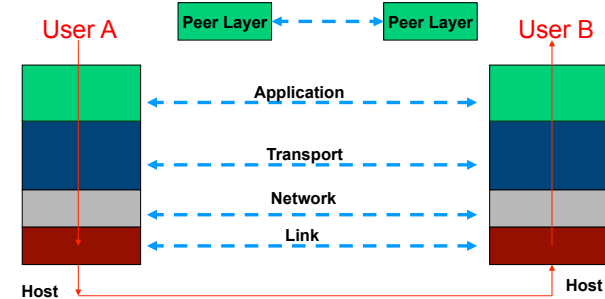




15-441 Computer Networking

Lecture 28 – Final Review

What is Layering?

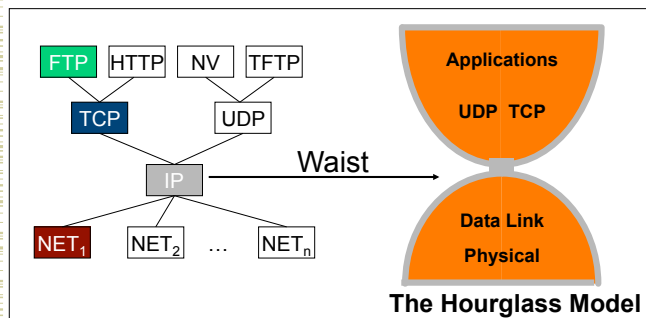


15-441 Fall 2011

© CMU 2005-2011

2

The Internet Protocol Suite



The waist facilitates interoperability

15-441 Fall 2011

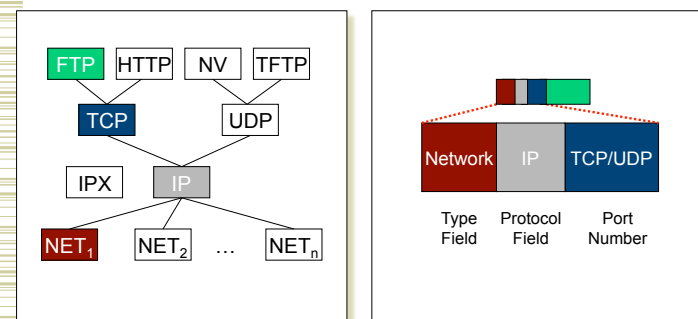
© CMU 2005-2011

3

Protocol Demultiplexing



- Multiple choices at each layer



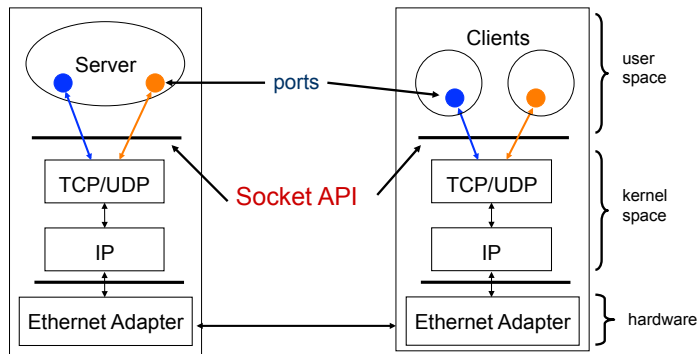
15-441 Fall 2011

© CMU 2005-2011

4

Server and Client

Server and Client exchange messages over the network through a common **Socket API**



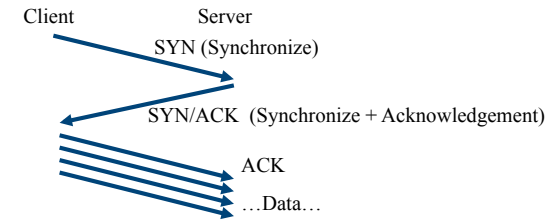
15-441 Fall 2011

© CMU 2005-2011

5

One more detail: TCP

- TCP connections need to be set up
 - "Three Way Handshake":



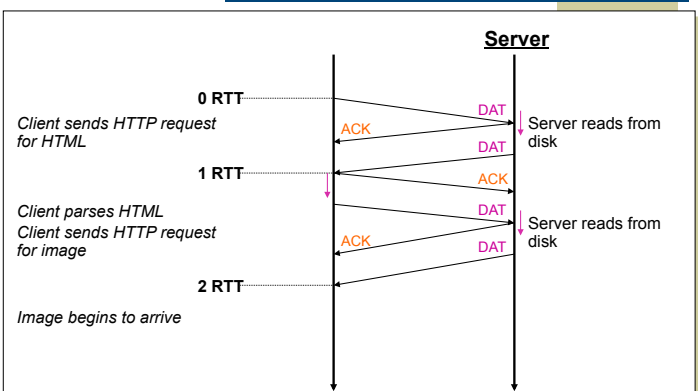
2: TCP transfers start slowly and then ramp up the bandwidth used (so they don't use too much)

15-441 Fall 2011

© CMU 2005-2011

6

Persistent Connection Solution

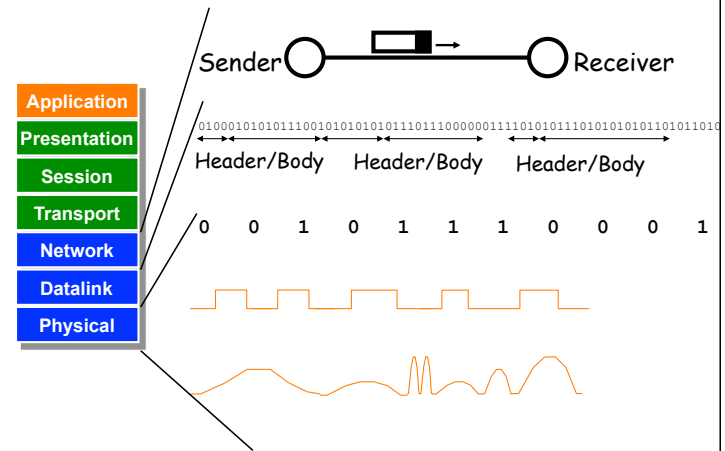


15-441 Fall 2011

© CMU 2005-2011

7

From Signals to Packets



15-441 Fall 2011

© CMU 2005-2011

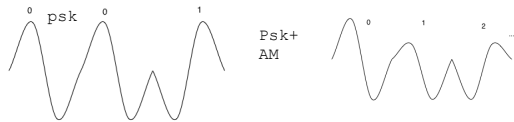
8

Past the Nyquist Limit

- More aggressive encoding can increase the channel bandwidth.

- Example: modems

- Same *frequency* - number of symbols per second
- Symbols have more possible values



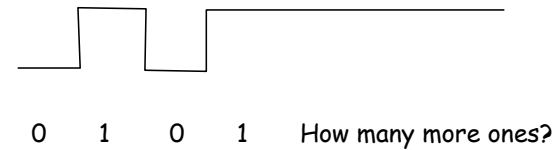
- Every transmission medium supports transmission in a certain frequency range.
 - The channel bandwidth is determined by the transmission medium and the quality of the transmitter and receivers
 - Channel capacity increases over time

15-441 Fall 2011

© CMU 2005-2011

9

Why Encode?



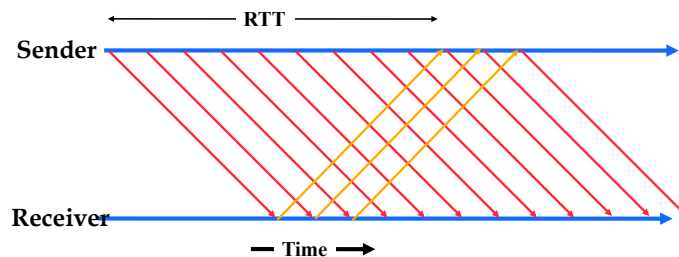
NRZ
NRZI
Manchester

15-441 Fall 2011

© CMU 2005-2011

10

Bandwidth-Delay Product



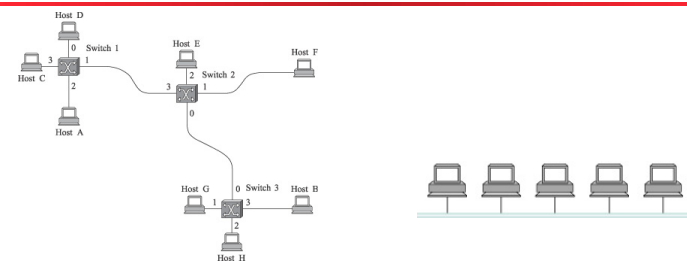
$$\text{Max Throughput} = \frac{\text{Window Size}}{\text{Roundtrip Time}}$$

15-441 Fall 2011

© CMU 2005-2011

11

Datalink Architectures



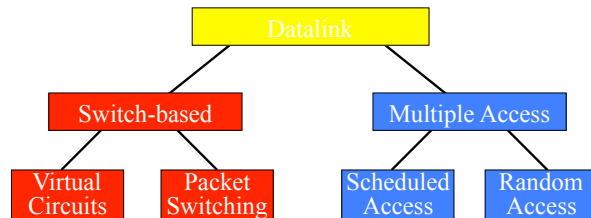
- Point-Point with switches
- Media access control.

15-441 Fall 2011

© CMU 2005-2011

12

Datalink Classification



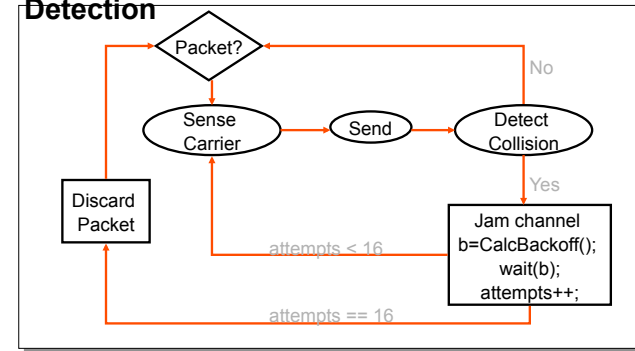
15-441 Fall 2011

© CMU 2005-2011

13

Ethernet MAC (CSMA/CD)

Carrier Sense Multiple Access/Collision Detection



15-441 Fall 2011

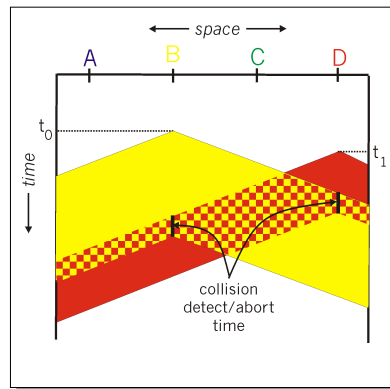
© CMU 2005-2011

14

Minimum Packet Size

What if two people sent really small packets

» How do you find collision?



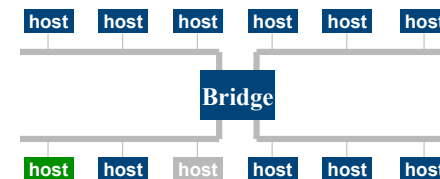
15-441 Fall 2011

© CMU 2005-2011

15

Learning Bridges

- Manually filling in bridge tables
 - Time consuming, error-prone
- Keep track of source address of packets arriving on every link, showing what segment hosts are on
 - Fill in the forwarding table based on this information



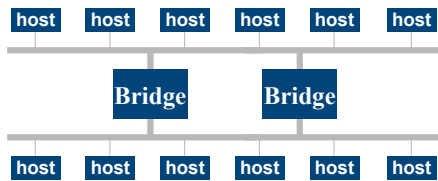
15-441 Fall 2011

© CMU 2005-2011

16

Spanning Tree Bridges

- More complex topologies can provide redundancy.
 - But can also create loops.
- What is the problem with loops?
- Solution: spanning tree

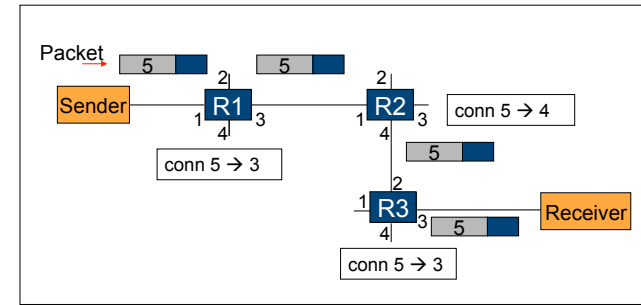


15-441 Fall 2011

© CMU 2005-2011

17

Simplified Virtual Circuits Example

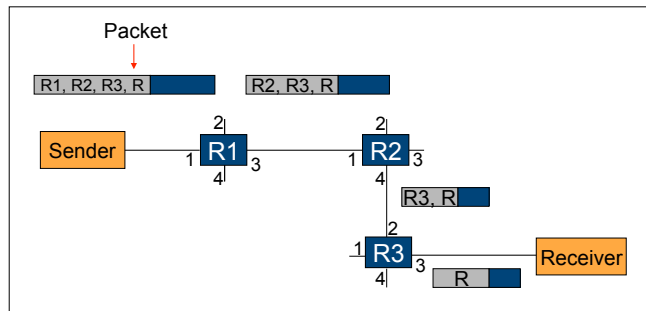


15-441 Fall 2011

© CMU 2005-2011

18

Source Routing Example

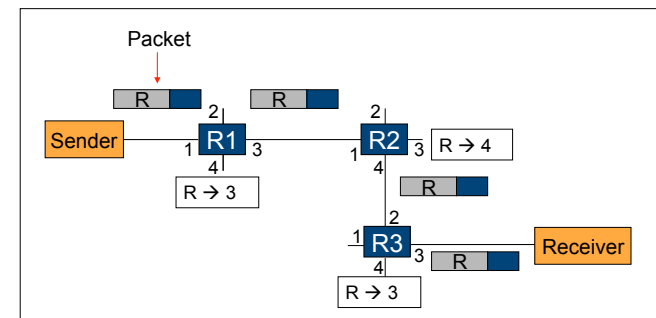


15-441 Fall 2011

© CMU 2005-2011

19

Global Address Example

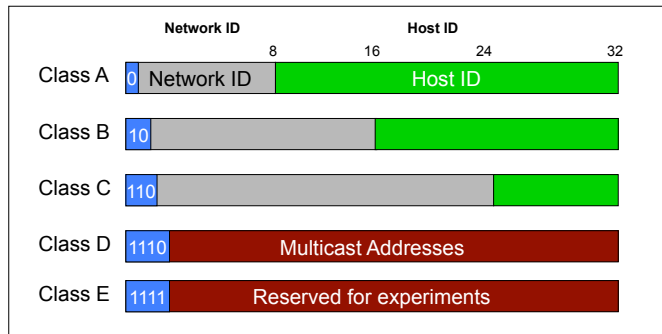


15-441 Fall 2011

© CMU 2005-2011

20

IP Address Classes (Some are Obsolete)



15-441 Fall 2011

© CMU 2005-2011

21

ARP Cache Example

- Show using command “arp -a”

Interface: 128.2.222.198 on Interface 0x1000003

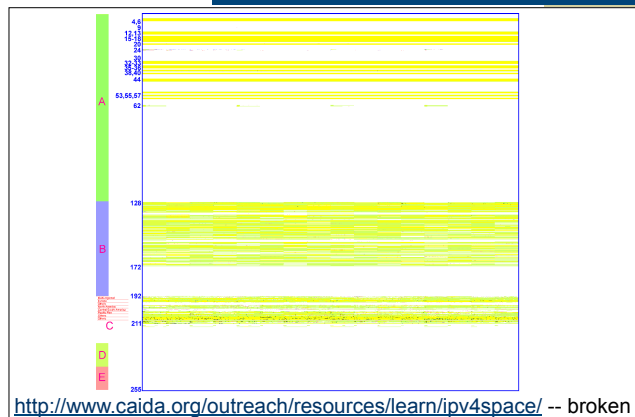
Internet Address	Physical Address	Type
128.2.20.218	00-b0-8e-83-df-50	dynamic
128.2.102.129	00-b0-8e-83-df-50	dynamic
128.2.194.66	00-02-b3-8a-35-bf	dynamic
128.2.198.34	00-06-5b-f3-5f-42	dynamic
128.2.203.3	00-90-27-3c-41-11	dynamic
128.2.203.61	08-00-20-a6-ba-2b	dynamic
128.2.205.192	00-60-08-1e-9b-fd	dynamic
128.2.206.125	00-d0-b7-c5-b3-f3	dynamic
128.2.206.139	00-a0-c9-98-2c-46	dynamic
128.2.222.180	08-00-20-a6-ba-c3	dynamic
128.2.242.182	08-00-20-a7-19-73	dynamic
128.2.254.36	00-b0-8e-83-df-50	dynamic

15-441 Fall 2011

© CMU 2005-2011

22

IP Address Utilization ('97)



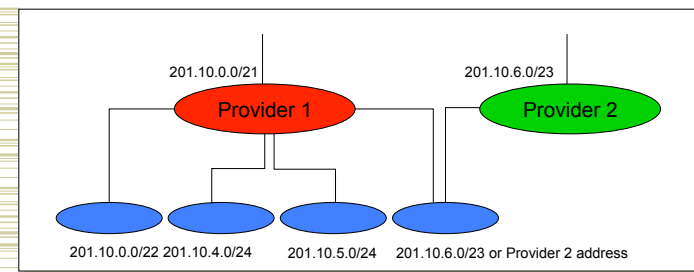
15-441 Fall 2011

© CMU 2005-2011

23

CIDR Implications

- Longest prefix match!!



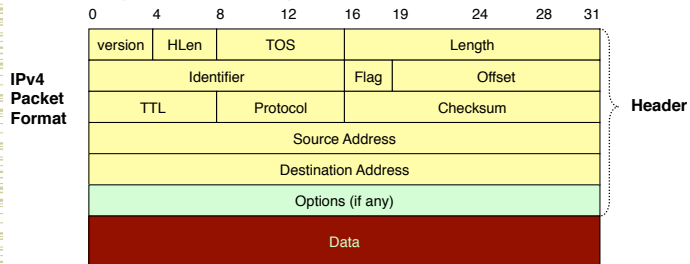
15-441 Fall 2011

© CMU 2005-2011

24

IP Service Model

- Low-level communication model provided by Internet
- Datagram
 - Each packet self-contained
 - All information needed to get to destination
 - No advance setup or connection maintenance
 - Analogous to letter or telegram

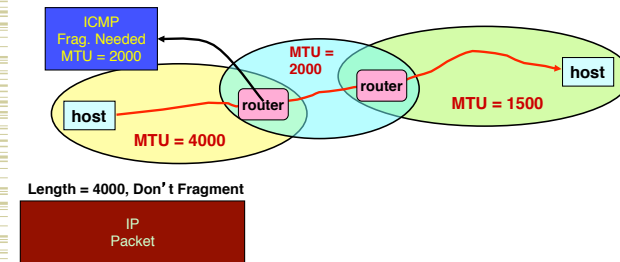


15-441 Fall 2011

© CMU 2005-2011

25

IP MTU Discovery with ICMP



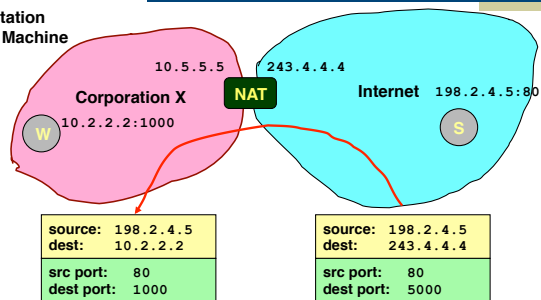
15-441 Fall 2011

© CMU 2005-2011

26

NAT: Server Response

W: Workstation
S: Server Machine



Int Addr	Int Port	NAT Port
10.2.2.2	1000	5000

- Firewall acts as proxy for client
 - Acts as destination for server messages
 - Relabels destination to local addresses

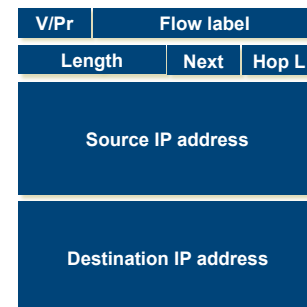
15-441 Fall 2011

© CMU 2005-2011

27

IPv6

- “Next generation” IP.
- Most urgent issue: increasing address space.
 - 128 bit addresses
- Simplified header for faster processing:
 - No checksum (why not?)
 - No fragmentation (?)
- Support for guaranteed services: priority and flow id
- Options handled as “next header”
 - reduces overhead of handling options

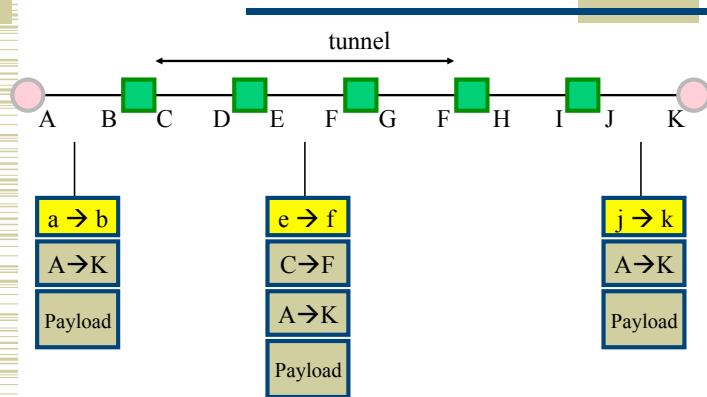


15-441 Fall 2011

© CMU 2005-2011

28

Tunneling Example

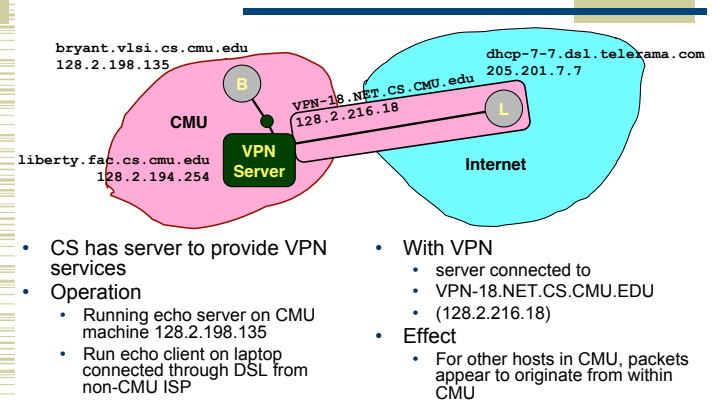


15-441 Fall 2011

© CMU 2005-2011

29

CMU CS VPN Example



- CS has server to provide VPN services
- Operation
 - Running echo server on CMU machine 128.2.198.135
 - Run echo client on laptop connected through DSL from non-CMU ISP

- With VPN
 - server connected to VPN-18.NET.CS.CMU.EDU (128.2.216.18)
- Effect
 - For other hosts in CMU, packets appear to originate from within CMU

15-441 Fall 2011

© CMU 2005-2011

30

Comparison of LS and DV Algorithms

Message complexity

- **LS**: with n nodes, E links, $O(nE)$ messages
- **DV**: exchange between neighbors only

Speed of Convergence

- **LS**: Relatively fast
 - Complex computation, but can forward before computation
 - may have transient loops
- **DV**: convergence time varies
 - may have routing loops
 - count-to-infinity problem
 - faster with triggered updates

Space requirements:

- LS maintains entire topology
- DV maintains only neighbor state

Robustness: router malfunctions

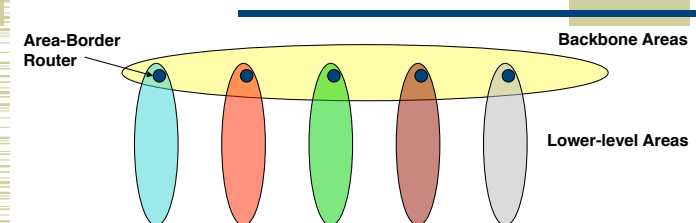
- **LS**: Node can advertise incorrect link cost
 - Each node computes its own table
- **DV**: Node can advertise incorrect path cost
 - Each node's table used by others (error propagates)

15-441 Fall 2011

© CMU 2005-2011

31

Routing Hierarchy



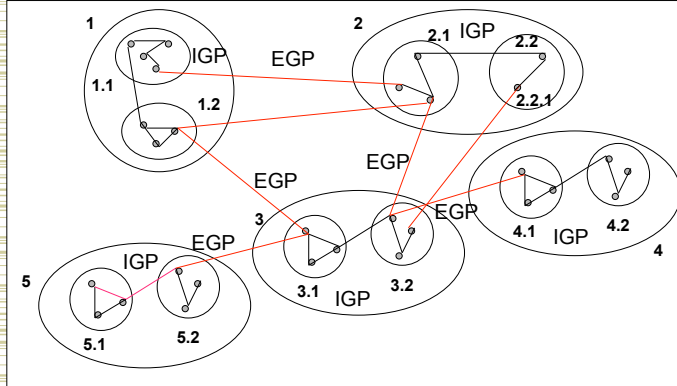
- Partition Network into "Areas"
 - Within area
 - Each node has routes to every other node
 - Outside area
 - Each node has routes for other top-level areas only
 - Inter-area packets are routed to nearest appropriate border router
- Constraint: no path between two sub-areas of an area can exit that area

15-441 Fall 2011

© CMU 2005-2011

32

Example

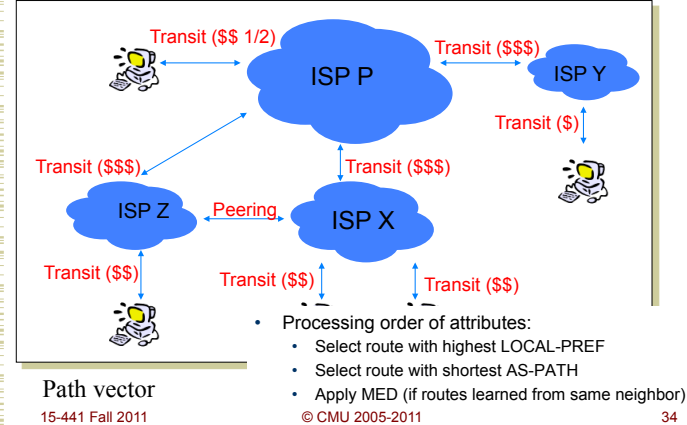


15-441 Fall 2011

© CMU 2005-2011

33

Transit vs. Peering



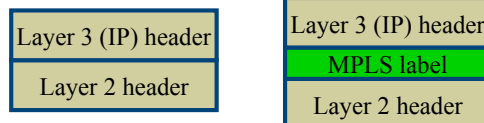
15-441 Fall 2011

© CMU 2005-2011

34

Multi Protocol Label Switching - MPLS

- Selective combination of VCs + IP
 - Today: MPLS useful for traffic engineering, reducing core complexity, and VPNs
- Core idea: Layer 2 carries VC label
 - Could be ATM (which has its own tag)
 - Could be a "shim" on top of Ethernet/etc.:
 - Existing routers could act as MPLS switches just by examining that shim -- no radical re-design. Gets flexibility benefits, though not cell switching advantages



15-441 Fall 2011

© CMU 2005-2011

35

DNS Records

RR format: (class, name, value, type, ttl)

- DB contains tuples called resource records (RRs)
 - Classes = Internet (IN), Chaosnet (CH), etc.
 - Each class defines value associated with type

FOR IN class:

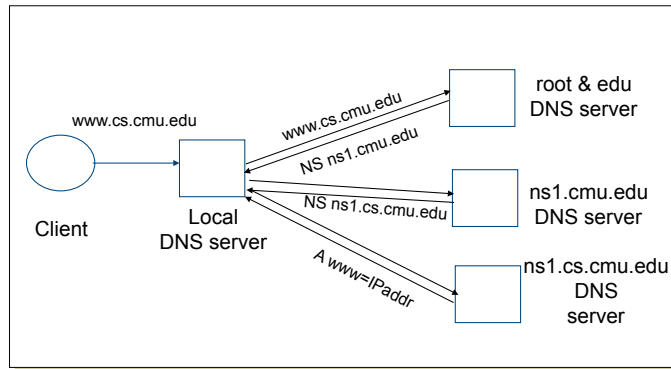
- Type=A
 - name** is hostname
 - value** is IP address
- Type=NS
 - name** is domain (e.g. foo.com)
 - value** is name of authoritative name server for this domain
- Type=CNAME
 - name** is an alias name for some "canonical" (the real) name
 - value** is canonical name
- Type=MX
 - value** is hostname of mailserver associated with **name**

15-441 Fall 2011

© CMU 2005-2011

36

Typical Resolution

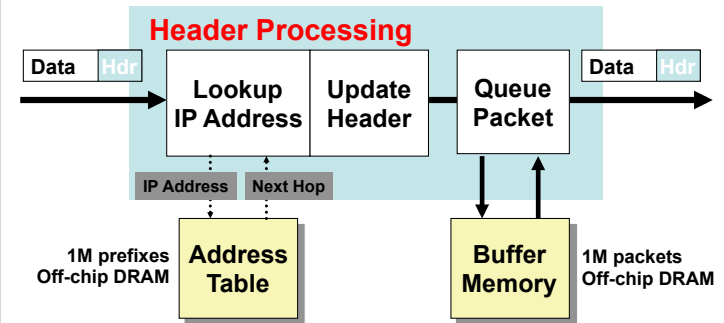


15-441 Fall 2011

© CMU 2005-2011

37

Generic Router Architecture

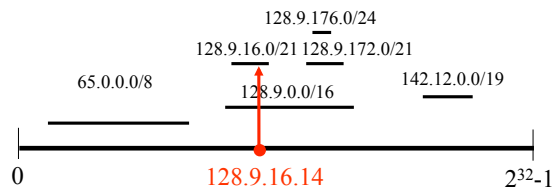


15-441 Fall 2011

© CMU 2005-2011

38

IP Lookups find Longest Prefixes



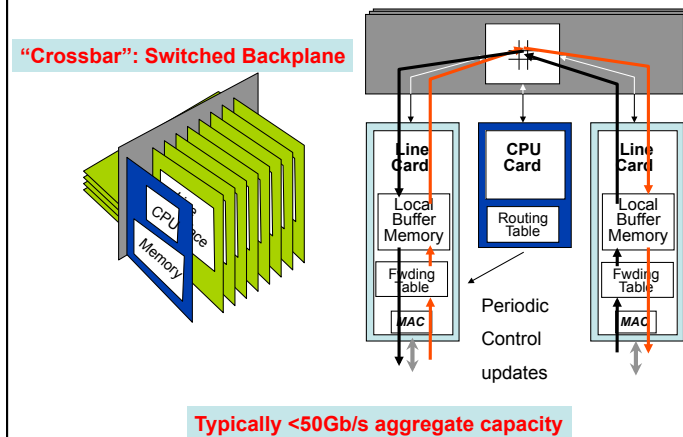
Routing lookup: Find the longest matching prefix (aka the most specific route) among all prefixes that match the destination address.

15-441 Fall 2011

© CMU 2005-2011

39

Third Generation Routers



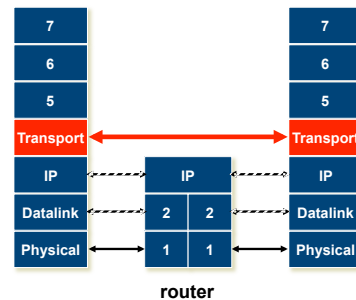
15-441 Fall 2011

© CMU 2005-2011

40

Transport Protocols

- Lowest level end-to-end protocol.
 - Header generated by sender is interpreted only by the destination
 - Routers view transport header as part of the payload

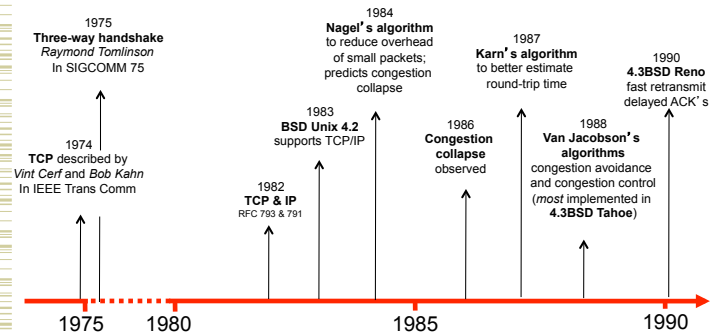


15-441 Fall 2011

© CMU 2005-2011

41

Evolution of TCP

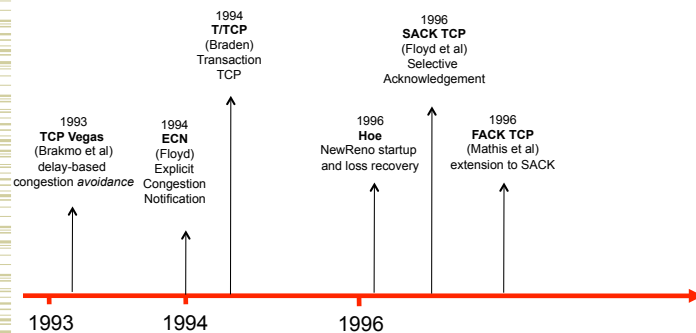


15-441 Fall 2011

© CMU 2005-2011

42

TCP Through the 1990s

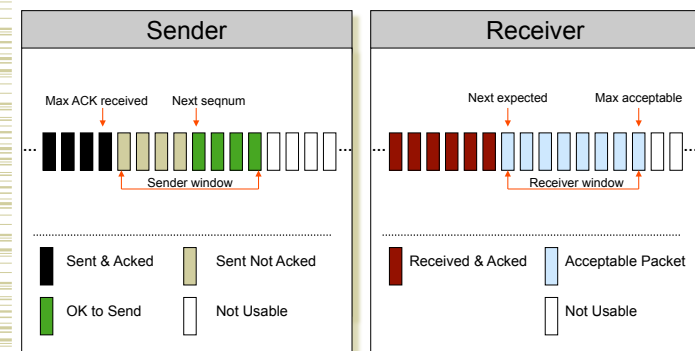


15-441 Fall 2011

© CMU 2005-2011

43

Sender/Receiver State

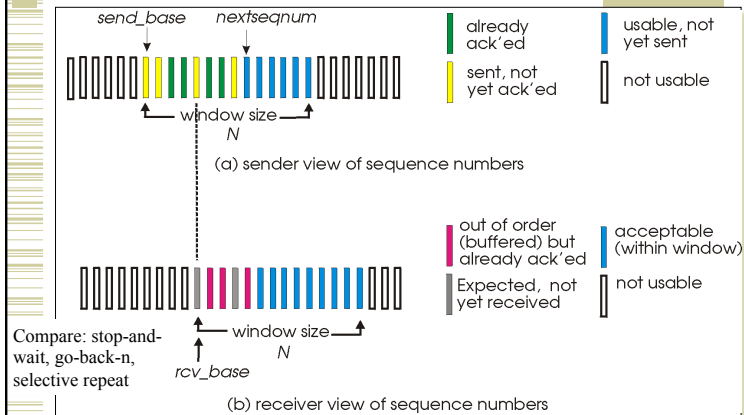


15-441 Fall 2011

© CMU 2005-2011

44

Selective Repeat: Sender, Receiver Windows



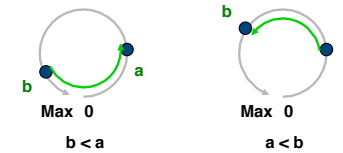
15-441 Fall 2011

© CMU 2005-2011

45

Sequence Numbers

- 32 Bits, Unsigned \rightarrow for bytes not packets!
 - Circular Comparison



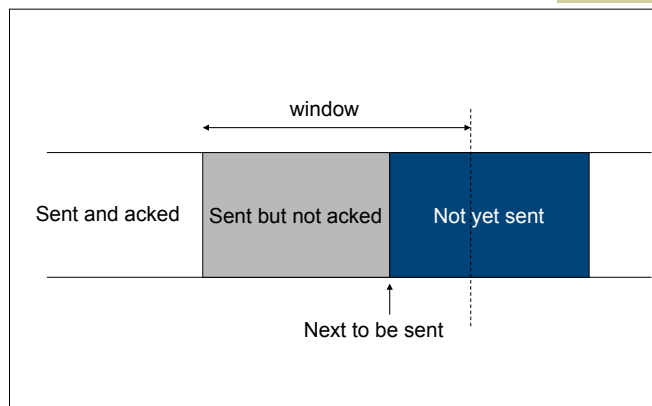
- Why So Big?
 - For sliding window, must have $|\text{Sequence Space}| > |\text{Sending Window}| + |\text{Receiving Window}|$
 - No problem
 - Also, want to guard against stray packets
 - With IP, packets have maximum lifetime of 120s
 - Sequence number would wrap around in this time at 286MB/s

15-441 Fall 2011

© CMU 2005-2011

46

Window Flow Control: Send Side

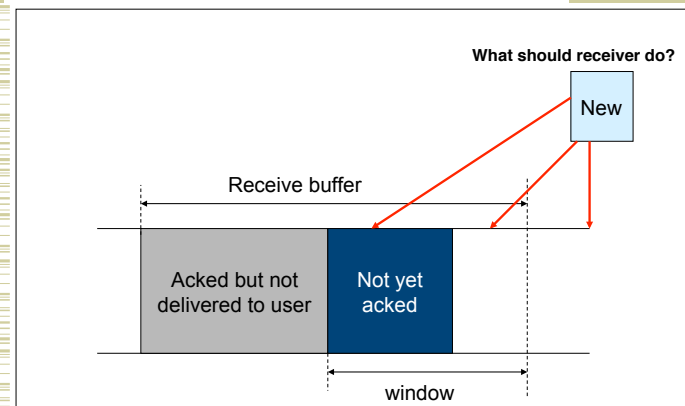


15-441 Fall 2011

© CMU 2005-2011

47

Window Flow Control: Receive Side

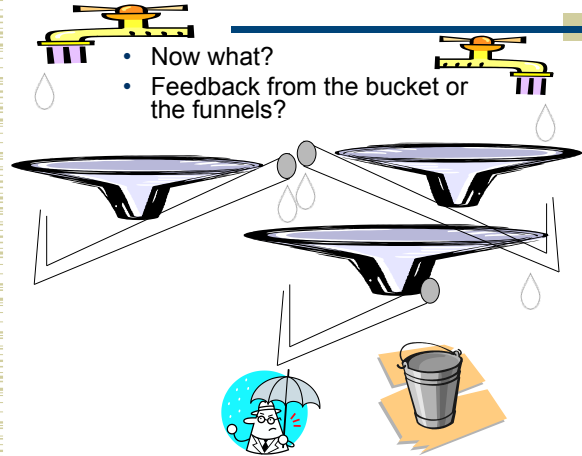


15-441 Fall 2011

© CMU 2005-2011

48

Plumbers Gone Wild 2!



- Now what?
- Feedback from the bucket or the funnels?

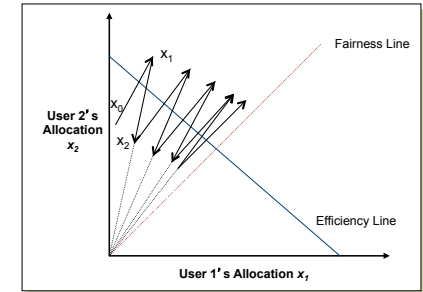
15-441 Fall 2011

© CMU 2005-2011

49

What is the Right Choice?

- Constraints limit us to AIMD
 - Can have multiplicative term in increase (MAIMD)
 - AIMD moves towards optimal point



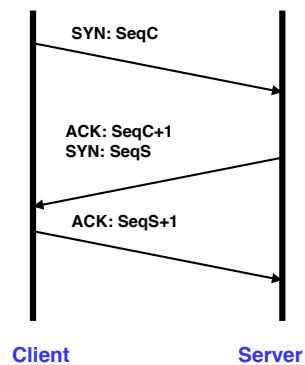
15-441 Fall 2011

© CMU 2005-2011

50

Establishing Connection: Three-Way handshake

- Each side notifies other of starting sequence number it will use for sending
 - Why not simply chose 0?
 - Must avoid overlap with earlier incarnation
 - Security issues
- Each side acknowledges other's sequence number
 - SYN-ACK: Acknowledge sequence number + 1
- Can combine second SYN with first ACK

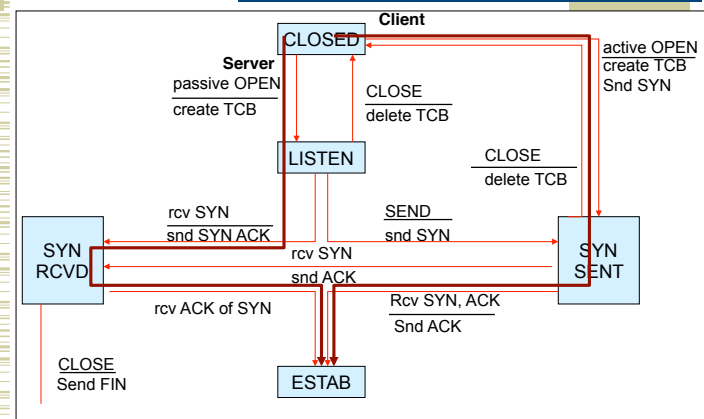


15-441 Fall 2011

© CMU 2005-2011

51

TCP State Diagram: Connection Setup

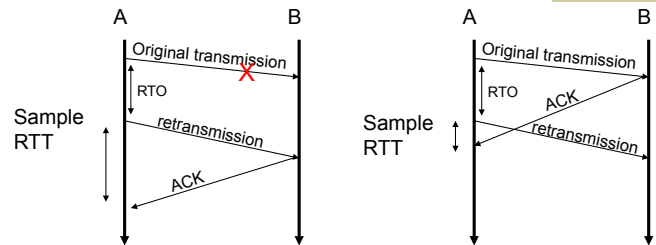


15-441 Fall 2011

© CMU 2005-2011

52

RTT Sample Ambiguity



Karn's RTT Estimator

- If a segment has been retransmitted:
 - Don't count RTT sample on ACKs for this segment
 - Keep backed off time-out for next packet
 - Reuse RTT estimate only after one successful transmission

15-441 Fall 2011

© CMU 2005-2011

53

Jacobson's Retransmission Timeout

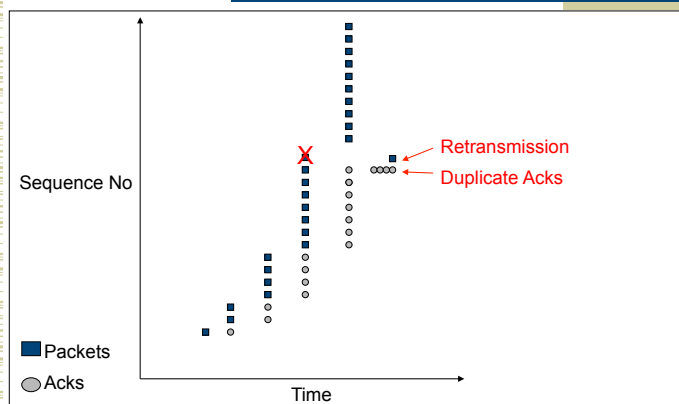
- Key observation:
 - At high loads, round trip variance is high
- Solution:
 - Base RTO on RTT and standard deviation
 - $RTO = RTT + 4 * rttvar$
 - $new_rttvar = \beta * dev + (1 - \beta) old_rttvar$
 - Dev = linear deviation
 - Inappropriately named – actually smoothed linear deviation

15-441 Fall 2011

© CMU 2005-2011

54

Fast Retransmit

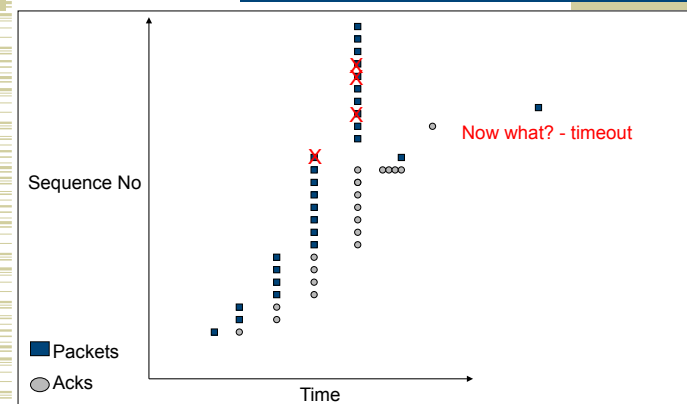


15-441 Fall 2011

© CMU 2005-2011

55

TCP (Reno variant)

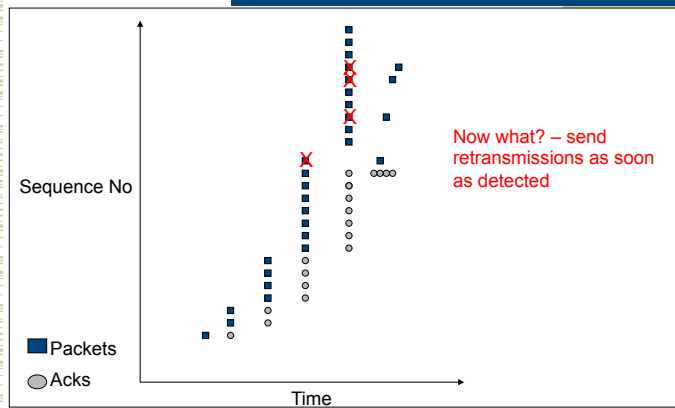


15-441 Fall 2011

© CMU 2005-2011

56

SACK



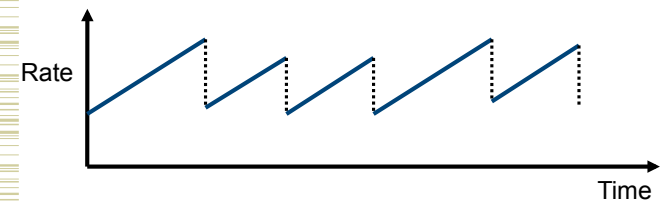
15-441 Fall 2011

© CMU 2005-2011

57

AIMD

- Distributed, fair and efficient
- Packet loss is seen as sign of congestion and results in a multiplicative rate decrease
 - Factor of 2
- TCP periodically probes for available bandwidth by increasing its rate



15-441 Fall 2011

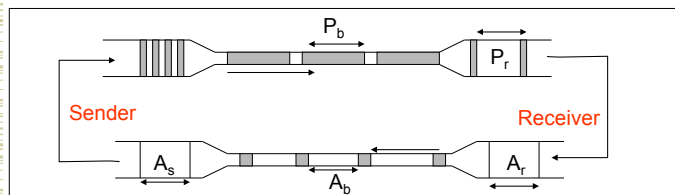
© CMU 2005-2011

58

TCP Packet Pacing

- Congestion window helps to “pace” the transmission of data packets
- In steady state, a packet is sent when an ack is received
 - Data transmission remains smooth, once it is smooth
 - Self-clocking behavior

Packet Conservation

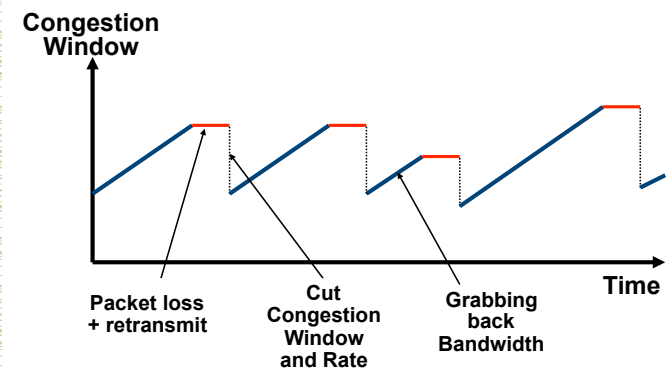


15-441 Fall 2011

© CMU 2005-2011

59

Congestion Avoidance Behavior



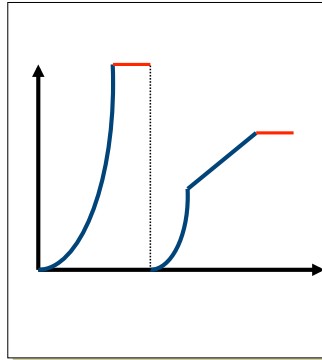
15-441 Fall 2011

© CMU 2005-2011

60

Slow Start Packet Pacing

- How do we get this clocking behavior to start?
 - Initialize $cwnd = 1$
 - Upon receipt of every ack, $cwnd = cwnd + 1$
- Implications
 - Window actually increases to W in $RTT * \log_2(W)$
 - Can overshoot window and cause packet loss

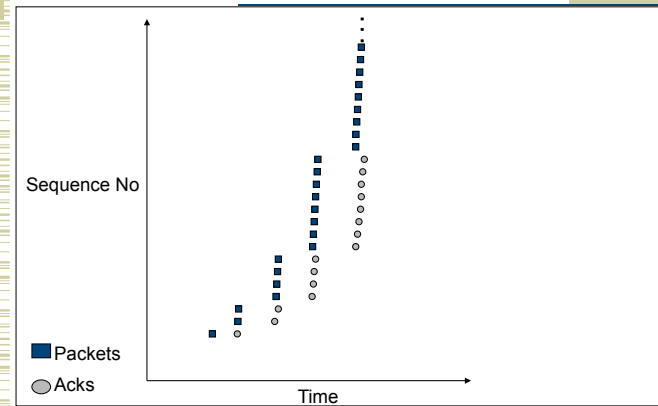


15-441 Fall 2011

© CMU 2005-2011

61

Slow Start Sequence Plot

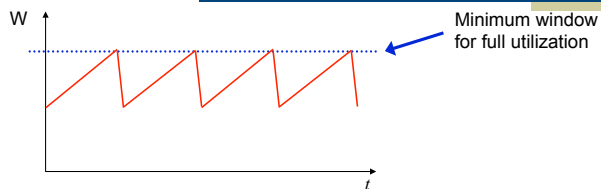


15-441 Fall 2011

© CMU 2005-2011

62

Summary Unbuffered Link



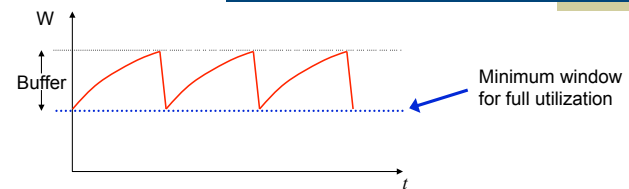
- The router can't fully utilize the link
 - If the window is too small, link is not full
 - If the link is full, next window increase causes drop
 - With no buffer it still achieves 75% utilization

15-441 Fall 2011

© CMU 2005-2011

63

Summary Buffered Link



- With sufficient buffering we achieve full link utilization
 - The window is always above the critical threshold
 - Buffer absorbs changes in window size
 - Buffer Size = Height of TCP Sawtooth
 - Minimum buffer size needed is $RTT * BW$
 - Delay? Between RTT and $2 * RTT$

15-441 Fall 2011

© CMU 2005-2011

64

TCP (Summary)

- General loss recovery
 - Stop and wait
 - Selective repeat
- TCP sliding window flow control
- TCP state machine
- TCP loss recovery
 - Timeout-based
 - RTT estimation
 - Fast retransmit
 - Selective acknowledgements

15-441 Fall 2011

© CMU 2005-2011

65

TCP (Summary)

- Congestion collapse
 - Definition & causes
- Congestion control
 - Why AIMD?
 - Slow start & congestion avoidance modes
 - ACK clocking
- Packet conservation
- TCP performance modeling
 - How does TCP fully utilize a link?
 - Role of router buffers

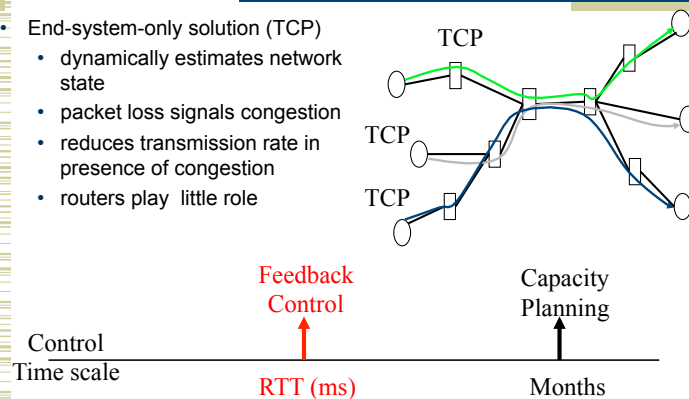
15-441 Fall 2011

© CMU 2005-2011

66

Congestion Control in Today's Internet

- End-system-only solution (TCP)
 - dynamically estimates network state
 - packet loss signals congestion
 - reduces transmission rate in presence of congestion
 - routers play little role



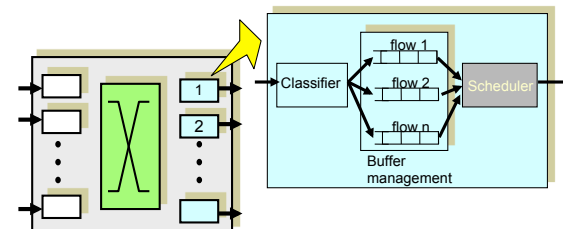
15-441 Fall 2011

© CMU 2005-2011

67

Router Mechanisms

- Buffer management: when and which packet to drop?
- Scheduling: which packet to transmit next?



15-441 Fall 2011

© CMU 2005-2011

68

Typical Internet Queueing

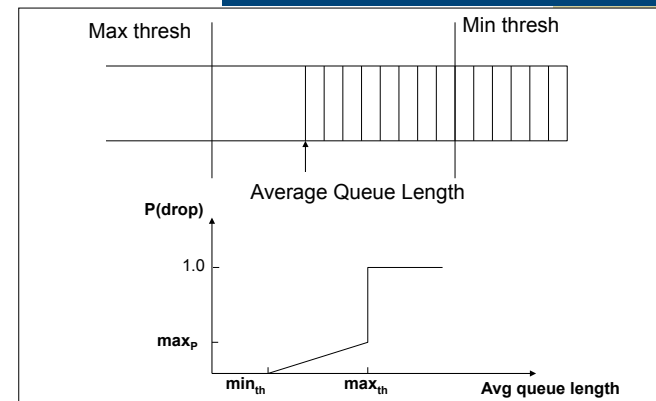
- FIFO (scheduling discipline) + drop-tail (drop policy)
 - Cong control at edges
 - No flow differentiation
 - Lock out
 - Random drop
 - Drop front
 - Full queues
 - Early random drop (RED)
 - Explicit congestion notification
 - decbit

15-441 Fall 2011

© CMU 2005-2011

69

RED Operation



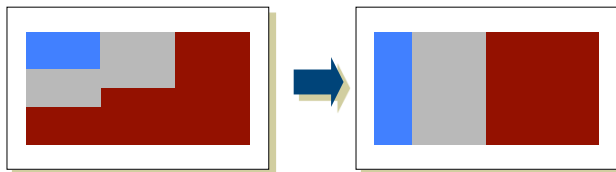
15-441 Fall 2011

© CMU 2005-2011

70

Fair Queuing

- Mapping bit-by-bit schedule onto packet transmission schedule
- Transmit packet with the lowest F_i at any given time
 - How do you compute F_i ?

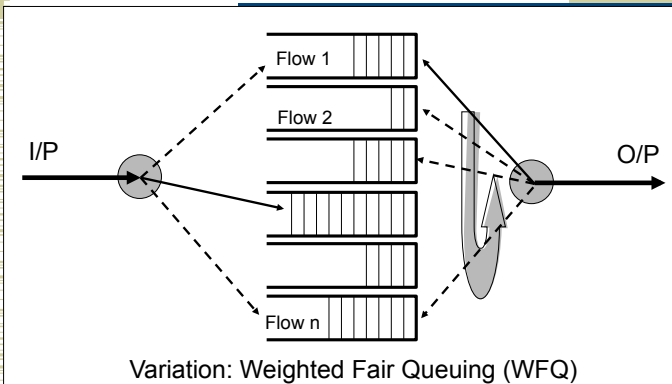


15-441 Fall 2011

© CMU 2005-2011

71

FQ Illustration

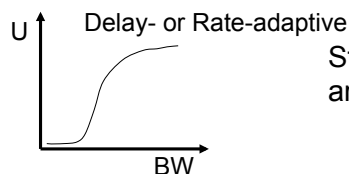
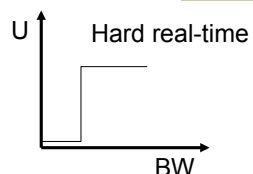
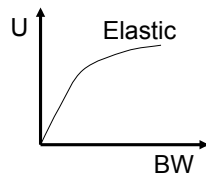


15-441 Fall 2011

© CMU 2005-2011

72

Utility Curve Shapes



Stay to the right and you are fine for all curves

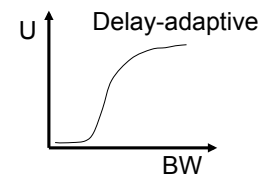
15-441 Fall 2011

© CMU 2005-2011

73

Admission Control

- If U is convex \rightarrow inelastic applications
 - $U(\text{number of flows})$ is no longer monotonically increasing
 - Need admission control to maximize total utility
- Admission control** \rightarrow deciding when adding more people would reduce overall utility
 - Basically avoids overload



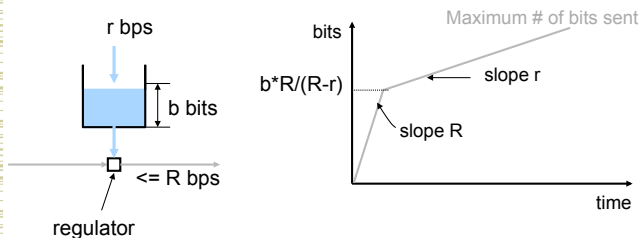
15-441 Fall 2011

© CMU 2005-2011

74

Token Bucket

- Parameters
 - r – average rate, i.e., rate at which tokens fill the bucket
 - b – bucket depth
 - R – maximum link capacity or peak rate (optional parameter)
- A bit is transmitted only when there is an available token



15-441 Fall 2011

© CMU 2005-2011

75

Guarantee Proven by Parekh

- Given:
 - Flow i shaped with token bucket and leaky bucket rate control (depth b and rate r)
 - Network nodes do WFQ
- Cumulative queuing delay D_i suffered by flow i has upper bound
 - $D_i < b/r$, (where r may be much larger than average rate)
 - Assumes that $\sum r <$ link speed at any router
 - All sources limiting themselves to r will result in no network queuing

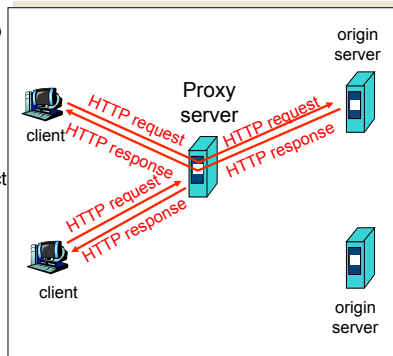
15-441 Fall 2011

© CMU 2005-2011

76

Web Proxy Caches

- User configures 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



15-441 Fall 2011

© CMU 2005-2011

77

W/Caching Example (3)

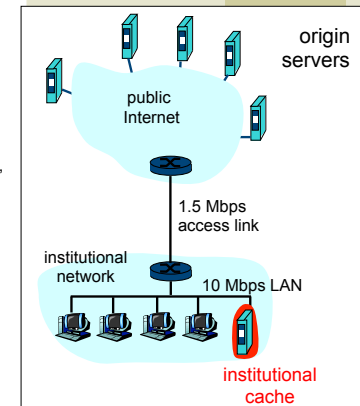
Install cache

- Suppose hit rate is .4

Consequence

- 40% requests will be satisfied almost immediately (say 10 msec)
- 60% requests satisfied by origin server
- Utilization of access link reduced to 60%, resulting in negligible delays
- Weighted average of delays

$$= .6 * 2 \text{ sec} + .4 * 10 \text{ msec} < 1.3 \text{ secs}$$



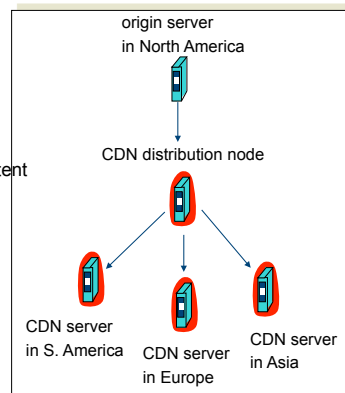
15-441 Fall 2011

© CMU 2005-2011

78

Content Distribution Networks (CDNs)

- The content providers are the CDN customers.
- Content replication**
 - CDN company installs hundreds of CDN servers throughout Internet
 - Close to users
- CDN replicates its customers' content in CDN servers. When provider updates content, CDN updates servers

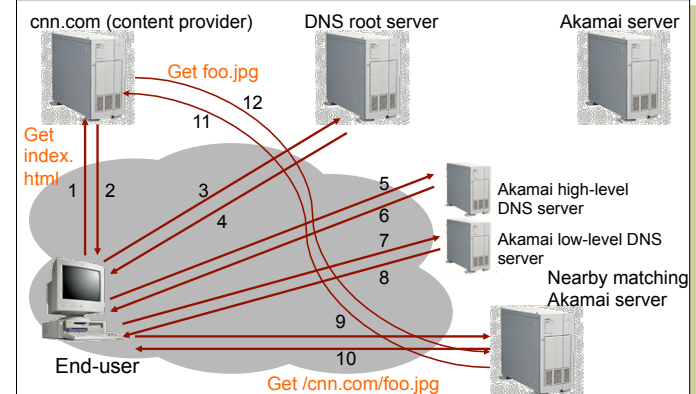


15-441 Fall 2011

© CMU 2005-2011

79

How Akamai Works

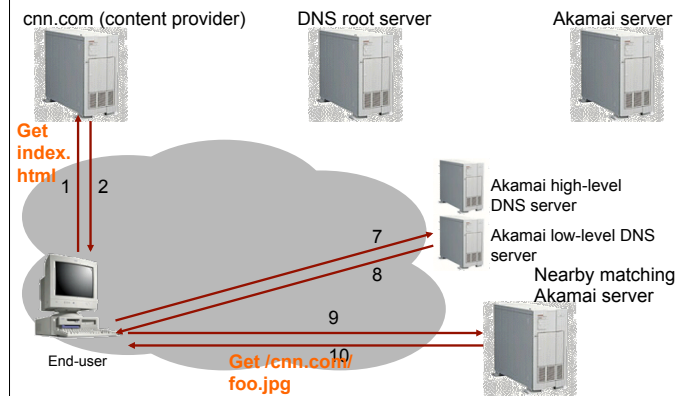


15-441 Fall 2011

© CMU 2005-2011

80

Akamai – Subsequent Requests



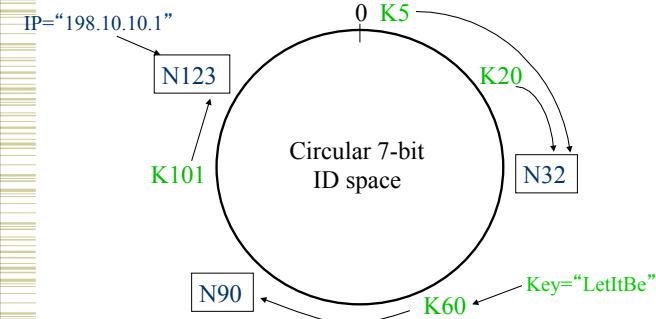
15-441 Fall 2011

© CMU 2005-2011

81

Consistent Hashing Example

Rule: A key is stored at its **successor** node with next higher or equal ID



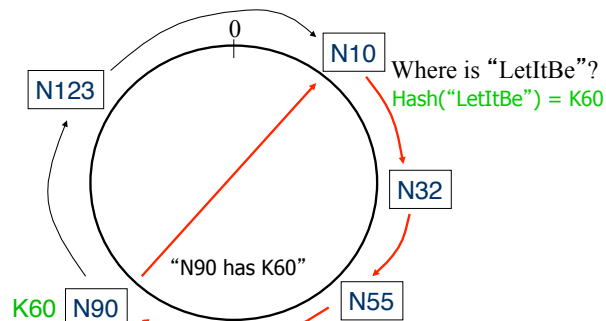
15-441 Fall 2011

© CMU 2005-2011

82

Lookups strategies

- Every node knows its successor in the ring
- Requires $O(N)$ lookups



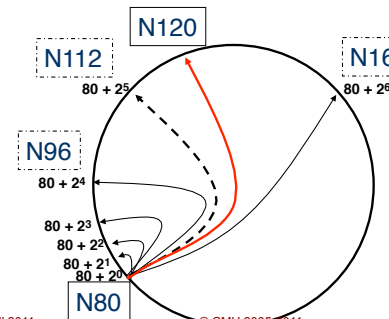
15-441 Fall 2011

© CMU 2005-2011

83

Reducing Lookups: Finger Tables

- Each node knows m other nodes in the ring (it has m fingers)
- Increase distance exponentially
- Finger i points to **successor** of $n + 2^{i-1}$ $i=1..m$



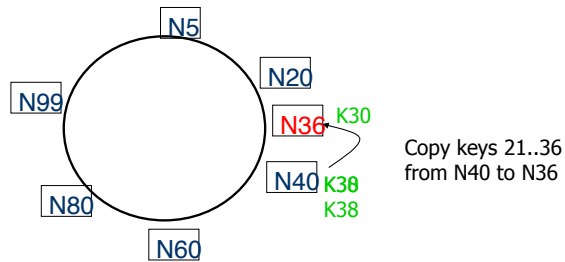
15-441 Fall 2011

© CMU 2005-2011

84

Join: Transfer Keys

- Only keys in the range are transferred



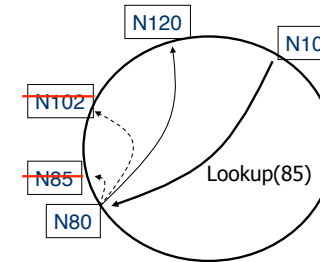
15-441 Fall 2011

© CMU 2005-2011

85

Handling Failures

- Problem:** Failures could cause incorrect lookup
- Solution:** *Fallback:* keep track of a list of immediate successors



15-441 Fall 2011

© CMU 2005-2011

86

Approaches to P2P

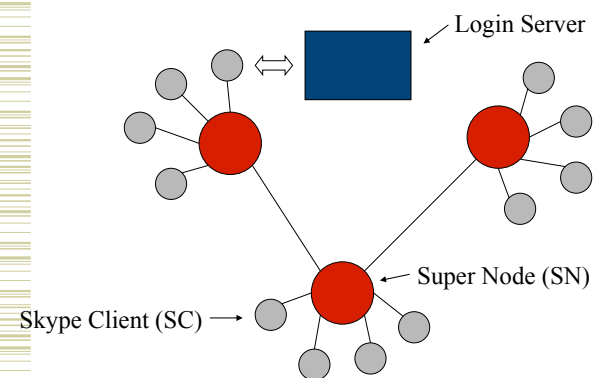
- Centralized
- Flooding
- Supernodes
- Routing
 - Structured
 - Un-structured

15-441 Fall 2011

© CMU 2005-2011

87

Skype Architecture



15-441 Fall 2011

© CMU 2005-2011

88

Routing Queries in Freenet

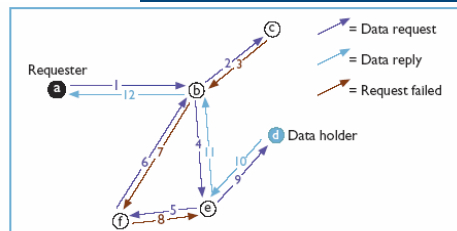


Figure 1. Typical request sequence. The request moves through the network from node to node, backing out of a dead-end (step 3) and a loop (step 7) before locating the desired file.

After success, node a creates a link in its routing table for the key to node d .

Note: alternatively, any node on path from d to a , e. g., e , can name itself as originator of data.

15-441 Fall 2011

© CMU 2005-2011

89



Routing to Mobile Nodes

- Obvious solution: have mobile nodes advertise route to mobile address/32??
- What are some possible solutions?
 - DHCP? (changing IP?)
 - TCP?
 - Learning bridges (e.g., at CMU)
 - Encapsulated PPP
 - Interception & forwarding

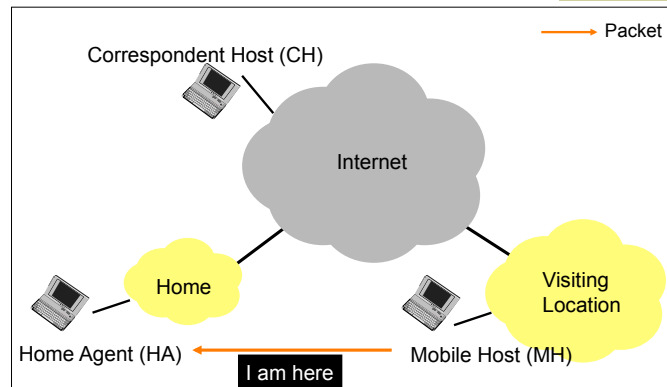
15-441 Fall 2011

© CMU 2005-2011

90



Mobile IP (MH Moving)



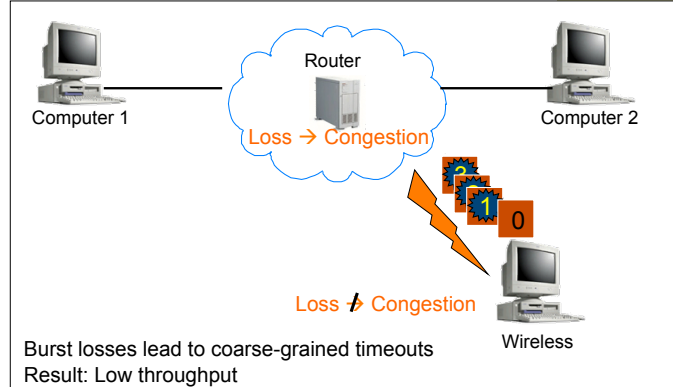
15-441 Fall 2011

© CMU 2005-2011

91



Wireless Bit-Errors



15-441 Fall 2011

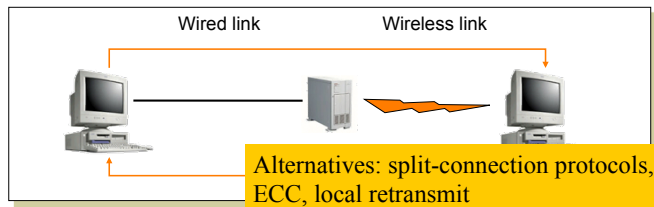
© CMU 2005-2011

92



Approach Styles (End-to-End)

- Improve TCP implementations
 - Not incrementally deployable
 - Improve loss recovery (SACK, NewReno)
 - Help it identify congestion (ELN, ECN)
 - ACKs include flag indicating wireless loss
 - Trick TCP into doing right thing → E.g. send extra dupacks



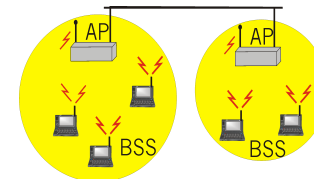
15-441 Fall 2011

© CMU 2005-2011

93

IEEE 802.11 Wireless LAN

- Wireless host communicates with a base station
 - Base station = access point (AP)
- **Basic Service Set (BSS)** (a.k.a. “cell”) contains:
 - **Wireless hosts**
 - **Access point (AP):** base station
- BSS's combined to form distribution system (DS)



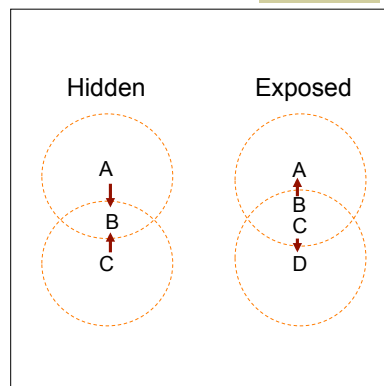
15-441 Fall 2011

© CMU 2005-2011

94

CSMA/CD Does Not Work

- Collision detection problems
 - Relevant contention at the **receiver**, not sender
 - Hidden terminal
 - Exposed terminal
 - Hard to build a radio that can transmit and receive at same time



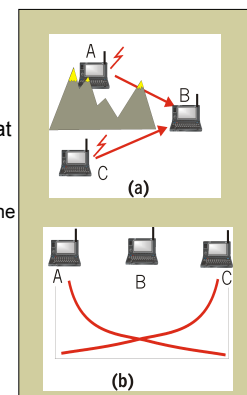
15-441 Fall 2011

© CMU 2005-2011

95

Hidden Terminal Effect

- **Hidden terminals:** A, C cannot hear each other
 - Obstacles, signal attenuation
 - Collisions at B
 - Collision if 2 or more nodes transmit at same time
- CSMA makes sense:
 - Get all the bandwidth if you're the only one transmitting
 - Shouldn't cause a collision if you sense another transmission
- Collision detection doesn't work
- **CSMA/CA: CSMA with Collision Avoidance**



15-441 Fall 2011

© CMU 2005-2011

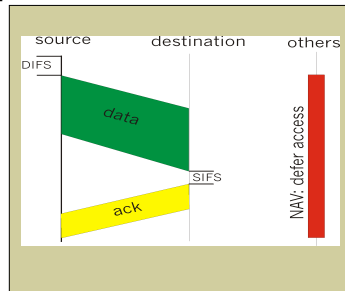
96

IEEE 802.11 MAC Protocol: CSMA/CA



802.11 CSMA: sender

- If sense channel idle for **DIFS** (Distributed Inter Frame Space) then transmit entire frame (no collision detection)
- If sense channel busy then binary backoff



802.11 CSMA receiver:

- If received OK return ACK after **SIFS** (Short IFS) (ACK is needed due to lack of collision detection)

15-441 Fall 2011

© CMU 2005-2011

97

Important Lessons



- Many assumptions built into Internet design
 - Wireless forces reconsideration of issues
- Link-layer
 - Spatial reuse (cellular) vs wires
 - Hidden/exposed terminal
 - CSMA/CA (why CA?) and RTS/CTS
- Network
 - Mobile endpoints – how to route with fixed identifier?
 - Link layer, naming, addressing and routing solutions
 - What are the +/- of each?
- Transport
 - Losses can occur due to corruption as well as congestion
 - Impact on TCP?
 - How to fix this → hide it from TCP or change TCP

15-441 Fall 2011

© CMU 2005-2011

98

Ad Hoc Networks



- All the challenges of wireless, plus:
 - No fixed infrastructure
 - Mobility (on short time scales)
 - Chaotically decentralized
 - Multi-hop!
- Nodes are both traffic sources/sinks and forwarders, no specialized routers
- The biggest challenge: routing

15-441 Fall 2011

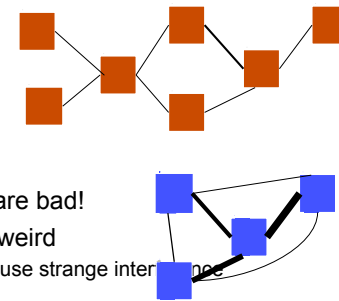
© CMU 2005-2011

99

Traditional Routing vs Ad Hoc



- Traditional network:
 - Well-structured
 - $\sim O(N)$ nodes & links
 - All links work \sim well
- Ad Hoc network
 - $O(N^2)$ links - but most are bad!
 - Topology may be really weird
 - Reflections & multipath cause strange interference
 - Change is frequent



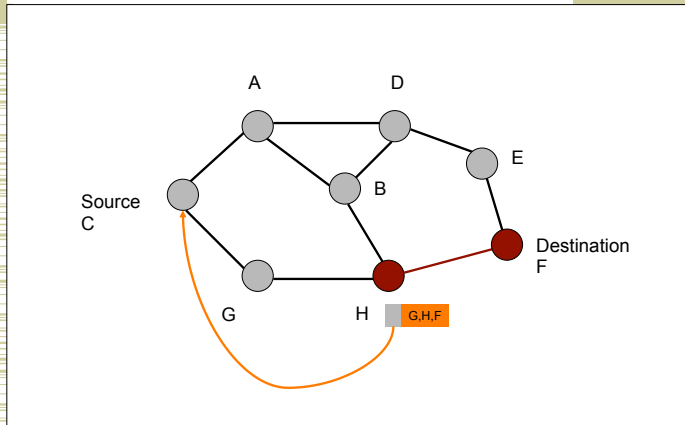
Traditional routing fails: DV loops, LS overhead, updates are power hungry, N^2 links: Instead proposed are: DSDV, AODV, DSR

15-441 Fall 2011

© CMU 2005-2011

100

H Responds to Route Request



15-441 Fall 2011

© CMU 2005-2011

101

Important Lessons

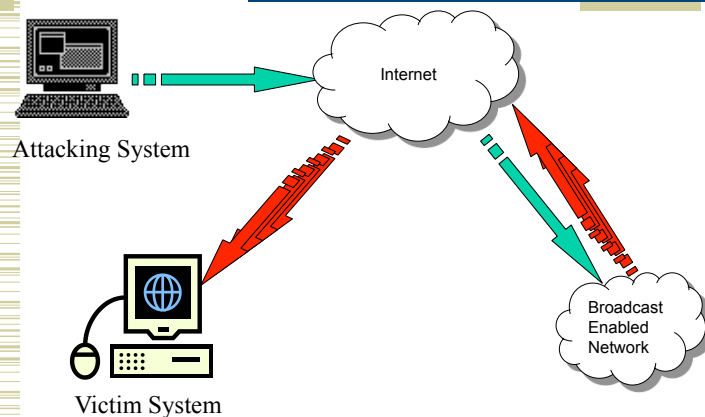
- Wireless is challenging
 - Assumptions made for the wired world don't hold
- Ad-hoc wireless networks
 - Need routing protocol but mobility and limited capacity are problems
 - On demand can reduce load; broadcast reduces overhead
- Special case 1 – Sensor networks
 - Power is key concern
 - Trade communication for computation
- Special case 2 – Vehicular networks
 - No power constraints but high mobility makes routing even harder, geographical routing

15-441 Fall 2011

© CMU 2005-2011

102

Smurf Attack

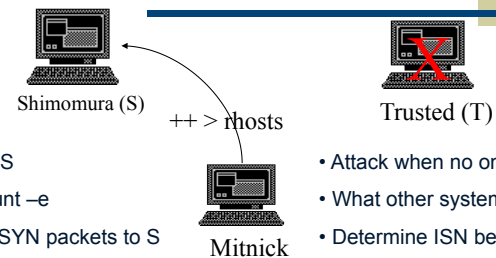


15-441 Fall 2011

© CMU 2005-2011

103

An Example



- Finger @S
- showmount -e
- Send 20 SYN packets to S
- SYN flood T
- Send SYN to S spoofing as T
- Send ACK to S with a guessed number
- Send "echo ++ > ~/.rhosts"

- Attack when no one is around
- What other systems it trusts?
- Determine ISN behavior
- T won't respond to packets
- S assumes that it has a session with T
- Give permission to anyone from anywhere

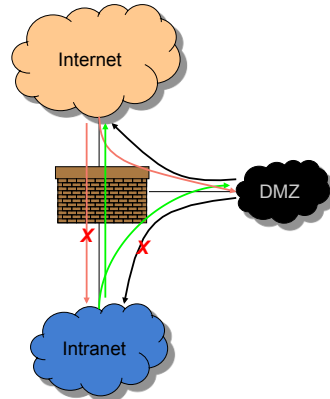
15-441 Fall 2011

© CMU 2005-2011

104

Typical Firewall Configuration

- Internal hosts can access DMZ and Internet
- External hosts can access DMZ only, not Intranet
- DMZ hosts can access Internet only
- Advantages?
 - If a service gets compromised in DMZ it cannot affect internal hosts



15-441 Fall 2011

© CMU 2005-2011

105

Sample Firewall Rule

Allow SSH from external hosts to internal hosts

Two rules

Inbound and out

How to know a p

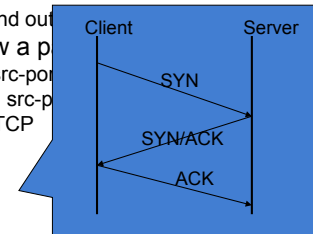
Inbound: src-po

Outbound: src-p

Protocol=TCP

Ack Set?

Problems?



Rule	Dir	Src Addr	Src Port	Dst Addr	Dst Port	Proto	Ack Set?	Action
SSH-1	In	Ext	> 1023	Int	22	TCP	Any	Allow
SSH-2	Out	Int	22	Ext	> 1023	TCP	Yes	Allow

15-441 Fall 2011

© CMU 2005-2011

106

What do we need for a secure comm channel?

- Authentication (Who am I talking to?)
- Confidentiality (Is my data hidden?)
- Integrity (Has my data been modified?)
- Availability (Can I reach the destination?)

15-441 Fall 2011

© CMU 2005-2011

107

The Great Divide

Symmetric Crypto
(Private key)
(E.g., AES)

Asymmetric Crypto
(Public key)
(E.g., RSA)

Shared secret
between parties?

Yes

No

Speed of crypto
operations

Fast

Slow

15-441 Fall 2011

© CMU 2005-2011

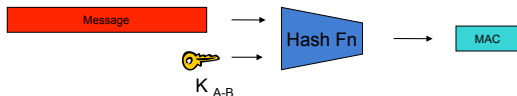
108

Symmetric Key: Integrity

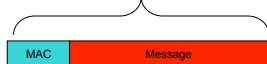
- Hash Message Authentication Code (HMAC)

Step #1:

Alice creates MAC



Step #2 Alice Transmits Message & MAC



Step #3

Bob computes MAC with message and K_{A-B} to verify.

Why is this secure?
How do properties of a hash function help us?

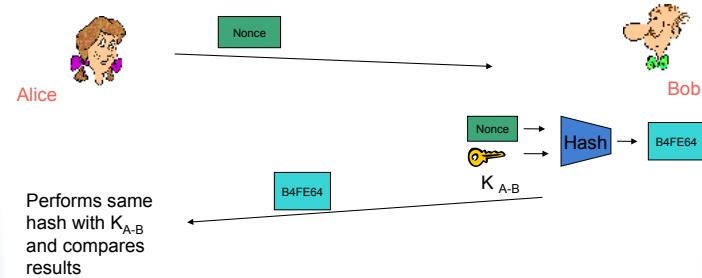
15-441 Fall 2011

© CMU 2005-2011

109

Symmetric Key: Authentication

- A "Nonce"
 - A random bitstring used only once. Alice sends nonce to Bob as a "challenge". Bob Replies with "fresh" MAC result.

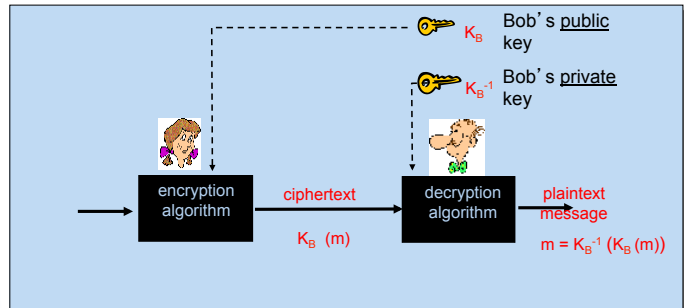


15-441 Fall 2011

© CMU 2005-2011

110

Asymmetric Key: Confidentiality



15-441 Fall 2011

© CMU 2005-2011

111

Asymmetric Key: Integrity & Authentication

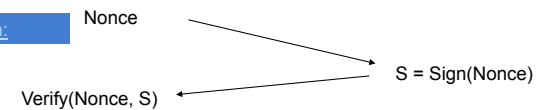
- We can use Sign() and Verify() in a similar manner as our HMAC in symmetric schemes.

Integrity:



Receiver must only check Verify(M, S)

Authentication:



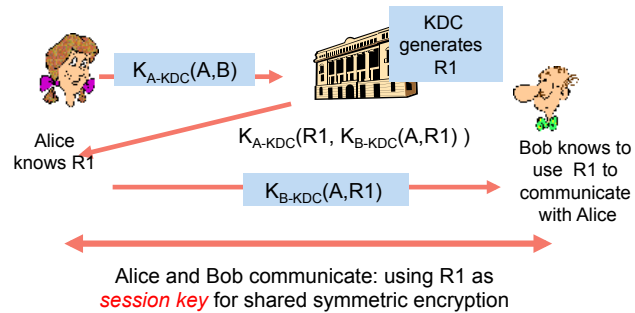
15-441 Fall 2011

© CMU 2005-2011

112

Key Distribution Center (KDC)

Q: How does KDC allow Bob, Alice to determine shared symmetric secret key to communicate with each other?



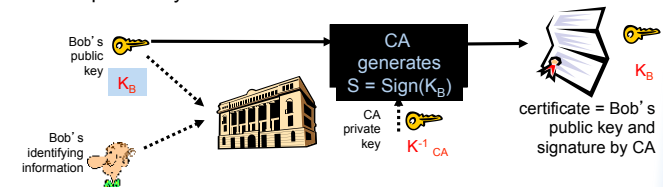
15-441 Fall 2011

© CMU 2005-2011

113

Certification Authorities

- **Certification authority (CA):** binds public key to particular entity, E.
- An entity E registers its public key with CA.
 - ♦ E provides "proof of identity" to CA.
 - ♦ CA creates certificate binding E to its public key.
 - ♦ Certificate contains E's public key AND the CA's signature of E's public key.



15-441 Fall 2011

© CMU 2005-2011

114