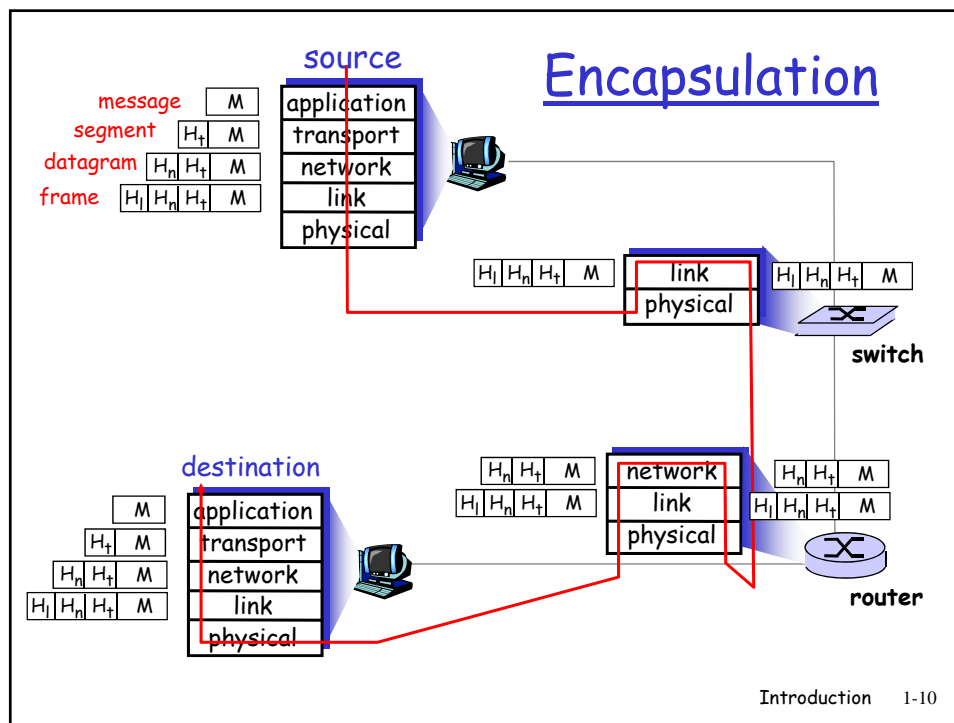
Introduction 1-8

Why layering?

Dealing with complex systems:

- ❑ explicit structure allows identification, relationship of complex system's pieces
 - layered **reference model** for discussion
- ❑ modularization eases maintenance, updating of system
 - change of implementation of layer's service transparent to rest of system
 - e.g., change in gate procedure doesn't affect rest of system
- ❑ layering considered harmful?

Introduction 1-9



Course Review: Chapter 2

- Application Architectures
- Application Layer Protocols
- Caching
- Response Time Modeling
- Pop3 vs. IMAP
- Domain Name Server
- P2P file sharing

Introduction 1-11

Application architectures

- Client-server
- Peer-to-peer (P2P)
- Hybrid of client-server and P2P

Introduction 1-12

Internet transport protocols services

TCP service:

- ❑ *connection-oriented*: setup required between client and server processes
- ❑ *reliable transport* between sending and receiving process
- ❑ *flow control*: sender won't overwhelm receiver
- ❑ *congestion control*: throttle sender when network overloaded
- ❑ *does not provide*: timing, minimum bandwidth guarantees

UDP service:

- ❑ unreliable data transfer between sending and receiving process
- ❑ does not provide: connection setup, reliability, flow control, congestion control, timing, or bandwidth guarantee

Q: why bother? Why is there a UDP?

Introduction 1-13

Internet apps: application, transport protocols


Application	Application layer protocol	Underlying transport protocol
e-mail	SMTP [RFC 2821]	TCP
remote terminal access	Telnet [RFC 854]	TCP
Web	HTTP [RFC 2616]	TCP
file transfer	FTP [RFC 959]	TCP
streaming multimedia	proprietary (e.g. RealNetworks)	TCP or UDP
Internet telephony	proprietary (e.g., Dialpad)	typically UDP

Introduction 1-14

Web caches (proxy server)

Goal: satisfy client request without involving origin server

- ❑ user sets browser: Web accesses via cache
- ❑ browser sends all HTTP requests to cache
 - object in cache: cache returns object
 - else cache requests object from origin server, then returns object to client



Introduction 1-15

More about Web caching

- ❑ Cache acts as both client and server
- ❑ Typically cache is installed by ISP (university, company, residential ISP)

Why Web caching?

- ❑ Reduce response time for client request.
- ❑ Reduce traffic on an institution's access link.
- ❑ Internet dense with caches enables "poor" content providers to effectively deliver content (but so does P2P file sharing)

Introduction 1-16


Response time modeling

Definition of RTT: time to send a small packet to travel from client to server and back.

Response time:

- ❑ one RTT to initiate TCP connection
- ❑ one RTT for HTTP request and first few bytes of HTTP response to return
- ❑ file transmission time

total = 2RTT + transmit time



Introduction 1-17

POP3 (more) and IMAP

More about POP3

- ❑ Previous example uses "download and delete" mode.
- ❑ Bob cannot re-read e-mail if he changes client
- ❑ "Download-and-keep": copies of messages on different clients
- ❑ POP3 is stateless across sessions

IMAP

- ❑ Keep all messages in one place: the server
- ❑ Allows user to organize messages in folders
- ❑ IMAP keeps user state across sessions:
 - names of folders and mappings between message IDs and folder name

Introduction 1-18

DNS: Domain Name System

People: many identifiers:

- SSN, name, passport #

Internet hosts, routers:

- IP address (32 bit) - used for addressing datagrams
- "name", e.g.,
ww.yahoo.com - used by humans

Q: map between IP addresses and name ?

Domain Name System:

- *distributed database*
implemented in hierarchy of many *name servers*
- *application-layer protocol*
host, routers, name servers to communicate to *resolve* names (address/name translation)
 - note: core Internet function, implemented as application-layer protocol
 - complexity at network's "edge"

Introduction 1-19

DNS

DNS services

- Hostname to IP address translation
- Host aliasing
 - Canonical and alias names
- Mail server aliasing
- Load distribution
 - Replicated Web servers: set of IP addresses for one canonical name

Why not centralize DNS?

- single point of failure
- traffic volume
- distant centralized database
- maintenance

doesn't *scale*!

Introduction 1-20

P2P: problems with centralized directory

- ❑ Single point of failure
- ❑ Performance bottleneck
- ❑ Copyright infringement

file transfer is decentralized, but locating content is highly decentralized

Introduction 1-21

Query flooding: Gnutella

- ❑ fully distributed
 - no central server
- ❑ public domain protocol
- ❑ many Gnutella clients implementing protocol


overlay network: graph

- ❑ edge between peer X and Y if there's a TCP connection
- ❑ all active peers and edges is overlay net
- ❑ Edge is not a physical link
- ❑ Given peer will typically be connected with < 10 overlay neighbors

Introduction 1-22

Exploiting heterogeneity: KaZaA

- ❑ Each peer is either a group leader or assigned to a group leader.
 - TCP connection between peer and its group leader.
 - TCP connections between some pairs of group leaders.
- ❑ Group leader tracks the content in all its children.



Introduction 1-23


Course Review: Chapter 3

- Transport Services and Protocols
- Transport Layer vs. Network Layer
- UDP
- GBN and Selective Repeat
- TCP
- Flow Control vs. Congestion Control
- TCP Fairness

Introduction 1-24

Transport services and protocols

- ❑ provide *logical communication* between app processes running on different hosts
- ❑ transport protocols run in end systems
 - send side: breaks app messages into *segments*, passes to network layer
 - rcv side: reassembles segments into messages, passes to app layer
- ❑ more than one transport protocol available to apps
 - Internet: TCP and UDP



Introduction 1-25

Transport vs. network layer

- ❑ *network layer*: logical communication between hosts
- ❑ *transport layer*: logical communication between processes
 - relies on, enhances, network layer services

Household analogy:


12 kids sending letters to 12 kids

- ❑ processes = kids
- ❑ app messages = letters in envelopes
- ❑ hosts = houses
- ❑ transport protocol = Ann and Bill
- ❑ network-layer protocol = postal service

Introduction 1-26

Internet transport-layer protocols

- ❑ reliable, in-order delivery (TCP)
 - congestion control
 - flow control
 - connection setup
- ❑ unreliable, unordered delivery: UDP
 - no-frills extension of "best-effort" IP
- ❑ services not available:
 - delay guarantees
 - bandwidth guarantees



Introduction 1-27

UDP: User Datagram Protocol [RFC 768]

- ❑ "no frills," "bare bones" Internet transport protocol
- ❑ "best effort" service, UDP segments may be:
 - lost
 - delivered out of order to app
- ❑ **connectionless**:
 - no handshaking between UDP sender, receiver
 - each UDP segment handled independently of others

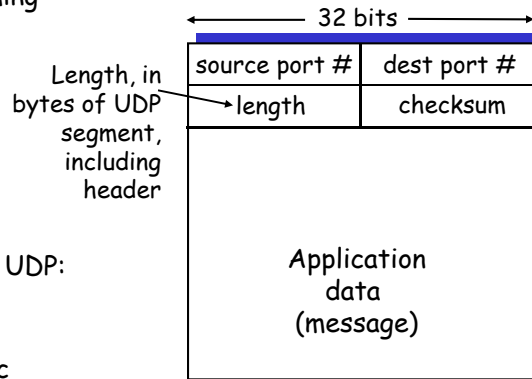
Why is there a UDP?

- ❑ no connection establishment (which can add delay)
- ❑ simple: no connection state at sender, receiver
- ❑ small segment header
- ❑ no congestion control: UDP can blast away as fast as desired

Introduction 1-28

UDP: more

- ❑ often used for streaming multimedia apps
 - loss tolerant
 - rate sensitive
- ❑ other UDP uses
 - DNS
 - SNMP
- ❑ reliable transfer over UDP: add reliability at application layer
 - application-specific error recovery!




UDP segment format

Introduction 1-29

TCP: Overview

RFCs: 793, 1122, 1323, 2018, 2581

- ❑ **point-to-point:**
 - one sender, one receiver
- ❑ **reliable, in-order byte stream:**
 - no "message boundaries"
- ❑ **pipelined:**
 - TCP congestion and flow control set window size
- ❑ **send & receive buffers**
- ❑ **full duplex data:**
 - bi-directional data flow in same connection
 - MSS: maximum segment size
- ❑ **connection-oriented:**
 - handshaking (exchange of control msgs) init's sender, receiver state before data exchange
- ❑ **flow controlled:**
 - sender will not overwhelm receiver



Introduction 1-30

TCP Round Trip Time and Timeout

$$\text{EstimatedRTT} = (1 - \alpha) * \text{EstimatedRTT} + \alpha * \text{SampleRTT}$$

- ❑ Exponential weighted moving average
- ❑ influence of past sample decreases exponentially fast
- ❑ typical value: $\alpha = 0.125$

Introduction 1-31

TCP: Overview

RFCs: 793, 1122, 1323, 2018, 2581

- ❑ **point-to-point:**
 - one sender, one receiver
- ❑ **reliable, in-order byte stream:**
 - no "message boundaries"
- ❑ **pipelined:**
 - TCP congestion and flow control set window size
- ❑ **send & receive buffers**
- ❑ **full duplex data:**
 - bi-directional data flow in same connection
 - MSS: maximum segment size
- ❑ **connection-oriented:**
 - handshaking (exchange of control msgs) init's sender, receiver state before data exchange
- ❑ **flow controlled:**
 - sender will not overwhelm receiver



Introduction 1-32

TCP Connection Management

Recall: TCP sender, receiver establish "connection" before exchanging data segments

- ❑ initialize TCP variables:
 - seq. #s
 - buffers, flow control info (e.g. RcvWindow)
- ❑ *client*: connection initiator

```
Socket clientSocket = new
Socket("hostname", "port
number");
```
- ❑ *server*: contacted by client

```
Socket connectionSocket =
welcomeSocket.accept();
```

SYN flood attack!!!

Three way handshake:

Step 1: client host sends TCP SYN segment to server

- specifies initial seq #
- no data

Step 2: server host receives SYN, replies with SYNACK segment


- server allocates buffers
- specifies server initial seq. #

Step 3: client receives SYNACK, replies with ACK segment, which may contain data

Introduction 1-33

TCP Flow Control

- ❑ receive side of TCP connection has a receive buffer:



- ❑ app process may be slow at reading from buffer

flow control

sender won't overflow receiver's buffer by transmitting too much, too fast

- ❑ speed-matching service: matching the send rate to the receiving app's drain rate

Introduction 1-34

Principles of Congestion Control

Congestion:

- ❑ informally: "too many sources sending too much data too fast for *network* to handle"
- ❑ different from flow control!
- ❑ manifestations:
 - lost packets (buffer overflow at routers)
 - long delays (queueing in router buffers)
- ❑ a top-10 problem!

Introduction 1-35

TCP Congestion Control

three mechanisms:

- AIMD
- slow start
- conservative after timeout events

Introduction 1-36


TCP AIMD

multiplicative decrease:

cut CongWin in half
after loss event

additive increase:

increase CongWin by
1 MSS every RTT in
the absence of loss
events: *probing*



Long-lived TCP connection

Introduction 1-37


TCP Slow Start

- ❑ When connection begins, CongWin = 1 MSS
 - Example: MSS = 500 bytes & RTT = 200 msec
 - initial rate = 20 kbps
- ❑ available bandwidth may be \gg MSS/RTT
 - desirable to quickly ramp up to respectable rate
- ❑ When connection begins, increase rate exponentially fast until first loss event

Introduction 1-38

TCP Fairness

Fairness goal: if K TCP sessions share same bottleneck link of bandwidth R , each should have average rate of R/K



Introduction 1-39

Fairness (more)

Fairness and UDP

- ❑ Multimedia apps often do not use TCP
 - do not want rate throttled by congestion control
- ❑ Instead use UDP:
 - pump audio/video at constant rate, tolerate packet loss
- ❑ Research area: TCP friendly

Fairness and parallel TCP connections

- ❑ nothing prevents app from opening parallel connections between 2 hosts.
- ❑ Web browsers do this
- ❑ Example: link of rate R supporting 9 connections:
 - new app asks for 1 TCP, gets rate $R/10$
 - new app asks for 11 TCPs, gets $R/2$!

Introduction 1-40

Key Network-Layer Functions

- *forwarding*: move packets from router's input to appropriate router output
- *routing*: determine route taken by packets from source to dest.
 - *Routing algorithms*
- analogy:
- *routing*: process of planning trip from source to dest
- *forwarding*: process of getting through single interchange

Introduction 1-41

Course Review: Chapter 4

- Forwarding vs. Routing
- Network layer connection and connection-less service
- Router Architecture
- IP Address/CIDR
- Network Address Translation
- IPv6

Introduction 1-42


Network layer connection and connection-less service

- ❑ Datagram network provides network-layer connectionless service
- ❑ VC network provides network-layer connection service
- ❑ Analogous to the transport-layer services, but:
 - **Service:** host-to-host
 - **No choice:** network provides one or the other
 - **Implementation:** in the core

Introduction 1-43

Datagram networks

- ❑ no call setup at network layer
- ❑ routers: no state about end-to-end connections
 - no network-level concept of "connection"
- ❑ packets forwarded using destination host address
 - packets between same source-dest pair may take different paths



Introduction 1-44

Longest prefix matching

<u>Prefix Match</u>	<u>Link Interface</u>
11001000 00010111 00010	0
11001000 00010111 00011000	1
11001000 00010111 00011	2
otherwise	3

Examples

DA: 11001000 00010111 00010110 10100001 Which interface?


DA: 11001000 00010111 00011000 10101010 Which interface?

Introduction 1-45

Router Architecture Overview

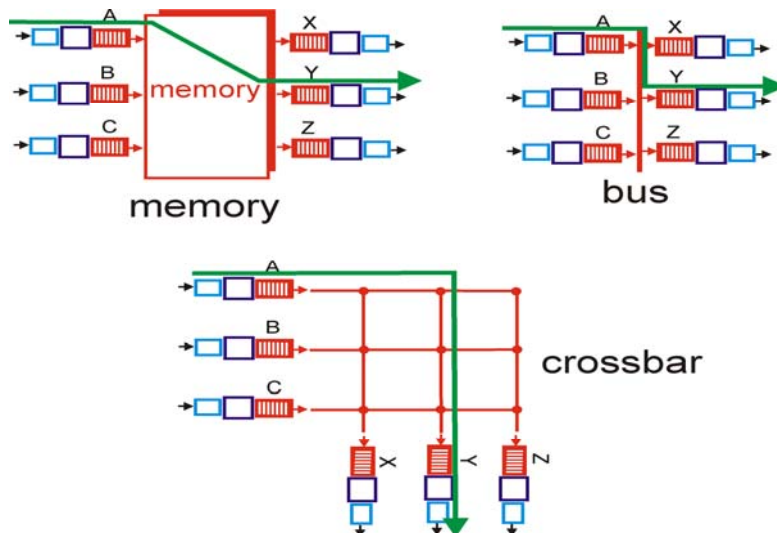
Two key router functions:

- ❑ run routing algorithms/protocol (RIP, OSPF, BGP)
- ❑ *forwarding* datagrams from incoming to outgoing link



Introduction 1-46

Three types of switching fabrics




Introduction 1-47

IP addressing: CIDR

CIDR: Classless InterDomain Routing

- subnet portion of address of arbitrary length
- address format: **a.b.c.d/x**, where x is # bits in subnet portion of address



Introduction 1-48