


1 Introduction



Internet: Billions of connected computing devices,
Protocol: Protocols define format, order of messages sent and received among network entities, and actions taken on message transmission, receipt, (control sending, receiving of messages e.g. TCP, IP, HTTP, Skype, 802.11)
Network edge: have 1) hosts: clients and servers. 2) access networks, physical media: wired, wireless communication links, (residential access nets, institutional access networks, mobile access networks)
Digital Subscriber Line: voice, data transmitted at different frequencies over **dedicated** line to central office. Use existing telephone line to central office DSLAM, data over DSL phone line goes to Internet, voice over DSL phone line goes to telephone net, < 2.5 Mbps upstream transmission rate (typically < 1 Mbps), < 24 Mbps downstream transmission rate (typically < 10 Mbps)
Cable Network: Uses HFC: hybrid fiber coax, asymmetric, up to 30Mbps downstream transmission rate, 2Mbps upstream transmission rate, Network of cable, fiber attaches homes to ISP router, homes share access network to cable headend, unlike DSL which has dedicated access to central office different channels transmitted in different frequency bands (FDM)
Enterprise Access Networks (Ethernet)




- Typically used in companies, universities, etc.
- 10 Mbps, 100Mbps, 1Gbps, 10Gbps transmission rates
- Today, end systems typically connect into Ethernet switch

Wireless Access Networks: Shared wireless access network connects end system to router via base station aka "access point"
Wireless LANs: within building (100 ft.), 802.11b/g/n (WiFi): 11.54, 450 Mbps transmission rate

Wireless access: provided by telco (cellular) operator, 10's km, between 1 and 10 Mbps, 3G, 4G: LTE
Host: Host sending function: takes application message, breaks into smaller chunks known as packets, of length L bits, transmits packet into access network at transmission rate R, link transmission rate, aka link capacity, aka link bandwidth
Physical Media: bit propagation between transmitter/receiver pairs, physical link: what lies between transmitter & receiver, guided media: signals propagate in solid media: copper, fiber, coax, Fiber optic cable, unguided media: signals propagate freely (Terrestrial microwave, LAN, Wide-area, Satellite), twisted pair (TP) two insulated copper wires
Network core: interconnected routers, network of networks
Packet-switching: takes L/R seconds to transmit (push out) L-bit packet into link at R bps, store and forward: entire packet must arrive at router before it can be transmitted on next link, end-end delay = 2L/R (assuming zero propagation delay)
Queueing and Loss: if arrival rate (n bits) to link exceeds transmission rate of link for a period of time: packets will queue, wait to be transmitted on link, packets can be dropped (lost) if memory (buffer) fills up 2nd width (bps), L: packet length (bits), a: average packet arrival rate, La/R < 0: avg. queueing delay small, La/R > 1: avg. queueing delay large, La/R > 1: more "work" arriving than can be serviced, average delay infinite!
Circuit-switching: end-end resources allocated to, reserved for "call" between source & destination: dedicated resources: no sharing, circuit segment idle if not used by call (no sharing), commonly used in traditional telephone networks



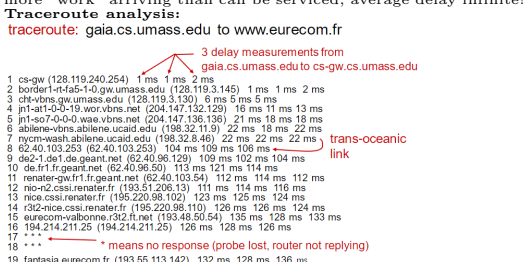
Delay, Loss, Throughput in networks:



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

1) cars "propagate" at 100 km/hr 2) toll booth takes 12 sec to service car (bit transmission time) 3) car bit; caravan + time to "push" entire caravan through toll booth onto highway: 12x10 = 120 sec, time for last car to propagate from 1st to 2nd toll both: 100km/(100km/hr) = 1 hr, i.e. 60 + 2 = 62 min


Queueing Delay: R: link bandwidth (bps), L: packet length (bits), a: average packet arrival rate, La/R < 0: avg. queueing delay small, La/R > 1: avg. queueing delay large, La/R > 1: more "work" arriving than can be serviced, average delay infinite!
Traceroute analysis:
 traceroute: gaia.cs.umass.edu to www.eurocom.fr



Packet Loss: queue (aka buffer) preceding link in buffer has finite capacity, packet arriving to full queue dropped (aka lost), lost packet may be retransmitted by previous node, by source end system, or not at all
Throughput: rate (bits/time unit) at which bits transferred between sender/receiver, **instantaneous:** rate at given point in time, **average:** rate over longer period of time
 per-connection end-end throughput:

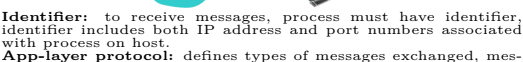
$$\min(R_s, R_r, R_{link})$$
 in practice: R_c or R_s is often bottleneck

Internet protocol stack: application: supporting network applications (FTP, SMTP, HTTP) transport: process-process data delivery (TCP, UDP), network: routing of packets from source to destination (IP, routing protocols), link: data transfer between neighboring network elements (Ethernet, 802.11 (WiFi), PPP), physical: bits "on the wire"



2 Application Layer

Client-server: server: always on host, permanent IP address, data centers for scaling; clients: communicate with server may be intermittently connected, may have dynamic IP addresses, do not communicate directly with each other.
P2P architecture: no always-on server, arbitrary end systems directly communicate, peers request service from other peers, provide service in return to other peers, self scalability - new peers bring new service capacity, as well as new service demands. Peers are intermittently connected and change IP addresses -> Complex Management.
Processes communicating: processes in different hosts communicate by exchanging messages, same process can be both a client (initiates communication and asks server waits to be contacted) for different connections; e.g., in P2P networks
Sockets: process sends/receives messages to/from its socket



application	data loss	throughput	time sensitive
file transfer	no loss	elastic	no
e-mail	no loss	elastic	no
Web documents	no loss	elastic	no
real-time audio/video	loss-tolerant	audio: 5kbps-1Mbps video: 10kbps-5Mbps	yes, 100's msec yes, few secs
stored audio/video	loss-tolerant	same as above	yes, 100's msec
interactive games	loss-tolerant	few kbps up	yes and no
text messaging	no loss	elastic	no

Socket programming: Two socket types for two transport services: UDP: unreliable datagram (User Datagram Protocol) TCP: reliable, byte stream-oriented (Transmission Control Protocol)

Socket programming with UDP
 UDP: no "connection" between client & server
 • no handshaking before sending data
 • sender explicitly attaches IP destination address and port # to each packet
 • receiver extracts sender IP address and port# from received packet
 UDP: transmitted data may be lost or received out-of-order

Socket programming with TCP
 client must contact server
 • server process must first be running
 • server must have created socket (door) that welcomes client's contact
 • allows server to talk with multiple clients
 • source port numbers used to distinguish clients (more in Chap 3)
 client contacts server by:
 • Creating TCP socket, specifying IP address, port number of server process
 • when client creates socket: client TCP establishes connection to server TCP
 application viewpoint:
 TCP provides reliable, in-order byte-stream transfer ("pipe") between client and server

HTTP: hypertext transfer protocol: client/server model client: browser that requests, receives, (using HTTP protocol) and "displays" Web objects, server: Web server sends (using HTTP protocol) objects in response to requests.
HTTP Steps: client initiates TCP connection (creates socket) to server, port 80, server accepts TCP connection from client. HTTP messages (application-layer protocol messages) exchanged between browser (HTTP client) and Web server (HTTP server). TCP connection closed.
HTTP is "stateless": server maintains no information about past client requests.


Non-persistent HTTP
 suppose user enters URL: www.someschool.edu/home/department/home.index (contains text, references to 10 pag images)
 1. HTTP client initiates TCP connection to HTTP server (process) at www.someschool.edu on port 80
 2. HTTP client sends HTTP request message (containing URL) into TCP connection socket. Message indicates that client wants object someDepartment/home.index
 3. HTTP server receives request message, forms response message containing requested object, and sends message into its socket
 4. HTTP server closes TCP connection. REPEAT TO TIMES FOR 10 IMG
 5. HTTP client receives response message containing HTML file, displays HTML file, finds 10 referenced jpeg objects

HTTP 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 = non-persistent HTTP response time = 2RTT + file transmission time
Persistent HTTP
 • server leaves connection open after sending response
 • subsequent HTTP messages between same client/server sent over open connection
 • client sends requests as soon as it encounters a referenced object
 • as little as one RTT for all the referenced objects
Pipelining
 • Send several requests at once
 • Push resources
QUIC
 • Eliminate first RTT


HTTP request message: two types of HTTP messages: request, response




Method types: HTTP 1.0: GET, POST, HEAD (asks server to leave requested object out of response) HTTP 1.1: GET, POST, HEAD, PUT, (uploads file in entity body to path specified in URL field), DELETE (deletes file specified in the URL field)




Cookies: four components: 1) C cookie header line of HTTP response message 2) Cookie header line in next HTTP request message 3) Cookie file kept on user's host, managed by user's browser 4) Back-end database at Website




Web caches (proxy 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




Electronic mail Three major components: 1) user agents 2) mail servers 3) SMTP: Simple Mail Transfer Protocol
User Agent: a.k.a. "mail reader", composing, editing, reading mail messages, e.g., Outlook, Thunderbird, iPhone mail client. outgoing, incoming messages stored on server Mail Servers: 1) mailbox contains incoming messages for user 2) message queue of outgoing (to be sent) mail messages (SMTP protocol between mail servers to send email messages, client: sending mail server, "server": receiving mail server).
SMTP Example: 1) Alice uses UA to compose message "to" bob@homeschool.edu 2) Alice's UA sends message to her mail server; message placed in message queue 3) client side of SMTP opens TCP connection with Bob's mail server 4) SMTP client sends Alice's message over the TCP connection 5) Bob's mail server places the message in Bob's mailbox 6) Bob invokes his user agent to read message



SMTP requires message header & body to be in 7-bit ASCII, SMTP server uses CRLF.CRLF to determine end of message
comparison with HTTP: HTTP: pull, SMTP: push. both have ASCII command/response interaction, status codes. HTTP: each object encapsulated in its own response message, SMTP: multiple objects sent in multipart message
 SMTP: protocol for exchanging email messages



different from SMTP MAIL FROM, RCPT TO: commands!
 • Body: the "message"
 • ASCII characters only
Mail access protocols retrieval from server: POP: Post Office Protocol [RFC 1939]: authorization, download. IMAP: Internet Mail Access Protocol [RFC 1730]: more features, including manipulation of stored messages on server. HTTP: gmail, Hotmail, Yahoo! Mail, etc.
DNS: domain name system: DNS services: hostname to IP address translation, host aliasing, canonical, alias names, mail server aliasing, load distribution, replicated Web servers: many IP addresses correspond to one name client wants IP for www.amazon.com; 1st approximation: 1) client queries root server to find com DNS server 2) client queries com DNS server to get amazon.com DNS server 3) client queries amazon.com DNS server to get IP address for www.amazon.com




Root Name Servers: Contacts authoritative name server if name mapping not known Gets mapping, Returns mapping to local name server.


Top-level domain (TLD) servers: responsible for com, org, net, edu, aero, jobs, museums, and all top-level country domains, e.g.: uk, fr, ca, jp, Network Solutions maintains servers for .com TLD, Educause for .edu TLD

Authoritative DNS servers: Organization's own DNS server(s), providing authoritative hostname to IP mappings for organization's named hosts, can be maintained by organization or service provider

- type=A**
 - name is hostname
 - value is IP address
- type=CNAME**
 - name is alias name for some "canonical" (the real name)
 - www.ibm.com is really serverest.backup2.ibm.com
 - value is canonical name
- type=NS**
 - name is domain (e.g., foo.com)
 - value is hostname of authoritative name server for this domain
- type=MX**
 - value is name of mail server associated with name



BitTorrent: File divided into 256Kb chunks. Peers in torrent send/receive file chunks, peer joining torrent has no chunks but will accumulate them over time from other peers, registers with tracker to get list of peers, connects to subset of peers ("neighbors") while downloading, peer uploads chunks to other peers, peer may change peers with whom it exchanges chunks, churn: peers may come and go, once peer has entire file, it may (selfishly) leave or (altruistically) remain in torrent.

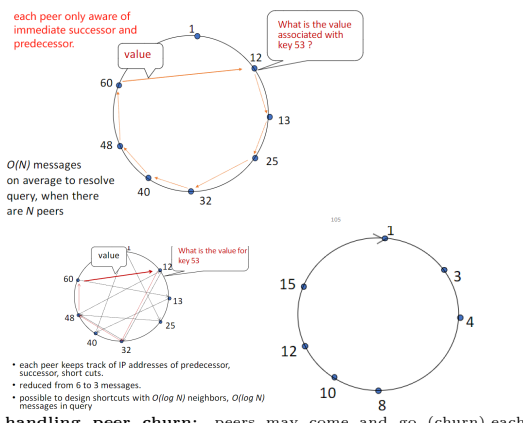


Distributed Hash Table (DHT): each peer only aware of immediate successor and predecessor.

What is the value associated with key 53?

What is the value associated with key 53?

What is the value associated with key 53?




Content Distribution Networks (CDNs): stores copies of content at CDN nodes, subscriber requests content from CDN (directed to nearby copy, retrieves content may choose different copy if network path congested)

3 Transport Layer

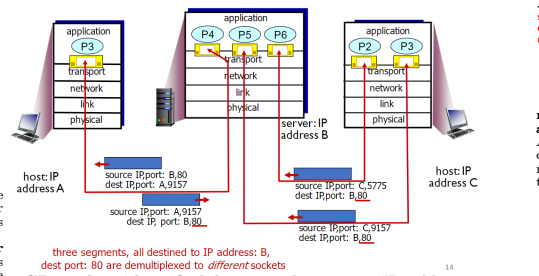
Transport services and protocols provide logical communication between app processes running on different hosts. **send side:** breaks app messages into segments, passes to network layer, **rcv side:** reassembles segments into messages, passes to app layer, more than one transport protocol available to apps (TCP and UDP)

Multiplexing/Demultiplexing




Demultiplexing: host receives IP datagrams, each datagram has source IP address, destination IP address, each datagram carries one transport-layer segment, each segment has **source, destination port number, host uses IP addresses & port numbers to direct segment to appropriate socket**

Connection-oriented demux:

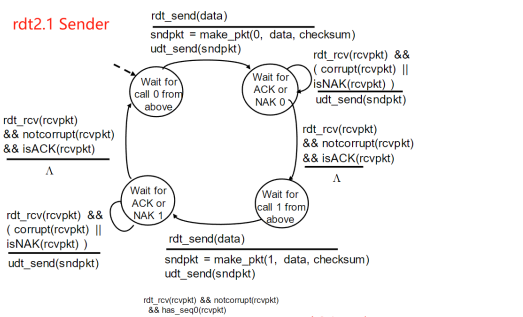
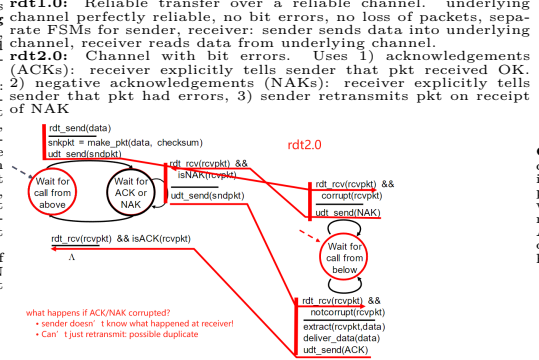



TCP socket identified by 4-tuple: source IP address, source port number, dest IP address, dest port number TCP socket identified by 4-tuple, source IP address, source port number, dest IP address, dest port number **Connectionless UDP:** "no frills," "bare bones" Internet transport protocol, "best effort" service, UDP segments may be: lost, delivered out-of-order to app, no handshaking between UDP sender, receiver, each UDP segment handled independently of others

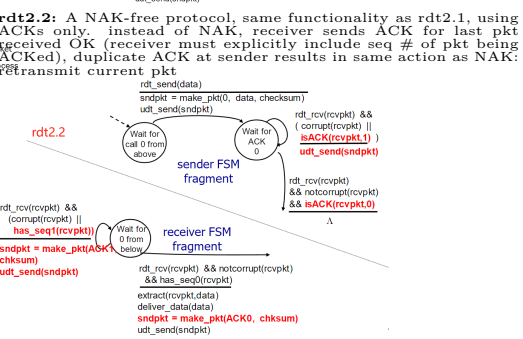


UDP checksum: detect "errors" in transmitted segment. **sender:** treat segment contents, including header fields, as sequence of 16-bit integers, checksum: addition (one's complement sum) of segment contents, sender puts checksum value into UDP checksum field. **receiver:** compute checksum of received segment, check if computed checksum equals checksum field value.

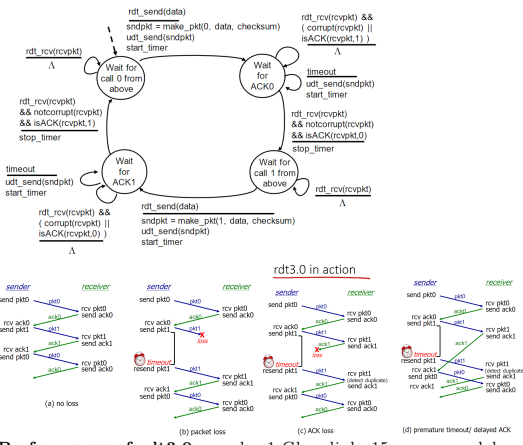
Reliable Data Transfer:



rdt2.2: A NAK-free protocol, same functionality as rdt2.1, using ACKs only, instead of NAK, receiver sends ACK for last pkt received OK (receiver must explicitly include seq # of pkt being ACKed), duplicate ACK at sender results in same action as NAK: retransmit current pkt



rdt3.0: Channels with errors and loss of packets (data, ACKs). **approach:** 1) sender waits "reasonable" amount of time for ACK, 2) retransmits if no ACK received in this time, 3) requires countdown timer, 4) if pkt (or ACK) just delayed (not lost): retransmission will be duplicate, but seq. # have already handles this, receiver must specify seq # of pkt being ACKed

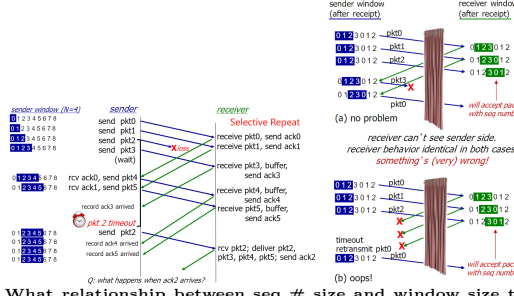


Performance of rdt3.0: under 1 Gbps link, 15 ms prop. delay, 8000 bit packet, pipeline to improve performance




GBN: Sender: k-bit seq # in pkt header, "window" of up to N, consecutive unack'd pkts allowed, ACK(n): ACKs all pkts up to, including seq # n "cumulative ACK" timer for oldest in-flight pkt, timeout(n): retransmit packet n and all higher seq # pkts in window. **Receiver:** ACK-only: always send ACK for correctly-received pkt with highest in-order seq #, may generate duplicate ACKs, need only remember expectedseqnum, out-of-order pkt; discard (don't buffer); no receiver buffering! re-ACK pkt with highest in-order seq #

rdt2.1: Sender, handles garbled ACK/NAKs, Handles bit corruptions that are detected by checksum, Uses a 1-bit sequence number to detect retransmission at receiver


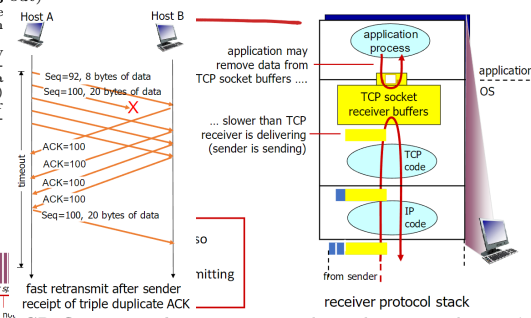
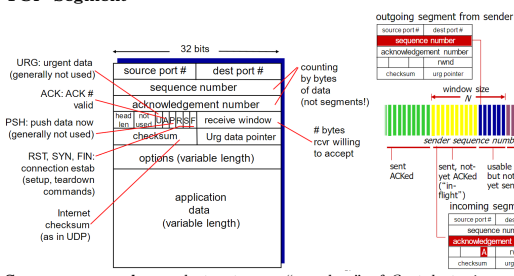


event at receiver	TCP receiver action
arrival of in-order segment with expected seq #. All data up to expected seq # already ACKed	delayed ACK. Wait up to 500ms for next segment. If no next segment, send ACK
arrival of in-order segment with expected seq #. One other segment has ACK pending	immediately send single cumulative ACK, ACKing both in-order segments
arrival of out-of-order segment higher-than-expected seq. #. Gap detected	immediately send duplicate ACK , indicating seq. # of next expected byte
arrival of segment that partially or completely fills gap	immediate send ACK, provided that segment starts at lower end of gap



What relationship between seq # size and window size to avoid problem in (b)? In technical terms, if N is the size of the sequence number space and W is the window size, the condition to avoid such problems is: $N > 2W$

TCP Overview: 1) full duplex data (bi-directional data flow in same connection, MSS: maximum segment size) 2) connection-oriented (handshaking initiates sender, receiver state before data exchange) 3) flow controlled (sender will not overwhelm receiver) 4) point-to-point (one sender, one receiver) 5) reliable, in-order byte stream (no "message boundaries"), 6) pipelined (TCP congestion and flow control set window size)



TCP Congestion Control: CWND sender limits transmission to *cwnd*. *cwnd* is dynamic, function of perceived network congestion. Thus: TCP sending rate roughly: $\text{send cwnd bytes, wait RTT for ACKs, then send more bytes rate} = \frac{\text{cwnd}}{\text{RTT}}$ bytes/sec

TCP Slow Start: 1) initially *cwnd* = 1 MSS, 2) double *cwnd* every RTT, (done by incrementing *cwnd* for every ACK received up to some threshold) 3) AIMD: Additive Increase Multiplicative Decrease.


Sequence numbers: byte stream "number" of first byte in segment's data acknowledgment: (seq # of next byte expected from other side, cumulative ACK).

TCP timeout interval: EstimatedRTT plus "safety margin"

$\text{DevRTT} = (1 - \beta) \times \text{DevRTT} + \beta \times |\text{SampleRTT} - \text{EstimatedRTT}|$

(typically, $\beta = 0.25$)

TCP flow control: receiver controls sender, so sender won't overflow receiver's buffer by transmitting too much, too fast. receiver "advertises" free buffer space by including *rwnd* value in TCP header of receiver-to-sender segments, (**RecvBuffer** size set via socket options, many operating systems autoadjust *cvBuffer*) 2) sender limits amount of unacked ("in-flight") data to receiver's *rwnd* value.

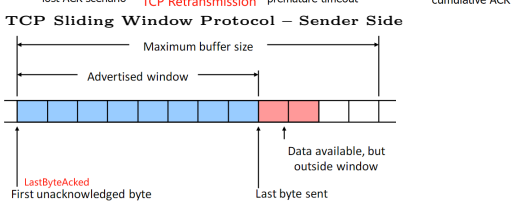
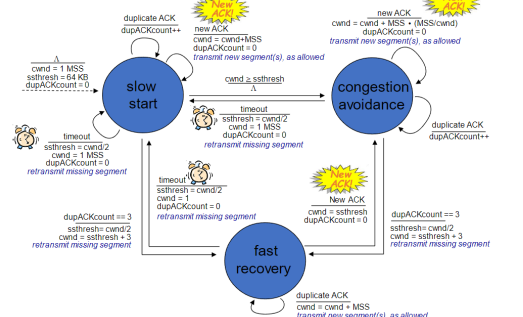

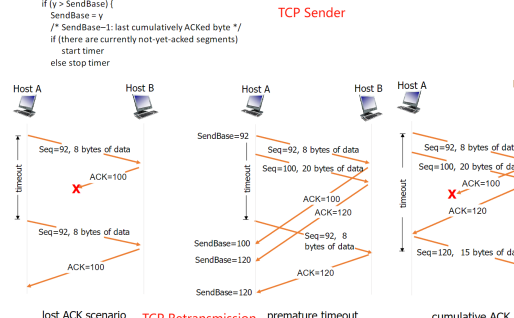
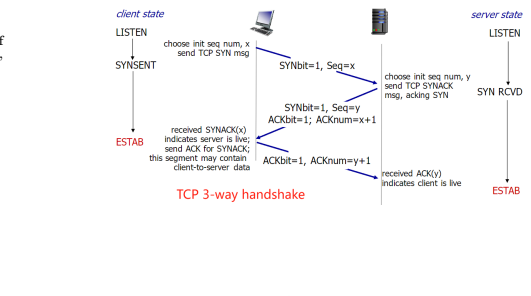




cwnd when detecting, reacting to loss: Loss indicated by timeout: *cwnd* set to 1 MSS; *rwnd* window the grows exponentially (as in slow start) to threshold, then grows linearly (as in additive increase), Loss indicated by 3 duplicate ACKs; **TCP Tahoe:** always sets *cwnd* to 1 (timeout or 3 duplicate acks), **TCP Reno:** (Fast Recovery), dup ACKs indicate network capable of delivering some segments, *cwnd* is cut in half window then grows linearly (as in additive increase)


TimeoutInterval = EstimatedRTT + 4 * DevRTT

estimated RTT "safety margin"

TCP reliable data transfer TCP creates *rdt* service on top of IP's unreliable service using pipelined segments, cumulative acks, single retransmission timer



Congestion control: "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 (queuing in router buffers)





TCP throughput: avg. TCP throughput as function of window size, RTT? ignore slow start, assume always data to send, W: window size (measured in bytes) where loss occurs, avg. window size (# in-flight bytes) is $3/4 W$, avg. TCP throughput = $\frac{3}{4W} \text{RTT}$

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

Explicit Congestion Notification: network-assisted congestion control: two bits in IP header (ToS field) marked by network router to indicate congestion, congestion indication carried to receiving host, receiver (seeing congestion indication in IP datagram) sets ECE bit on receiver-to-sender ACK segment to notify sender of congestion

TCP ACK generation:

if $y > \text{SendBase}$ {
 SendBase = y
 /* SendBase=1: last cumulatively ACKed byte */
 if there are currently not-yet-acked segments)
 start timer
 else stop timer




4 Network Layer

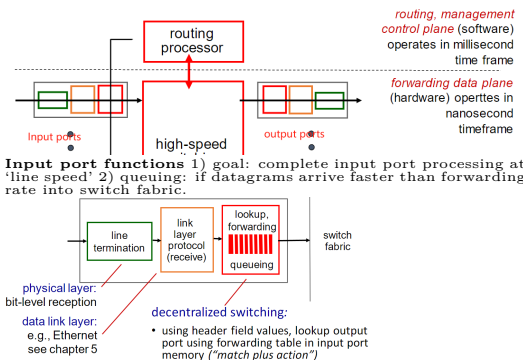
Network-layer functions: 1) forwarding: move packets from router's input to appropriate router output. 2) routing: determine route taken by packets from source to destination, routing algorithms

Data plane: 1) local, per-router function, 2) determines how datagram arriving on router input port is forwarded to router output port 3) forwarding function

Control plane: 1) network-wide logic 2) determines how datagram is routed among routers along end-end path from source host to destination host 3) two control-plane approaches: (traditional routing algorithms: implemented in routers, software-defined networking (SDN): implemented in (remote) servers)



Network service model:
example services for individual datagrams: guaranteed delivery, guaranteed delivery with less than 40 msec delay, **example services for a flow of datagrams:** in-order datagram delivery, guaranteed minimum bandwidth to flow, restrictions on changes in inter-packet spacing
Router architecture overview




Input port functions 1) goal: complete input port processing at 'line speed' 2) queuing: if datagrams arrive faster than forwarding rate into switch fabric.

Destination Address Range	Link interface
11001000 00010111 00010***	0
11001000 00010111 00011000	1
11001000 00010111 00011***	2
otherwise	3


Longest prefix matching: when looking for forwarding table entry for given destination address, use longest address prefix that matches destination address.

examples:
 DA: 11001000 00010111 000101010001 which interface? 0
 DA: 11001000 00010111 00011000 10101010 which interface? 1

Switching fabrics: transfer packet from input buffer to appropriate output buffer, switching rate: rate at which packets can be transferred from inputs to outputs, (often measured as multiple of input/output line rate, N inputs: switching rate N times line rate desirable)

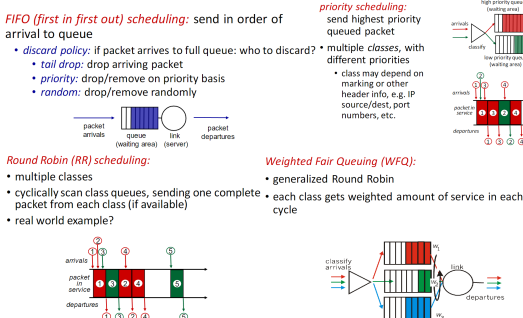


Switching via a bus, datagram from input port memory to output port memory via a shared bus, bus contention: switching speed limited by bus bandwidth, 30 Gbps.
Switching via interconnection network, overcome bus bandwidth limitations, banyan networks, crossbar, other interconnection nets initially developed to connect processors in multiprocessor, advanced design: fragmenting datagram into fixed length cells, switch cells through the fabric, 60 Gbps.
Input Port Queuing: Caused by: 1) Slow switching fabric, 2) Output port contention

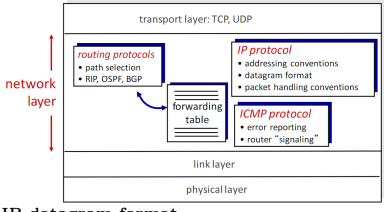


Reducing Input Queuing: Why? Reduce HOL blocking, Avoid packet drops at input queues, Save on queue memory, **How?** Increase switch fabric speed, Increase inbound capacity of output ports

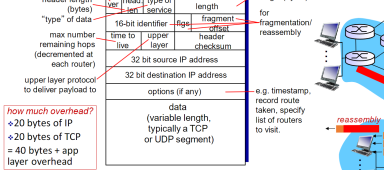
Output ports buffer: required when datagrams arrive from fabric faster than the transmission rate (How much buffering? RFC 3439 rule of thumb: average buffering equal to "typical" RTT (say 250 msec) times link capacity C, e.g., C = 10 Gbps link: 2.5 Gbit buffer, recent recommendation [Appenzeller'04]: with N flows, buffering equal to: $\frac{RTT \times C}{\sqrt{N}}$)
Output port scheduling:




Internet network layer



IP datagram format





IP fragmentation, reassembly: MTU (max transfer size). So large IP datagram divided ("fragmented") within net.




IP address: 32-bit identifier for host, router interface, interface: connection between host/router and physical link, Router's typically have multiple interfaces, host typically has one or two interfaces (e.g., wired Ethernet, wireless 802.11), IP addresses associated with each interface


Subnets: IP address: subnet part - high order bits, host part - low order bits, What's a subnet?, device interfaces with same subnet part of IP address, can physically reach each other without intervening router



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




DHCP: Dynamic Host Configuration Protocol dynamically get ip address from a server
DHCP overview: 1) host broadcasts "DHCP discover" msg [optional], 2) DHCP server responds with "DHCP offer" msg [optional], 3) host requests IP address: "DHCP request" msg, 4) DHCP server sends address: "DHCP ack"

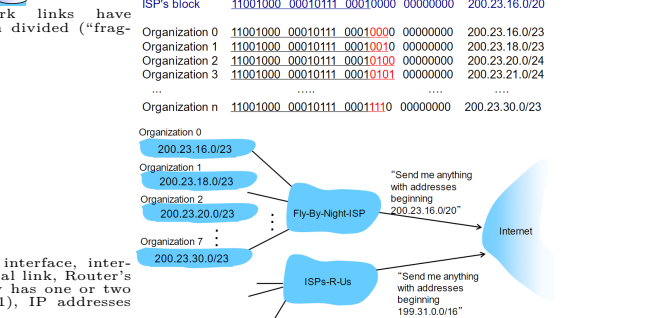


DHCP: example 1) connecting laptop needs its IP address, addr of first-hop router, addr of DNS server: use DHCP router with DHCP server built into router. 2) DHCP request encapsulated in UDP, encapsulated in IP, encapsulated in 802.11 Ethernet. 3) Ethernet frame broadcast (dest: FFFFFFFF) on LAN, received at router running DHCP server. 4) Ethernet demuxed to


IP demuxed, UDP demuxed to DHCP. 5) DHCP server formulates DHCP ACK containing client's IP address, IP address of first-hop router for client, name & IP address of DNS server 6) encapsulation of DHCP server, frame forwarded to client, demuxing up to DHCP at client. 7) client now knows its IP address, name and IP address of DNS server, IP address of its first-hop router




Hierarchical addressing: allows efficient advertisement of routing information. (network get subnet part of IP addr from allocated portion of its provider ISP's address space)



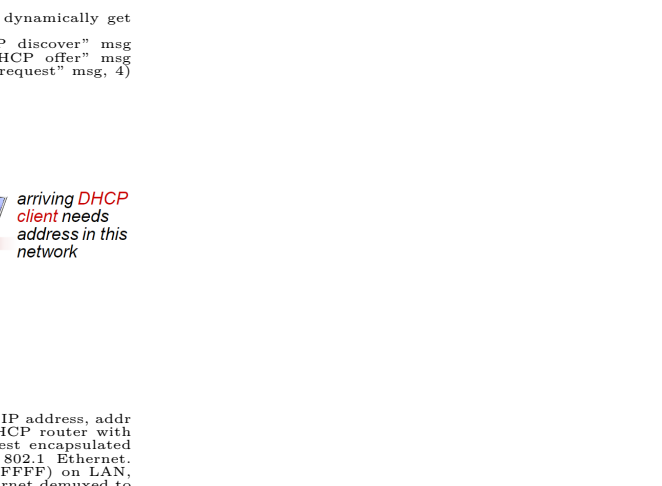
NAT: network address translation local network uses just one IP address as far as outside world is concerned. **Advantages:** 1) range of addresses not needed from ISP: just one IP address for all devices 2) can change addresses of devices in local network without notifying outside world 3) can change ISP without changing addresses of devices in local network 4) devices inside local net not explicitly addressable, visible by outside world (a security plus)



NAT router must: 1) outgoing datagrams: replace (source IP address, port #) of every outgoing datagram to (NAT IP address, new port #) 2) remember (in NAT translation table) every (source IP address, port #) to (NAT IP address, new port #) translation pair 3) incoming datagrams: replace (NAT IP address, new port #) in dest fields of every incoming datagram with corresponding (source IP address, port #) stored in NAT table




2: NAT router changes datagram source addr from 10.0.0.1, 3345 to 138.76.29.7, 5001, updates table
 3: reply arrives dest. address: 138.76.29.7, 5001
 4: NAT router changes datagram dest addr from 138.76.29.7, 5001 to 10.0.0.1, 3345




IPv6: initial motivation: 32-bit address space soon to be completely allocated. IPv6 datagram format: fixed-length 40 byte header, no fragmentation allowed

IPv6 datagram format

priority: identify priority among datagrams in flow
flow label: identify datagrams in same "flow."
 (concept of "flow" not well defined).
next header: identify upper layer protocol for data




Routing Protocols Classification: global ("link state") algorithms or decentralized ("distance vector") information? static (routes change slowly over time) or dynamic (routes change more quickly, periodic update, in response to link cost changes)?



Distance vector algorithm: $D_x(y)$ = estimate of least cost from x to y . x maintains distance vector $D_x = [D_x(y), y \in N]$. node x , knows cost to each neighbor v : $c(x,v)$, maintains its neighbors' distance vectors. For each neighbor v , x maintains $D_v = [D_v(y), y \in N]$

clearly, $d_1(x) = 5, d_1(z) = 3, d_1(x) = 3$
 B-F equation says:
 $d_i(x) = \min_v (c(i,v) + d_v(x))$
 $= \min(2+5, 1+3, 1+3) = 4$
 node achieving minimum is next hop in shortest path, used in forwarding table

key idea:
 • From time-to-time, each node sends its own distance vector estimate to neighbors
 • when it receives new DV estimate from neighbor, it updates its own DV using B-F equation
 $D_i(y) \leftarrow \min_v (c(i,v) + D_v(y))$ for each node $v \in N$
 • under certain conditions, the estimate $D_i(y)$ converge to the actual least cost $d_i(y)$



link cost changes:
 • node detects local link cost change
 • updates routing info, recalculates distance vector
 • if DV changes, notify neighbors

poisoned reverses:
 • if Z routes through Y to get to X:
 • 2 hops Y to Z (Z is distance to Y is infinite (so Y won't route to X via Z))
 • will this completely solve count to infinity problem?

Comparison of LS and DV algorithms

message complexity: LS: with n nodes, E links, $O(nE)$ msgs sent
 DV: exchange between neighbors only
 • convergence time varies

speed of convergence: LS: $O(n^2)$ algorithm requires $O(nE)$ msgs
 DV: convergence time varies
 • may have oscillations
 • may be routing loops
 • count-to-infinity problem

Scalable Routing: forwarding table configured by both intra- and inter-AS routing algorithms, intra-AS routing determine entries for destinations within AS, inter-AS & intra-AS determine entries for external destinations

Inter-AS tasks


AS1 must:
 1. learn which dests are reachable through AS2, which through AS3
 2. propagate this reachability info to all routers in AS1

Job of inter-AS routing!

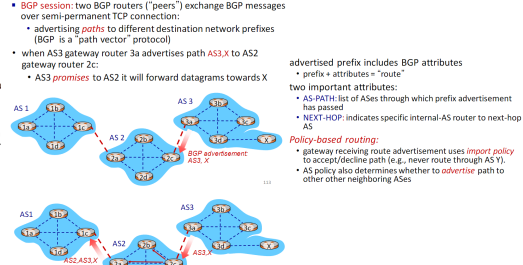
Intra-AS Routing in the internet: also known as interior gateway protocols (IGP), most common intra-AS routing protocols: RIP: Routing Information Protocol, OSPF: Open Shortest Path First, IGRP: Interior Gateway Routing Protocol (Cisco proprietary for decades, until 2016)
OSPF (Open Shortest Path First)

"open": publicly available
 uses link-state algorithm
 • link state packet dissemination
 • topology map at each node
 • route computation using Dijkstra's algorithm
 router floods OSPF link-state advertisements to all other routers in **entire AS**
 • carried in OSPF messages directly over IP (rather than TCP or UDP)
 • link state: for each attached link


Hierarchical OSPF



Internet inter-AS routing: BGP (Border Gateway Protocol): the de facto inter-domain routing protocol, "glue that holds the Internet together". BGP provides each AS a means to: **eBGP:** obtain subset reachability information from neighboring ASes; **iBGP:** propagate reachability information to all AS-internal routers, determine "good" routes to other networks based on reachability information and policy, allows subnet to advertise its existence to rest of Internet



AS2 router 2c receives path advertisement AS2,X (via eBGP) from AS3 router 3a
 • Based on AS2 policy, AS2 router 2c accepts path AS2,X, propagates (via iBGP) to all AS2 routers
 • Based on AS2 policy, AS2 router 2a advertises (via eBGP) path AS2, AS2,X to AS1 router 1c



gateway router may learn about multiple paths to destination:
 • AS1 gateway router 1c learns path AS2,AS3,X from 2a
 • AS1 gateway router 1c learns path AS3,X from 3a
 • Based on policy, AS1 gateway router 1c chooses path AS2,X, X and advertises path within AS1 via iBGP


BGP, OSPF, forwarding table entries
 Q: how does router set forwarding table entry to distant prefix?



BGP route selection:
 • router may learn about more than one route to destination AS, selects route based on:
 1. local preference value attribute: policy decision
 2. shortest AS-PATH
 3. closest NEXT-HOP router: hot potato routing
 4. additional criteria

hot potato routing: choose local gateway that has least intra-domain cost (e.g., 2d chooses 2a, even though more AS hops to X); don't worry about inter-domain cost!

BGP: achieving policy via advertisements



Suppose an ISP only wants to route traffic to/from its customer networks (does not want to carry transit traffic between other ISPs)
 • A advertises path A-w to B and to C
 • B chooses not to advertise B-A-w to C
 • B gets no "revenue" for routing C-B-A-w, since none of C, A, or A's B's customers
 • C does not learn about C-B-A-w path
 • C will route C-A-w (not using B) to get to w


Why different Intra-, Inter-AS routing?

policy:
 • inter-AS: admin wants control over how its traffic routed, who routes through its net.
 • intra-AS: single admin, so no policy decisions needed

scale:
 • hierarchical routing saves table size, reduced update traffic


performance:
 • intra-AS: can focus on performance
 • inter-AS: policy may dominate over performance

Internet network layer: historically has been implemented via distributed, per-router approach, monolithic router contains switching hardware, runs proprietary implementation of Internet standard protocols (IP, RIP, IS-IS, OSPF, BGP) in proprietary router OS (e.g., Cisco IOS), different "middleboxes" for different network layer functions: firewalls, load balancers, NAT boxes, ...




Software defined networking (SDN) a logically centralized control plane?, easier network management: avoid router mis-configurations, greater flexibility of traffic flows, table-based forwarding allows "programming" routers, centralized "programming" easier: compute tables centrally and distribute, distributed "programming" more difficult: compute tables as result of distributed algorithm (protocol) implemented in each and every router, open (non-proprietary) implementation of control plane
Traffic engineering: difficult traditional routing what if network operator wants u-to-z traffic to flow along uvwx, x-to-z traffic


to flow xyzw? (need to define link weights so traffic routing algorithm computes routes accordingly)



Software defined networking (SDN)
 1. generalized flow-based forwarding (e.g., OpenFlow)
 2. control plane, data plane separation
 3. control plane functions external to data plane switches




Flow table in a router (computed and distributed by controller) define router's match-action rules
 Each router contains a flow table that is computed and distributed by a logically centralized routing controller





flow: defined by header fields
 • generalized forwarding: simple packet-handling rules
 • Pattern: match values in packet header fields
 • Action: for matched packet: drop, forward, modify, matched packet or send matched packet to controller
 • Priority: disambiguate overlapping patterns
 • Counters: Bytes and packets

Destination-based forwarding:
 1. src=1.2.*.* , dest=3.4.5.* → drop
 2. src=*.*.*.* , dest=3.4.*.* → forward(2)
 3. src=10.1.2.3, dest=*.*.*.* → send to controller




SDNData plane switches: 1) fast, simple, commodity switches implementing generalized data-plane forwarding in hardware, 2) switch flow table computed, installed by controller, 3) API for table-based switch control (e.g., OpenFlow), defines what is controllable and what is not, 4) protocol for communicating with controller (e.g., OpenFlow)

SDN controller (network OS): 1) maintain network state information 2) interacts with network control applications "above" and northbound API interacts with network switches "below" 3) southbound API 3) implemented as distributed system for performance, scalability, fault-tolerance, robustness

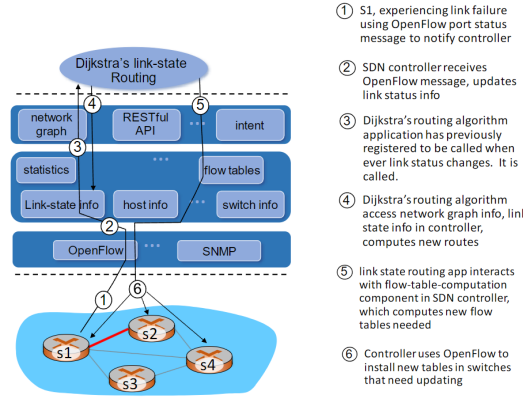


OpenFlow protocol
 • operates between controller, switch
 • TCP used to exchange messages
 • optional encryption
 • three classes of OpenFlow messages:
 • controller-to-switch
 • asynchronous (switch to controller)
 • symmetric (misc)




Key controller-to-switch messages
 • features: controller queries switch features, switch replies
 • configure: controller queries/sets switch configuration parameters
 • modify-table: add, delete, modify flow entries in the OpenFlow tables
 • packet-out: controller can send this packet out of specific switch port

Key switch-to-controller messages
 • packet-in: transfer packet (and its control) to controller. See packet-out message from controller
 • flow-removed: flow table entry deleted at switch
 • port status: inform controller of a change on a port.



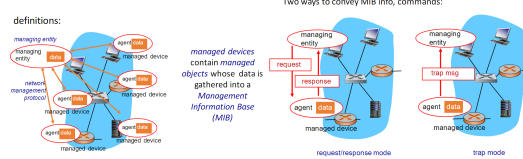
SDN network-control apps: 1) "brains" of control: implement control functions using lower-level services, API provided by SDN controller 2) unbundled: can be provided by 3rd party: distinct from routing vendor, or SDN controller
OpenDaylight (ODL) controller



SDN challenges, hardening the control plane: dependable, reliable, performance-scalable, secure, distributed system, robustness to failures: leverage strong theory of reliable distributed system for control plane, dependability, security: "baked in" from day one, networks, protocols meeting mission-specific requirements, e.g., real-time, ultra-reliable, ultra-secure, Internet-scaling
ICMP: internet control message protocol: used by hosts & routers to communicate network level information, error reporting: unreachable host, network, port, protocol, echo request/reply (used by ping), network-layer "above" IP; ICMP msgs carried in IP datagrams,
ICMP message: type, code plus first 8 bytes of IP datagram causing error

Type	Code	Description
0	0	echo reply (ping)
0	1	dest. network unreachable
0	2	dest. host unreachable
0	3	dest. protocol unreachable
0	3	dest. port unreachable
0	6	dest. network unknown
0	7	dest. host unknown
0	4	source quench (congestion control - not used)
8	0	echo request (ping)
9	0	route advertisement
10	0	router discovery
11	0	TTL expired
12	0	bad IP header

Network management includes the deployment, integration and coordination of the hardware, software, and human elements to monitor, test, poll, configure, analyze, evaluate, and control the network and element resources to meet the real-time, operational performance, and Quality of Service requirements at a reasonable cost.




Message type	Function
GetRequest	manager-to-agent: 'get me data' (data instance, net data in list of data)
GetNextRequest	manager-to-agent: get next MIB value
InformRequest	manager-to-manager: here's MIB value
SetRequest	manager-to-agent: set MIB value
Response	Agent-to-manager: value, response to Request
Trap	Agent-to-manager: inform manager of exceptional event

Link layer: introduction has responsibility of transferring datagram from one node to physically adjacent node over a link (terminology: hosts and routers: **nodes**, communication channels that connect adjacent nodes along communication path: **links**, wired links, wireless links, LANs, layer-2 packet: **frame**, encapsulates datagram)

Link layer services

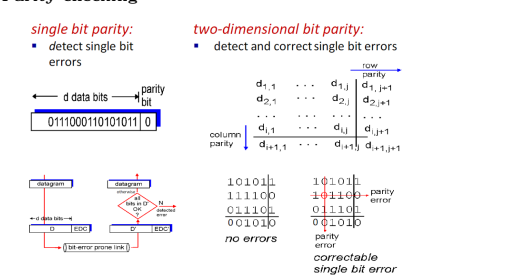
- framing, link access:**
 - encapsulate datagram into frame, adding header, trailer
 - channel access if shared medium
 - "MAC" addresses used in frame headers to identify source, destination
 - different from IP address
- reliable delivery between adjacent nodes**
 - we learned how to do this already (chapter 3)
 - sdlem used on low bit-error link (fiber, some twisted pair)
 - wireless links: high error rates
 - Q: why both link-level and end-end reliability?
- flow control:**
 - padding between adjacent sending and receiving nodes
- error detection:**
 - errors caused by signal attenuation, noise
 - receiver detects presence of errors:
 - signals sender for retransmission or drops frame
- error correction:**
 - receiver identifies and corrects bit error(s) without resorting to retransmission
- half-duplex and full-duplex**
 - with half duplex, nodes at both ends of link can transmit, but not at same time

Where is the link layer implemented? 1) in each and every host, 2) link layer implemented in "adapter" (aka network interface card NIC) or on a chip, Ethernet card, 802.11 card; Ethernet chipset, implements link, physical layer, attaches into host's system buses, combination of hardware, software, firmware




Error detection EDC = Error Detection and Correction bits (redundancy) D = Data protected by error checking, may include

header fields, Error detection not 100% reliable, protocol may miss some errors, but rarely, larger EDC field yields better detection and correction



Cyclic redundancy check more powerful error-detection coding using data bits, D, as a binary number, choose r-1 bit pattern (generator), G, goal: choose r CRC bits, R, such that, D/R is exactly divisible by G (modulo 2), receiver knows G, divides D/R, by G. If non-zero remainder: error detected, can detect all burst errors less than r+1 bits, widely used in practice (Ethernet, 802.11 WiFi, ATM)



Multiple access links, protocols 1) point-to-point, PPP for dial-up access, point-to-point link between Ethernet switch, host, 2) broadcast (shared wire or medium), old-fashioned Ethernet, upstream HFC, 802.11 wireless LAN

Multiple access protocol Given: single shared broadcast channel, two or more simultaneous transmissions by nodes: interference, collision if node receives two or more signals at the same time. **Multiple access protocols distributed algorithm** that determines how nodes share channel, i.e., determine when node can transmit, communication about channel sharing must use channel itself, no out-of-band channel for coordination.

An ideal multiple access protocol given: broadcast channel of rate R bps desiderata: 1. when one node wants to transmit, it can send at rate R. 2. when M nodes want to transmit, each can send at average rate R/M 3. fully decentralized; no special node to coordinate transmissions, no synchronization of clocks, slots 4. simple

MAC three broad classes: 1. channel partitioning, divide channel into smaller "pieces" (time slots, frequency, code), allocate piece to node for exclusive use, 2. random access, channel not divided, avoid collisions, "recover" from collisions, "taking turns" 3. nodes take turns, but nodes with more to send can take longer turns


TDMA: time division multiple access

- access to channel in "rounds"
- each station gets fixed length slot (length = packet transmission time) in each round
- unused slots go idle
- example: 6-station LAN, 1,3,4 have packets to send, slots 2,5,6 idle

Random access protocol: when node has packet to send, transmit at full channel data rate R, no a priori coordination among nodes, two or more transmitting nodes - "collision", random access MAC protocol specifies: how to detect collisions, how to recover from collisions (e.g., via delayed retransmissions), examples of random access MAC protocols: slotted ALOHA, ALOHA, CSMA, CSMA/CD, CSMA/CA

Slotted ALOHA assumptions: 1) all frames same size, 2) time divided into equal size slots (time to transmit 1 frame), 3) nodes start to transmit only slot beginning, 4) nodes are synchronized, if 2 or more nodes transmit in slot, these nodes detect collision

Slotted ALOHA operation: when node obtains fresh frame, transmits in next slot, if no collision: node can send new frame in next slot, if collision: node retransmits frame in each subsequent slot with prob. p until success



CSMA (carrier sense multiple access): listen before transmit; if channel sensed idle: transmit entire frame, if channel sensed busy, defer transmission. collisions can still occur: propagation delay means two nodes may not hear each other's transmission collision: entire packet transmission time wasted, distance & propagation delay play role in determining collision probability spatial layout of nodes

CSMA/CD: carrier sensing, deferral as in CSMA

- collisions detected with short time
- colliding transmissions aborted, reducing channel wastage

collision detection:

- easy in wired LANs: measure signal strengths, compare transmitted, received signals
- difficult in wireless LANs: received signal strength overwhelmed by local transmission strength

human analogy: the polite conversationalist

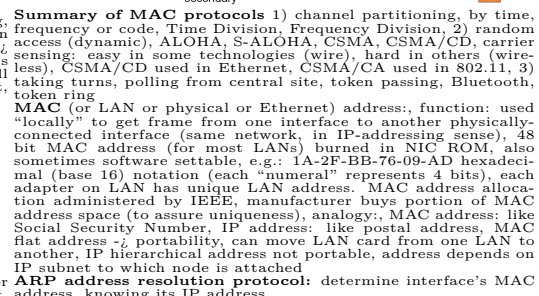


Ethernet CSMA/CD algorithm

- NIC receives datagram from network layer, creates frame
- If NIC senses channel idle, starts frame transmission. If NIC senses busy, it waits until channel idle, then transmits.
- If NIC transmits entire frame without detecting another transmission, NIC is done with frame 1
- If NIC detects another transmission while transmitting, aborts and sends jam signal
- After aborting, NIC enters binary (exponential) backoff:
 - after mth collision, NIC chooses K at random from {0,1,2,...,2^m-1}. NIC waits K*512 bit times, returns to Step 2
 - longer backoff interval with more collisions

"Taking turns" MAC protocols polling:


- primary node "invites" secondary nodes to transmit in turn
- concerns:
 - polling overhead
 - latency
 - single point of failure (primary)



Summary of MAC protocols 1) channel partitioning, by time, frequency or code, Time Division, Frequency Division, 2) random access (dynamic), ALOHA, S-ALOHA, CSMA, CSMA/CD, carrier sensing: easy in some technologies (wire), hard in others (wireless), CSMA/CD used in Ethernet, CSMA/CA used in 802.11, 3) taking turns, polling from central site, token passing, Bluetooth, token ring

MAC (or LAN or physical or Ethernet) address: function: used "locally" to get frame from one interface to another physically-connected interface (same network, in IP-addressing sense), 48 bit MAC address (for most LANs) burned in NIC ROM, also sometimes software settable, e.g.: IA-2F-BB-76-09-AD hexadecimal (base 16) notation (each "numeral" represents 4 bits), each adapter on LAN has unique LAN address. MAC address allocation administered by IEEE, manufacturer buys portion of MAC address space (to assure uniqueness), analogy; MAC address: like Social Security Number, IP address: like postal address, MAC flat address - portability, can move LAN card from one LAN to another, IP hierarchical address not portable, address depends on IP subnet to which node is attached

ARP address resolution protocol: determine interface's MAC address, knowing its IP address



Addressing: routing to another LAN

walkthrough: send datagram from A to B via R


- focus on addressing - at IP (datagram) and MAC layer (frame)
- assume A knows B's IP address
- assume A knows IP address of first hop router, R
- assume A knows R's MAC address

A creates IP datagram with IP source A, destination B

A creates link-layer frame with R's MAC address as destination address, frame contains A-to-B IP datagram

R forwards datagram with IP source A, destination B

R creates link-layer frame with B's MAC address as destination address, frame contains A-to-B IP datagram



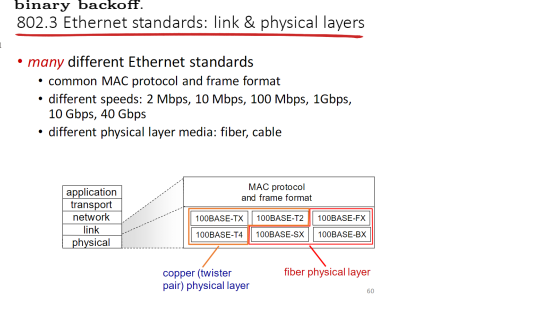
Ethernet "dominant" wired LAN technology: single chip, multiple speeds (e.g., Broadcom BCM5761), first widely used LAN technology, simple, cheap, kept up with speed race: 10 Mbps - 10 Gbps. bus: popular through mid 90s, all nodes in same collision domain (can collide with each other), star: prevails today, active switch in center, each "spoke" runs a (separate) Ethernet protocol (nodes do not collide with each other)

preamble: 7 bytes with pattern 10101010 followed by one byte with pattern 10101011, used to synchronize receiver, sender clock rates

addresses: 6 byte source, destination MAC addresses, if adapter receives frame with matching destination address, or with broadcast address (e.g. ARP packet), it passes data in frame to network layer protocol, otherwise, adapter discards frame, type: indicates higher layer protocol (mostly IP but others possible, e.g., Novell IPX, AppleTalk).

CRC: cyclic redundancy check at receiver, error detected: frame is dropped


Ethernet connectionless: no handshaking between sending and receiving NICs, unreliable: receiving NIC doesn't send acks or nacks to sending NIC, data in dropped frames recovered only if initial sender uses higher layer rdt (e.g., TCP), otherwise dropped data lost. Ethernet's MAC protocol: **unslothed CSMA/CD with binary backoff.**



Ethernet switch, link-layer device: takes an active role, store, forward Ethernet frames, examine incoming frame's MAC address, selectively forward frame to one-or-more outgoing links when frame is to be forwarded on segment, uses CSMA/CD to access segment, transparent, hosts are unaware of presence of switches, **plug-and-play, self-learning**, switches do not need to be configured.

Switch: multiple simultaneous transmissions


- hosts have dedicated, direct connection to switch
- switches buffer packets
- Ethernet protocol used on each incoming link, but no collisions; full duplex
- each link is its own collision domain



switching: A-to-A' and B-to-B' can transmit simultaneously, without collisions

Switch: self-learning: switch learns which hosts can be reached through which interfaces, when frame received, switch "learns" location of sender; incoming LAN segment, records sender/location pair in switch table


Self-learning, forwarding: example



Problem: If there is a loop in the extended LAN, a packet could circulate forever, Side question: Why would we have loops?, **Solution:** Select which switches should actively forward, Create a spanning tree to eliminate unnecessary edges, Adds robustness, Complicates learning/forwarding


Defining a Spanning Tree

- Basic Rules
 - Switch with the lowest ID is the root
 - For a given switch
 - A port in the direction of the root switch is the root port
 - For a given LAN
 - The switch closest to the root (or the switch with the lowest ID to break ties) is the designated switch for a LAN
 - The corresponding port is the designated port
 - Switches with no designated ports and ports that are neither a root port nor a designated port are not part of the tree.




Switches vs. routers

- both are store-and-forward:
- routers: network-layer devices (examine network-layer headers)
- switches: link-layer devices (examine link-layer headers)
- both have forwarding tables:
- routers: compute tables using routing algorithms, IP addresses
- switches: learn forwarding table using flooding, learning, MAC addresses



VLANs Virtual Local Area Network switch(es) supporting VLAN capabilities can be configured to define multiple virtual LANs over single physical LAN infrastructure.

port-based VLAN: switch ports grouped (by switch management software) so that single physical switch



operates as multiple virtual switches

traffic isolation: frames to/from ports 1-8 can only reach ports 1-8

- can also define VLAN based on MAC addresses of endpoints, rather than switch port

dynamic membership: ports can be dynamically assigned among VLANs

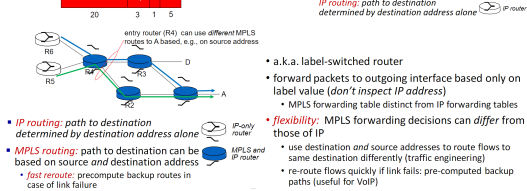
forwarding between VLANs: done via routing (just as with separate switches)

- in practice vendors sell combined switches plus routers

trunk port: carries frames between VLANs defined over multiple physical switches

- frames forwarded within VLAN between switches can't be vanilla 802.1 frames (must carry VLAN ID info)
- 802.1q protocol adds/removed additional header fields for frames forwarded between trunk ports

Multiprotocol label switching (MPLS): initial goal: high-speed IP forwarding using fixed length label (instead of IP address), fast lookup using fixed length identifier (rather than shortest prefix matching), borrowing ideas from Virtual Circuit (VC) approach, but IP datagram still keeps IP address!




Data center networks: 10's to 100's of thousands of hosts, often closely coupled, in close proximity: e-business (e.g. Amazon) content-servers (e.g., YouTube, Akamai, Apple, Microsoft), search engines, data mining (e.g., Google)

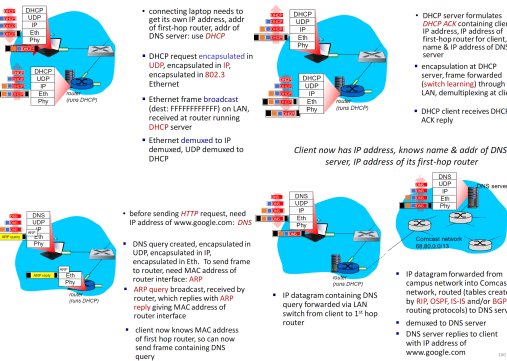
Challenges: 1) multiple applications, each serving massive numbers of clients 2) managing/balancing load, avoiding processing, networking, data bottlenecks

load balance: application-layer routing

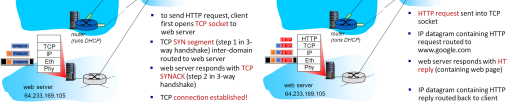
- receives external client requests
- directs workload within data center
- returns results to external client (hiding data center internals from client)



A day in the life of a web request, journey down protocol stack complete!, application, transport, network, link, putting-it-all-together: synthesis!, goal: identify, review, understand protocols (at all layers) involved in seemingly simple scenario: requesting www page, scenario: student attaches laptop to campus network, requests/receives www.google.com




Client now has IP address, knows name & addr of DNS server, IP addr of its first-hop router



Elements of a wireless network


- base station: typically connected to wired network, also used as backbone link, multiple access protocol coordinates link access, various data rates, transmission distance
- wireless hosts: laptop, smartphone, run applications, may be stationary (non-mobile) or mobile, wireless does not always mean mobility



Wireless network taxonomy

	single hop	multiple hops
infrastructure (e.g., APs)	host connects to base station (WiFi, WiMAX, cellular) which connects to larger Internet	host may have to rely through several wireless nodes to connect to larger Internet; mesh-net
no infrastructure	no base station, no connection to larger Internet (Bluetooth, ad-hoc nets)	no base station, no connection to larger Internet; may have to rely to reach either a given wireless network or MANET/WANET

Wireless Link Characteristics: important differences from wired link **decreased signal strength:** radio signal attenuates as it propagates through matter (path loss) **interference from other sources:** standardized wireless network frequencies (e.g., 2.4 GHz) shared by other devices (e.g., phone); devices (motors) interfere as well **multipath propagation:** radio signal reflects off objects ground, arriving at destination at slightly different times




$$P_{Rx} = \frac{G_{Tx} G_{Rx} \lambda^2}{(4\pi d)^2} P_{Tx}$$

$$\text{Signal-to-Noise Ratio: } SNR = \frac{|h|^2 \times |x(t)|^2}{\ln(t)^2} = \frac{|h|^2 P_{Tx}}{N}$$

$$\text{Path Loss (dB)} = 10 \log_{10} P_{Tx}/P_{Rx}$$


SNR: signal-to-noise ratio

- High SNR - easier to extract signal from noise (a "good thing")
- Low SNR - hard to extract signal from noise (a "bad thing")



SNR versus BER tradeoffs

- given physical layer modulation: **Higher SNR -> Low BER**
- given SNR: choose physical layer that meets BER requirement, gives highest throughput
- SNR may change with mobility: dynamically adapt physical layer modulation technique, coding/2-rate adaptation



Shannon Capacity Theorem: Capacity = Bandwidth x log2(1 + SNR)

multipath propagation: radio signal reflects off objects ground, arriving at destination at slightly different times




important differences from wired link ...

- **interference from other sources:** standardized wireless network frequencies (e.g., 2.4 GHz) shared by other devices (e.g., phone); devices (motors) interfere as well

MAC Protocols necessary to avoid interference

Problem of radio transmission: frequency dependent fading can wipe out narrow band signals for duration of the interference, **Solution:** spread the narrow band signal into a broad band signal using a special code



DSSS (Direct Sequence Spread Spectrum)

- XOR the signal with pseudo noise (PN) sequence (chipping sequence)
- Advantages:
 - reduces frequency selective fading
 - Robust to interference
 - Multi-user



FHSS (Frequency Hopping Spread Spectrum)

- Discrete changes of carrier frequency
- Advantages:
 - Frequency selective fading and interference limited to short period
 - uses only small portion of spectrum at any time
 - Secure
 - Used in Bluetooth & military applications



Advantage of signal attenuation: Spatial Reuse

- To double transmission range, we need: 4x more overall power!
- To transmit over two hops, we need: 2x more overall power!




Multi-hop wireless networks:

- Increase TX power: increase transmission range by N times, need N^2 x more power
- Multi-hop links: increase transmission range by N times, need N x more power

Ad hoc multi-hop wireless networks!


Frequency multiplex

- Separation of spectrum into smaller frequency bands
- Channel gets band of the spectrum for the whole time
- Advantages:
 - no dynamic coordination needed
 - works also for analog signals
- Disadvantages:
 - waste of bandwidth if traffic distributed unevenly
 - inflexible
 - guard spaces



Time multiplex


- Channel gets the whole spectrum for a certain amount of time
- Advantages:
 - only one carrier in the medium at any time
 - throughput high even for many users
- Disadvantages:
 - precise synchronization necessary



Time and frequency multiplex

- A channel gets a certain frequency band for a certain amount of time (e.g. GSM)
- Advantages:
 - better protection against tapping
 - protection against frequency selective interference
 - higher data rates compared to code multiplex
- Disadvantages:
 - Precise coordination required

Code Division Multiple Access (CDMA) 1) unique "code" assigned to each user; i.e., code set partitioning, all users share same frequency, but each user has own "chipping" sequence (i.e., code) to encode data, allows multiple users to "coexist" and transmit simultaneously with minimal interference (if codes are "orthogonal"), encoded signal = (original data) X (chipping sequence), decoding: inner-product of encoded signal and chipping sequence, Example codes: Gold Codes, Walsh Codes


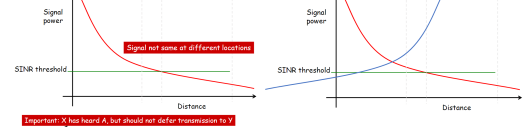


The Channel Access Problem: Multiple nodes share a channel, Pairwise communication desired, Simultaneous communication not possible, MAC Protocols, Suggests a scheme to schedule communication, Maximize number of communications, Ensure fairness among all transmitters

2 Observations on CSMA/CD 1) Transmitter can send/listen concurrently, If nothing is received while sending, then success 2) The signal is identical at Tx and Rx (Non-dispersive)

Both observations do not hold for wireless

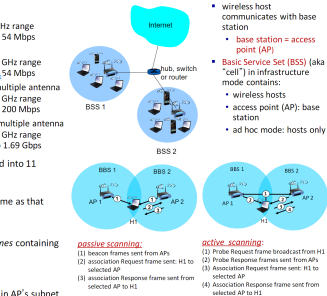
- Important: C has not heard A, but can interfere at receiver B
- Important: X has heard A, but should not defer transmission to Y



IEEE 802.11 Wireless LAN

802.11a

- 2.4 GHz unlicensed spectrum
 - up to 11 Mbps
 - direct sequence spread spectrum (DSSS) in physical layer
 - all hosts use same channel code
- 802.11a: multiple antenna
- 2.4 GHz range
 - up to 200 Mbps
- 802.11a: multiple antenna
- 2.4 GHz range
 - up to 1.69 Gbps
- 802.11b: 2.4GHz-2.485GHz spectrum divided into 11 channels at different frequencies
- AP admin chooses frequency for AP
 - interference possible: channel can be same as that chosen by neighboring AP!
 - host: must associate with an AP
 - scans channels, listening for beacon frames containing AP's name (SSID) and MAC address
 - selects AP to associate with
 - may perform authentication
 - will typically run DHCP to get IP address in AP's subnet



IEEE 802.11: multiple access:



- 1) avoid collisions: 2+ nodes transmitting at same time, 802.11: CSMA - sense before transmitting, don't collide with ongoing transmission by other node.
- 2) no collision detection!, difficult to receive (sense collisions) when transmitting due to weak received signals (fading), can't sense all collisions in any case: hidden terminal, fading. So, goal: avoid collisions: CSMA/C(ollision)A(voidance)

802.11 sender


- 1) If sense channel idle for BSS then transmit entire frame (no CTS)
- 2) If sense channel busy then start random backoff time timer count down while channel idle transmit when timer expires
- if no ACK, increase random backoff interval, repeat 2.

802.11 receiver


- if frame received OK
- return ACK after SIFS (ACK needed due to hidden terminal problem)



802.11 frame: addressing



4G/5G cellular networks, the solution for wide-area mobile Internet, widespread deployment/use: more mobile-broadband-connected devices than fixed- broadband-connected devices every (5-1 in 2019)!



Mobile device, smartphone, tablet, laptop, IoT, ... with 4G LTE radio, 64-bit International Mobile Subscriber Identity (IMSI), stored on SIM (Subscriber Identity Module) card, LTE jargon: User Equipment (UE)


Base station: at "edge" of carrier's network, manages wireless resources; mobile devices in its coverage area ("cell"), coordinates device authentication with other elements, similar to WiFi AP but: 1) active role in user mobility 2) coordinates with nearby base stations to optimize radio use, LTE jargon: eNode-B

Home Subscriber Service, stores info about mobile devices for which the HSS's network is their "home network", works with MME in device authentication


Serving Gateway (S-GW), PDN Gateway (P-GW), lie on data path from mobile to/from Internet. P-GW: gateway to mobile cellular network, Looks like any other internet gateway router, provides NAT services, other routers: extensive use of tunneling

Mobility Management Entity device authentication (device-to-network network-to-device) coordinated with mobile home network HSS, mobile device management: device handover between cells, tracking/paging device location, path (tunneling) setup from mobile device to P-GW


LTE: data plane control plane separation



LTE mobiles: sleep modes as in WiFi, Bluetooth: LTE mobile may put radio to "sleep" to conserve battery; light sleep: after 100's msec of inactivity, wake up periodically (100's msec) to check for downstream transmissions, deep sleep: after 5-10 secs of inactivity, mobile may change cells while deep sleeping - need to re-establish association

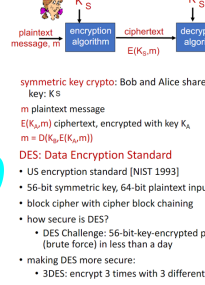


5G



Security: confidentiality: only sender, intended receiver should "understand" message contents, sender encrypts message, receiver decrypts message authentication: sender, receiver want to confirm identity of each other message integrity: sender, receiver want to ensure message not altered (in transit, or afterwards) without detection access and availability: services must be accessible and available to users

An adversary do 1) eavesdrop; intercept messages 2) actively insert message into connection 3) impersonation can fake (spoof) source address in packet (or any field in packet) 3) hijacking "take over" ongoing connection by removing sender or receiver, inserting himself in place 4) denial of service: prevent service from being used by others (e.g., by overloading resources).



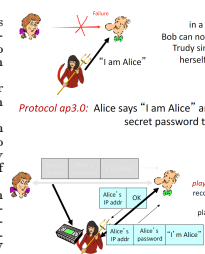
DES: Data Encryption Standard

- US encryption standard (NIST 1973)
- 56-bit symmetric key, 64-bit plaintext input
- block cipher with cipher block chaining
- how secure is DES?
 - DES Challenge: 56-bit key-encrypted phrase decrypted (brute force) in less than a day
- making DES more secure:
 - 3DES: encrypt 3 times with 3 different keys

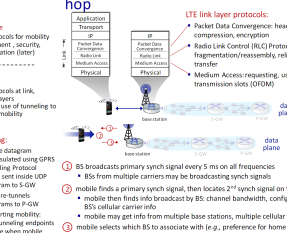
RSA: Creating public/private key pair

1. choose two large prime numbers p, q (e.g., 1024 bits each)
2. compute $n = pq, z = (p-1)(q-1)$
3. choose e (with e and n have no common factors with e, z are "relatively prime")
4. choose d such that $ed-1$ is exactly divisible by z . (in other words: private key is d)
5. public key is (n, e) , private key is (n, d)

RSA: encryption, decryption



LTE data plane protocol stack: first hop




home network HSS:

- identify & services info, while in home network and roaming
- all IP:
 - carriers interconnect with each other, and public internet at exchange points
 - may 3G, 2G, 3G; not all IP, handled otherwise

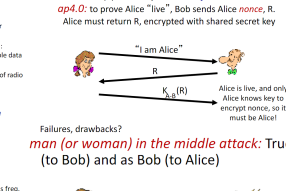
all IP:

- carriers interconnect with each other, and public internet at exchange points
- may 3G, 2G, 3G; not all IP, handled otherwise

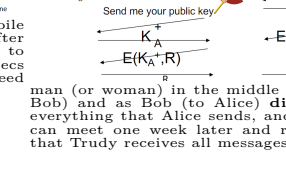
5G: microservice-like architecture



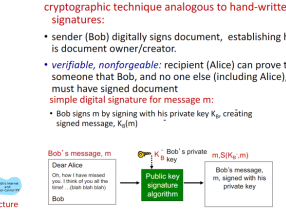
Goal: avoid playback attack



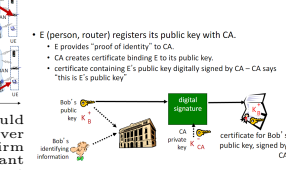
man (or woman) in the middle attack: Trudy poses as Alice (to Bob) and as Bob (to Alice)




Digital signatures



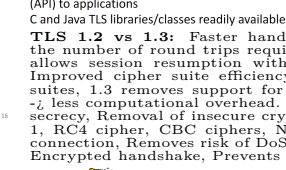
Certification authority (CA): binds public key to particular entity, E.




SSL/TLS: Transport Layer Security, Originally called Secure Sockets Layer (SSL) until 1996, widely deployed security protocol, supported by almost all browsers, web servers, https, billions \$/year over TLS, Current version: TLS1.3, TLS 1.3 is backwards compatible with TLS 1.2 but restricts usage on certain ciphers, provides confidentiality, integrity, authentication, original goals: Web e-commerce transactions, encryption (especially credit-card numbers), Web-server authentication, optional client authentication, minimum hassle in doing business with new merchant, available to all TCP applications, secure socket interface



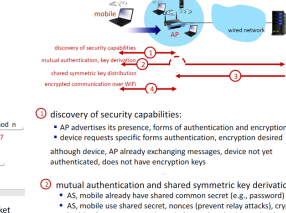
TLS 1.2 Handshake



TLS 1.3 Handshake



802.11: authentication, encryption



802.11: WPA3 handshake

