



# CSE 461

MIDTERM REVIEW

# MIDTERM OVERVIEW

- 50 minutes
- Closed book, closed notes
- Covers topics in lectures, projects and homeworks
- Not intended to test things you can easily look up
  - If something seems like you could Google it in a second, ask JZ
- Mixture of straightforward questions and conceptual thought questions
- **Bonus Question:** name both TAs (first and last name) for 461
  - Nat Guy and Paul Vines





# NETWORK LAYERS & ENCAPSULATION



# APPLICATION LAYER

Application

Application

- Used by applications
- Protocol is arbitrary



# TRANSPORT LAYER

Transport

Transport

- Involves packaging of data for transport
- UDP/TCP and ports

# NETWORK LAYERS & ENCAPSULATION

Network

Network

- Handles issues related to routing on the network
- Data treated as packets



# DATA LINK/PHYSICAL LAYERS



- Data link layer
  - Puts data onto the actual line
  - Error-correcting codes to account for line noise are in the data link layer
    - At this level, data consists of frames
- Physical layer
  - Actual electrical or wireless oscillations

# ADDRESSING

- MAC addresses
- IP addresses
- Ports
- Sockets





# MAC ADDRESSES

- 48-bit
- Identify instance of specific network interface hardware

# IP ADDRESSES

- 32-bit (in IPv4) or 128-bit (in IPv6)
- Identify a host on a network
- Can change dynamically
- **Bonus Question** : what does IP stand for?
  - Internet protocol



# PORTS

- 16-bit
- Identify communication channels on a specific host
- Often map to applications
- **Bonus Question** : what application uses port 21?
  - FTP

# SOCKETS

- Programming interface for networking
- Most common implementation is Berkeley sockets
- Allows data to be sent with file descriptor-like structures
- **Bonus Question** : name two valid *type* arguments you can specify for a socket
  - SOCK\_STREAM, SOCK\_DGRAM, SOCK\_SEQPACKET, SOCK\_RAW



# UDP VS. TCP

UDP	TCP
Unreliable	Reliable
Connection-less	Connection-oriented
No acknowledgements	Acknowledgements
No flow control	Sliding window
No sequence numbers	Sequence numbers

# TELNET VS. FTP

Telnet	FTP
Used for sending text data; originally for remote login into a server	Used for sending files
Data transfer and control on same channel	Separate channels for control and data transfer; control channel uses Telnet
Uses TCP	Also uses TCP
Once connected, server and client essentially the same	Server and client behave very differently
Not secure; largely replaced by SSH	Not secure; somewhat replaced by SFTP



# METRICS

- Bandwidth
- Latency
- Throughput, goodput
- Channel utilization
- Shannon's theorem
- Nyquist rate

# FREQUENCY & BANDWIDTH

- Frequency: rate of an oscillation
- Bandwidth: measures the width of a range of frequencies
- $\text{Bandwidth} = \text{freq}_{\text{upper}} - \text{freq}_{\text{lower}}$
- Human hearing bandwidth: ~20kHz (20kHz - 20 Hz)
- “Bandwidth” and “bitrate” are often used interchangeably; this is a different definition
- **Bonus Question:** what’s the frequency range and bandwidth of 802.11 b/g?
  - 2.4 GHz to 2.5 GHz; 100 MHz



# LATENCY

- Time between source and destination
- Shortest possible latency bounded by  $c$
- Ping can measure round-trip latency
- Why might latency vary between ping tests?

# THROUGHPUT & GOODPUT

- Throughput: measures how much data can be sent in a given time period (a.k.a. bitrate)
- E.g., 100 Gbps
- Bits that you can send (i.e., put onto the wire) per amount of time
- Goodput: excludes protocol bits and retransmitted data packets
- What factors might cause goodput < throughput?
  - Protocol overhead
  - Dropped or corrupted packets
  - Flow control



# CHANNEL UTILIZATION

- Calculates how much of the channel is being used
- Percent of the time the channel is in use
- $(\text{sent data size}) / ((\text{channel bitrate}) * (\text{round-trip latency}))$
- If the time to send data is non-negligible, you will want to add the  $(\text{sent data size}) *$
- If you're using stop-and-wait and only sending 1KB at a time over a 1MBps channel with round-trip latency of 10s (and negligible ACK size), what's the channel utilization?
  - $2^{10} / (2^{20} * 10) = \sim 0.01\%$






# SHANNON THEOREM

- Tells about maximum bitrate in the presence of noise
- Capacity = bandwidth \*  $\log_2(1 + \text{signal/noise})$
- $C = B \log_2(1 + S/N)$
- What are the implications of this?





# NYQUIST RATE



- To recover a waveform, the sampling rate must be at least two times the highest frequency
- Telephone sampling rate is 8kHz; what are the implications of this?
- What sampling rate would be required to recover all frequencies audible by humans? (Up to 20kHz)
  - Audio CDs use 44.1kHz sampling rates for this reason

# HTTP

- HTTP 1.0
  - Initial connection over TCP acts as a preamble
  - Content-length can designate payload end
    - Bad for streaming
    - Alternative: drop the connection!
  - Caching used heavily
- HTTP 1.1
  - Data comes as a stream, chunked into defined lengths
  - Connections are reused, reducing overhead
  - Some pipelining possible, but limited (HOL blocking)
- HTTP 2.0
  - Reduces latency through compression
  - Allows asynchronous sending/multiplexing



# ERROR HANDLING


- Parity bits
- Hamming codes
- Checksums
- CRCs
- Bursty errors and error locality





# PARITY BITS

- Bits check parity on a set of bits
- Even parity: bits add to 0
- Odd parity: bits add to 1
- Multiple parity bits (on odd bits/ on even bits, etc.) can increase effectiveness
- **Bonus Question** : What parity bit would need to go in the x to achieve even parity? 0010101x
  - 1





# HAMMING CODES

- An extension of bit parity, where parity check bits are in “powers of two” positions

Bit string:	0	0	1	0	0	0	0	1	0	0	1
Bit number:	1	2	3	4	5	6	7	8	9	10	11
Par/msg:	p	p	m	p	m	m	m	p	m	m	m

- Each data bit is checked (with even parity) by check bits that make up its “power of two” sum
  - E.g., for data bit 7, add to sum for parity bits  $4 + 2 + 1$
  - Possible to recover from single errors
  - See section notes for more details
- Hamming distance: minimum number of bit flips necessary to change one string into another

# CHECKSUMS & CRCS

- Checksums:
  - Adds all words in data as unsigned numbers, allowing to overflow
  - Sum was then compared to check data integrity
- CRCS:
  - Specific type of checksum that uses polynomial division
- Both are integrity checks using a fixed size of data





# IDEMPOTENCE


- Property that an operation can be applied multiple times without changing the result beyond the initial application
- Essential for servers where the same packet may arrive multiple times
- Limerick
  - There once was a misconfigured client
  - Its registrations were quite noncompliant.
  - It sent the same thing 20 times,
  - But the operation worked fine;
  - Idempotent protocols aren't state-reliant!

# CARRIER SENSE MULTIPLE ACCESS WITH COLLISION DETECTION (CSMA/CD)

- Nodes detect traffic on the line and wait to transmit until it's clear
- Collisions are detected, and a random amount of time is waited before a retransmit
- **Bonus question:** Why doesn't Ethernet use this anymore?
  - Hosts can communicate directly with Ethernet switches, completely avoiding collisions.



# HIDDEN TERMINAL PROBLEM



- Issue that can cause wireless interference
- **Bonus Question:** what are some ways to mitigate this?
  - Static
    - Increase both stations' power so they can hear each other
    - Relocate the stations or obstacles
  - Dynamic
    - RTS/CTS

# BINARY EXPONENTIAL BACKOFF

- Wait a random number of slots between 0 and  $n$
- $n$  starts at 1 and doubles each time there's a collision
- Useful for collision-heavy environments
- Stable, because it'll eventually get through






# CODE DIVISION MULTIPLEXING

- All stations send at the same time, with same frequencies
- Each station has a unique chip sequence, all orthogonal to each other
- E.g.,  $(1,1,1,1) \perp (1,1,-1,-1) \perp (1,-1,1,-1)$
- Each station sends that sequence for 1, and the negation of that sequence for 0
- Receiver decodes signal by taking the dot product of the received signal with the chip sequence for each station



# CODE DIVISION MULTIPLEXING (2)

- Signals on the right are all orthogonal (including each one's negation)
- If added on top of each other, they're always separable with the method on the previous slide





# OTHER TOPICS

- Protocol design
  - Data and control channels can be separate or mixed
  - Necessary reliability influences whether data needs to be checked and to what degree
- Sliding window protocol
  - Receiver must receive a certain minimum number of segments before sender can send new data
  - Used in TCP
- Cumulative ACKing
  - ACKing a sequence number means you've received **all** data preceding that sequence number

# PROJECT 1 QUESTIONS

- Could you use cumulative ACKs without breaking the protocol? Would it be useful?
- Why use headers?
- Why does the registration server need to be idempotent?
  - Idempotent: can be applied multiple times without changing the result beyond the initial application



# SAMPLE EXAM QUESTIONS



# SAMPLE QUESTION

Ethernet and 802.11 both support multiple data rates. When an Ethernet cable is plugged into the device, it communicates with the other end, chooses a rate to use, and then sticks with it. 802.11 devices, however, continuously talk to the AP to choose a specific rate to use for the next short while, adjusting that rate up and down as they please.

Why is it a good idea for 802.11 to repeatedly choose transmission rates? Why is it **not** a good idea for Ethernet to do this? (Discuss with someone near you.)



# ANSWER

The 802.11 signal-to-noise ratio can change dramatically over time, which strongly affects the possible transmission rates. Unless 802.11 adapted, it would have to choose between wide coverage at low rates and high rates at low coverage. Dynamic adaptation lets it try to achieve both. Ethernet operates in a much more constrained environment, with strict limits on signal quality imposed by the specification. The environment does not change dynamically.

# SAMPLE QUESTION

A sender and receiver are using the sliding window protocol, with cumulative ACKs. The receiver has a bug where it repeats the last ACK it sent whenever it hasn't received a data frame in the last 10 msec. The connection is bad enough that some data and ACK frames may be lost. Will the sender and receiver achieve reliable transmission? Will this protocol fail? ? (Discuss with someone near you.)



# ANSWER

They'll achieve reliable transmission, assuming no bugs other than those described. The extra ACKs will simply repeat a previously sent ACK. They won't affect the sender's window; the sender will view them as duplicate ACKs and drop them.

# SAMPLE QUESTION

In Project 1 your code sent UDP packets from a client to a server. If we were to look at the bits actually being carried on the wire (assuming we're using a wired network), we'd find the destination IP address and port were part of the bits being sent. Did the bits on the wire also carry a destination MAC address, or not?



# ANSWER

Yes. All data delivery happens at the link layer, which uses MAC addresses. Higher level protocols, like UDP, are encapsulated in link layer frames.

# SAMPLE QUESTION

Sliding window protocols specify a sending window and a receiving window. Can it ever be useful for the sending window to be larger than the receiving window? Briefly explain your answer.



# ANSWER

Yes, possibly. The sending window limits the amount of data that can be in flight but unacknowledged. The receiver window is a limit on the memory available for buffering and reordering. If the receiver is capable of consuming each incoming frame basically as it arrives, it could have a very small receive window. In that case, the send window could be larger than the receive, to allow frames to be in flight (on the wire).

# ADDITIONAL STUDY SUGGESTIONS

- Read through your project 0 & 1 code and diagram what it's doing
- Review HW problems; do similar problems
- Watch David Wetherall's Coursera course videos ([link](#))
- Review old midterms
  - Available on CSE site (but cover somewhat different material than what ours will)



ANY QUESTIONS?

