

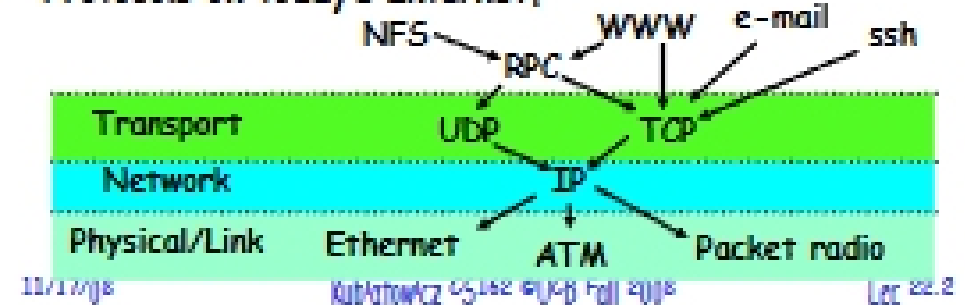
CS162
Operating Systems and
Systems Programming
Lecture 23

Networking III

November 19, 2008
Prof. John Kubiawicz
<http://inst.eecs.berkeley.edu/~cs162>

Review: Network Protocols

- **Protocol:** Agreement between two parties as to how information is to be transmitted
 - Example: system calls are the protocol between the operating system and application
 - Networking examples: many levels
 - Physical level: mechanical and electrical network (e.g. how are 0 and 1 represented)
 - Link level: packet formats/error control (for instance, the CSMA/CD protocol)
 - Network level: network routing, addressing
 - Transport Level: reliable message delivery
- Protocols on today's Internet;



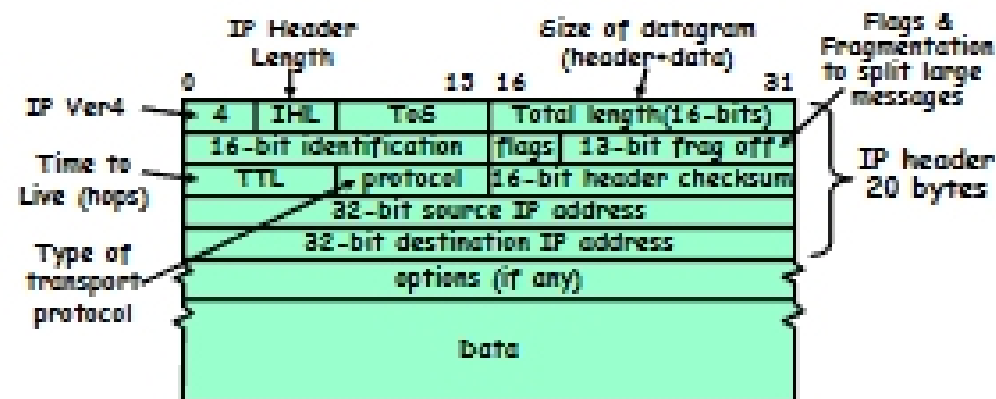
11/17/08

Kubiawicz CS162 @UCB Fall 2008

Lec 22.2

Review: IP Packet Format

- IP Packet Format:



11/17/08

Kubiawicz CS162 @UCB Fall 2008

Lec 22.3

Goals for Today

- Networking
 - Continue discussion of reliable messaging
 - Sequence numbers for ordering
 - Acknowledgments for reliability
- TCP windowing
- Sockets
- Messages
 - Send/receive
 - One vs. two-way communication

Note: Some slides and/or pictures in the following are adapted from slides ©2005 Silberschatz, Galvin, and Gagne. Many slides generated from my lecture notes by Kubiawicz.


11/17/08

Kubiawicz CS162 @UCB Fall 2008

Lec 22.4

Performance Considerations

- Before we continue, need some performance metrics
 - **Overhead:** CPU time to put packet on wire
 - **Throughput:** Maximum number of bytes per second
 - Depends on "wire speed", but limited by slowest router or by congestion
 - **Latency:** time until first bit of packet arrives at receiver
 - Raw transfer time + overhead at each routing hop



- Contributions to Latency
 - Wire latency: depends on speed of light on wire
 - about 1-1.5 ns/foot
 - Router latency: depends on internals of router
 - Could be < 1 ms (for a good router)
 - Question: can router handle full wire throughput?
- What is the End-to-end MTU? **Minimum across path**
- What is the End-to-end Throughput? **Minimum across path**

11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.5

Sample Computations

- E.g.: Ethernet within Soda
 - Latency: speed of light in wire is 1.5ns/foot, which implies latency in building < 1 μ s (if no routers in path)
 - Throughput: 10-1000Mb/s
 - Throughput delay: packet doesn't arrive until all bits
 - Ex: 4KB/100Mb/s = 0.3 milliseconds (same order as disk!)
- E.g.: ATM within Soda
 - Latency (same as above, assuming no routing)
 - Throughput: 155Mb/s
 - Throughput delay: 4KB/155Mb/s = 200 μ s
- E.g.: ATM cross-country
 - Latency (assuming no routing):
 - 3000miles * 5000ft/mile \Rightarrow 15 milliseconds
 - How many bits could be in transit at same time?
 - 15ms * 155Mb/s = 290KB
 - In fact, Berkeley \rightarrow MIT Latency \sim 45ms
 - 872KB in flight if routers have wire-speed throughput
- **Requirements for good performance:**
 - Local area: minimize overhead/improve bandwidth
 - Wide area: keep pipeline full!

11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.6

Sequence Numbers

- **Ordered Messages**
 - Several network services are best constructed by ordered messaging
 - Ask remote machine to first do x, then do y, etc.
 - Unfortunately, underlying network is packet based:
 - Packets are routed one at a time through the network
 - Can take different paths or be delayed individually
 - IP can reorder packets! P_0, P_1 might arrive as P_1, P_0
- Solution requires queuing at destination
 - Need to hold onto packets to undo misordering
 - Total degree of reordering impacts queue size
- Ordered messages on top of unordered ones:
 - Assign sequence numbers to packets
 - 0, 1, 2, 3, 4, ...
 - If packets arrive out of order, reorder before delivering to user application
 - For instance, hold onto #3 until #2 arrives, etc.
 - Sequence numbers are specific to particular connection
 - Reordering among connections normally doesn't matter
 - If restart connection, need to make sure use different range of sequence numbers than previously...

11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.7

Reliable Message Delivery: the Problem


- All physical networks can garble and/or drop packets
 - Physical media: packet not transmitted/received
 - If transmit close to maximum rate, get more throughput - even if some packets get lost
 - If transmit at lowest voltage such that error correction just starts correcting errors, get best power/bit
 - Congestion: no place to put incoming packet
 - Point-to-point network: insufficient queue at switch/router
 - Broadcast link: two host try to use same link
 - In any network: insufficient buffer space at destination
 - Rate mismatch: what if sender send faster than receiver can process?
- **Reliable Message Delivery on top of Unreliable Packets**
 - Need some way to make sure that packets actually make it to receiver
 - Every packet received at least once
 - Every packet received at most once
 - Can combine with ordering: every packet received by process at destination exactly once and in order

11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.8

Using Acknowledgements



- How to ensure transmission of packets?
 - Detect garbling at receiver via checksum, discard if bad
 - Receiver acknowledges (by sending "ack") when packet received properly at destination
 - Timeout at sender: if no ack, retransmit
- Some questions:
 - If the sender doesn't get an ack, does that mean the receiver didn't get the original message?
 - * No
 - What if ack gets dropped? Or if message gets delayed?
 - * Sender doesn't get ack, retransmits. Receiver gets message twice, acks each.


11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.0

How to deal with message duplication

- Solution: put sequence number in message to identify re-transmitted packets
 - Receiver checks for duplicate #'s; Discard if detected
- Requirements:
 - Sender keeps copy of unack'd messages
 - * Easy: only need to buffer messages
 - Receiver tracks possible duplicate messages
 - * Hard: when ok to forget about received message?
- Alternating-bit protocol:
 - Send one message at a time; don't send next message until ack received
 - Sender keeps last message; receiver tracks sequence # of last message received
- Pros: simple, small overhead
- Con: Poor performance
 - Wire can hold multiple messages; want to fill up at (wire latency \times throughput)
- Con: doesn't work if network can delay or duplicate messages arbitrarily



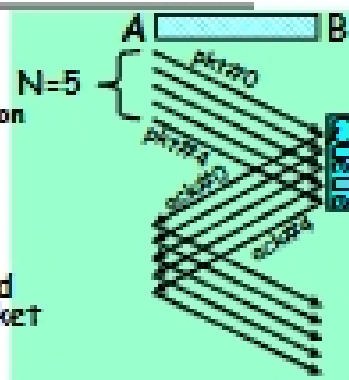
11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.10

Better messaging: Window-based acknowledgements

- Window based protocol:
 - Send up to N packets without ack
 - * Allows pipelining of packets
 - * Window size (N) $<$ queue at destination
 - Each packet has sequence number
 - * Receiver acknowledges each packet
 - * Ack says "received all packets up to sequence number X "/send more
- Acks serve dual purpose:
 - Reliability: Confirming packet received
 - Flow Control: Receiver ready for packet
 - * Remaining space in queue at receiver can be returned with ACK
- What if packet gets garbled/dropped?
 - Sender will timeout waiting for ack packet
 - * Resend missing packets \Rightarrow Receiver gets packets out of order!
 - Should receiver discard packets that arrive out of order?
 - * Simple, but poor performance
 - Alternative: Keep copy until sender fills in missing pieces?
 - * Reduces # of retransmits, but more complex
- What if ack gets garbled/dropped?
 - Timeout and resend just the un-acknowledged packets



11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.11

Administrivia

- Projects:
 - Project 4 design document due Monday, November 24th
- MIDTERM II; Wednesday Dec 3rd
 - Location: 10 Evans, 5:30pm - 8:30pm
 - Topics:
 - * All material from last midterm and up to Monday 12/1
 - * Lectures #13 - 26
 - * One cheat sheet (both sides)
- Final Exam
 - Thursday, Dec 18th, 8:00-11:00am
 - Topics: All Material except last lecture (freebie)
 - Two Cheat sheets.
- Final Topics: Any suggestions?
 - Please send them to me...

11/17/08

Kubiatowicz CS162 @UCB Fall 2008

Lec 22.12