CSE 461: Midterm Review

Spring 2021

Administrivia

- Midterm on Monday 5/3

- HW3 due Monday 5/10

- Project 2 due Wednesday 5/12

Midterm Review Section

- Today: A brief review of lecture materials
  - Concepts, Protocols, Algorithms, …
- What **YOU** should do after this section and before the exam:
  - Go through the lecture slides
  - Think about the **problems** that each protocol/algorithm tries to solve
    - Pros and cons of current approaches?
   - Any other possible solutions?
    - What has not been solved yet?

Networks Overview

- Parts of a network
- Types of links
- Key interfaces
- Sockets
- Traceroute
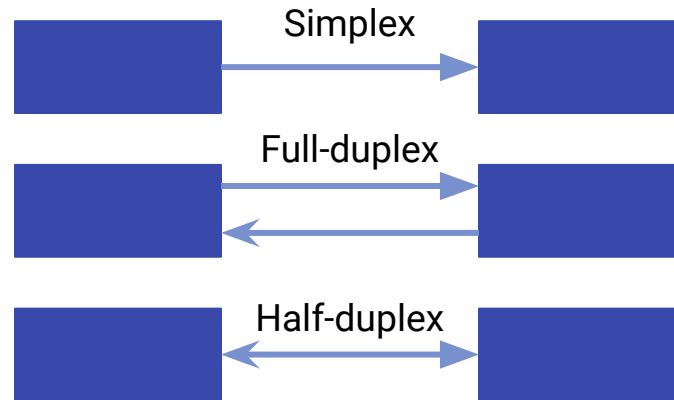- Protocols and layers
- Encapsulation

Parts of a Network

- Parts of a Network

host   router   link   app

- Types of Links

Simplex

Full-duplex

Half-duplex

Protocols and Layers

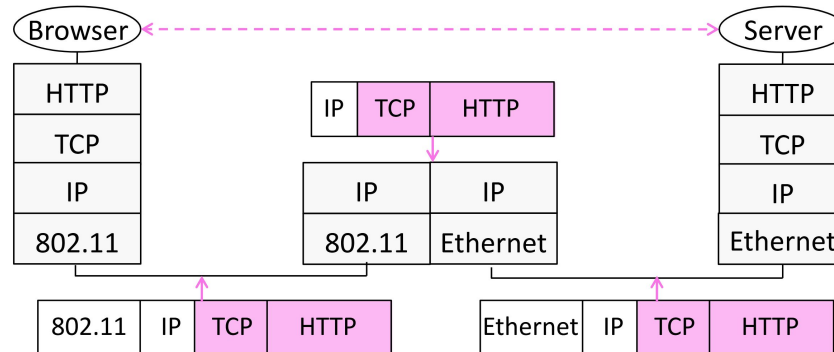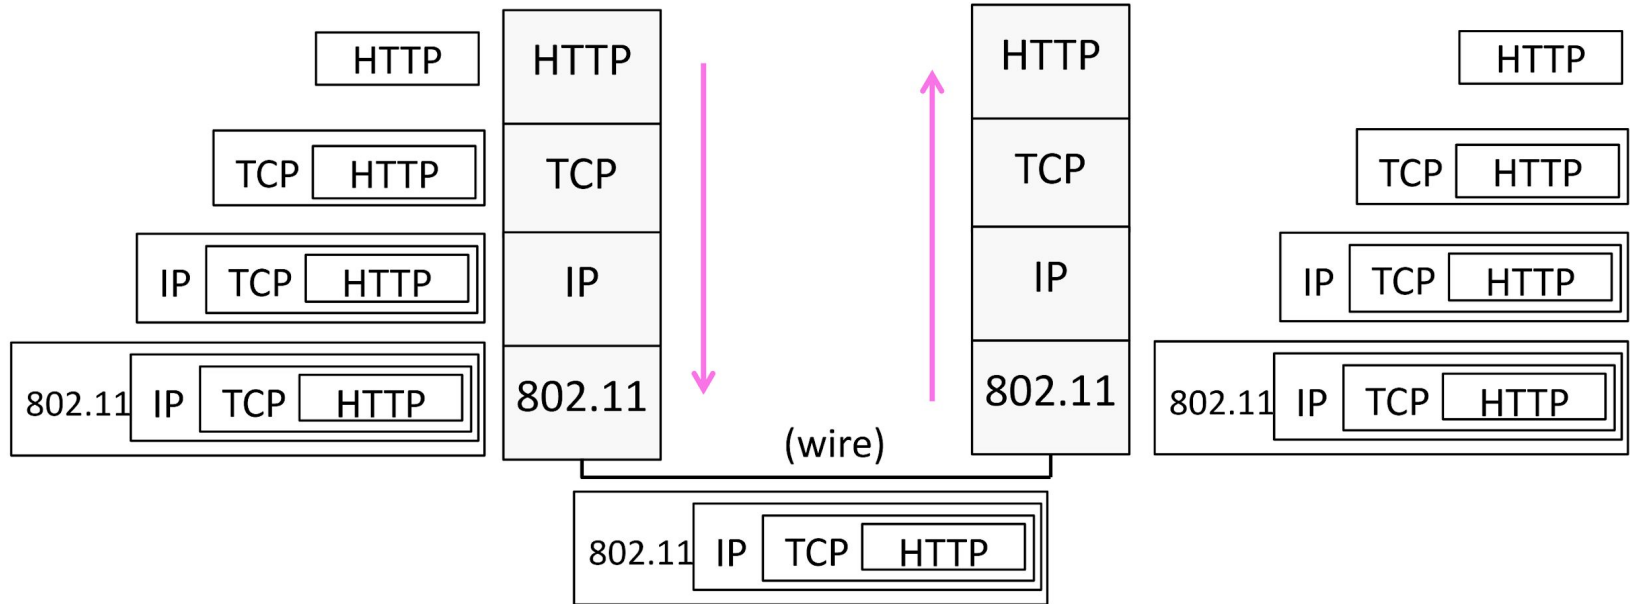|	Purpose	Example Protocols
Application	Programs that use network service	HTTP, DNS
Transport	Provides end-to-end data delivery	TCP, UDP
Network	Sends packets across multiple networks	IP
Link	Sends frames across a link	Ethernet, Wi-Fi
Physical	Transmit bits	--

Protocols and Layers

Advantages

- Use information hiding to connect different systems
- Information reuse to build new protocols

Disadvantages

- Adds overhead
- Hides information

Encapsulation

HTTP | HTTP

TCP HTTP | TCP

IP TCP HTTP | IP

802.11 IP TCP HTTP | 802.11

(wire)

HTTP | HTTP

TCP | TCP HTTP

IP | IP TCP HTTP

802.11 | 802.11 IP TCP HTTP

802.11 IP TCP HTTP

Network Layer

- Network Service Models
- IP Address and Forwarding
- DHCP, ARP, ICMP
- NAT, IPv6
- Routing Algorithms
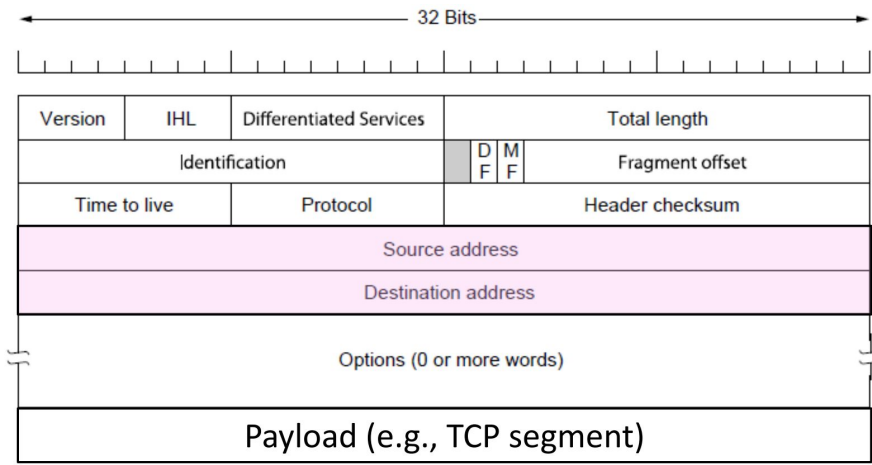- BGP

Motivation

- What does the network layer do?
    - Connect different networks (send packets over multiple networks)

- Why do we need the network layer?
    - Switches don't scale to large networks
    - Switches don't work across more than one link layer technology
    - Switches don't give much traffic control

Network Service Models

Datagrams	Virtual Circuits
Connectionless service	Connection-oriented service
No setup	Connection setup required
Packets contain destination address	Packets contain label for circuit
Routers looks up address in its forwarding table to determine next hop	Router looks up circuit in forwarding table to determine next hop
Example: IP	Example: MPLS

Internetworking - IP

- How do we connect different networks together?

- **IP - Internet Protocol**

- Lowest Common Denominator
  - Asks little of lower-layer networks
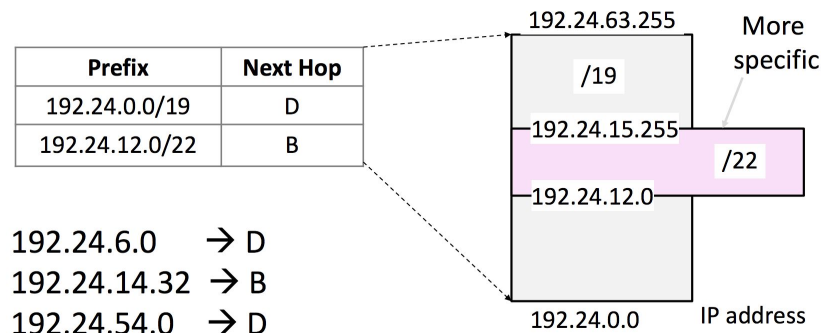 - Gives little as a higher layer service

4. Application	SMTP	HTTP	RTP	DNS
3. Transport		TCP	UDP
2. Internet		IP	
1. Link	Ethernet	3G
	Cable	DSL	802.11

32 Bits

Version	IHL	Differentiated Services		Total length
Identification			D F	M F	Fragment offset
Time to live		Protocol	Header checksum
Source address
Destination address
Options (0 or more words)

Payload (e.g., TCP segment)

IP Addresses Prefix and Forwarding

- IP prefix   a.b.c.d/**L**

  - Represents addresses that have the same first L bits

  - e.g. 128.13.0.0/16 -> all 65536 addresses between 128.13.0.0 to 128.13.255.255

 - e.g. 18.31.0.0/32 -> 18.31.0.0 (only one address)

- **Longest Matching Prefix**

  - Find the longest prefix that contains the destination address, i.e., the most specific entry

Prefix	Next Hop
192.24.0.0/19	D
192.24.12.0/22	B

192.24.6.0   → D
192.24.14.32 → B
192.24.54.0   → D

192.24.63.255

More specific

/19

192.24.15.255

/22

192.24.12.0

192.24.0.0   IP address

DHCP - Dynamic Host Configuration Protocol

- Bootstrapping problem

- Leases IP address to nodes

- UDP

- Also setup other parameters:
  - DNS server
  - IP address of local router
 - Network prefix

Client

Server

DISCOVER

Broadcast

OFFER

REQUEST

ACK
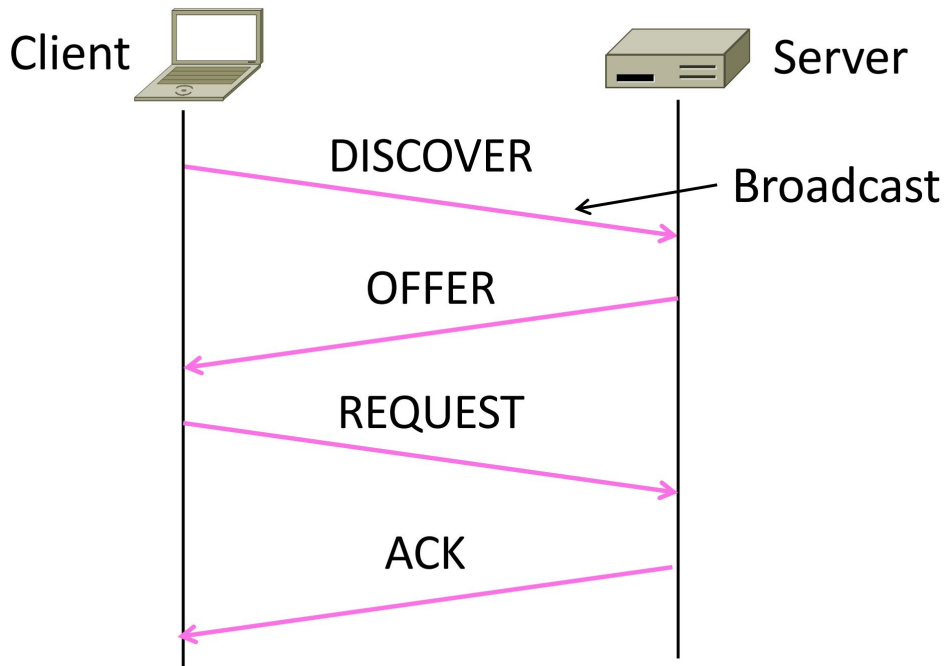
ARP - Address Resolution Protocol

- MAC is needed to send a frame over the local link

- ARP to map an IP to MAC

- Sits on top of link layer

Link layer

Source Ethernet	Dest. Ethernet	Source IP	Dest. IP	Payload ...

From NIC

From ARP

From DHCP

Node

Target

REQUEST
Who has IP 1.2.3.4?

Broadcast

REPLY
I do at 1:2:3:4:5:6

ICMP - Internet Control Message Protocol

- Provides error reporting and testing

- Companion protocol to IP

- Traceroute, Ping

NAT - Network Address Translation

- One solution to **IPv4 address exhaustion**
- Map many private IP to one public IP, with different port number
- Pros: useful functionality (firewall), easy to deploy, etc.
- Cons: Connectivity has been broken!
- Many other cons...

What host thinks	What ISP thinks
Internal IP:port	**External IP : port**
192.168.1.12 : 5523	44.25.80.3 : 1500
192.168.1.13 : 1234	44.25.80.3 : 1501
192.168.2.20 : 1234	44.25.80.3 : 1502

Internal source

Internal IP:port	External IP : port
192.168.1.12 : 5523	44.25.80.3 : 1500

External destination
IP=X, port=Y

Src = 192.168.1.12:5523

NAT

Src = 44.25.80.3:1500
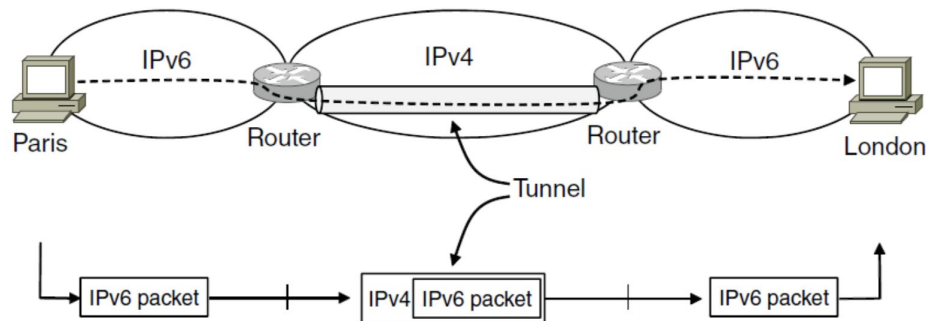
Dst = 123.1.1.10:5000

Dst = 123.1.1.10:5000

IPv6

- A much better solution to IPv4 address exhaustion

- Uses 128-bit addresses, with lots of other changes

- IPv6 version protocols: NDP -> ARP, SLAAC -> DHCP

- Problem: being incompatible with IPv4. Solution: Tunnelling

What's my IP

2601:602:8b00:5f0:30b3:2d19:3fe:db9e

Your public IP address

Routing

- The process of deciding in which direction to send traffic

- Delivery models: unicast, broadcast, multicast, anycast

- Goals: correctness, efficient paths, fair paths, fast convergence, scalability

- Rules: decentralized, distributed setting

Techniques to Scale Routing

Hierarchical Routing

- Route first to the region, then to the IP prefix within the region

IP Prefix Aggregation and Subnets

- Adjusting the size of IP prefixes
  - Internally split one large prefix
  - Externally join multiple IP prefixes

Best Path Routing

Distance Vector Routing

Each node maintains a vector of distances and next hops to all destinations.

Problems: Count-to-infinity scenario when removing links.

Link State Routing (widely used)

Phase 1. **Topology Dissemination**: Each node floods neighboring links. They learn full topology.

Phase 2. **Route Computation**: Each node runs Dijkstra algorithm (or equivalent)

Algorithm details available in lecture slides

BGP - Border Gateway Protocol

- Internet-wide routing between ISPs (ASes)
  - Each has their own policy decisions
- Peer and Transit (Customer) relationship
- Border routers of ISPs announce BGP routes only to other parties who may use those paths.
- Border routers of ISPs select the best path of the ones they hear in any, non-shortest way

BGP example

- Transit (Provider & Customer)
  - Provider announces everything it can reach to its customer
    - AS1 to AS2: you can send packet to AS4 through me
  - Customer only announces its customers to provider
    - AS2 to AS1: you can send packet to A through me
- Peer (ISP 1 & ISP 2)
  - ISP 1 only announces its customer to ISP 2 and vice versa
    - AS2 to AS3: you can send packet to A through me
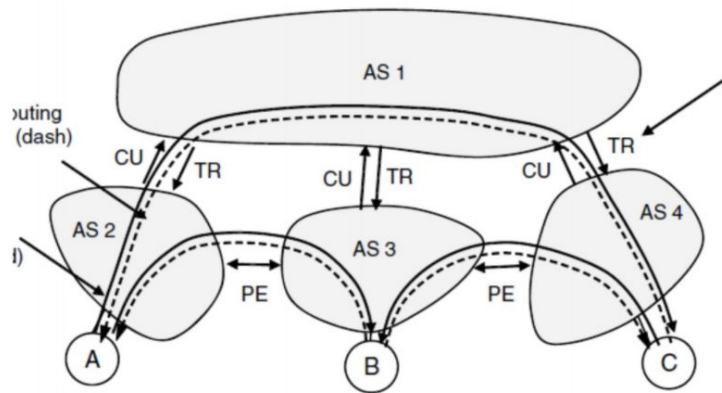
Transport Layer

- Service Models
- TCP vs UDP
- TCP Connections
- Flow Control and Sliding Window
- TCP Congestion Control
- Newer TCP Implementations
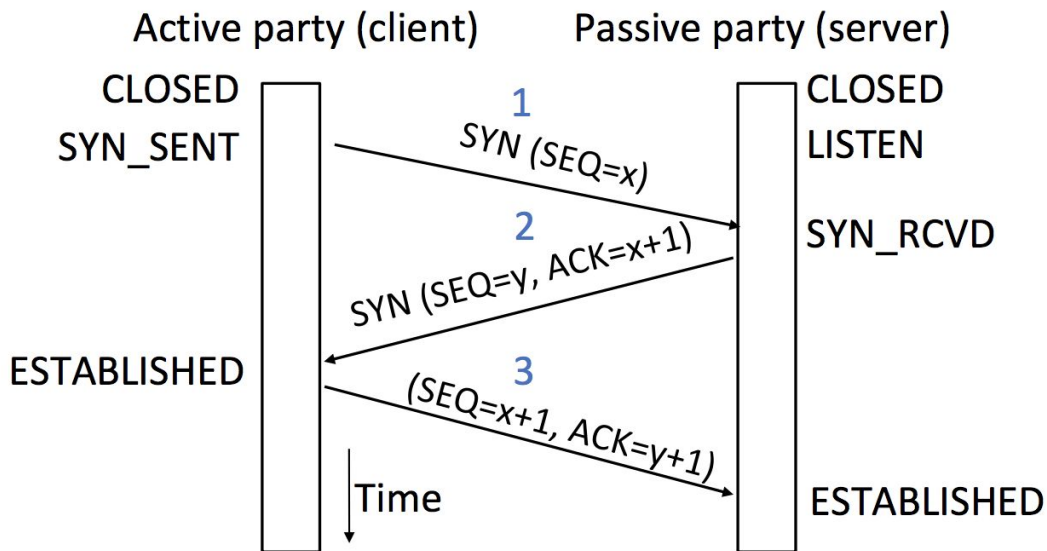
Service Models

- Transport Layer Services

  - Datagrams (UDP): Unreliable Messages

  - Streams (TCP): Reliable Bytestreams

- Socket API: simple abstraction to use the network

  - Port: Identify different applications / application layer protocols on a host
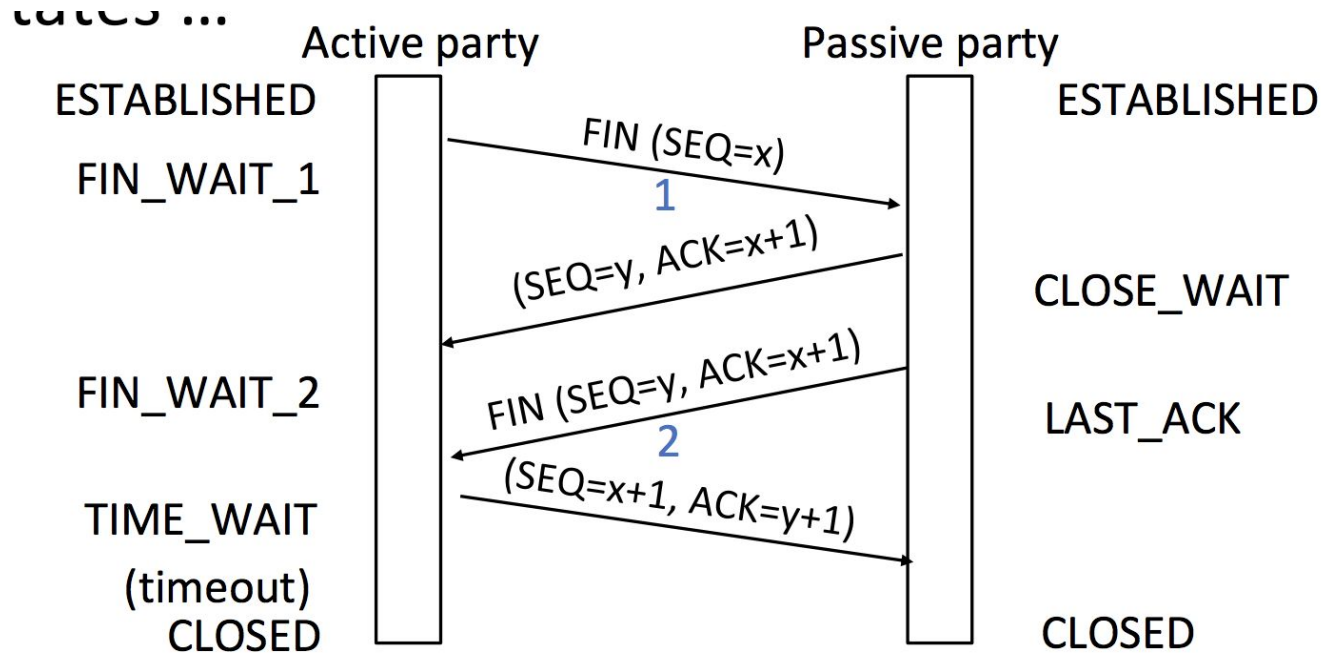
TCP vs UDP

| TCP (Streams)	UDP (Datagrams)
Connections	Datagrams
Bytes are delivered once, reliably, and in order	Messages may be lost, reordered, duplicated
Arbitrary length content	Limited message size
Flow control matches sender to receiver	Can send regardless of receiver state
Congestion control matches sender to network	Can send regardless of network state

TCP Connection Establishment

Three-way handshake

CSE 461 University of Washington

TCP Connection Release

tates ...

| Active party		Passive party
ESTABLISHED		ESTABLISHED
FIN_WAIT_1	FIN (SEQ=x) **1** →	
	← (SEQ=y, ACK=x+1)	CLOSE_WAIT
FIN_WAIT_2	← FIN (SEQ=y, ACK=x+1) **2**	LAST_ACK
TIME_WAIT	(SEQ=x+1, ACK=y+1) →	
(timeout) CLOSED		CLOSED

Flow Control - Sliding Window Protocol

- Instead of stop-and-wait, sends W packets per 1 RTT
    - To fill network path, W = B * RTT / packet_size
- Receiver sends ACK upon receiving packets
    - Go-Back-N (similar to project 1 stage b): not efficient
    - **Selective Repeat**
        - Receiver passes data to app in order, and buffers out-of-order segments to reduce retransmissions
        - ACK conveys highest in-order segment
            - As well as hints about out-of-order segments
- **Selective Retransmission** on sender's side

Flow Control - ACK Clock
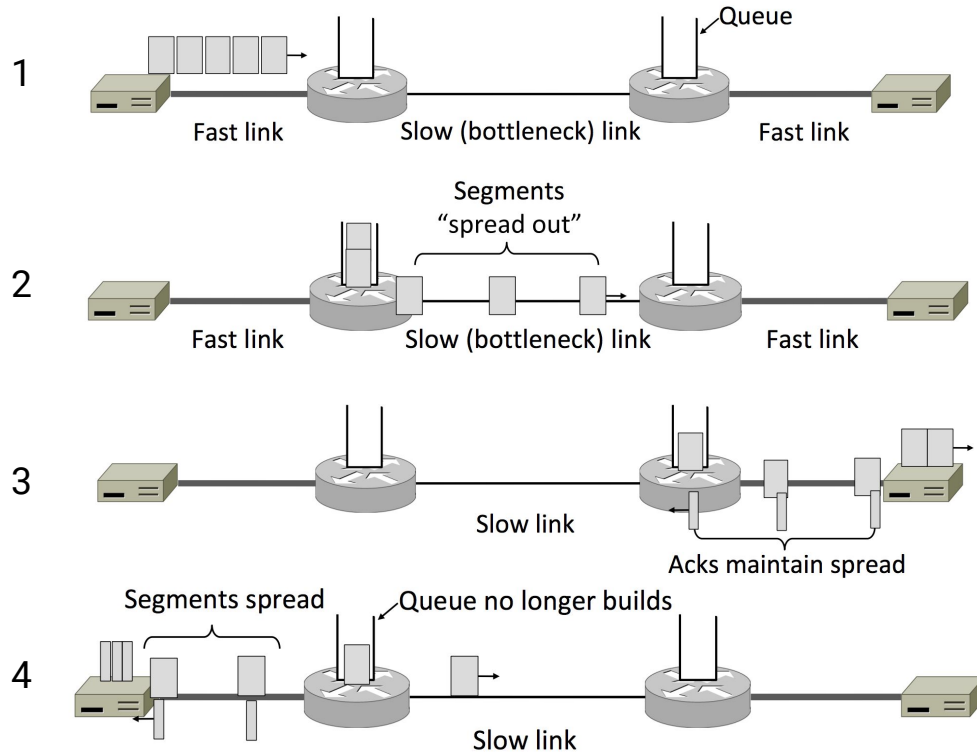
Flow Control - Sliding Window Protocol (2)

- Flow control on receiver's side

  - In order to avoid loss caused by user application not calling recv(), receiver tells sender its available buffer space (WIN)

  - Sender uses lower of the WIN and W as the effective window size

- How to set a **timeout** for retransmission on sender's side?

 - Adaptively determine timeout value based on smoothed estimate of RTT

$$SRTT_{N+1} = 0.9*SRTT_N + 0.1*RTT_{N+1}$$
$$Svar_{N+1} = 0.9*Svar_N + 0.1*|RTT_{N+1} - SRTT_{N+1}|$$
$$TCP\ Timeout_N = SRTT_N + 4*Svar_N$$

Max-Min Fair Allocation

1. Start with all flows at rate 0

2. Increase the flows until there is a new bottleneck in the network

3. Hold fixed the rate of the flows that are bottlenecked

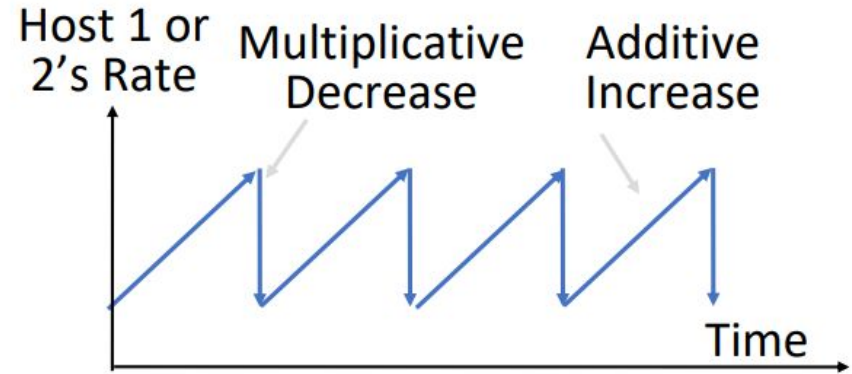4. Go to step 2 for any remaining flows

TCP Bandwidth Allocation

- Closed loop: use feedback to adjust rates
  - NOT open loop: reserve bandwidth before use
- Host driven: host sets/enforces allocations
  - NOT network driven
- Window based
  - NOT rate based
- Congestion signal
 - Packet loss, Packet delay, Router indication

AIMD!

AIMD - Additive Increase Multiplicative Decrease

AIMD

- **Slow-Start** (used in AI)
  - Double cwnd until packet timeout
  - Restart and double until cwnd/2, then AI
- **Fast-Retransmit** (used in MD)
  - Three duplicate ACKs = packet loss
  - Don't have to wait for TIMEOUT
- **Fast-Recovery** (used in MD)
  - MD after fast-retransmit
  - Then pretend further duplicate ACKs are the expected ACKs

TCP Reno

Network-Side Congestion Control

- Explicit Congestion Notification (**ECN**)
  - Router detects the onset of congestion via its queue. When congested, it marks affected packets in their IP headers
  - Marked packets arrive at receiver; treated as loss. TCP receiver reliably informs TCP sender of the congestion