

Ask test questions to Chris

2 polynomial generators

Scheme 1 G_1 $h_0 h_1 h_2 h_3$
 1101 G_2 1110

$x[0] \ x[1] \ \dots \ x[n] \ \dots$

$P_1(0) \ \dots \ P_1(n)$

$P_2(0) \ \dots \ P_2(n)$

$$P_1[n] = x[n] \cdot h_0 + x[n-1] \cdot h_1 + x[n-2] \cdot h_2 + x[n-3] \cdot h_3$$

$$= x[n] \cdot 1 + x[n-1] \cdot 1 + x[n-2] \cdot 0 + x[n-3] \cdot 0$$

addition
mod 2

$$P_2[n] = x[n] \oplus x[n-1] \oplus x[n-2]$$

- Uses G_2

$1+1 \text{ mod } 2$
 add normally
 then divide by 2
 take remainder

for each input bit/finestep transmit $P_1 P_2$

Scheme 2 G_1 110101 G_2 111011

$0+0=0$
 $0+1=1$
 $1+0=1$
 $1+1=0$
 XOR

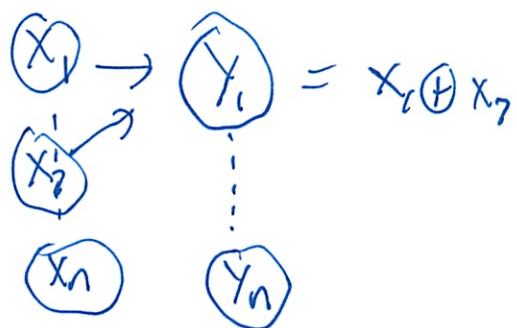
$$P_1[n] = x[n] \oplus x[n-1] \oplus x[n-3] \oplus x[n-5]$$

$$P_2[n] = x[n] \oplus x[n-1] \oplus x[n-2] \oplus x[n-4] \oplus x[n-5]$$

② Digression

Is found empirically by testing
for fixed length of code

How to find empirically



Can put all sorts of edges in

Art is deciding which edges to put in

"Expanders"

Won't go into details how to find them

What is the rate of the two codes?

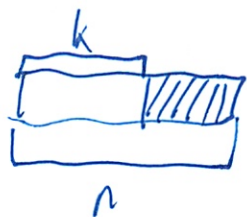
Scheme 1 $\left\{ \begin{array}{l} (n, k, d) \\ \text{rate} = k/n \end{array} \right.$

$$\begin{aligned} k &= k \\ n &= 2k \end{aligned}$$

rate $\frac{1}{2}$ since 2 bits for every message bit transmitted

③ Scheme 2

$$\text{rate} = \frac{k}{n}$$



$$k = k$$

$$n = 2k$$

$$(2k, k)$$

↑ can find

find non zero codeword w/ min # of 1s
see last citation

$$\text{rate} = \frac{k}{2k} = \frac{1}{2}$$

k = Constraint Length

The length of ^m polynomials

how much history are you involving

Scheme 1 $G_{01} = 4 \rightarrow$ since 0, 1, 2, 3

$$G_{02} = 4$$

↑ ^{for constraint length for scheme}
max of any polynomial
max = 4

Scheme 2 $G_{01} = 6$
 $G_{02} = 6$

$$L_{\text{max}} = 6$$

9

Which code leads to better performance (lower BER)

- not precise

- touchy feely

Scheme 2 has more history

- more error correction possibility

A good G w/ shorter constraint length is better than
bad G w/ longer " "

d) Someone says add G_3 . New rate:

$$\text{Rate} = \frac{1}{3}$$

Constraint length = 4

Certainly not weaker

Adding more structure - so more error correction

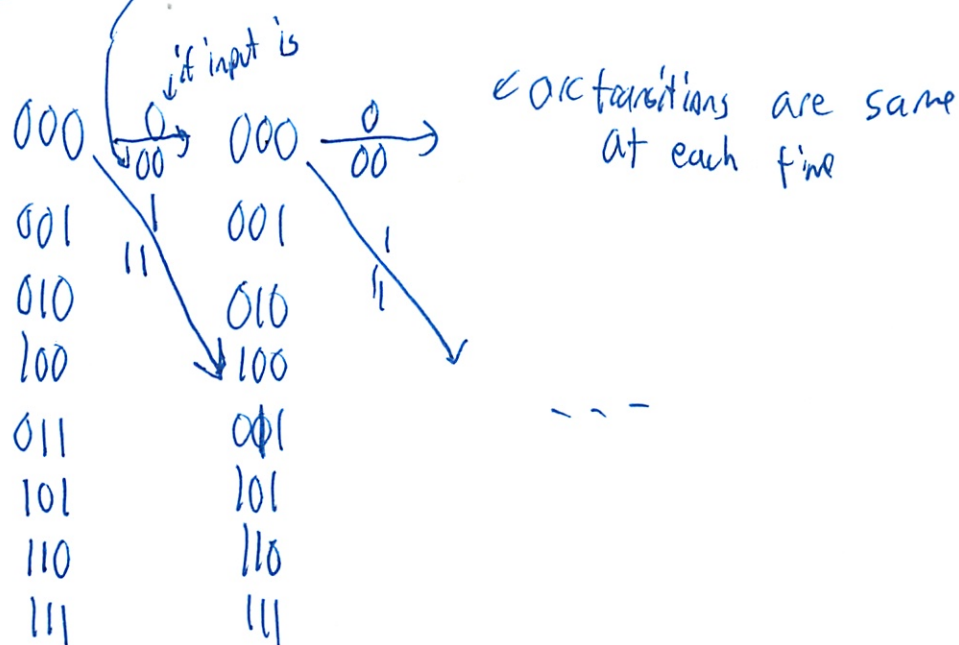
At cost of rate/bandwidth

Decoding



(5)

2^{k-1} Parity bits $x[n] + x[n-1] + x[n-2]$ for the state it is going towards



Verterbi - Dynamic Programming

On tree graph

Belief prediction

A different convolution code

$$G_0 = 111 \quad G_1 = 110$$

Construct state diagram

$$P_1[n] = x[n] \oplus x[n-1] + x[n-2]$$

$$P_2[n] = x[n] \oplus x[n-1]$$

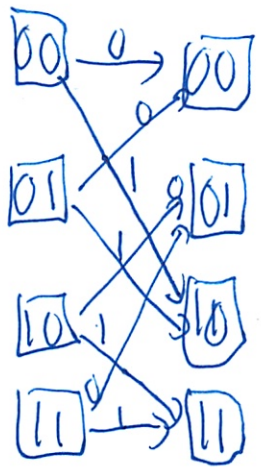
$$\text{Rate} = \frac{1}{2}$$

$$\text{Constraint length} = \text{---} = 3$$

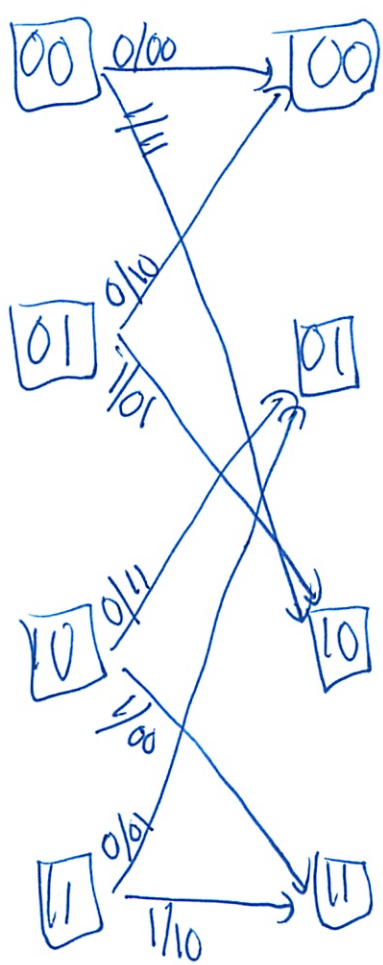
$$\text{History} = 2^{k-1} = 4 \text{ states possible}$$

6

Need 2 bits of history + current bit



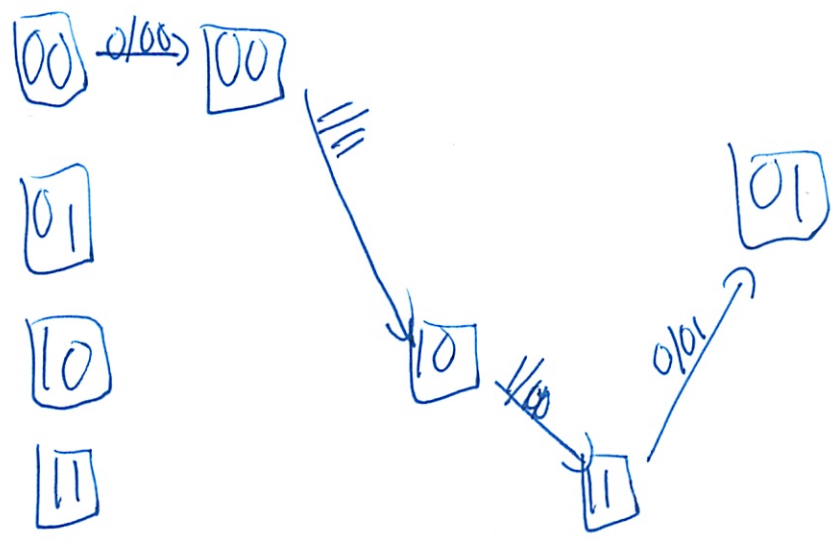
Rewrite better



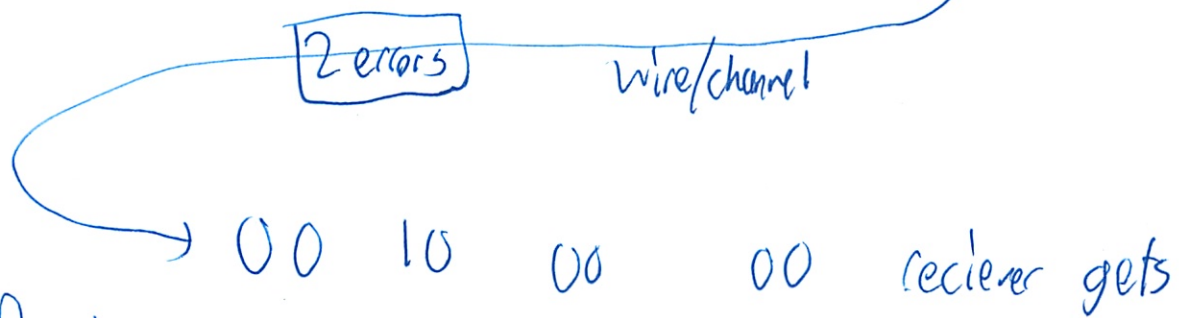
7

Encode
~~Remember~~ 0110 ← message

Remember start w/ 00 state



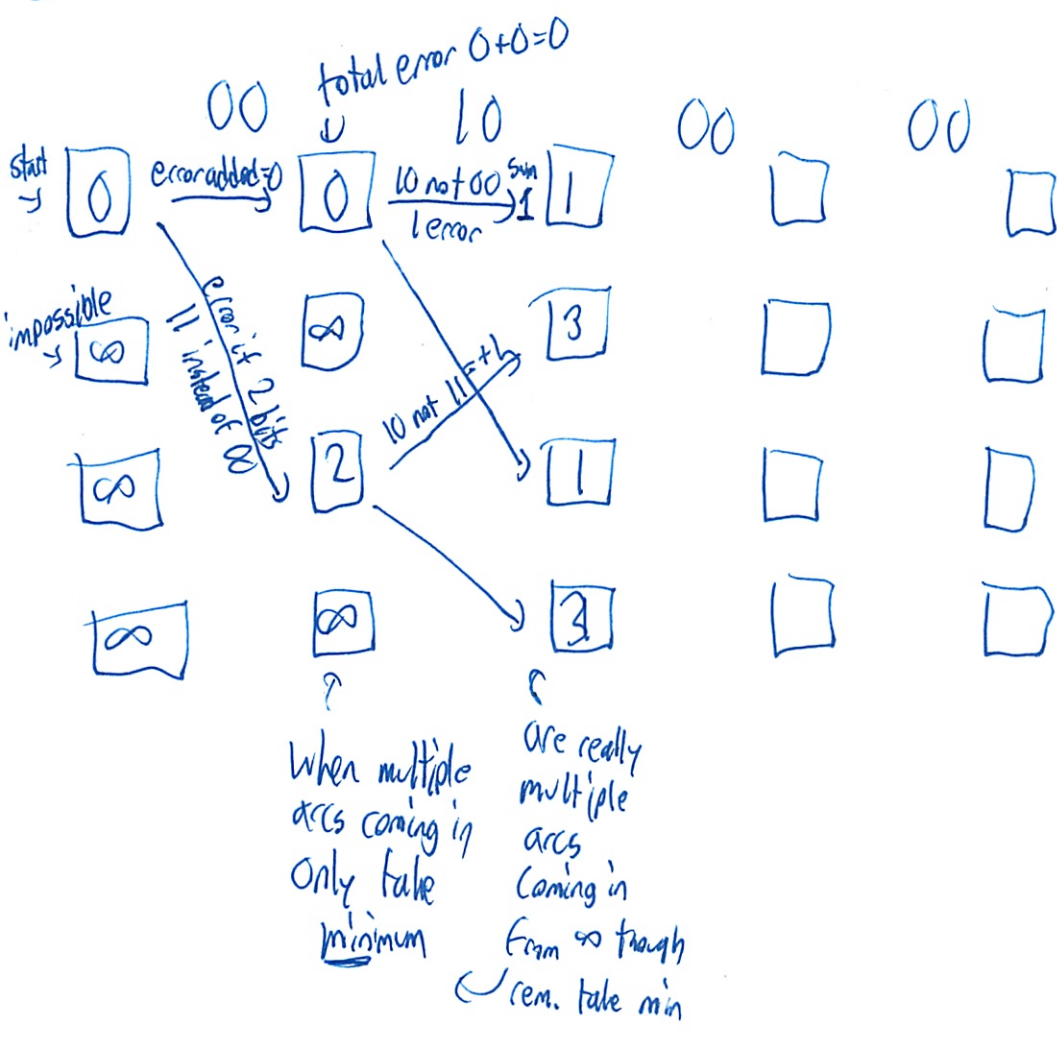
Tx	00	11	00	01
State	00	10	11	01



Decode (Shortest path algorithm)



8



Arcs tell you what bits were transmitted

Interview 4

3/11

St dev - not abs value

- otherwise only half dist

- not as large as normal since only half



Consider 3×3 or 4×4

- can only fix 1 error

- hamming $d = 3$

- or SEC 2

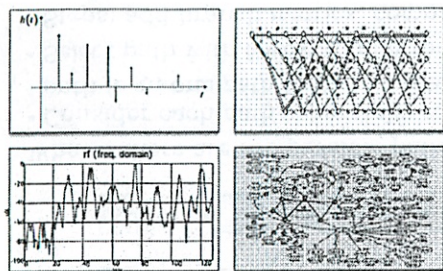


- can extend out to multiple errors

but harder to code

Lots of material

Classes not any more - (6.45) old videos



INTRODUCTION TO EECS II DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2011 Lecture #11

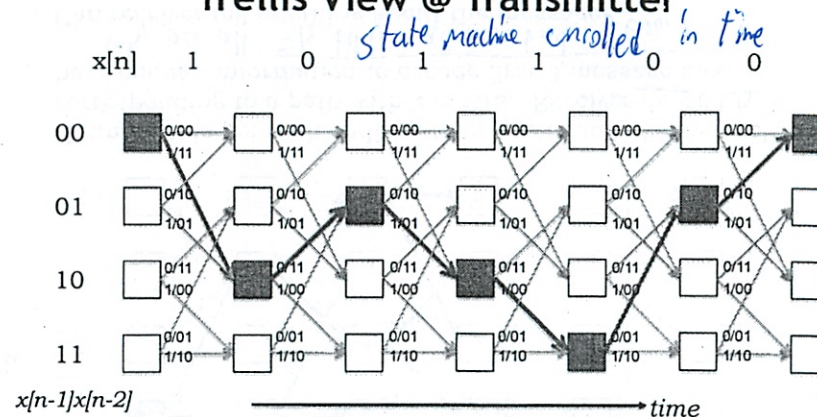
- state machines & trellises
- path and branch metrics
- Viterbi algorithm: add-compare-select
- hard decisions vs. soft decisions

6.02 Spring 2011

Lecture 11, Slide #1

Last time in 6.02...

Trellis View @ Transmitter



Send: 11 11 01 00 01 10

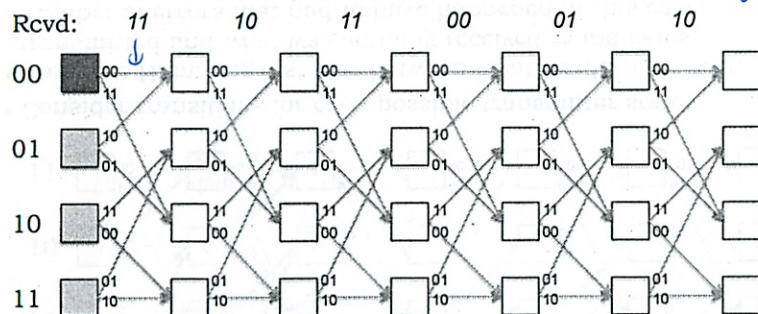
Current state \rightarrow
old 2 recent bits

6.02 Spring 2011

Lecture 11, Slide #2

Receiver

Finding the Most-likely Path

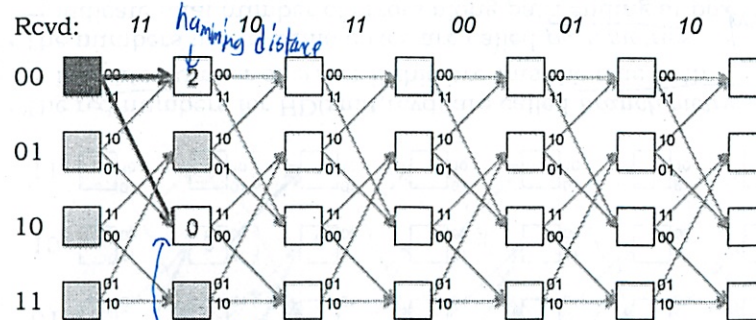


Given the received parity bits, the receiver must find the most-likely sequence of transmitter states, i.e., the path through the trellis that minimizes the Hamming distance between the received parity bits and the parity bits the transmitter would have sent had it followed that state sequence.

6.02 Spring 2011

Lecture 11, Slide #3

Processing 1st pair of parity bits



never reached

- Transmitter is known to be in state 00 at start of message.
- If next message bit 0, next state is 00, transmit 00
 - So receiving 11 means there have been two errors in the channel
- If next message bit is 1, next state is 10, transmit 11
 - So receiving 11 means there have been no errors in the channel

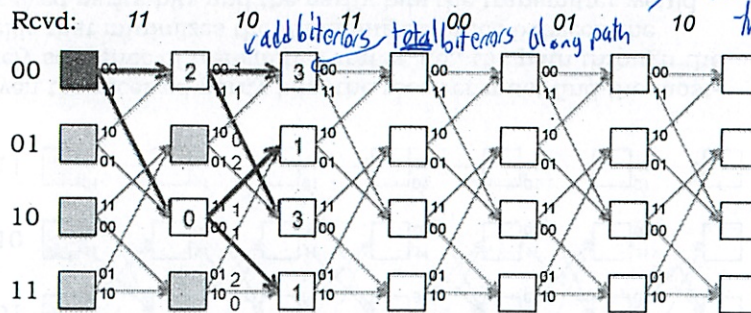
6.02 Spring 2011

but can't fully be sure here

Lecture 11, Slide #4

3/10

Processing 2nd pair of parity bits

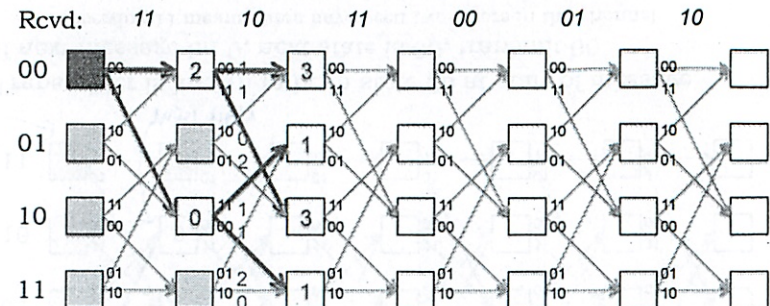


- Consider transitions for each possible transmitter state
- Compute Hamming distance between what would have been transmitted and what was actually received → indicates number of errors that had to have happened in this case.
- Enter total errors along path in each destination state in next column of the trellis, color transition arc red

6.02 Spring 2011

Lecture 11, Slide #5

Branch Metrics, Path Metrics



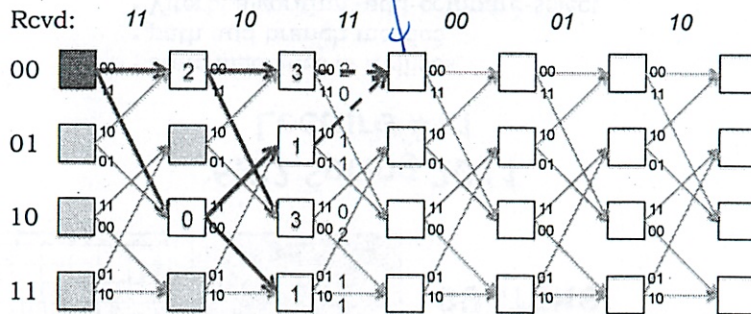
- The red numbers for $HD(x_{mit}, r_{cvd})$ are called branch metrics → indicate number of errors if this arc was the true path
- The numbers in the trellis boxes are called path metrics → indicate total number of errors along path ending at box
- The red arrows indicate sequence of transmitter states that end at a particular column and state.

6.02 Spring 2011

Lecture 11, Slide #6

say $HD(00, 01) = 1$

Processing 3rd pair of parity bits



What if there are two possible paths to a particular state?

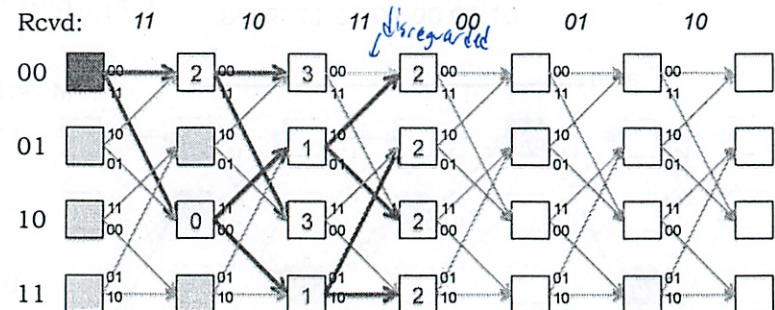
- Consider each path separately: compute total errors along each path (e.g., one path to 00 has 5 total errors, the other 2 errors)
- Select path with fewest total errors as the most-likely path
- Steps: add branch metrics, compare total errors, select path

6.02 Spring 2011

Lecture 11, Slide #7

add, compare, select

What Can We Tell About Message?



Hmm, at this point all ending states are equally likely, each corresponding to a path with 2 errors. Receiver doesn't (yet) have enough information to decode first 3 message bits.

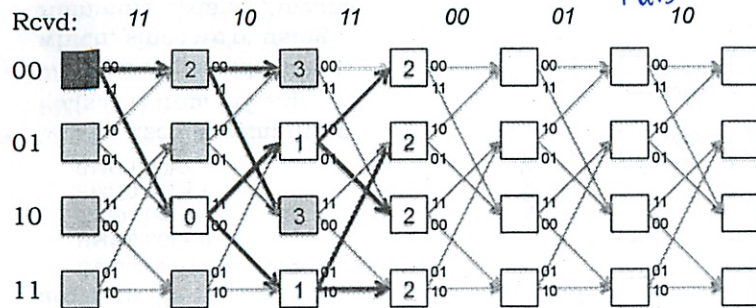
Can receiver tell anything about the message?

6.02 Spring 2011

Lecture 11, Slide #8

Survivor Paths

(Can start eliminating stuff that's not likely)



Color boxes grey and remove arrows (in next slide)

Receiver can make some deductions:

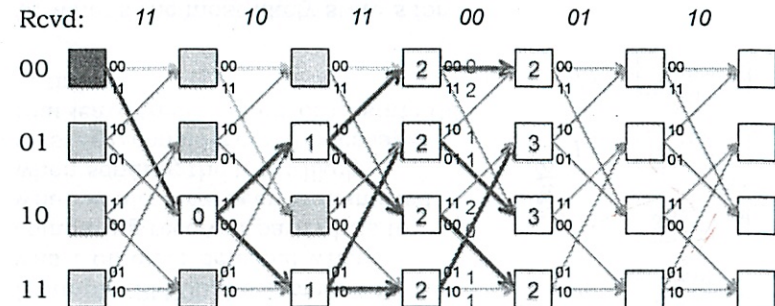
Some earlier states are no longer part of any most-likely path. We can eliminate partial paths leading to those states since they will not be part of the final most-likely path. Do this recursively...

6.02 Spring 2011

Lecture 11, Slide #9

that go to a higher than used/min state and anything before that if that is its exclusive arrow!

Processing 4th pair of parity bits



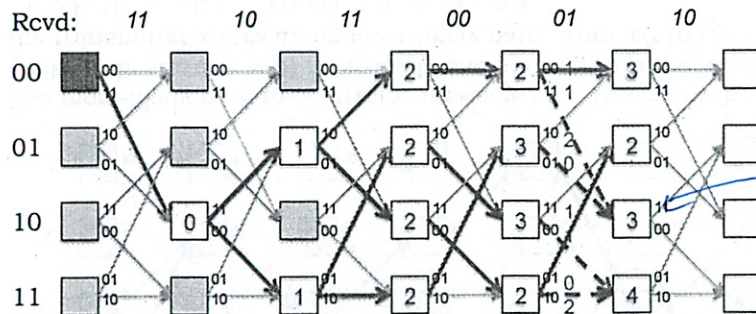
The usual: ↑ Starting to see message!

- Add branch metrics to previous total errors
- Compare paths arriving at destination state
- Select path with smallest total errors as most-likely path

6.02 Spring 2011

Lecture 11, Slide #10

Processing 5th pair of parity bits



what if tie?

When there are "ties" (sum of metrics are the same)

- Make an arbitrary choice about incoming path
- If state is not on most-likely path: choice doesn't matter
- If state is on most-likely path: choice may matter and error correction has failed

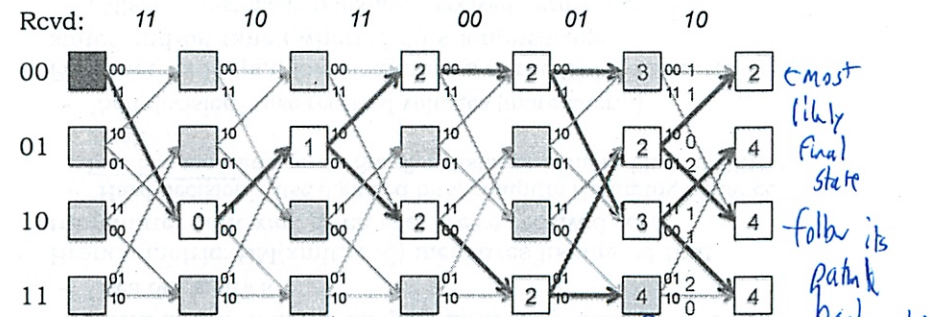
add, compare, select

6.02 Spring 2011

Lecture 11, Slide #11

pick one randomly - guessing

Processing 6th pair of parity bits



most likely final state follow its path backwards!

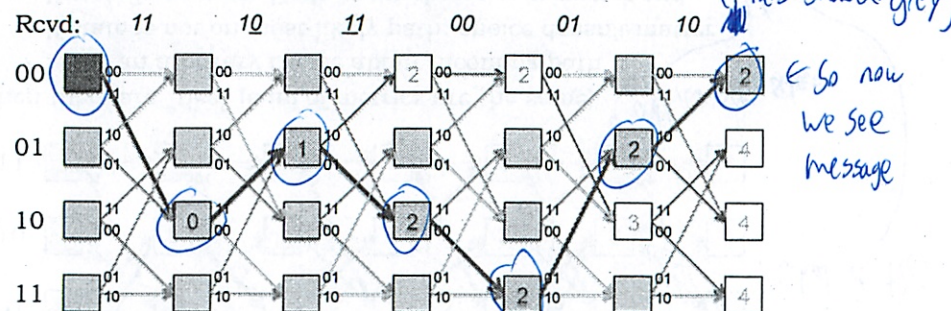
When we reach end of received parity bits:

- Each state's path metric indicates how many errors have happened on most-likely path to state
- Most-likely final state has smallest path metric
- Ties mean end of message uncertain (but survivor paths may merge to a unique path earlier in message)

6.02 Spring 2011

Lecture 11, Slide #12

Traceback



Use most-likely path to determine message bits, tracing path backwards starting with most-likely end state. In this example, the transmitter states along most-likely path, from left to right:

Write down states 10 01 10 11 01 00

Message bits from high-order state bit:

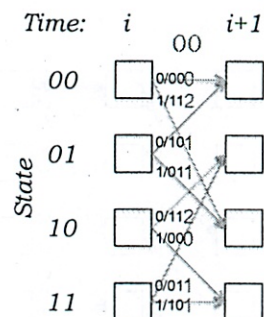
101100

Viterbi Algorithm

- Want: Most likely message sequence
- Have: (possibly corrupted) received parity sequences
- Viterbi algorithm for a given k and r :
 - Works incrementally to compute most likely message sequence
 - Uses two metrics
- Branch metric: $BM(xmit, rcvd)$ measures likelihood that transmitter sent $xmit$ given that we've received $rcvd$.
 - “Hard decision”: use digitized bits, compute Hamming distance between $xmit$ and $rcvd$. Smaller distance is more likely if BER is small
 - “Soft decision”: use received voltages (more later...)
- Path metric: $PM[s, i]$ for each state s of the 2^{k-1} transmitter states and bit time i where $0 \leq i < \text{len}(\text{message})$.
 - $PM[s, i] = \text{most likely } BM(xmit_m, \text{received parity}) \text{ over all message sequences } m \text{ that leave transmitter in state } s \text{ at time } i$
 - $PM[s, i+1]$ computed from $PM[s, i]$ and $p_0[i], \dots, p_{r-1}[i]$

Hard-decision Branch Metric

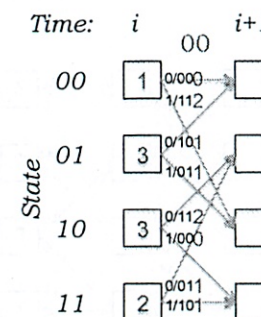
- BM = Hamming distance between expected parity bits and received parity bits
- Compute BM for each transition arc in trellis
 - Example: received parity = 00
 - $BM(00, 00) = 0$
 - $BM(01, 00) = 1$
 - $BM(10, 00) = 1$
 - $BM(11, 00) = 2$
- Will be used in computing $PM[s, i+1]$ from $PM[s, i]$.
- We'll want most likely BM, which, since we're using Hamming distance, means minimum BM.



Computing $PM[s, i+1]$

Starting point: we've computed $PM[s, i]$, shown graphically as label in trellis box for each state at time i .

Example: $PM[00, i] = 1$ means there was 1 bit error detected when comparing received parity bits to what would have been transmitted when sending the most likely message, considering all messages that leave the transmitter in state 00 at time i .



Q: What's the most likely state s for the transmitter at time i ?

A: state 00 (smallest $PM[s, i]$)

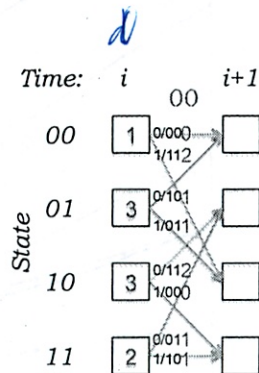
Computing $PM[s, i+1]$ cont'd.

Q: If the transmitter is in state s at time $i+1$, what state(s) could it have been in at time i ?

A: For each state s , there are two predecessor states α and β in the trellis diagram

Example: for state 01, $\alpha=10$ and $\beta=11$.

Any message sequence that leaves the transmitter in state s at time $i+1$ must have left the transmitter in state α or state β at time i .



4 diff possible last 2 bits (constraint level of 3)

$k=4$ 8 boxes

$k=15$ 2¹⁴ boxes

twice work
grows exponentially w/ constraint length

Computing $PM[s, i+1]$ cont'd.

Formalizing the computation:

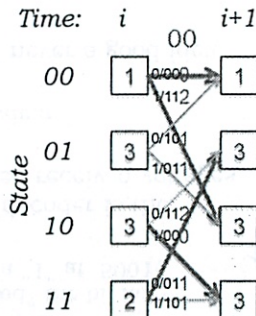
$$PM[s, i+1] = \min(PM[\alpha, i] + BM[\alpha \rightarrow s], PM[\beta, i] + BM[\beta \rightarrow s])$$

Example:

$$PM[01, i+1] = \min(PM[10, i] + 2, PM[11, i] + 1) = \min(3 + 2, 2 + 1) = 3$$

Notes:

- 1) Remember with arc was min; saved arcs will form a path through trellis
- 2) If both arcs have same sum, break tie arbitrarily (e.g., when computing $PM[11, i+1]$)

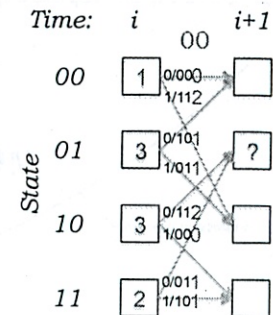


Computing $PM[s, i+1]$ cont'd.

Example cont'd: to arrive in state 01 at time $i+1$, either

- 1) The transmitter was in state 10 at time i and the i^{th} message bit was a 0. If that's the case, the transmitter sent 11 as the parity bits and there were 2 bit errors since we received 00. Total bit errors = $PM[10, i] + 2 = 5$ OR
- 2) The transmitter was in state 11 at time i and the i^{th} message bit was a 0. If that's the case, the transmitter sent 01 as the parity bits and there was 1 bit error since we received 00. Total bit errors = $PM[11, i] + 1 = 3$

Which is mostly likely?



Viterbi Algorithm Summary

- Branch metrics measure the likelihood by comparing received parity bits to possible transmitted parity bits computed from possible messages.
- Path metric $PM[s, i]$ measures the likelihood of the transmitter being in state s at time i assuming the mostly likely message of length i that leaves the transmitter in state s .
- Most likely message? The one that produces the most likely $PM[s, N]$.
- At any given time there are 2^{k-1} most-likely messages we're tracking \rightarrow time complexity of algorithm grows exponentially with constraint length k .

Hard Decisions

Hamming distance

- As we receive each bit it's immediately digitized to "0" or "1" by comparing it against a threshold voltage
 - We lose the information about how "good" the bit is: a "1" at .9999V is treated the same as a "1" at .5001V
- The branch metric used in the Viterbi decoder is the Hamming distance between the digitized received voltages and the expected parity bits
 - This is called *hard-decision Viterbi decoding*
- Throwing away information is (almost) never a good idea when making decisions
 - Can we come up with a better branch metric that uses more information about the received voltages?

throws away info vs

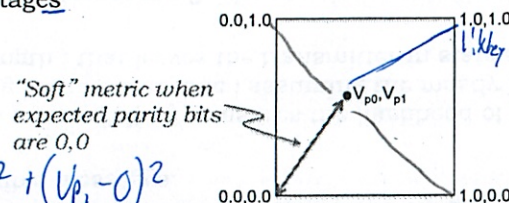
V_{th} ~~31~~ 30

Soft Decisions

looks closer at Voltage .50001

both "1" but different

- Let's limit the received voltage range to [0.0, 1.0]
 - $V_{eff} = \max(0.0, \min(1.0, V_{received}))$
 - Voltages outside this range are "good" 0's or 1's
- Define our "soft" branch metric as the square of the Euclidian distance between received V_{eff} and expected voltages



"Soft" metric when expected parity bits are 0,0

$$(V_{p0}-0)^2 + (V_{p1}-0)^2$$

No sq. root!

Soft-decision decoder chooses path that minimizes sum of the squares of the Euclidian distances between received and expected voltages

- Different branch metric but otherwise the same recipe

6.02 Spring 2011

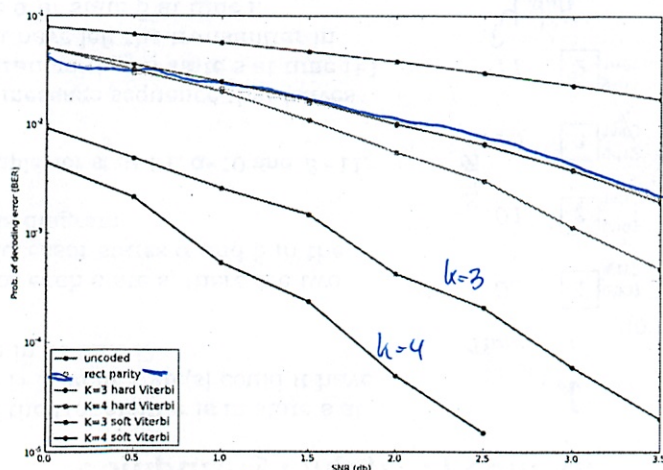
Lecture 11, Slide #22

Adding the folding pt # not hamming d'istance
just a few lines different

6.02 Spring 2011

Lecture 11, Slide #21

More Work, Better BER

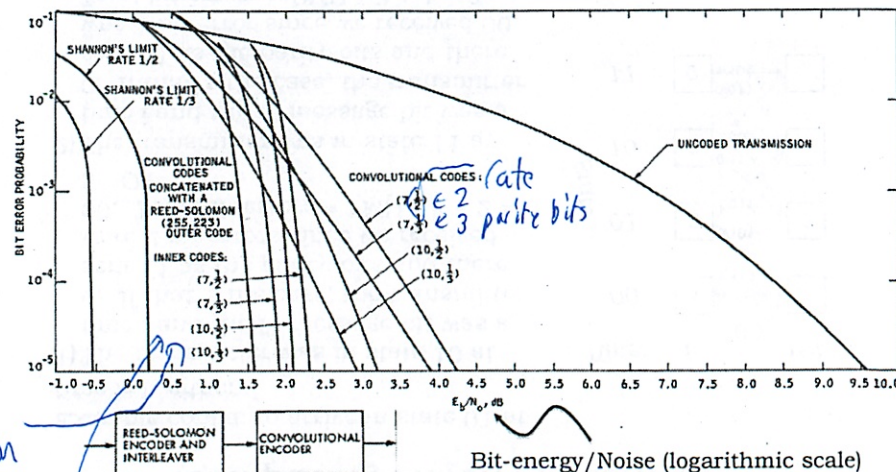


k = # of places using a bit

6.02 Spring 2011

Lecture 11, Slide #23

Code performance

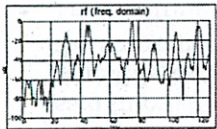
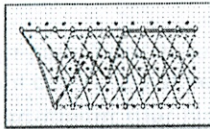
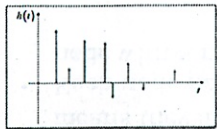


new generation here
turbo codes

Source: Butman, Deutsch, Miller, "Performance of Concatenated Codes for Deep Space Missions"
6.02 Spring 2011

Lecture 11, Slide #24

Conv. bad at burst errors



How to do sharing
of channel

INTRODUCTION TO SECS II DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2011 Lecture #12

- shared medium → media access protocol
- time division multiple access (TDMA)
- contention protocols, Alohanet

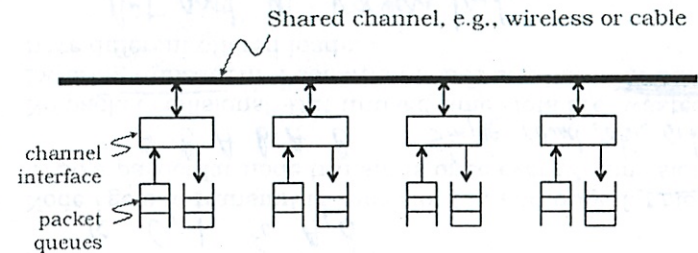
time
sharing

Ethernet was
originally shared
Now mostly
point to point
Wireless is
shared

6.02 Spring 2011

Lecture 12, Slide #1

Shared Communications Channels



- Basic constraint: avoid collisions between transmitters
 - Collisions can be detected from corrupted packets
- Wanted: a communications protocol ("rules of engagement") that ensures good performance overall

what qualifies as good?

6.02 Spring 2011

Lecture 12, Slide #2

"Good" Performance

- High utilization
 - Channel capacity is a limited resource → use it efficiently
 - Ideal: use 100% of channel capacity in transmitting packets
 - Waste: idle periods, collisions, protocol overhead
- Fairness
 - Divide capacity equally among requesters
 - But not every node is requesting all the time...
- Bounded wait
 - An upper bound on the wait before successful transmission
 - Important for isochronous communications (e.g., voice/video)
- Scalability
 - Accommodate changing number of nodes, hopefully without changing implementation of any given node

don't want
overhead of
"raise your hand
to transmit"

everyone can
talk

can't just divide evenly

6.02 Spring 2011

Lecture 12, Slide #3

Sharing Protocols

- Protocols ≡ "rules of engagement" for good performance
 - Known as media access control (MAC) or multiple access control
- Time division multiple
 - Share time "slots" between requesters
 - Prearranged: time division multiple access (TDMA) *rand robin*
 - Not prearranged: contention protocols (e.g., Alohanet). These are interesting because each node operates independently
- Frequency division
 - Give each transmitter its own frequency, receivers choose "station"
 - Our topic for after spring break
- Code division
 - Uses unique orthogonal pseudorandom code for each transmitter
 - Channel adds transmissions to create combined signal
 - Receiver listens to one "dimension" of combined signal using dot product of code with combined signal.

6.02 Spring 2011

Lecture 12, Slide #4

separate x, y, z channel dimension
separate out w/ dot product

3/14

Utilization

Useful

- Utilization measures the throughput of a channel:

$$U_{\text{channel}} = \frac{\text{total throughput over all nodes}}{\text{maximum data rate of channel}}$$

- Example: 10 Mbps channel, four nodes get throughputs of 1, 2, 2 and 3 Mbps. So utilization is $(1+2+2+3)/10 = 0.8$.
- $0 \leq U \leq 1$. Utilization can be less than 1 if
 - The nodes have packets to transmit (nodes with packets in their transmit queues are termed *backlogged*), but the protocol is inefficient.
 - There is insufficient *offered load*, i.e., there aren't enough packets to transmit to use the full capacity of the channel. *have stuff to send if turn*
- With backlogged nodes, perfect utilization is easy: just let one node transmit all the time! But that wouldn't be fair... *usually we assume people could make full use of channel*

Fairness

- Many plausible definitions. A standard recipe:
 - Measure throughput of nodes = x_i , over a given time interval
 - Say that a distribution with lower standard deviation is "fairer" than a distribution with higher standard deviation.
 - Given number of nodes, N , fairness F is defined as

index

$$F = \frac{\left(\sum_{i=1}^N x_i\right)^2}{N \sum_{i=1}^N x_i^2}$$

- $1/N \leq F \leq 1$, where $F=1/N$ implies single node gets all the throughput and $F=1$ implies perfect fairness.
- We'll see that there is often a tradeoff between fairness and utilization, i.e., fairness mechanisms often impose some overhead, reducing utilization.

6.02 Spring 2011

Lecture 12, Slide #5

6.02 Spring 2011

Lecture 12, Slide #6

if 3, 3, 3, 3

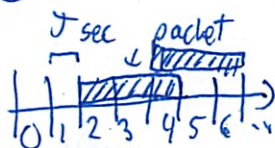
$$F_{\text{then}} = \frac{12^2}{4 \cdot (9+9+9+9)} = 1$$

if 3, 0, 0, 0

$$F = \frac{9}{4 \cdot 9} = \frac{1}{4}$$

Abstraction for Shared Medium

- Time is divided into *slots* of equal length
- Each node can start transmission of a packet only at the beginning of a time slot
- All packets are of the same size and hence take the same amount of time to transmit, equal to some integral multiple of time slots.
- If the transmissions of two or more nodes overlap, they are said to *collide* and none of the packets are received correctly. Note that even if the collision involves only part of the packet, the entire packet is assumed to be lost.
- Transmitting nodes can detect collisions, which usually means they'll retransmit that packet at some later time.
- Each node has a queue of packets awaiting transmission. A node with a non-empty queue is said to be *backlogged*.



Can only start at a start of time slot and packet is multiple of T

Overlap - so both packets failed

Nodes know they failed, retransmit

Time Division Multiple Access (TDMA)

- Suppose that there is a centralized resource allocator and a way to ensure time synchronization between the nodes - for example, a cellular base station.
- For N nodes, give each node a unique index in the range $[0, N-1]$. Assume each slot is numbered starting at 0. *gives it the time points*
- Node i gets to transmit in time slot t if, and only if, $t \bmod N = i$. So a particular node transmits once every N time slots. *ie 0, 1 for A, B*
- No packet collisions! But unused time slots are "wasted", lowering utilization. Poor when nodes send data in bursts or have different offered loads. *A B A B A B - simple round robin scheme*

Not good on varying loads

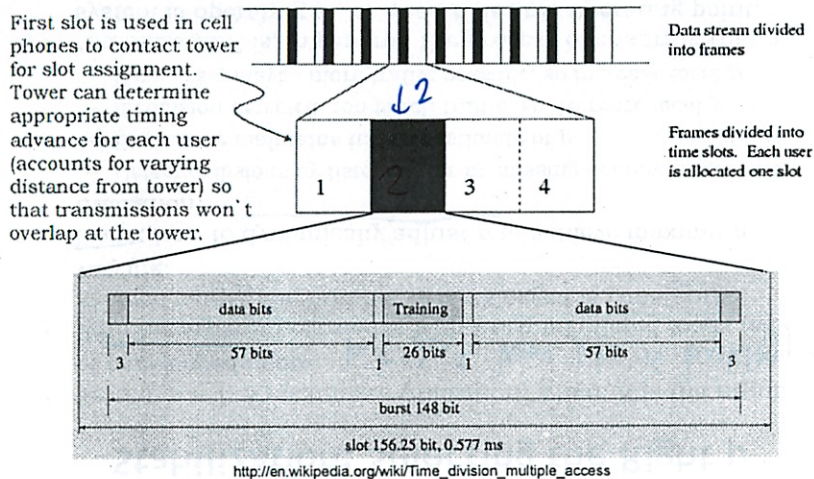
6.02 Spring 2011

Lecture 12, Slide #7

6.02 Spring 2011

Lecture 12, Slide #8

TDMA for GSM Phones

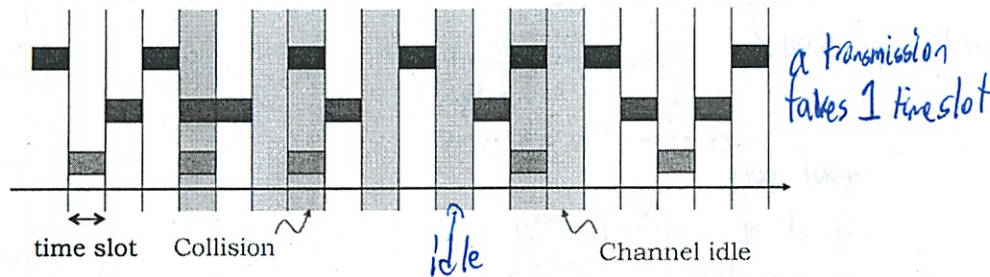


6.02 Spring 2011

1st slot = RATCH - request to tower
When to translate
tower gives you back timing instructions inc allowance for time of flight, a few bits left if moving

Success, Idleness, Collisions

Slotted Everything happens in slots



- Throughput = Uncollided packets per time interval
- Utilization = Throughput / Channel Rate = 12/20 = .6

* transmitters can detect collisions

successful transmission

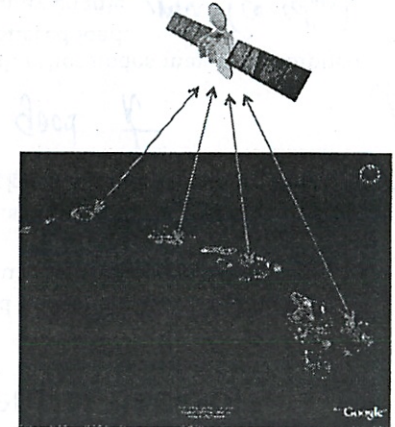
Lecture 12, Slide #11

6.02 Spring 2011

Contention Protocols: Alohanet

To improve performance when there are burst data patterns or skewed loads, use a contention protocol where allocation is not predetermined.

Alohanet was a satellite-based data network connecting computers on the Hawaiian islands. One frequency was used to send data to the satellite, which rebroadcast it on a different frequency to be received by all stations.



Stations could only "hear" the satellite, so had to decide independently when it was their turn to transmit.

6.02 Spring 2011

uplink freq to sat
downlink freq to everyone
But when to transmit? Did not want sat to decide

Lecture 12, Slide #10

Slotted Aloha

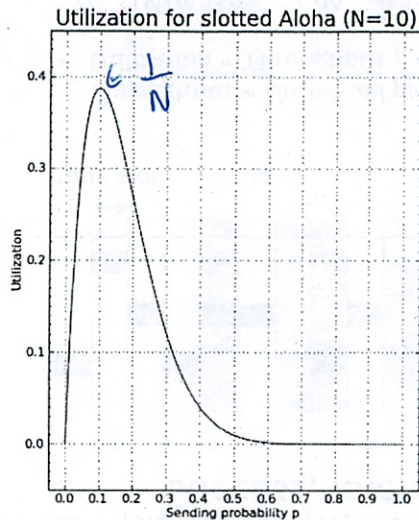
- Aloha protocol followed by each of N nodes:
 - if a node is backlogged, it sends a packet in the next time slot with probability p . *b/w 0 and 1*
random.random() < p
- Assume (for now) each packet takes exactly one time slot to transmit (slotted Aloha)
- Utilization when all nodes backlogged? The probability that exactly one node sends a packet.
 - prob(send a packet) = p
 - prob(don't send a packet) = $1-p$
 - prob(only one sender) = $p(1-p)^{N-1}$
 - There are $(N \text{ choose } 1) = N$ ways to choose the one sender

$$U_{\text{slotted Aloha}} = Np(1-p)^{N-1}$$

6.02 Spring 2011

Lecture 12, Slide #12

Maximizing Utilization



To determine maximum:
set $dU/dp = 0$, solve for p .

Result: $p = 1/N$, so

$$U_{\max} = \left(1 - \frac{1}{N}\right)^{N-1}$$

differentiate w/ resp to p to find max

As $N \rightarrow \infty$, $U_{\max} \rightarrow \frac{1}{e} \approx 37\%$

why e?

$$\ln\left(1 - \frac{1}{N}\right)^{N-1} = (N-1)\ln\left(1 - \frac{1}{N}\right)$$

taylor series expansion

$$= (N-1)\left(-\frac{1}{N} - \frac{1}{2N^2} - \frac{1}{3N^3} - \dots\right)$$

$$= -1 + \frac{1}{N} + \frac{1}{6N^2} + \frac{1}{12N^3} + \dots$$

Should remember in this class

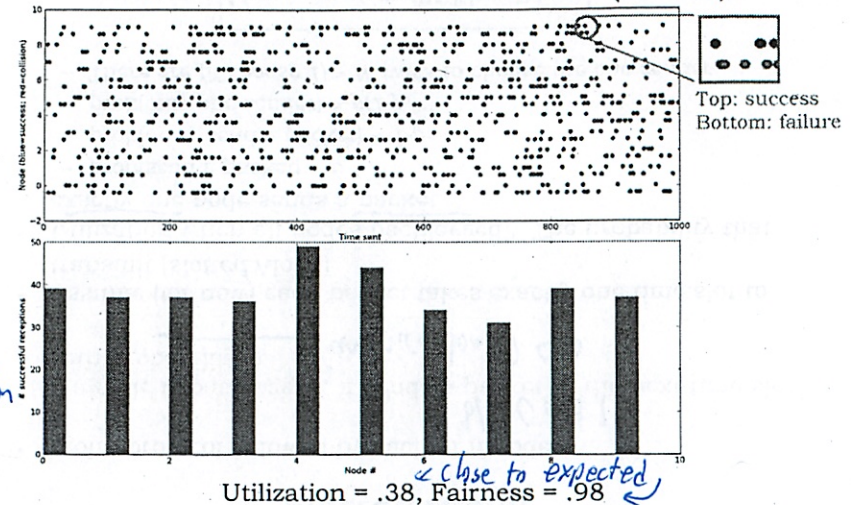
$$= -1 \text{ as } N \rightarrow \infty$$

Lecture 12, Slide #13

$$U_{\max} = \left(1 - \frac{1}{10}\right)^9 = .99$$

What is this - manually, way to take limit?
Can you do 37% utilization?
→ It's very simple

Simulation of Slotted Aloha (N=10)



python PS6_2.py -r -n 10 -p .1 -t 1000

Lecture 12, Slide #14

Stabilization: Selecting the Right p

- Setting $p = 1/N$ maximizes utilization, where N is the number of backlogged nodes. *but don't know who is backlogged*
- With bursty traffic or nodes with unequal offered loads (aka skewed loads), the number of backlogged is constantly varying.
- Issue: how to dynamically adjust p to achieve maximum utilization?
 - Detect collisions by listening, or by missing acknowledgement
 - Each node maintains its own estimate of p
 - If collision detected, too much traffic, so decrease local p
 - If success, maybe more traffic possible, so increase local p
- "Stabilization" is, in general, the process of ensuring that a system is operating at, or near, a desired operating point.
 - Stabilizing Aloha: finding a p that maximizes utilization as loading changes.

Simple strategy

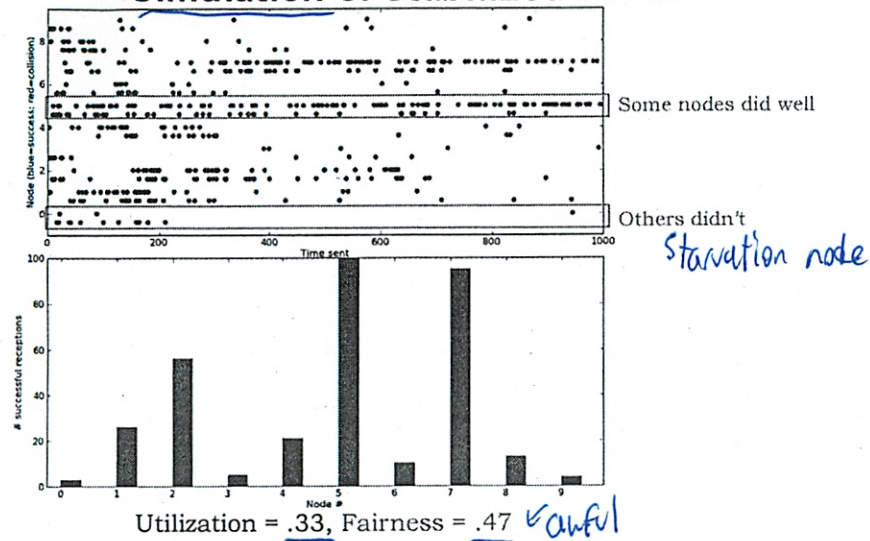
Binary Exponential Back-off

- Decreasing p on collision
 - Estimate of N (# of backlogged nodes) too low, p too high
 - To quickly find correct value use multiplicative decrease: $p \leftarrow p/2$ *halving p*
 - k collisions in a row: p decreased by factor of 2^k
 - Binary: 2, exponential: k , back-off: smaller $p \rightarrow$ more time between tries
- Increasing p on success
 - While we were waiting to send, other nodes may have emptied their queues, reducing their offered load.
 - If increase is too small, slots may go idle *min (correction)*
 - Try multiplicative increase: $p \leftarrow \max(2^k p, 1)$
 - Or maybe just: $p \leftarrow 1$ to ensure no slots go idle

*exponential - if k things in a row
 p is changing exponentially over time*

switch b/w

Simulation of Stabilized Aloha



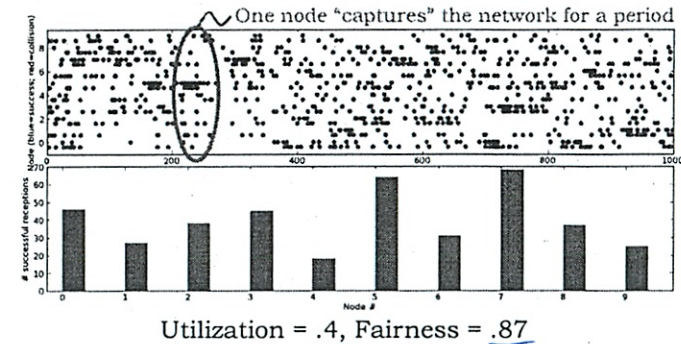
6.02 Spring 2011 python PS6_3.py -r -n 10 -t 1000

Lecture 12, Slide #17

What Went Wrong?

- Starvation
 - Too many successive failures $\rightarrow p$ very small \rightarrow no xmit attempts
 - Result: significant long-term unfairness
 - Try a reduction rule with a lower bound: $p \leftarrow \max(p_{\min}, p/2)$
 - Choosing $p_{\min} \ll 1/\max(N)$ seems to work best

add lower bound

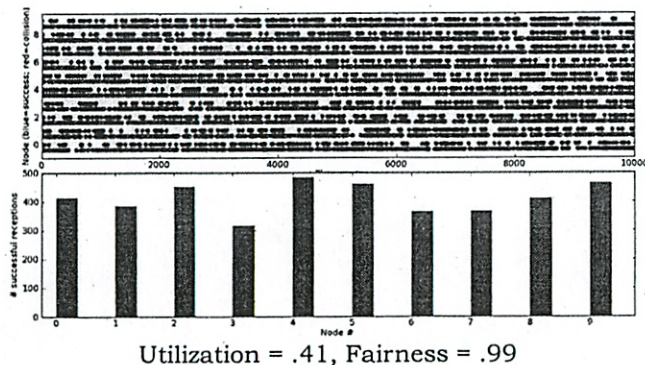


6.02 Spring 2011 python PS6_3.py -r -n 10 -t 1000 --pmin=.05 Lecture 12, Slide #18

better

Limiting the Capture Effect

- Capture effect
 - A successful node maintains a high p (avg. near 1)
 - Starves out other nodes for short periods
 - Try an increase rule with an upper bound: $p \leftarrow \min(p_{\max}, 2*p)$



6.02 Spring 2011 python PS6_3.py -r -n 10 -t 1000 --pmin=.05 --pmax=0.8 Lecture 12, Slide #19

*So much
4/ upper
bound*



Media Access Control (MAC)

Machines interested in total throughput

Not just their throughput

So backs off so others can talk

IEEE 802.11xx standards

Prof. Shah's topic

Want

- simple - utilize only local info, little computation
- efficient - try to always make useful transmissions
- fair - how to judge that?

- proportional fair

- max-min fair

"Socialist"

- "Capitalist"

"Cake Cutting Algorithms"

$$T \sim o(T)$$

$$\lim_{T \rightarrow \infty} \frac{g(T)}{T} \rightarrow 0$$

Example $\rightarrow \frac{T - \sqrt{T}}{T} = 1 - \frac{1}{\sqrt{T}}$

loss should be much smaller than total transmission

Cake Cutting

For $n=2$

- person A cuts
- B chooses piece

$n=3$

move knife $L \rightarrow R$ at ∞ by slow rate

Anyone can stop the knife by saying "I want it"

(2)

Read book "Cake cutting algorithms"
"Stable marriage algorithms"

Aloha algorithm

$p(T)$

local information only
- each node on own

$$p(0) = 1$$

$$p(t+1) = \begin{cases} p(t)/2 & \text{if raised at } T + \text{failed since conflict} \\ p(t) & \text{Didn't raise hand} \end{cases}$$

\downarrow Raised hand + successful reset p to 1

- greedy
- monopolizes channel?

- does not do well when channel becomes congested

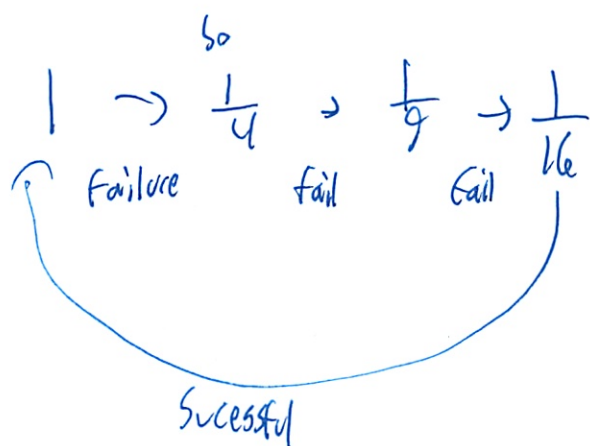
So then Polynomial backoff

$$p(k) \approx 2^{-k} \quad k = \# \text{ of unsuccessful attempts}$$

$$\rightarrow \frac{1}{(k+1)^\alpha} \quad \alpha > 1$$

always works ~~and~~

3



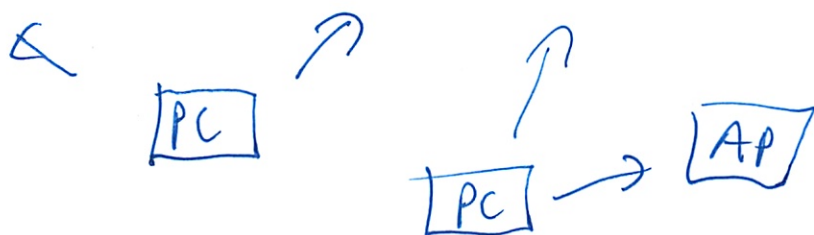
p of transmitting next round $\alpha = 2$

Doesn't work when more than one source everyone is transmitting

Have bunch of APs to connect to

AP

AP



But if each PC talks to diff AP - you are actually talking to both

- interference
- talking loudly allows other people to hear

#2 ^{Binary} Exponential Backoff

Will it ensure 2 nodes will never collide

False

One person A
B

2^{k_1} $2^{-(k_1+1)}$
 2^{k_2} collide, so $2^{-(k_2+1)}$
 at each time step, p of trans.

4

$$P(\text{collision}) = 2^{-(k_1+1)} \cdot 2^{-(k_2+1)}$$

Which means still small chance of collision

Polynomial - since it gets smaller ~~he~~ may have collisions - but it will eventually work

Since $P() \xrightarrow{\text{mm}} 0$
tends

b) correct (did not here)

c) skip

d) Over timescale it improves throughput vs not using _____

No - since starvation can happen

(Need to see on paper)

Why did e appear?

1 each transmits at every time step w/ $P()$

2 ↘

3 →

3 →

⋮ ↗

N ↗

(M)

$P = \frac{1}{N}$ is best case (from simulations)

$$P_{\text{success}} = \left(1 - \frac{1}{N}\right)^{N-1}$$

$$\hookrightarrow P(\text{Node 1 trans success}) + P(\text{Node 2 trans success}) + \dots +$$

$$P(\text{Node N trans success})$$

call synthetic $\lambda_0 = N \cdot P(\text{Node 1 success trans})$

⑤

For ~~trans~~ $P(\text{Node } i \text{ trans success}) = P(\text{Node } i \text{ trans}) \text{ AND } P(\text{No one else trans})$
 $= \frac{1}{N} \left(1 - \frac{1}{N}\right)^{N-1}$

So in total

$$= N \cdot \frac{1}{N} \left(1 - \frac{1}{N}\right)^{N-1}$$
$$= \left(1 - \frac{1}{N}\right)^{N-1}$$

Seq decreases to $\frac{1}{e} \approx .3762 \dots$

$$\lim_{x \rightarrow \infty} \left(1 - \frac{1}{x}\right)^x$$

P_i comes from circle

e comes from i

3/13

To save your work, click the SAVE button at the bottom of this page. You can revisit this page, revise your answers and SAVE as often as you like.

To submit the assignment, click the SUBMIT button at the bottom of this page. YOU CAN SUBMIT ONLY ONCE. Once the assignment has been submitted, you can continue to view this page but will no longer be able to make any changes to your answers.

6.02 Spring 2011: Plasmeier, Michael E.

PSet PS5

Dates & Deadlines

issued: Mar-09-2011 at 00:00

due: Mar-17-2011 at 06:00

checkoff due: Mar-22-2011 at 06:00

Help is available from the staff in the 6.02 lab (38-530) during lab hours -- for the staffing schedule please see the [Lab Hours](#) page on the course website. We recommend coming to the lab if you want help debugging your code.

For other questions, please try the 6.02 on-line Q&A forum at [Piazza](#).

Your answers will be graded by actual human beings, so your answers aren't limited to machine-gradable responses. Some of the questions ask for explanations and it's always good to provide a short explanation of your answer.

Problem 1.

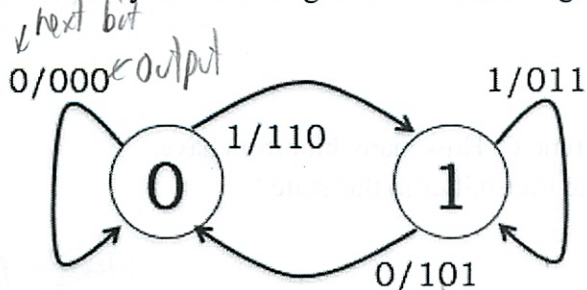
Consider a convolutional code with three generator polynomials

$$p0[n] = (x[n] + x[n-1]) \bmod 2$$

$$p1[n] = x[n]$$

$$p2[n] = x[n-1]$$

as shown by the following state transition diagram:



A. What are the constraint length (K) and code rate (r) for this code?

(n, k, d) but its $n=0$ $k=3$ 1013

$r = \text{rate} = \frac{k}{n}$ $n = 1$ $r = \frac{1}{3}$ 3 bits trans. for every message bit

$k = \text{non zero codewords w/ min \# of 1s}$

(points: 0.5)

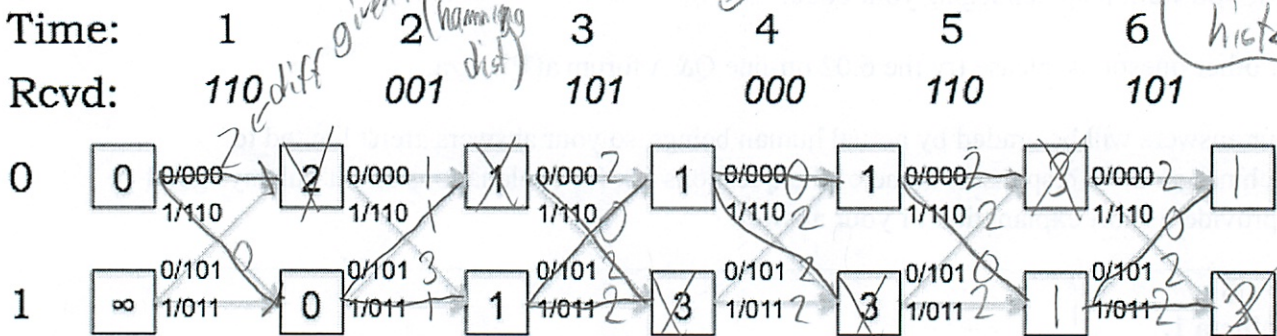
B. What sequence of bits would be transmitted for the message "0110"? Assume the transmitter starts in the state labeled "0".

Try SM

000 110 011 101

(points: 0.5)

The Viterbi algorithm is used to decode a sequence of received parity bits, as shown in the following trellis diagram:



C. Please determine the path metrics for the last two columns of the trellis.

3 1

1 3

(points: 0.5)

D. What is the most-likely state of the transmitter after time 6? How many bit errors have been detected for the most likely message that left the transmitter in this state?

State 0

1 possible error

should try crossing at needed

(points: 0.5)

E. What is the most-likely message sequence as decoded using the trellis above?

how exactly do you ~~at~~ rock off - just following 0 error path
is that right? - or want to go backward
0 1 1 0 0 1 0

(points: 0.5)

F. Given the answer to part E, at what time during the transmission did the first bit error occur?

2 The only

(points: 0.5)

Draw arrows darker that are used in the min test

Problem 2.

(ok this is the error following path)
- wrote in

A technique called **puncturing** can be used to obtain a convolutional code with a rate higher than $1/r$, when r generators are used. The idea is to produce the r parity streams as before and then carefully omit sending some of the parity bits that are produced on one or more of the parity streams. One can express which bits are sent and which are not on each stream as a vector of 1's and 0's: for example, on a given stream, if the normal convolutional encoder produces 4 bits on a given parity stream, the vector (1 1 0 1) says that the first, second, and fourth of these bits are sent, but not the third, in the punctured code.

Punctured codes are useful in communication systems that adjust their code rate in response to the observed BER of the channel (in packet systems, this would be measured by the number of packets with CRC failures). If the BER is low, using a punctured code may still yield an acceptable BER after decoding, while enjoying the increased throughput of a higher code rate.

How does it communicate puncturing decisions?

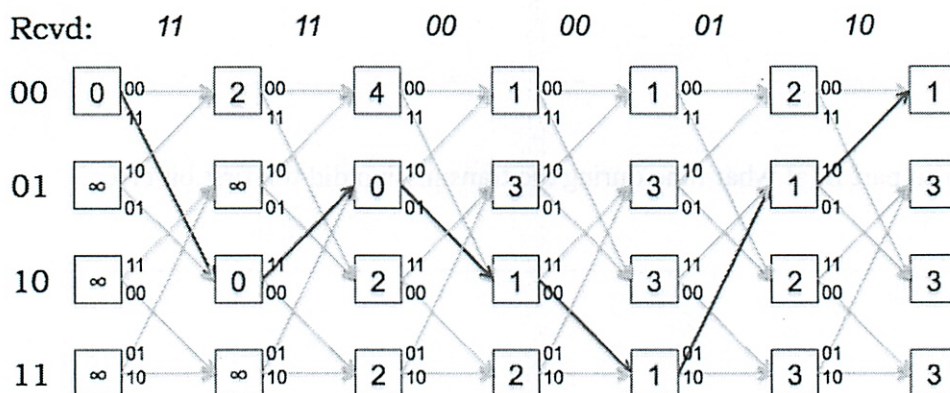
A. Suppose we start with a rate $1/2$ convolutional code without puncturing. The encoder then uses the vector (1 0) on the first parity stream and the vector (1 1) on the second one: that is, it sends every other parity bit produced on the first parity stream and every parity bit produced on the second one. What is the rate of the resulting convolutional code?

same as previous

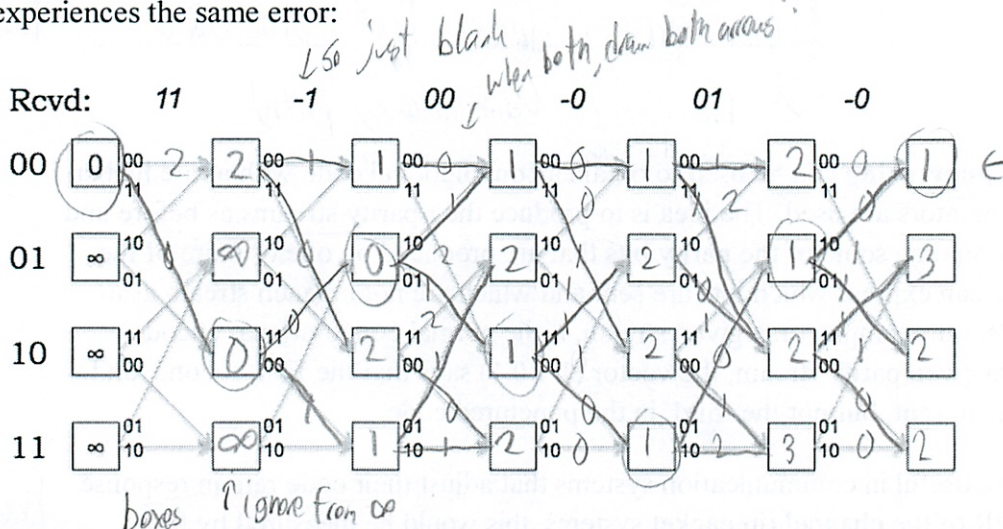
So diff code for each generator/stream
says for every timestep

$r = \frac{k}{n}$ What is new k ? guessing $k = \frac{n}{2}$ for $\frac{1}{2}$
 $\frac{1}{2}$ means take 2x as many bits to send same message
 So now $\frac{1}{2} \cdot 2 + \frac{3}{2} \cdot 1 = \frac{2+3}{2} = \frac{5}{2}$

- B. When using a punctured code, missing parity bits don't participate in the calculation of branch metrics. Consider the following trellis from a transmission using $k=3, r=1/2$ convolutional code that experiences a single bit error:



Here's the trellis for the same transmission using the punctured code from part (A) that experiences the same error:



Fill in the path metrics, remembering that with the punctured code from part (A) that the branch metrics are calculated as follows:

BM(-0, 00)=0 BM(-0, 01)=1 BM(-0, 10)=0 BM(-0, 11)=1
 BM(-1, 00)=1 BM(-1, 01)=0 BM(-1, 10)=1 BM(-1, 11)=0

where the missing incoming parity bit has been marked with a "-". Use a text representation of the matrix (a 2-D array of numbers) when typing in the answer.

See above

(points: 1)

C. Did using the punctured code result in a different decoding of the message?

No

(points: 0.5)

Python Task 1: Viterbi decoder for convolutional codes -- Hard decision decoding

Useful download links:

[PS5_tests.py](#) -- test jigs for this assignment

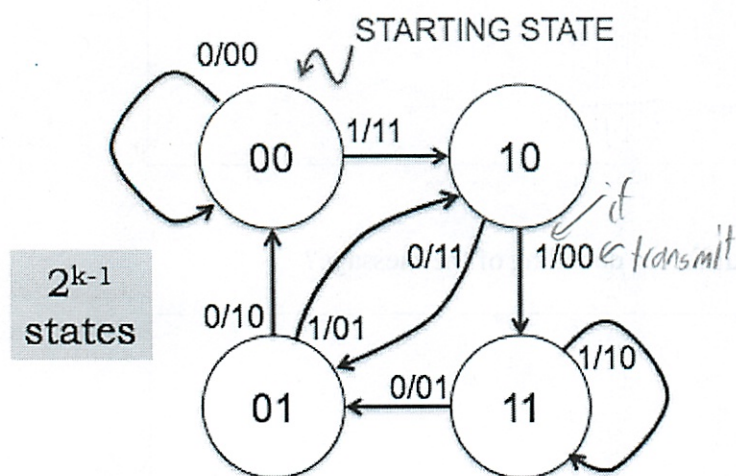
[PS5_1.py](#) -- template file for this task

As explained in lecture, a convolutional encoder is characterized by two parameters: a constraint length K and a set of r generator functions $\{G_0, G_1, \dots\}$. The encoder processes the message one bit at a time, generating a set of r parity bits $\{p_0, p_1, \dots\}$ by applying the generator functions to the current message bit, $x[n]$, and $K-1$ of the previous message bits, $x[n-1]$, $x[n-2]$, \dots , $x[n-(K-1)]$. The r parity bits are then transmitted and the encoder moves on to the next message bit.

The operation of the encoder may be described as a state machine, as you know (we hope). The figure below is the state transition diagram for a rate $1/2$ encoder with $K=3$ using the following two generator functions:

$G_0: 111$, i.e., $p_0[n] = (1 \cdot x[n] + 1 \cdot x[n-1] + 1 \cdot x[n-2]) \bmod 2$
 $G_1: 110$, i.e., $p_1[n] = (1 \cdot x[n] + 1 \cdot x[n-1] + 0 \cdot x[n-2]) \bmod 2$

A generator function can be described compactly by simply listing its K coefficients as a K -bit binary sequence, or even more compactly (for a human reader) if we construe the K -bit sequence as an integer, e.g., $G_0: 7$ and $G_1: 6$. We will use the latter (integer) representation in this task.



In this diagram the states are labeled with the two previous message bits, $x[n-1]$ and $x[n-2]$, in left-to-right order. The arcs -- representing transitions between states as the encoder completes the processing of the current message bit -- are labeled with $x[n] / p_0p_1$.

You can read the transition diagram as follows: "If the encoder is currently in state 11 and if the current message bit is a 0, transmit the parity bits 01 and move to state 01 before processing the next message bit." And so on, for each combination of current states and message bits. The encoder starts in state 00 and you should assume that all the bits before the start were 0.

A stream of parity bits corrupted by additive noise arrives at the receiver. Using the information in these parity bits, the decoder attempts to deduce the actual sequence of states traversed by the encoder and recover the transmitted message. Because of errors, the receiver cannot know exactly what that sequence of states was, but it can determine the most likely sequence using the Viterbi decoding algorithm.

The Viterbi decoder works by determining a **path metric** $PM[s, n]$ for each state s and bit time n . Consider all possible encoder state sequences that cause the encoder to be in state s at time n . In hard decision decoding, the most-likely state sequence is the one that produced the parity bit sequence nearest in Hamming distance to the sequence of received parity bits. Each increment in Hamming distance corresponds to a bit error. $PM[s, n]$ records this smallest Hamming distance for each state at the specified time.

The Viterbi algorithm computes $PM[... , n]$ incrementally. Initially

$$PM[s, 0] = 0 \text{ if } s == \text{starting_state} \text{ else } \infty$$

The decoding algorithm uses the first set of r parity bits to compute $PM[... , 1]$ from $PM[... , 0]$. Then, it uses the next set of r parity bits to compute $PM[... , 2]$ from $PM[... , 1]$. It continues in this fashion until it has processed all the received parity bits.

Here are the steps for computing $PM[... , n]$ from $PM[... , n-1]$ using the next set of r parity bits to be processed:

For each state s :

1. Looking at the encoder's state transition diagram, determine the two predecessor states α and β , which have transition arcs that arrive at state s .
if how do programmatically?
 Using the rate $1/2$, $k=3$ encoder above, if s is 10 then, in no particular order, $\alpha = 00$ and $\beta = 01$.
2. For the state transition $\alpha \rightarrow s$ determine the r parity bits that the encoder would have transmitted; call this r -bit sequence p_α . Similarly, for the state transition $\beta \rightarrow s$ determine the r parity bits the encoder would have transmitted; call this r -bit sequence p_β .

Continuing the example from Step 1: $p_\alpha = 11$ and $p_\beta = 01$.

3. Call the next set of r received parity bits p_{received} . Compute the Hamming distance between p_α and p_{received} . This Hamming distance is a **branch metric** for the state transition $\alpha \rightarrow s$, so we'll label it BM_α . Similarly, compute the Hamming distance between p_β and p_{received} ; we'll call it BM_β .

Continuing the example from Step 2: assuming the received parity bits $p_{\text{received}} = 00$, then $BM_\alpha = \text{hamming}(11, 00) = 2$ and $BM_\beta = \text{hamming}(01, 00) = 1$.
do compare separate so can do soft

4. Compute two trial path metrics that correspond to the two possible paths leading to state s :

$$\begin{aligned} PM_\alpha &= PM[\alpha, n-1] + BM_\alpha \\ PM_\beta &= PM[\beta, n-1] + BM_\beta \end{aligned}$$

PM_α is the Hamming distance between the transmitted bits and the parity bits received so far, assuming that we arrive at state s via state α . Similarly, PM_β is the Hamming distance between the transmitted and parity bits received so far, assuming we arrive at state s via state β .

Continuing the example from Step 3: assuming $PM[\alpha, n-1] = 5$ and $PM[\beta, n-1] = 3$, then $PM_\alpha = 7$ and $PM_\beta = 4$.

5. Now compute $PM[s, n]$ by picking the smaller of the two trial path metrics. Also record which state we chose to be the most-likely predecessor state for s :

```
if  $PM_\alpha \leq PM_\beta$ :
     $PM[s, n] = PM_\alpha$ 
     $\text{Predecessor}[s, n] = \alpha$ 
else
     $PM[s, n] = PM_\beta$ 
     $\text{Predecessor}[s, n] = \beta$ 
```

Completing the example from step 4: $PM[s, n] = 4$ and $\text{Predecessor}[s, n] = \beta$.

so never dump steps

We can use the following recipe at the receiver to determine the transmitted message:

1. Initialize $PM[\dots, 0]$ as described above.

2. Use the Viterbi algorithm to compute $PM[... , n]$ from $PM[... , n-1]$ and the next r parity bits; repeat until all received parity bits have been consumed. Call the last time point N .
3. Identify the most-likely ending state of the encoder by finding the state s that has the minimum value of $PM[s, N]$. If there are several states with the same minimum value, choose one arbitrarily.
4. Determine message bit that caused the transition to state s . You can make this determination simply from knowing s ; you don't need to know the predecessor state. Set $s = \text{Predecessor}[s, N]$, decrement N , and repeat this step as long as $N > 0$, accumulating the message bits in reverse order.

So never "dump" unneeded steps

In this task we'll write the code for some methods of the `ViterbiDecoder` class. One can make an instance of the class, supplying K and the parity generator functions, and then use the instance to decode messages transmitted by the matching encoder.

The decoder will operate on a sequence of received voltage samples; the choice of which sample to digitize to determine the message bit has already been made, so there's one voltage sample for each bit. The transmitter has sent a 0 Volt sample for a "0" and a 1 Volt sample for a "1", **but those nominal voltages have been corrupted by additive Gaussian noise zero mean and non-zero variance**. Assume that 0's and 1's appear with equal probability in the transmitted message.

*glist = parity bits
- so they calc when you call it*

PS5_1.py is the template file for this task (not shown inline here).

You need to write the following functions:

so just small parts

`number = branch_metric(self, expected, received)`
`expected` is an r -element list of the expected parity bits (or you can also think of them as voltages given how we send bits down the channel). `received` is an r -element list of actual sampled voltages for the incoming parity bits. In this task, we will do hard decision decoding, so digitize the received voltages to get bits and then compute the Hamming distance between the expected sequence and the received sequences. That's the branch metric.

*what is this
- oh for the branch
- oh given!
- don't worry how they got it*

Consider using `PS5_tests.hamming(seq1, seq2)` which computes the Hamming distance between two binary sequences.

Oh done already

`viterbi_step(self, n, received_voltages)`
 compute `self.PM[... , n]` from the batch of r parity bits and the path metrics for `self.PM[... , n-1]` computed on the previous iteration. Follow the algorithm described above. In addition to making an entry for `self.PM[s, n]` for each state s , keep track of the most-likely predecessor for each state in the `self.Predecessor` array. You'll find the following instance variables and methods useful (scan the code to understand how to use them):

*Can't write yet!
- they should do in steps*

```
self.predecessor_states
self.expected_parity
self.branch_metric()
```

*called each time step
n = time*

```
s, n = most_likely_state(self, n)
```

Identify the most-likely ending state of the encoder by finding the state s that has the minimum value of $PM[s, n]$, where n points to the last column of the trellis. If there are several states with the same minimum value, choose one arbitrarily. Return the info as a tuple (s, n) .

```
message = traceback(self, s, n)
```

Starting at state s at time n , use the trellis along the most-likely path back to time step 0. Note that you're tracing backwards through the trellis, so you'll be collecting the message bits in reverse order before returning the final message. Hint: you can determine the message state itself -- think about how the state changes when the transmitter makes a transition. Your code will work for different values of $K=3$ and $K=4$ codes.

already known

0	000
1	001
2	010
3	011
4	100
5	101
6	110
7	111

rough the
each
at you'll
right
e. Hint:
n the
state
that
at least

do step
by step

Get from
state to
message

no just added how?

transition
diagram

w/o having
to regenerate

Oh expected bits!

Predecessor

IR

$x[n-1] \times [n-2] \times [n-3]$

Oh that is not etc

What I was
interested in

Oh only
use first!

The testing code at the end of the template reports whether the decoding was successful or not.

File to upload for Task 1:

Browse...

(points: 8)

Here's the debugging printout generated when Alyssa P. Hacker ran the example from lecture.

```
1 [ 2.000000000e+00  1.000000000e+06  0.000000000e+00  1.00000100e+06] [0 2 0 3]
2 [ 3.  1.  3.  1.] [0 2 0 2]
3 [ 2.  2.  2.  2.] [1 3 1 3] error
4 [ 2.  3.  3.  2.] [0 3 1 2]
5 [ 3.  2.  3.  4.] [0 3 0 2]
6 [ 2.  4.  4.  4.] [1 2 0 2]
```

PM
error
error

Each line shows the new column added to the PM and Predecessor matrices after each call to `viterbi_step`. The last line indicates that the most-likely final state is 0 and that there were two bit errors on the most-likely path leading to that state.

- A. Consider the most-likely path back through the trellis using the Predecessor matrix. At which processing steps did the decoder detect a transmission error? Call the first line of the debugging printout Step 1; the next, Step 2; etc.

even though wrong dir

Step 6

0

1

3

2 (error)

4 (error)

initial not set up w/ A

(points: 0.5)

Ben Bitdiddle ran Task #1 and found the following path metrics at the end for the four states:

[6200. 6197. 6199. 6199.]

after long message; w/ lots of errors

- B. What is the most likely ending state? And what can you say about the number of errors the decoder detected? Were they all corrected?

too big!??

(points: 0.5)

Here's an example printout from running decoder tests with a 500000-bit message using two different convolutional codes, one with a constraint length of 3 (with generators 111 and 110), the other with a constraint length of 4 (with generators 1101 and 1110).

**** ViterbiDecoder: k=3, rate=1/2 ****

BER: without coding = 0.006200, with coding = 0.000218

Trellis state at end of decoding:

[6200. 6197. 6199. 6199.]

**** ViterbiDecoder: k=4, rate=1/2 ****

BER: without coding = 0.006200, with coding = 0.000022

Trellis state at end of decoding:

[6203. 6204. 6205. 6200. 6203. 6202. 6203. 6202.]

Please answer the following questions based on the printout from the k=3, rate=1/2 decoder.

- C. What was the most-likely final state of the transmitter? And what were the most-likely final two bits of the message? You can answer both questions from the final trellis state shown in the printout.

2 bits for every message

(points: 0.5)

- D. The test message used was 500,000 bits long. Since we're using a rate=1/2 convolutional code, there were 1,000,000 parity bits transmitted over the channel. The BER "without coding" was calculated by comparing the transmitted bit stream that went into the noisy channel with the bit stream that came out of the noisy channel. How many transmission errors occurred?

$$\text{BER} = \frac{\text{\# of bits wrong}}{\text{\# total bits}}$$

Just
without

$$1006200 = \frac{x}{1,000,000} \quad x = 6200$$

$$1000218 = \frac{x}{1,000,000} \quad x = 218$$

(points: 0.5)

- E. After passing the received parity bits through the decoder and calculating the most-likely transmitted message, the BER "with coding" was calculated by comparing the original message with the decoded message. In this example, the density of bit errors was high enough that the decoder was unable to find and correct all the transmission errors. How many uncorrected errors remained after the decoder did the best it could at correcting errors?

Just 218

(points: 0.5)

Now consider what happened when testing the the $k=4$, rate= $1/2$ decoder.

- F. What was the most-likely final state of the transmitter? And what were the most-likely final *three* bits of the message? You can answer both questions from the final trellis state shown in the printout.

3
011

(points: 0.5)

- G. How many uncorrected errors remained after the decoder did the best it could at correcting errors?

22

(points: 0.5)

- H. Consider a given message bit -- how many transmit bits are influenced by that message bit in both the $k=3$ and $k=4$ cases? How does that difference improve the error correction capability of the $k=4$ code?

Oh earlier I confused message + transmit bits

easy qv!

(points: 0.5)

Python Task 2: Decoding with a soft-decision branch metric

Useful download link:

[PS5_2.py](#) -- template file for this task

Let's change the branch metric to use soft decision decoding, as described in lecture.

The soft metric we'll use is the **square of the Euclidean distance** between the received vector of dimension r , the number of parity bits produced per message bit, of voltage samples and the expected r -bit vector of parities. Just treat them like a pair of r -dimensional vectors and compute the squared distance.

Oh just

[PS5_2.py](#) is the template file for this task (not shown inline here).

$$[rec[0] - exp[0]]^2 + [rec[1] - exp[1]]^2$$

Complete the `branch_metric` method for the `SoftViterbiDecoder` class. Note that other than changing the branch metric calculation, the hard decision and soft decision decoders are identical. The code we have provided runs a simple test of your implementation. It also tests whether the branch metrics are as expected.

have to limit first

File to upload for Task 2:

- Totally wrong
- float?

(points: 2)

not using!

Python Task 3: Comparing the performance of error correcting codes

Useful download link:

[PS5_3.py](#) -- template file for this task

In this task we'll run some experiments to measure the probability of a decoding error (i.e., the bit error rate) using the codes you implemented in Tasks 2, 3, and 4 above. There's no additional code for you to write, just some results to analyze. The codes all have rate 1/2:

Oh in P-set I don't convert to int - well do in Hamming En

And no limiting (ins changed later)

1. The stream code using rectangular parity from Task 2. What should `nrows` and `ncols` be if we want a rate 1/2 code?
2. A convolutional code with hard decision decoding and constraint length 3.
3. A convolutional code with hard decision decoding and constraint length 4.
4. A convolutional code with soft decision decoding and constraint length 3.

45 min!

5. A convolutional code with soft decision decoding and constraint length 4.

PS5_3.py is the template file for this task.

Run the code -- please be patient, it takes a while; you can do something else while this script does its job. In the end you should see a set of plots of the probability of decoding error (the bit error rate) as a function of the amount of noise on the channel (the x axis is proportional to $\log(1/\text{noise})$), tracking the "signal to noise ratio" (SNR) of the channel. Larger values of x correspond to **lower** noise, and a lower error probability for a given coding scheme. Note that the y axis is also on a log scale.

The Task 3 script saves the graph in a .png file called PS5_plots.png -- please upload the result:

Plot to upload for Task 3: <input type="text"/>	<input type="button" value="Browse..."/>	
-------------------------------------------------	------------------------------------------	-------------------------------------------------------------------------------------

(points: 1)

Looking the generated graph, for an SNR of 2 db list the codes in order of performance, from highest to lowest (better performing codes have a lower BER for a given SNR). Considering the Viterbi codes, briefly comment on the how much extra work the receiver has to perform as the code performance improves.

(points: 1)

You can save your work at any time by clicking the Save button below. You can revisit this page, revise your answers and SAVE as often as you like.

To submit the assignment, click on the Submit button below. YOU CAN SUBMIT ONLY ONCE after which you will not be able to make any further changes to your answers. Once an assignment is submitted, solutions will be visible after the due date and the graders will have access to your answers. When the grading is complete, points and grader comments will be shown on this page.

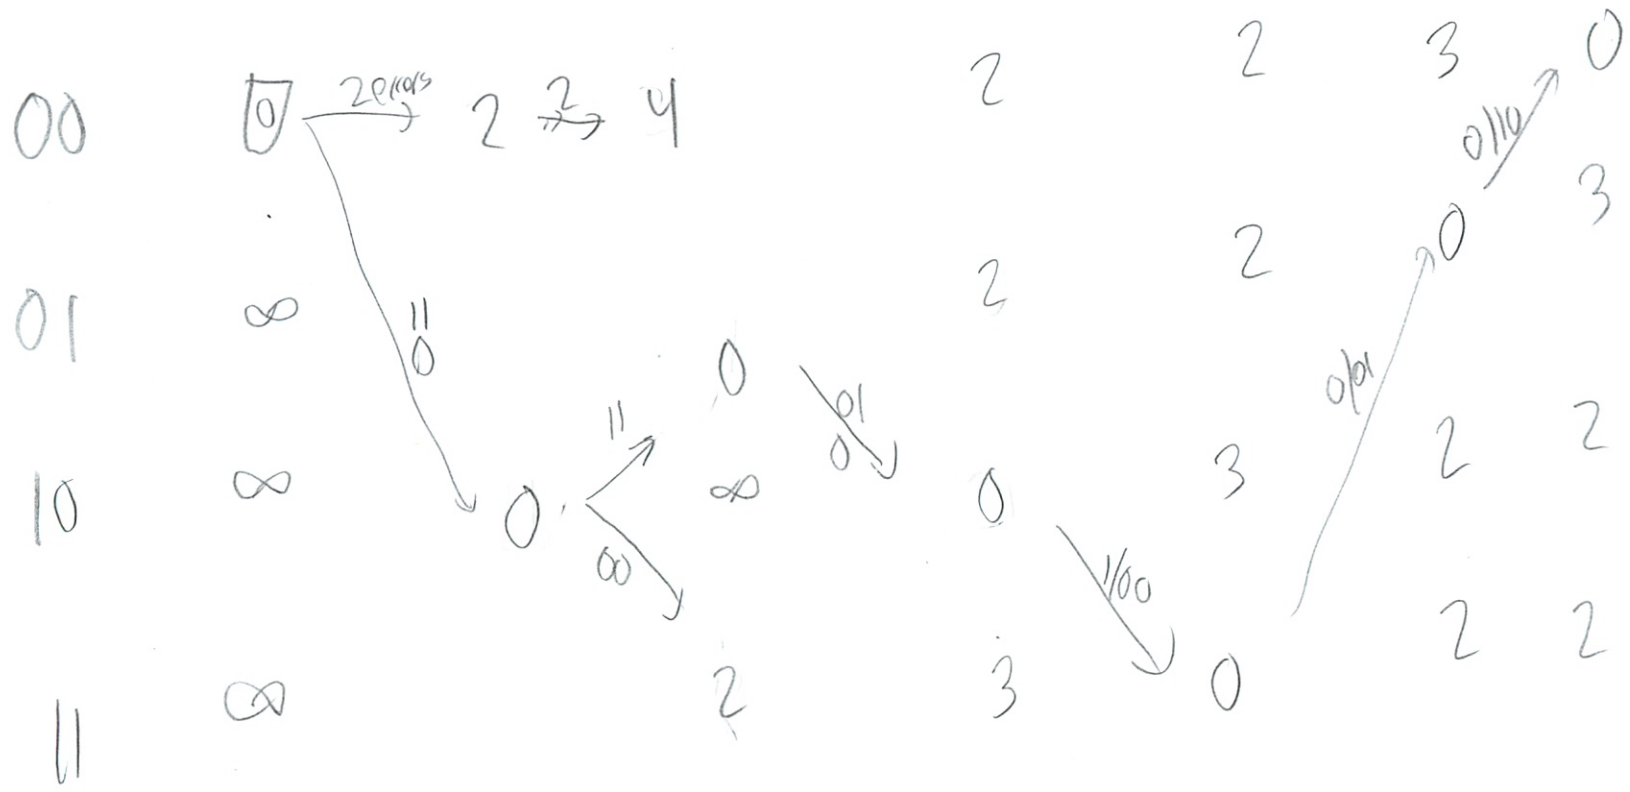
Message ———— 1 ———— 0 ———— 1 ———— 1 ———— 0 ———— 0

Genie Tx ———— 11 ———— 11 ———— 01 ———— 00 ———— 01 ———— 10



Message
 Tx 1 0 1 1 0 0

 11 11 01 00 01 10



So past states 0 1 3 2 1 2 | 0

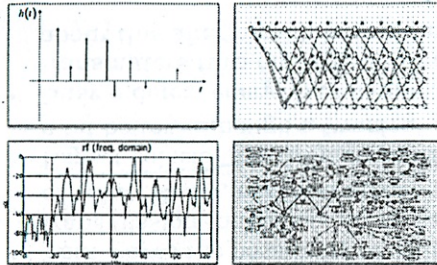
 ms 0 0 1 1 0 1 | 0

← backwards

10 11 00 ✓

Oh never
 did first
 msg hit!

Stabilized Slotted Aloha Summary



INTRODUCTION TO BECS II DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2011 Lecture #13

- unslotted Aloha
- carrier sense, contention windows
- code division glimpse

Contention protocol
Who gets to talk?
Not everyone always using bandwidth

best fit

6.02 Spring 2011

Lecture 13, Slide #1

6.02 Spring 2011

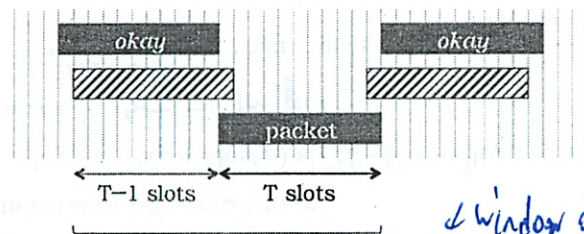
Lecture 13, Slide #2

any overlap collides

Unslotted Aloha

- Packets take T time slots to transmit
 - As slots get smaller and T grows, approximates transmission at arbitrary times. *continuous*
- Collisions are no longer "perfect"
 - Any overlap between multi-slot packets is a collision
 - Larger window of vulnerability to other transmissions

no more nice slots



Any other packet transmitted in these $2T-1$ time slots will collide with target packet

window of vulnerability

6.02 Spring 2011

Lecture 13, Slide #3

6.02 Spring 2011

Lecture 13, Slide #4

- Assumptions
 - Time is divided into slots
 - Transmissions begin on slot boundaries
 - Packets are 1 slot long
- Choose to transmit packet with probability p
 - Each node uses collisions/success to adjust p
 - Adjust p between p_{min} (avoid starvation) and p_{max} (avoid capture)
 - On collision $p \leftarrow \max(p_{min}, p/2)$, on success $p \leftarrow \min(p_{max}, 2*p)$ *but upper + lower bounds to prevent starvation*
- Utilization = throughput achieved / maximum data rate
 - $U = \text{prob}(\text{exactly one transmission in a slot}) = N*p*(1-p)^{N-1}$, where N is number of backlogged nodes
 - Maximized when $p = 1/N$, $U_{max} = (1 - 1/N)^{N-1}$
 - As $N \rightarrow \infty$, $U_{max} \rightarrow 1/e \approx 37\%$
 - While each node's p is never exactly $1/N$, the goal is to have its average value over modest intervals be approximately $1/N$.

Stabilized

~37% - looks low but pretty good

and ~~very~~ very simple - ~~don't~~ but can't get exactly .125 : 8 : 4 : 2 : .05 but on avg will get

Utilization in Unslotted Aloha

Probability of no transmission for $2T-1$ slots:

$$(1-p)^{2T-1}$$

Probability of a sender experiencing no collisions:

$$\begin{aligned} \text{Exactly = } & \text{for OOD nodes } \rightarrow \geq [p(1-p)^{2T-2}] [(1-p)^{2T-1}]^{N-1} = p(1-p)^{(2T-1)N-1} \\ & \text{= for nodes that try to send new packet while busy with last one!} \end{aligned}$$

Utilization = throughput / maximum rate:

$$U_{\text{unslotted Aloha}} \geq \frac{Np(1-p)^{(2T-1)N-1}}{1/T} = TNp(1-p)^{(2T-1)N-1}$$

Self-interference

lower bound

3/16

U_{\max} for Unslotted Aloha

Maximization with respect to p :

$$\log(\dots) = \text{const} + \log(p) + [(2T-1)N-1]\log(1-p)$$

Derivative:

$$\frac{1}{p} + \frac{(2T-1)N-1}{1-p}, \text{ which equals 0 at } p = \frac{1}{(2T-1)N}$$

Plugging back into U :

$$U_{\max} = \frac{T}{2T-1} \left(1 - \frac{1}{(2T-1)N} \right)^{(2T-1)N-1}$$

For large N : $U_{\max} \approx \left(\frac{T}{2T-1} \right) \frac{1}{e}$

For large N , T : $U_{\max} \approx \frac{1}{2e}$

Half the utilization of slotted Aloha:
makes sense: twice the window of vulnerability

Lecture 13, Slide #5

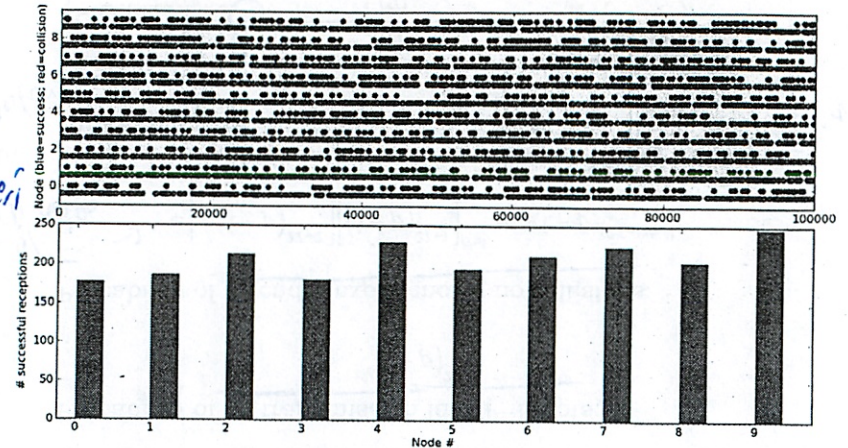
interested in asymptotic behavior

small time slots

pair to max.
log of expression

is it bigger or smaller?

Simulation of Unslotted Aloha



Utilization = .21, Fairness = .99

python PS6_3.py -r -n 10 -t 100000 --pmin=.005 --pmax=0.8 -s 10

6.02 Spring 2011

Lecture 13, Slide #6

had to drop down

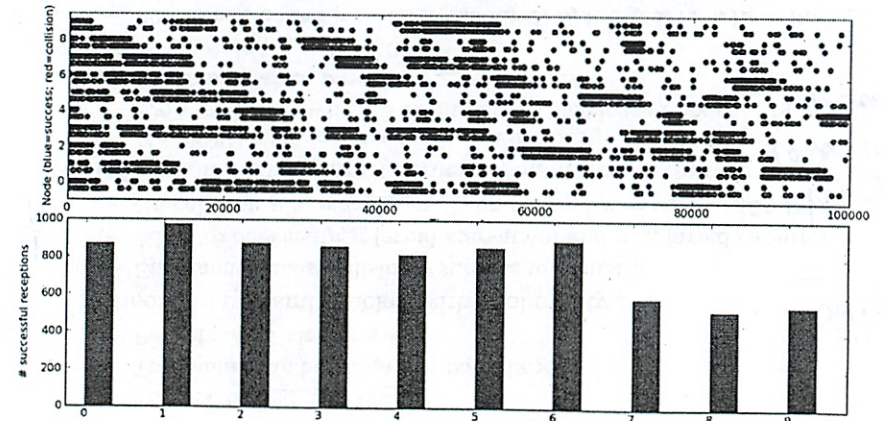
size of 10

Carrier Sense to see if someone else

is trans

- Reduce collisions with on-going transmissions by transmitting only if channel appears not to be busy.
- For large T (slots/packet) if channel is busy this cycle, the same sender will probably be transmitting more of their packet next cycle
- When the channel is idle, there's no chance of interrupting an on-going transmission.
- That leaves the possibility of colliding with another transmission that starts at the same time – a one slot window of vulnerability, not $2T-1$ slots.
- Expect collisions to drop dramatically, utilization to be quite a bit better, although a “wasted” slot is now necessary
- Busy = detect energy on channel. On wireless channels, transmitters turn on carrier to transmit (we'll learn more about this after break), hence the term “carrier sense”.

Simulation of Carrier Sense



Utilization = .78, Fairness = .96

python PS6_4.py -r -n 10 -t 100000 --pmin=.005 --pmax=0.8 -s 10

6.02 Spring 2011

Lecture 13, Slide #8

only 15 characters to change

detecting business is hard, slow

Lecture 13, Slide #7

6.02 Spring 2011

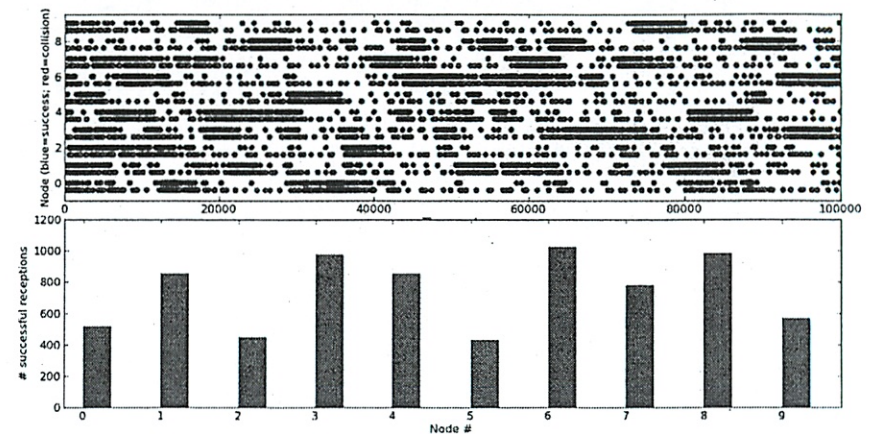
just look to see that channel no busy wait for idle period

Request to send
Clear to send
but spread out requests

Contention Windows

- Contention Window: parameter is some integer CW
- When node wants to transmit, it picks a random number r uniformly in $[1, CW]$ and sends after the r^{th} idle slot from the current time. *size = # time slots*
- If transmission collides: $CW \leftarrow \max(CW_{min}, CW/2)$
If transmission succeeds: $CW \leftarrow \min(CW_{max}, CW*2)$ *not much contention*
- Node is guaranteed to attempt a transmission within CW slots. With the earlier scheme, there was always the chance (though exponentially decreasing) that a node may not transmit within some fixed number of time slots.

Simulation of Contention Windows



Utilization = .74, Fairness = .92

python PS6_5.py -r -n 10 -t 100000 -W 256 -s 10

Summary of MAC Protocols

- Goal of MAC protocols is to maximize utilization and fairness
- TDMA is a good choice when nodes are all (or mostly) backlogged *like Cell Phones*
 - Round-robin sharing provides known communication capacity and bounded wait
 - It's precisely fair, 100% utilization if all nodes have packets
 - Poor choice when traffic is bursty or if some nodes have a higher offered load than others
- Contention protocols dynamically adapt to changing traffic *data transfer*
 - Distributed protocol (each node makes its own decisions based on transmission experience) avoids cost of centralized controller having to know which nodes have packets to send
 - Parameter (p , CW) that controls when packets are sent is adjusted so that prob(sending packet) is lowered when collisions are detected and raised when transmissions are successful.

CW dynamically updated

Summary (cont'd.)

- Slotted Aloha – based on very simple rule: transmit with probability p .
 - Dynamic adjustment of p “stabilizes” the protocol.
 - Use binary exponential backoff to adjust p downward
 - Utilization maximized when $p = 1/(\text{number of backlogged nodes})$
 - For large numbers of backlogged nodes $U \approx 1/e \approx 37\%$
 - For fairness: $p_{min} \leq p \leq p_{max}$
- Unslotted Aloha – packets take multiple time slots to send, models transmissions at arbitrary times
 - Gets half of the max utilization of slotted Aloha due to doubled window of vulnerability to collisions
 - Carrier sense avoids collisions from packets once transmission has started \rightarrow much better utilization
 - Fairness still requires bounds on p

Code Division Multiple Access (CDMA)

- Two vectors are orthogonal if their dot products are 0. Here's a set of 4 mutually orthogonal vectors:

$V1: (1, 1, 1, 1)$
 $V2: (1, 1, -1, -1)$
 $V3: (1, -1, 1, -1)$
 $V4: (1, -1, -1, 1)$

- Assign each transmitter a particular orthogonal vector (V_i) it will use to encode its transmissions (called the "chip code"). With vectors shown above we can support 4 transmitters.

- If message bit is 0, transmit $-V_i$
 - If message bit is 1, transmit V_i
- 1 message bit \rightarrow $\text{len}(V_i)$ "chips"

- Channel will sum the transmitted values:

send 00 using $V1$: $-1 -1 -1 -1 -1 -1 -1 -1$
 send 01 using $V2$: $-1 -1 1 1 1 1 -1 -1$
 send 11 using $V3$: $1 -1 1 -1 1 -1 1 -1$
 send 10 using $V4$: $1 -1 -1 1 -1 1 1 -1$
 channel: $0 -4 0 0 0 0 0 -4$

Sums up in air

6.02 Spring 2011

Lecture 13, Slide #13

can get large

- biggest prob: don't make message so large the receiver can't sense

Asynchronous CDMA

- Use N orthogonal vectors to multiplex N transmitters (e.g., use a $N \times N$ Walsh/Hadamard Matrix) *all rows + columns are orthogonal*
- Scheme described above works for synchronous CDMA when all symbols are transmitted starting at same moment. For example this works fine for a cell tower transmitting to mobile phones.
- But hard to synchronize mobile phone transmissions, so use asynchronous CDMA: *not accurate enough*
 - Can't create transmissions that are truly orthogonal if they start at different times
 - Approximate orthogonality with longer uncorrelated pseudo-random sequences (called pseudo-noise or PN). "pseudo" implies that sequence can be reconstructed at receiver given a known starting point. *Overlaps uncorrelated*
 - Assuming equal signal strengths from each transmitter at receiver, if we decode bits using a particular PN sequence synchronized with desired transmitter, we'll get desired signal plus some uncorrelated noise from other transmitters.

6.02 Spring 2011

Lecture 13, Slide #15

CDMA Receiver

- At receiver take groups of $\text{len}(V)$ bits and form dot product with V_i for desired channel.
 - If result is negative, message bit is 0
 - If result is positive, message bit is 1

channel: $0 -4 0 0 0 0 0 -4$

receive using $V1$: $1 1 1 1 1 1 1 1$ *← chip codes for each V_i*
 dot product: $-4 -4$
 message bits: $0 0$

receive using $V2$: $1 1 -1 -1 1 1 -1 -1$
 dot product: $-4 4$
 message bits: $0 1$

receive using $V4$: $1 -1 -1 1 1 -1 -1 1$
 dot product: $4 -4$
 message bits: $1 0$

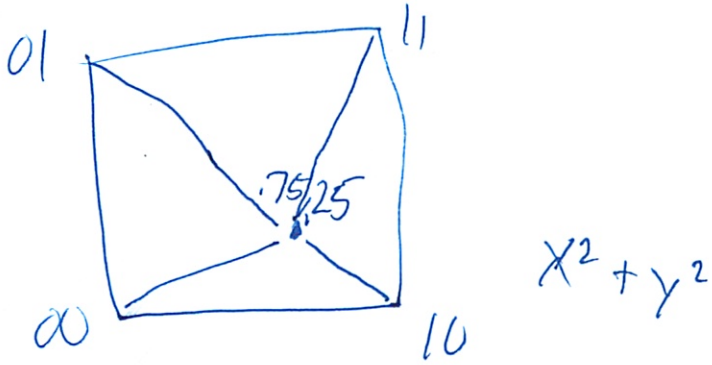
6.02 Spring 2011

Lecture 13, Slide #14

\uparrow dot product for each V_i assigned
 $V3$ no problem either - just out of space on slide

P-Set 5 Interview

3/6



Cases when close to Threshold

SNR - ratio $\frac{\text{signal}}{\text{noise}}$

~~Lower~~ higher better

~~usually~~ usually log
Characteristic of channel

$k=3, 4$ both have gain $\frac{1}{2}$
 \uparrow history \uparrow # generators

red parity - last pset

1	0
1	0

4 8
Cate $\frac{1}{2}$

all have same rate

(15 min late)

Unslotted Aloha

- finding the probabilities of collision
(skipping old stuff on board)

$$U(T, N, p) = T \cdot p (1-p)^{N(2T-1)-1}$$

$$\frac{\partial U}{\partial p} = T(1-p)^{N(2T-1)-1} + T p (1-p)^{N(2T-1)-2} \cdot (-1) \cdot N(2T-1)$$

$$= T(1-p)^{N(2T-1)-2} \left[(1-p) - p(N(2T-1) - 1) \right]$$

≈ 0 Set = to 0 and solve for p

$$\frac{p}{1-p} = \frac{1}{N(2T-1)-1} \rightarrow p^* = \frac{1}{N(2T-1)}$$

$$U^* = \frac{T}{N(2T-1)} \cdot \left(1 - \frac{1}{N(2T-1)} \right)^{N(2T-1)-1} \quad \text{for one node}$$

$$U_{\text{total}} = U^* N$$

$$= \frac{T}{2T-1} \left(1 - \frac{1}{N(2T-1)} \right)^{N(2T-1)} \left(1 - \frac{1}{N(2T-1)} \right)^{-1}$$

When large \uparrow T, N \uparrow \uparrow

$$\underbrace{\frac{1}{2}}_{\frac{1}{2}} \underbrace{\frac{1}{e}}_{\frac{1}{e}} \underbrace{1}_1$$

$$= \frac{1}{2e} \quad \text{cost of lack of sync.}$$

(2)

These formulas were with self transmission (OCD nodes)

Keep for math

Does not really matter at large N, T

$$P_{\text{success}} = p \underbrace{(1-p)^{2T-2}}_{\text{self trans}} ((1-p)^{2T-1})^{N-1}$$

$$= p(1-p)^{2T-1}(N-1) + 2T-2$$

With Carrier Sense

- other ~~people~~ transmitters can listen to channel and listen if someone transmitting
- if they hear a packet on network, wait for ~~it to end first~~ idle time
- Only really works on short ping time

2 types of slots

- waste
 - idle since no one wants to talk
 - collision
 - idle since people want to transmit, but p is low
- Useful

What fraction of slots are useful?

Slot can be in useful or wasteful state

(w)

(u)

5

CDMA

Each transmitter gets a "signature"

T_x	1	1	1	1
$\{-1, 1\} \Rightarrow b$	b_1	b_1	b_1	b_1
0	1			
	-1	-1	-1	-1

Multiply bit by signature value each time

T_x	v_1	v_2	...	v_{1n}
0	-1	b_1	b_1	b_1
1	1			
	s_1	s_2	...	s_{1n}
	+	+		+
	s'_1	s'_2		s'_{1n}
	"	"		"
	y_1	y_2		y_n

③

Making small change to protocol: At end of each timeslot

{ Stop w/ prob $\frac{1}{T}$
 { continue otherwise

So transmit T slots at a time on average

for the
 system



After stop transmitting will be 1 wasteful slot

Next w/ ~~prob~~ another transmitters will start successfully

Q
 \rightarrow prob. that a node becomes successful

Or with ~~at~~ $1-Q$ there is a ~~conflict~~ collision and still useless

$Q = N \cdot p \cdot (1-p)^{N-1}$ 1 of the people (or of them) must trans - the rest must be silent

Now solve Markov Chain to find steady state (remember G.042)

$$P_W = P_W(1-Q) + P_U\left(\frac{1}{T}\right)$$

$P_U = \dots$ see next pg

$$P_W + P_U = 1$$

9

$$P_w \cdot Q = \frac{1}{T} P_v$$

$$P_w = \frac{1}{QT} P_v$$

$$P_v + \frac{1}{QT} P_v = 1$$

$$P_v = \left(\frac{1}{1 + \frac{1}{QT}} \right)$$

Best Q ~~was~~ $\left\{ \begin{array}{l} p = \frac{1}{N-1} \end{array} \right.$ larger Q is better
 $\left\{ \begin{array}{l} Q = Np(1-p)^{N-1} \end{array} \right. \sim \frac{1}{e}$
Greater if N is small

$$P_v = \frac{1}{1 + \frac{e}{T}}$$

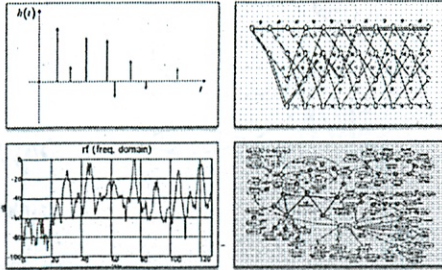
Say $T = 10e$ ^{related to} (length of packet)

$$P_v = \frac{1}{1 + \frac{1}{10}}$$

$$= \frac{10}{11} \approx 90\%$$

If you knew N

Re watch 4/4



INTRODUCTION TO BECS II DIGITAL COMMUNICATION SYSTEMS

6.02 Spring 2011 Lecture #14

- complex exponentials
- discrete-time Fourier series
- spectral coefficients
- band-limited signals

Freq division multiplexing

each gets small section of spectrum

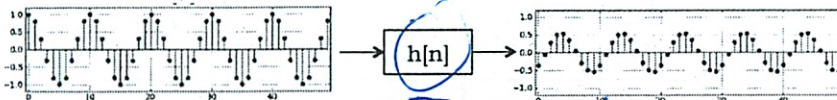
- radio

- cable - internet and TV

6.02 Spring 2011

Lecture 14, Slide #1

Sinusoids and LTI Systems



Frequency division multiplexing depends on an interesting property of LTI channels:

if the channel input $x[n]$ is a sinusoid of a given amplitude, frequency and phase, the response will be a sinusoid at the same frequency, although the amplitude and phase may be altered. As we'll see, the change in amplitude and phase may depend on the frequency of the input.

The same property holds when the inputs are complex exponentials, which are closely related to sines and cosines (and, perhaps surprisingly, are much easier to analyze!).

tool used to analyze

6.02 Spring 2011

Lecture 14, Slide #3

Frequency Division Multiplexing

To engineer the sharing of a channel through frequency division multiplexing we'll need a new set of tools that will let us understand the behavior of signals and systems in the frequency domain. Plan:

- This week

- Analyze the frequency content of signals using the discrete-time Fourier series
- Determine what happens when we band-limit a signal
- Characterize LTI systems by their frequency response
- Introduce filters: LTI systems that eliminate a region of frequencies from a signal *Remove energy from bands*

- Next week

- Using modulation to position band-limited signals in different regions of the frequency spectrum
- Receiving a particular signal from a shared spectrum

6.02 Spring 2011

Lecture 14, Slide #2

need new tools

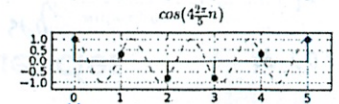
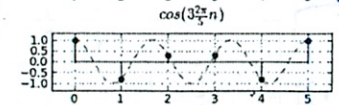
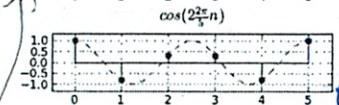
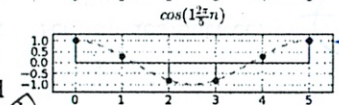
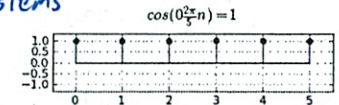
Periodic Sequences *all w/ period of 5 in discrete systems*

A sequence $x[n]$ is said to be periodic with a period of N samples ("periodic with period N ") if

$$x[n] = x[n + N]$$

A sequence that is periodic with period N is also periodic with period $2N, 3N, \dots$, and so on.

A sinusoidal sequence that is periodic with period N can only have one of a finite number of frequencies: all the harmonics of the fundamental frequency $2\pi/N$ radians/sample, i.e., frequencies of the form $k \cdot (2\pi/N)$ for some integer k .



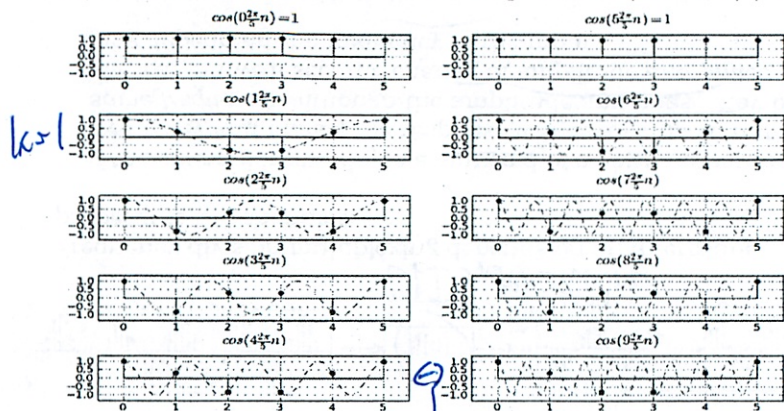
6.02 Spring 2011

waveforms = value of seq

all periodic w/ period 5 are no more possible

Frequencies $k \cdot (2\pi/N)$ when $k \geq N$

Frequency $k \cdot (2\pi/N)$ yields the same sequence as $(k \bmod N) \cdot (2\pi/N)$



Unique sequences: $k = 0, 1, 2, \dots, N-1$

Highest frequency: $(N-1) \cdot (2\pi/N) < 2\pi$

6.02 Spring 2011

Lecture 14, Slide #5

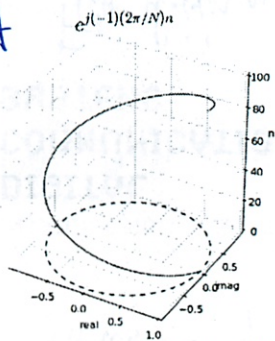
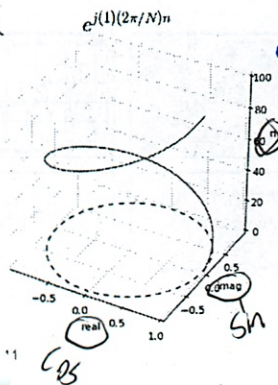
its the same blue dots.
Discrete system can't tell diff.
so only N possible freq

Complex Exponentials

A complex exponential is a complex-valued function of a single argument – an angle measured in radians. Euler's formula shows the relation between complex exponentials and our usual trig functions:

$$e^{j\varphi} = \cos(\varphi) + j\sin(\varphi)$$

Produces a complex value



easier to work in exponentials
 $\cos(\phi) \cdot \cos(\theta)$
 $e^{j\phi} \cdot e^{j\theta}$

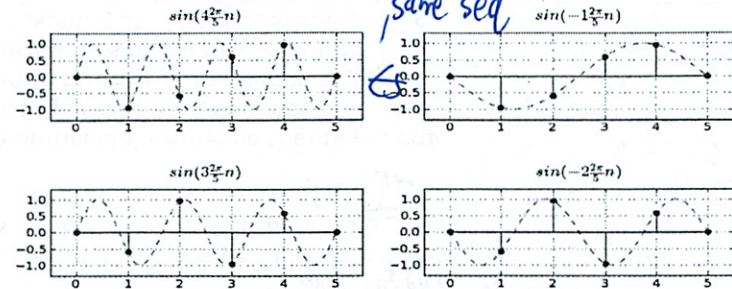
add exponents

Lecture 14, Slide #7

freq = # radians/sample

- simply book keeping
radians/sample

Negative Frequencies



Sequences of frequency $k \cdot (2\pi/N)$, i.e., frequencies between 0 and 2π , are identical to sequences of frequency $-(N-k) \cdot (2\pi/N)$, i.e., between -2π and 0.

In 6.02, our convention will be to specify frequencies in the range $-\pi$ and π , corresponding to k 's in the range $-(N/2)$ to $(N/2)$.

6.02 Spring 2011

$$\frac{7\pi}{4} = -\frac{\pi}{4}$$

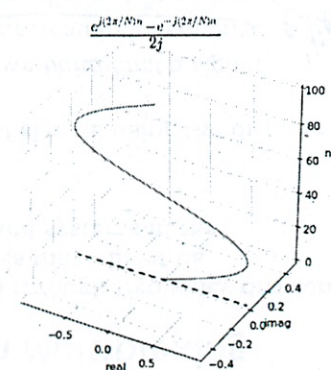
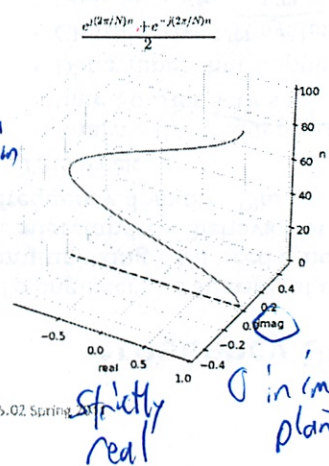
(did a lot of this in Hs)

Lecture 14, Slide #5

will go $-\pi$ \rightarrow π just packing up, arrive at same pts
Sine and Cosine and $e^{j\varphi}$

$$\cos(\varphi) = \frac{1}{2}e^{j\varphi} + \frac{1}{2}e^{-j\varphi} \quad \text{Can write complex \# as sin + cos}$$

$$\sin(\varphi) = \frac{1}{2j}e^{j\varphi} - \frac{1}{2j}e^{-j\varphi}$$



Lecture 14, Slide #8

Useful Properties of $e^{j\varphi}$

When $\varphi = 0$:

$$e^{j0} = 1 = (\cos 0) + j \sin 0$$

When $\varphi = \pm\pi$:

$$e^{j\pi} = e^{-j\pi} = -1$$

$$e^{j\pi n} = e^{-j\pi n} = (-1)^n$$

alternates

Summing samples over one period:

3rd interesting property

$$\sum_{n=-\infty}^{\infty} e^{jk\frac{2\pi}{N}n} = \begin{cases} N & k = 0, \pm N, \pm 2N, \dots \\ 0 & \text{otherwise} \end{cases}$$

n ranges over any N consecutive integers, e.g., $n = 0, 1, \dots, N-1$

6.02 Spring 2011

when add  = 0

Lecture 14, Slide #9

prove in tutorial problems

$$e^{j\pi} + 1 = 0$$

One of magic formulas of all time

Discrete-time Fourier Series

If $x[n]$ is periodic with period N , it can be expressed as the sum of scaled periodic complex exponentials:

(can express any periodic fn as sum of scaled, complex exponentials)

Complex exponential with period N and fundamental frequency $2\pi/N$.

$$x[n] = \sum_{k=-\infty}^{\infty} a_k e^{jk\left(\frac{2\pi}{N}\right)n}$$

constant that changes complex

k ranges over any N consecutive integers. Two common choices:

- k starts at 0 ($0 \leq \text{freq} \leq 2\pi$) to $N-1$
- k starts at $-N/2$ ($-\pi \leq \text{freq} \leq \pi$)

The spectral coefficients a_k for each of the discrete frequencies are, in general, complex, changing both the amplitude and phase of the associated complex exponential. If $x[n]$ is real, $a_k = a_k^*$.

does not matter where start/stop

Use to go back/forth time + freq terms need to know a_k solve for

6.02 Spring 2011

Lecture 14, Slide #10

Solving for the a_k

Start with:

$$x[n] = \sum_{k=-\infty}^{\infty} a_k e^{jk\left(\frac{2\pi}{N}\right)n}$$

Multiply both sides by $e^{-jr(2\pi/N)n}$ and sum over N terms:

$$\begin{aligned} \sum_{n=-\infty}^{\infty} x[n] e^{-jr\left(\frac{2\pi}{N}\right)n} &= \sum_{n=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} a_k e^{jk\left(\frac{2\pi}{N}\right)n} e^{-jr\left(\frac{2\pi}{N}\right)n} \\ &= \sum_{k=-\infty}^{\infty} a_k \sum_{n=-\infty}^{\infty} e^{j(k-r)\left(\frac{2\pi}{N}\right)n} \end{aligned}$$

sum over N consecutive integers

this whole thing

From slide 9: N if $k=r$, 0 otherwise

trying diff ks - mostly 0 - except when $k=r$ get N

$$a_k = \frac{1}{N} \sum_{n=-\infty}^{\infty} x[n] e^{-jk\left(\frac{2\pi}{N}\right)n}$$

6.02 Spring 2011

Lecture 14, Slide #11

Discrete-time Fourier Series Pair

$$x[n] = \sum_{k=-\infty}^{\infty} a_k e^{jk\left(\frac{2\pi}{N}\right)n}$$

freq to time
Synthesis equation

$$a_k = \frac{1}{N} \sum_{n=-\infty}^{\infty} x[n] e^{-jk\left(\frac{2\pi}{N}\right)n}$$

spectral coeff.

time to freq
Analysis equation
fourier series pair

If we have N samples of a periodic waveform of period N , we can find the waveform's spectral coefficients using the analysis equation.

If we have the spectral coefficients, we can reconstruct the original time-domain waveform using the synthesis equation.

6.02 Spring 2011

Lecture 14, Slide #12

avg value sinusoid = 0

What are the a_k - this is the
 $x[n] = \cos(r \frac{2\pi}{N} n)$ crank the math way

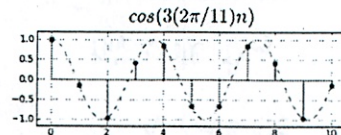
$$a_k = \frac{1}{N} \sum_{n=0}^{N-1} \cos(k \frac{2\pi}{N} n) e^{-jk \frac{2\pi}{N} n}$$

$$= \frac{1}{2N} \sum_{n=0}^{N-1} \left[e^{j(r-k) \frac{2\pi}{N} n} + e^{-j(r+k) \frac{2\pi}{N} n} \right]$$

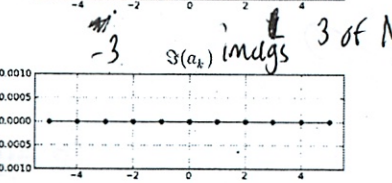
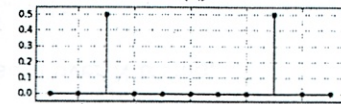
$$= \frac{1}{2N} \sum_{n=0}^{N-1} e^{j(r-k) \frac{2\pi}{N} n} + \frac{1}{2N} \sum_{n=0}^{N-1} e^{-j(r+k) \frac{2\pi}{N} n}$$

$$a_k = \begin{cases} \frac{1}{2} & k = \pm r \\ 0 & \text{otherwise} \end{cases}$$

1st spectral coefficient = $1/2$



Reals



Lecture 14, Slide #13

will write on p-set

r = index
 = # cycles in N samples

$$x[n] = 1 + 2 \cos(3 \frac{2\pi}{11} n) - 3 \sin(5 \frac{2\pi}{11} n)$$

Again, by inspection: since the cos and sin are different frequencies, we can analyze them separately.

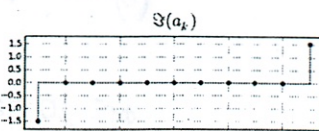
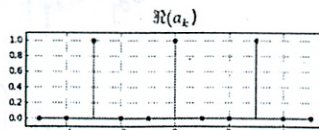
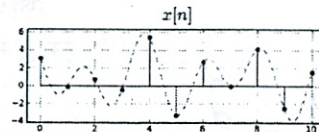
a_0 = average value = 1

$a_{\pm 3} = 2(1/2) = 1$ [from cos term]

$a_{-5} = -3(j/2) = -1.5j$ [from sin term]

$a_5 = -3(-j/2) = 1.5j$

$a_k = 0$ otherwise



Lecture 14, Slide #15

use superposition to solve
 behaving the way we want - getting energy at freq we want.

$$x[n] = \sin(r \frac{2\pi}{N} n)$$

This time let's do it "by inspection".
 First rewrite $x[n]$ (see slide #8):

$$x[n] = \frac{1}{2j} e^{jr \frac{2\pi}{N} n} - \frac{1}{2j} e^{-jr \frac{2\pi}{N} n}$$

Now $x[n]$ is a sum of complex exponentials and we can determine the a_k directly from the equation;

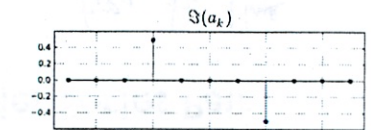
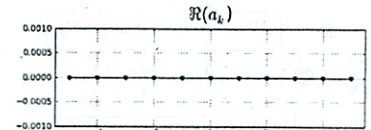
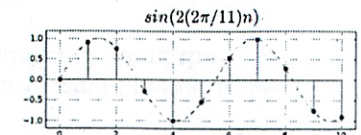
$$a_r = \frac{1}{2j} = -\frac{j}{2}$$

$$a_{-r} = -\frac{1}{2j} = \frac{j}{2}$$

$$a_k = 0 \text{ otherwise}$$

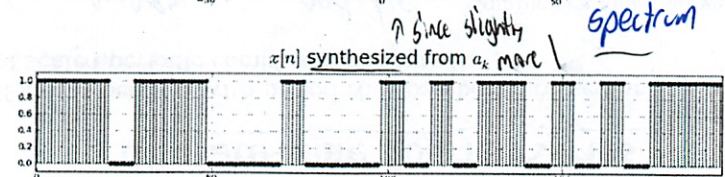
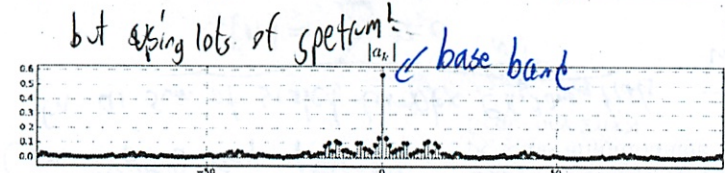
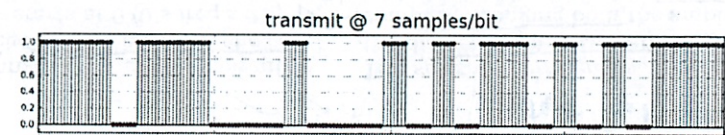
cleverer way

avg will be 0



Lecture 14, Slide #14

Spectrum of Digital Transmissions



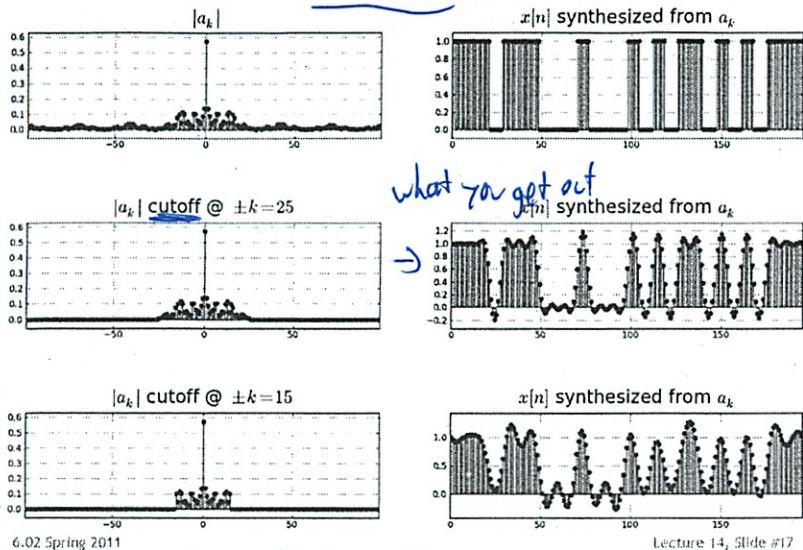
Lecture 14, Slide #16

the spectral coefficients for the seq

try to
limit
freqs

to stay on spectrum

Effect of Band-limiting a Transmission



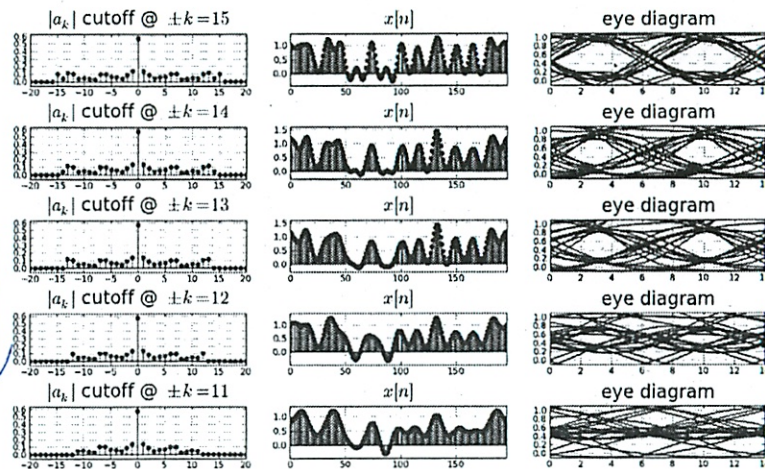
what you get out

→

limit freq used

square wave turning to sin wave - like if sin
since high freq responsible for the limited transitions
- like when it was slow channel

How Low Can We Go?



7 samples/bit → 14 samples/period → $k = (N/14) = (196/14) = 14$

6.02 Spring 2011

Lecture 14, Slide #18

Where do we expect eye to disappear?

|||||.....|||||
14 = period

How many 14s in 196

$$k = \frac{196}{14} = 14$$

↑
Coincidence

throw
away
more &
more

~~0/0~~

$$e^{j\theta} = \cos \theta + j \sin \theta$$

$$j = i$$

↑ the (S) notation

$$\underbrace{x}_{\text{real}} + j \underbrace{y}_{\text{complex}}$$

modulus $|x + jy| = \sqrt{x^2 + y^2}$

Remember $\cos^2 \theta + \sin^2 \theta = 1$

Can write \cos, \sin in form of complex #

$$e^{j\theta} = \cos \theta + j \sin \theta$$

(a)

$$e^{-j\theta} = \cos(-\theta) + j \sin(-\theta)$$

$$= \cos \theta - j \sin \theta$$

(b)

If add - will cancel - get formula for \cos

$$\frac{e^{j\theta} + e^{-j\theta}}{2} = \cos \theta$$

$$\frac{(a) + (b)}{2}$$

$$\frac{e^{j\theta} - e^{-j\theta}}{2j} = \sin \theta$$

$$\frac{(a) - (b)}{2j}$$

②

Represent time series in summation of complex exponentials

Look at periodic \rightarrow it repeats

$x[n]$ period N

$$x[n] = x[n+N] \text{ for all } n$$

If want can represent as

$$x[n] = \sum_{k \in \langle N \rangle} a_k \cdot e^{j \cdot \frac{2\pi}{N} \cdot k \cdot n}$$

\hookrightarrow summation of N functions,
each function is nice

(can think of

$$-\frac{N}{2} \quad \frac{N}{2} - 1$$

or
$$0 \quad N-1$$

(choose N consecutive things

In class defined

$$a_k = \frac{1}{N} \sum_{n \in \langle N \rangle} x[n] e^{-j \cdot \frac{2\pi}{N} \cdot k \cdot n}$$

(d)
Inverse
DFT

*

Lots of devices use

Say get signal w/ period N

- get $\hat{x}[n]$

(c)
Distributed
Fourier \leftarrow spelling
Transform

③

- Calculate a_k
How many calculations?

Each a_k has $x[n]$ terms

- each term 2 multiplications + a few additions

If have to do N things then

$$\cancel{N} \cdot N + \underbrace{c}_{\text{forget}}$$

Fast Fourier Transform

- does $c + N \log N$

- is very, very important

To get (d) out of (c)
(c) $\times e^{j \cdot r \cdot \frac{2\pi}{N} \cdot n}$

$$\sum_{n \in \mathbb{Z}/N\mathbb{Z}} x[n] e^{j \cdot r \cdot \frac{2\pi}{N} \cdot n} = \sum_{n \in \mathbb{Z}/N\mathbb{Z}} \sum_{k \in \mathbb{Z}/N\mathbb{Z}} a_k e^{j \cdot \frac{2\pi}{N} \cdot n \cdot (k-r)}$$

Exchange index k, n

$$= \sum_{k \in \mathbb{Z}/N\mathbb{Z}} a_k \left(\sum_{n \in \mathbb{Z}/N\mathbb{Z}} \underbrace{\left(e^{j \cdot \frac{2\pi}{N} \cdot (k-r)} \right)^n}_{\alpha} \right)$$

$$= \sum_{n \in \mathbb{Z}/N\mathbb{Z}} \alpha^n$$

(4)

$$= \sum_{n=0}^{N-1} \alpha^n$$

$$= \begin{cases} N & \text{if } \alpha = 1 \\ 0 & \text{if } \alpha \neq 1 \end{cases}$$

α will be 1

when $k-r=0$

$$k=r$$

$$\text{term} = N$$

entire thing reduces to

$$= a_r \cdot N$$

What we were looking for to find (d)

When $\alpha \neq 1$ - $k-r \neq 0$ - some complex #

- modulus = 1

but does ~~not~~ $\neq 1$

$$\sum_{n=0}^{N-1} \frac{(1-\alpha)\alpha^n}{(1-\alpha)} = \frac{1}{1-\alpha} \sum_{n=0}^{N-1} (\alpha^n - \alpha^{n+1})$$

$$= \frac{1}{1-\alpha} (1 - \alpha + \alpha - \alpha^2 + \alpha^2 - \alpha^3 + \alpha^3 - \dots + \alpha^{N-1} - \alpha^N)$$

5

Caroles So

$$= \frac{1}{1-\alpha} (1-\alpha^N)$$

$$\alpha^N = \left(e^{j \cdot l \cdot \left(\frac{2\pi}{N} \right)} \right)^N$$

$$= e^{j \cdot l \cdot (2\pi)}$$

$$= \cos(2\pi l) + j \sin(2\pi l)$$

$$= 1$$

$$x[n] = \cos \left(\frac{2\pi}{N} \cdot r \right)$$

$$0 \leq r \leq N-1$$

Using first formula

$$= \frac{1}{2} e^{j \cdot \frac{2\pi}{N} \cdot r \cdot n} + \frac{1}{2} e^{j \cdot \frac{2\pi}{N} \cdot (-r) \cdot n}$$

Imp

Implies that

$$a_k = \begin{cases} \frac{1}{2} & \text{if } k=r, k=-r \\ 0 & \text{otherwise} \end{cases}$$

cos only has this real

same

6

Same for \sin

$$x[n] = \sin\left(\frac{2\pi}{N} \cdot r \cdot n\right)$$

$$= \frac{1}{2j} \cdot e^{j \cdot \frac{2\pi}{N} \cdot r \cdot n} - \frac{1}{2j} \cdot e^{j \cdot \frac{2\pi}{N} \cdot (-r) \cdot n}$$

$$\hookrightarrow a_k = \begin{cases} \cancel{\frac{1}{2j}} - \frac{j}{2} & \text{if } k=r \\ \frac{j}{2} & k=-r \\ 0 & \text{otherwise} \end{cases}$$

Compositions

$$x[n] = x_1[n] + x_2[n]$$

both period N

Super imposition

Say did FT of x_1 and x_2

$\hookrightarrow \text{got } (a_k^1)_{k \in \mathbb{Z}N}$

$\hookrightarrow \text{got } (a_k^2)_{k \in \mathbb{Z}N}$

So $x[n]$

- could just compute

$-(a_k)_{k \in \mathbb{Z}N}$

⑦

Or just add them

$$a^k = a^1_k + a^2_k$$

$$= \sum_k a^1_k e^{j \frac{2\pi}{N} \cdot k \cdot n} + \sum_k a^2_k \cdot e^{j \frac{2\pi}{N} \cdot k \cdot n}$$

$$= \sum_k (a^1_k + a^2_k) e^{j \frac{2\pi}{N} \cdot k \cdot n}$$

by def
 a_k

Now have a toolkit for FT

Find each w/ formula

Then add them up

Other tutorial problems

#3 $x[n] = 1 + \sin\left(\frac{2\pi}{N} \cdot n\right) + 3 \cos\left(\frac{2\pi}{N} \cdot n\right) + \cos\left(2 \cdot \frac{2\pi}{N} \cdot n + \frac{\pi}{2}\right)$

~~signal~~
signal fn

What FT.

8)

First - Calc F. coefficients for each term
 Then - add them up

Constant Term

- exactly $\cos(0)$

$$x[n] = 3 y[n]$$

$$a_k = 3 a_k$$

So write 1 as

$$\cos\left(\frac{2\pi}{N} \cdot 0 \cdot n\right)$$

So from above

$$a_0 = 1$$

$$a_{\text{other}} = 0$$

$$\left\{ \begin{array}{l} \text{Technically} \\ a_k = 0 \text{ for all } k \neq 0 \end{array} \right.$$

2nd Term

F. coefficients

$$r = 1$$

$$\text{So } a_k = \begin{cases} -\frac{1}{2} & \text{if } k=1 \\ \frac{1}{2} & k=-1 \\ 0 & \text{otherwise} \end{cases}$$

9

Third Term

$$a_k = \begin{cases} \frac{3}{2} & k = \pm 1 \\ 0 & \text{otherwise} \end{cases}$$

← multiplied by 3

Fourth Term

not standard term

So use trig to re-jigger

$$\cos\left(\theta + \frac{\pi}{2}\right) = -\sin\theta$$

$$\text{So } -\sin\left(2 \cdot \frac{2\pi}{N} \cdot n\right)$$

$$\text{So } a_k = \begin{cases} +\frac{j}{2} & k=2 \\ -\frac{j}{2} & k=-2 \\ 0 & \text{otherwise} \end{cases}$$

Now add

each k }
$$\begin{aligned} a_0 &= 1 + 0 + 0 + 0 = 1 \\ a_1 &= 0 + -\frac{j}{2} + \frac{3}{2} + 0 = \frac{3}{2} - \frac{j}{2} \\ a_2 &= 0 + 0 + 0 + \frac{j}{2} \\ a_{-1} &= 0 + \frac{j}{2} + \frac{3}{2} + 0 = \frac{3}{2} + \frac{j}{2} \\ a_{-2} &= 0 + 0 + 0 - \frac{j}{2} = -\frac{j}{2} \\ a_k \text{ for } k \neq -2, -1, 0, 1, 2 &= 0 + 0 + 0 + 0 = 0 \end{aligned}$$

10

B, C, D pretty much same thing

F) $x[n]$ has exactly 1 non 0 value

$$x[0] \dots x[m] \dots x[N-1]$$

\uparrow
 $\neq 0$

Can you compute modulus of all F coefficients

$$a_k = \frac{1}{N} \sum_{n \in \mathbb{Z}_N} x[n] \cdot e^{-j \cdot \frac{2\pi}{N} \cdot k \cdot n}$$

Only term that matters is when it is not 0

$$= \frac{1}{N} \cdot x[m] \cdot e^{-j \cdot \frac{2\pi}{N} \cdot k \cdot m}$$

Compute modulus of that thing

$$|a_k| = \frac{1}{N} \cdot |x[m]| \cdot |e^{j \dots}|$$

$$= \frac{|x[m]|}{N} \underbrace{1}_{\text{all } k \text{ s are } = \text{ modulus}}$$

Something sparse in time domain - becomes full in freq domain

Heisenberg uncertainty principle is essentially this

- You can not measure 2 qv exactly correctly, at same time

- moment
- position

⑪

If you knew time exactly

- freq would be spread out

Do 4, 5 on own.

(Don't really get, so should do)

To save your work, click the SAVE button at the bottom of this page. You can revisit this page, revise your answers and SAVE as often as you like.

To submit the assignment, click the SUBMIT button at the bottom of this page. YOU CAN SUBMIT ONLY ONCE. Once the assignment has been submitted, you can continue to view this page but will no longer be able to make any changes to your answers.

6.02 Spring 2011: Plasmeier, Michael E.

PSet PS6

Dates & Deadlines

issued: Mar-16-2011 at 00:00

due: Mar-31-2011 at 06:00

checkoff due: Apr-05-2011 at 06:00

Help is available from the staff in the 6.02 lab (38-530) during lab hours -- for the staffing schedule please see the [Lab Hours](#) page on the course website. We recommend coming to the lab if you want help debugging your code.

For other questions, please try the 6.02 on-line Q&A forum at [Piazza](#).

Your answers will be graded by actual human beings, so your answers aren't limited to machine-gradable responses. Some of the questions ask for explanations and it's always good to provide a short explanation of your answer.

Aloha

Problem 1.

Consider a shared medium with N backlogged nodes running the slotted Aloha MAC protocol without any backoffs. A "wasted slot" is one in which no node sends data. The fraction of time during which no node uses the medium is the "waste" of the protocol. Each packet is 1 time slot long.

A. If the sending probability is p , what is the waste?

$$\text{Waste} = 1 - \text{utilization}$$

$$U = N p (1-p)^{N-1}$$

(points: 0.5)

B. Suppose N is large. If the Aloha sending probability, p , for each node is picked so as to

maximize the utilization, what is the probability of a collision?

have waste $1 - \frac{1}{e}$
 but what is $p(\text{collision})$? - don't see anywhere!
 i assume it is all waste!
 Or think fresh - though remember seeing
 (points: 0.5) $P(\text{collision}) = 1 - P(\text{noe}) - P(\text{one 1 ie success})$

Problem 2.

Alyssa and Ben are the only two users on a shared medium broadcast network running a variant of slotted Aloha. Their computers are configured such that Alyssa is 1.5 times as likely to send a packet as Ben. Assume that both computers are backlogged and that each packet is one slot long.

- A. For Alyssa and Ben, what is their probability of transmission such that the utilization of their network is maximized? Please give numeric answers.

$U_{\max} = \frac{1}{N}$
 - but computers set wrong
 $A = \frac{2}{3}$ $B = \frac{1}{3}$ no not right!
 (points: 0.5) $A = .6$ $B = .4$

- B. What is the maximum utilization? Please give a numeric answer.

$U = Np(1-p)^{N-1}$ think about it - 6.042!
 but w/ diff p 's
 $\frac{1}{2}(1 - \frac{1}{3}) + \frac{1}{3}(1 - \frac{1}{2}) = .52$
 (points: 0.5)

Problem 3.

You have two computers, A and B, sharing a wireless network in your room. The network runs the slotted Aloha protocol with equal-sized packets; each packet is 1 slot long. You want B to get twice the throughput over the wireless network as A whenever both nodes are backlogged. You configure A to send packets with probability $p_a = 0.25$. What should you set the transmission probability of B (p_b) to, in order to achieve your throughput goal? Assume that if exactly one node sends a packet in a time slot, it will be received successfully, but if both nodes do so, then neither packet will be received successfully.

Want B to get twice throughput
 What does that mean? B is ~~more~~ successful more often? Just says x;
 Or b transmits most of the time? 1.75?

(points: 1) How calc throughput? ^{successful} transmissions over time

Problem 4. B successful when $p_B(1-p_a)$ so $2p_a(1-p_b) = p_b(1-p_a)$

Ben Bitdiddle sets up a shared medium wireless network with one access point and N client nodes. Assume that both the access point and the N client nodes are backlogged. Each of the N clients wants to send its packets to the access point; the access point's packets are destined to various clients. The network uses slotted Aloha with each packet fitting exactly in one slot. Ben sets the transmission probability, p , of each client node to $1/N$ and sets the transmission probability of the access point to a value p_a .

- A. Determine the utilization of the network in terms of N and p_a . Hint: when $N=4$ and $p_a=0.25$, the utilization is approximately 40%.

$$U = Np(1-p)^{N-1} + p_a(1-p)^N$$

just access point
 Oh client $\frac{1}{N}$ ap p_a

$$4 \cdot 0.25(1-0.25)^{4-1} \approx 42\%$$

$$Np(1-p)^{N-1}(1-p_a) + p_a(1-p)^N \approx 0.395$$

(points: 0.5)

- B. What is the utilization of the network when N is large?

$$\lim_{N \rightarrow \infty} \text{of above}$$

$$= 0.367$$

(points: 0.25)

- C. Suppose N is large. What value of p_a ensures that the aggregate throughput of packets received successfully by the N clients is the same as the throughput of the packets received successfully by the access point?

So first find client success

$$Np(1-p)^{N-1}(1-p_a) \stackrel{\text{AP success}}{=} p_a(1-p)^N$$

$p = \frac{1}{N}$ solve for p_a

- which is actually just utilization

(points: 0.25)

complicated thing

$$a = \left(\frac{N-1}{N}\right)^{-N} - Na\left(\frac{N-1}{N}\right)^N + a\left(\frac{N-1}{N}\right)^N + \left(\frac{N-1}{N}\right)^N$$

take $\lim_{N \rightarrow \infty}$

- no result found!

$$p_a = 0.25$$

$$p_a + p_b = 1$$

no!

$$1.5 - 0.5 = 1.75b$$

$$1.5 = 1.25b$$

$$b = 1.2$$

$$a = \frac{1}{2}$$

$$\frac{1-a}{e} = \frac{a}{e}$$

try other side again w/ x sign

$$\text{got } \frac{a}{e}$$

add sign

0 = not found

Try one side at a time - but has a!

Try limit first

From here on, only the client nodes are backlogged -- the access point has no packets to send. *do clients no that*
Each client node sends with probability p (don't assume it is $1/N$).

Ben Bitdiddle comes up with a cool improvement to the receiver at the access point. If exactly one node transmits, then the receiver works as usual and is able to correctly decode the packet. If exactly two nodes transmit, he uses a method to cancel the interference caused by each packet on the other, and is (quite remarkably) able to decode both packets correctly.

- C. What is the probability, P_2 , of *exactly* two of the N nodes transmitting in a slot? Note that we want the probability of *any* two nodes sending in a given slot.

with no AP

$$N^2 \cdot p \cdot p \cdot (1-p)^{N-2}$$

2 trans *all except 2 transmit*
in possible combos

(points: 0.5)

- D. What is the utilization of slotted Aloha with Ben's receiver modification? Write your answer in terms of N , p , and P_2 , where P_2 is defined in the problem above.

am I mixing up p and v?

$$U = 2P_2$$

that might be more correct (ie correct instead)

(points: 0.5)

Problem 5.

what is this again
Media access control - guessing general term

Alyssa P. Hacker is designing a MAC protocol for a network used by people who: live on a large island, never sleep, never have guests, and are always on-line. Suppose the island's network has N nodes and the island dwellers always keep *exactly some four of these nodes backlogged*. The nodes communicate with each other by beaming their data to a satellite in the sky, which in turn broadcasts the data down. If two or more nodes transmit in the same slot, their transmissions collide; however, the satellite uplink doesn't interfere with the downlink. The nodes on the ground cannot hear each other and each node's packet transmission probability is non-zero. Alyssa uses a slotted protocol with all packets equal to one slot in length.

- just normal?*
- A. For the slotted Aloha protocol with a fixed per-node transmission probability, p , what is the maximum utilization of this network? (Note that although there are N nodes in all, only some four of them are constantly backlogged.)

just normal $\frac{1}{N} = \frac{1}{4}$
ignoring how we know which of the 4

(points: 0.5)

- B. In this network, as mentioned above, four of the N nodes are constantly backlogged, but the set of backlogged nodes is not constant. Suppose Alyssa must decide between slotted Aloha with a transmission probability of $1/5$ or time division multiple access (TDMA) among the N nodes. For what N does the expected utilization of this slotted Aloha protocol exceed that of TDMA?

Strict TDMA

why this

$$U \text{ w/ } p = \frac{1}{5} = N \frac{1}{5} (1 - \frac{1}{5})^{N-1}$$
 TDMA U - depends how many want to transmit
 $N=4$ does not matter - or no it does
 $U_{TDMA} = \frac{4}{N} = \frac{4}{4} = 1$

(points: 0.5)

Problem 6.

Token-passing is a variant of a TDMA MAC protocol. Here, the N nodes sharing the medium are numbered $0, 1, \dots, (N-1)$. The token starts at node 0 . A node can send a packet if, and only if, it has the token. When node i with the token has a packet to send, it sends the packet and then passes the token to node $(i+1) \bmod N$. If node i with the token does not have a packet to send, it passes the token to node $(i+1) \bmod N$. To pass the token, a node broadcasts a token packet on the channel and all other nodes hear it correctly.

A data packet occupies the channel for time T_d . A token packet occupies the channel for time T_k . If s of the N nodes in the network have data to send when they get the token, calculate the utilization of the channel in terms of the parameters above. Note that the bandwidth used to send tokens is pure overhead; the throughput we want corresponds to the rate at which data packets are sent.

Hint: When 20% of the nodes have data to send (i.e., $s/N = 0.2$) and $T_d = 10 \cdot T_k$, the utilization is $2/3$.

$U = \frac{T_d}{T_d + T_k}$ but how to include which nodes active?
 $T_k = \frac{1}{N-s}$ Oh # to send = s
 $(N-s)$ don't send

(points: 1)

stumper!
 $\frac{1}{N-s}$ should really be able to figure out!

$$\frac{T_d}{T_d + [(N-s) + 1] T_k}$$

$$T_d = 10 T_k$$

$$\frac{10 T_k}{10 T_k + (N-s+1) T_k}$$

test $\frac{10}{10 + 9.1} = 52\%$
not close

Introduction to this week's Python tasks.

This lab uses **WSim**, a simple packet-level network simulator for a shared medium network. You will be writing a small amount of code to develop various MAC protocols and measure how they perform under different conditions. Much of your work will be on experimenting with various parameters and explaining what you observe.

In each experiment, all the nodes run the same MAC protocol. The simulator executes a set of steps every time slot; time increments by 1 each slot.

You can run the python programs for this lab using python from the command line, e.g., `python PS6_.py`. **This lab does not work well in IDLE (you can use IDLE to edit files, but running them may not work as expected).**

To understand the different parameters one can set in WSim, go to a shell (i.e., command line prompt) and enter:

`python PS6_1.py -h`

might need to go to lab

This command prints out the various options; the important ones are:

1. **Retry:** If two or more nodes are actively sending a packet in the same time slot, they collide and both packets are considered lost. The `-r` option decides whether the node should *retry* the packet or not upon failure. In WSim, the feedback about whether a packet succeeded or not (i.e., collided) is instantaneous, with the sending node discovering it in the same time slot as the transmission. By default, the retry option is "off". When it is turned on, at an *offered* load of 100% (which is what we will use in this lab), the **actual** load presented to the system **exceeds** the channel's maximum rate. That is, we would expect most queues to be backlogged most of the time with these settings.
2. **Packet size:** In any given experiment, the size of a packet is fixed. It has to be an integral number of time slots in size (1 or more). To set the packet size, use the `-s` option; the default is 1. Notice that setting a large packet size (say, 10) emulates an "unslotted" network.
3. **Skew:** The `-k` option specifies whether the load is skewed or not. The load itself is generated according to a random process, whose details aren't important for this lab. By default, the skew is off, so all nodes generate the same load on average. When the `-k` option is set, then the total offered load, L , is divided in geometrically-spaced amounts. Node 0 presents a load of $L/2$, node 1 $L/4$, and so on. The last two nodes each present the same load, $L/2^{N-1}$, where N is the total number of nodes.
4. **Number of nodes:** The number of nodes in a run of WSim; default is 16. Set using `-n`.
5. **GUI:** The `-g` option turns on the graphical user interface, which may be of some use in debugging your code. **We recommend that you set the parameters for the simulation from the command line and NOT from the GUI, as the GUI's**

parameter setting code may not port well across different python installations.

Experimental method

WSim runs for a specified number of time slots (settable using the `-t` option, with a default of 10000) and prints out some performance numbers (and possibly displays some graphs) at the end. Of interest to us are the **utilization** and **fairness** numbers, which are defined in the lecture notes. The code reports a "weighted" fairness number as well, which is the fairness index calculated over the ratio of the observed throughputs to offered loads. The (unweighted) inter-node fairness is calculated over the throughputs alone, without regard to the offered load.

In this lab, you will implement the core of three MAC protocols---**TDMA**, **stabilized Aloha** (with backoffs), and **CSMA**---and understand their performance. Each node is an object, which has three methods that you can use to implement the core of the MAC protocol:

```
boolean = node.channel_access(time,ptime,numnodes)
```

This method is called by WSim every time slot when the node has a packet waiting to be sent in its queue. This method should return `True` if the MAC protocol you're implementing would like a packet sent in the current time slot, and `False` otherwise. `time` is current time, `ptime` is the packet size in time slots, and (for TDMA) `numnodes` is the number of nodes in the network.

```
node.on_collision(packet)
```

Called every time slot in which the node has experienced a collision.

```
node.on_xmit_success(packet)
```

Called every time slot in which the node has successfully sent a packet (i.e., no collisions occurred during transmission).

You can modify any per-node state that you want to in these methods (e.g., the state maintained by the backoff scheme, statistics of interest, etc.). Make sure to add the code to initialize this state in the Node object's `__init__` function, whose body is included in the lab task files.

In many of our experiments, we will use the `-r` option, which cause the nodes to retry upon experiencing a collision. (Of course, the MAC protocol's `channel_access()` method will determine when the retry actually occurs.)

Python Task #1: TDMA

Useful download links:

[PS6_wsim.py](#) -- packet-level simulator

[PS6_1.py](#) -- template file for this task

Implement a simple TDMA scheme by suitably filling in the `channel_access()` method in

the template file `PS6_1.py`. Recall that in a TDMA scheme, time is divided into `numnodes` equal-size slots, each long enough to accommodate the transmission of a single packet and that each node is allocated one of the slots to use when it has a packet to transmit. Note that a node can determine its unique node number (an integer between 0 and `numnodes-1`) by calling `self.get_id()`.

The slightly tricky part of this function is to correctly handle packet sizes that are larger than 1 time slot. We want the TDMA scheme to treat each packet as an atomic unit of transmission; when the protocol determines that a given node can send, that node should send the complete packet. Put another way, we want the effective size of a time slot in the scheme to be equal to the packet size. You will probably find it easier to first write the function and run it for a packet size of 1 slot, then modify your code to correctly handle larger packet sizes. Note that when the packet size is set to some value using the `-s` option, all nodes will use that value.

Run the following (after you test it for various packet sizes to ensure that there are no collisions):

- `python PS6_1.py -t 2000`
(16 nodes, packet size = 1 slot, simulation time = 2000 slots)
- `python PS6_1.py -s 7 -t 14000`
(16 nodes, packet size = 7 slots, simulation time = 14000 slots)
- `python PS6_1.py -k -n 20`
(20 nodes, skewed load, packet size = 1 slot, simulation time = 10000, the default value)

When you're ready, please submit the file with your code using the field below.

File to upload for Task 1:

Browse...

(points: 1)

Questions:

- A. Please run the following experiments using a skewed load (`-k`) where the load offered by a node decreases with the node number, i.e., high-numbered nodes have a packet to transmit much less frequently than low-numbered nodes.

```
python PS6_1.py -k -n 10
python PS6_1.py -k -n 20
python PS6_1.py -k -n 40
python PS6_1.py -k -n 80
```

Please report the utilization from each experiment (look for `util` in the printout). With a skewed load, as one increases the number of nodes, what happens to the utilization? Why?

↓ b/c share slot

(points: 0.5)

- B. What is the number of nodes at which the network utilization is smaller than 0.25 for the skewed workload? (Because each run is randomized, run it a few times to be confident of your answer.)

Just guess + check

(points: 0.5)

Python Task #2: Aloha with Fixed Probability of Sending

Useful download link:

PS6_2.py -- template file for this task

In this task, we will implement slotted Aloha with a fixed transmission probability and measure its performance under different conditions. The program allows you to set the transmission probability using the -p option. Looking at PS6_2.py, note that the __init__ function of AlohaWirelessNetwork sets each node's "p" to the configured value of the transmission probability. Your task is to implement the Aloha MAC protocol by providing the correct code in channel_access(), so that the node will transmit packets with the configured probability. Your scheme should work when a packet is 1 slot long, but also for longer packet sizes (of course, the utilization may be different for different packet sizes).

Please note: For this task as well as the subsequent ones, do not use numnodes in your code, even though it is accessible. The reason is that we want your algorithm to work for arbitrarily distributed offered loads, and using numnodes will not help achieve that.

Test your code by running the following for different values of p and observe the resulting utilization. These tests use the -r option to force multiple nodes to usually be backlogged. With this option, each collision causes a retry. Because the (default) offered load is 100%, the retries ensure that the total offered load exceeds the channel rate. You can also observe this backlog if you run the following tests with the -g option and pay attention to the queue lengths at the nodes.

`python PS6_2.py -n 16 -r -p value`

What is this?
GVI

When you're ready, please submit the file with your code using the field below.

File to upload for Task 2:

Browse...

(points: 1)

Questions:

- A. Experiment with different values of p and observe the resulting utilization. What value of p in the range (0,1) maximizes the utilization in your experiments? What is the maximum utilization?

$1/16$

(points: 0.33)

- B. We will now explore a skewed load to see whether that affects the best choice of p . Run the following for a few values of p .

```
python PS6_2.py -n 16 -r -k -p value
```

What value(s) of p in the range (0,1) maximizes the utilization in your experiments in this case? What is the maximum utilization?

.0625 = .25	.2 = .40	.23 = .4
.1 = .32	.25 = .39	.21 = .41
.15 = .37	.22 = .41	

(points: 0.33)

- C. How do the optimal values of p and the corresponding utilization compare with the no-skew case? Why are they different?

(points: 0.34)

Python Task #3: Stabilizing Aloha (with Backoff)

Useful download link:

PS6 3.py -- template file for this task

In this task, we will develop a stablization method for Aloha using randomized backoffs to replace the fixed probability of Task #2. Our goal is to adaptively select the transmission attempt probability, p , used in the `channel_access` method. To do that, write your code in `PS6_3.py` to adjust p in the `on_collision` and `on_xmit_success` methods, which are called when a packet transmission fails and succeeds, respectively.

We will use two parameters, p_{\max} and p_{\min} . These correspond to the maximum and minimum values of the transmission attempt probability, p . The values of these parameters can be set from the command line when you run the program, and are available as `self.network.pmax` and `self.network.pmin` respectively (see the `__init__` function of `AlohaWirelessNetwork`). In your code, ensure that $p_{\min} \leq p \leq p_{\max}$.

You can use any algorithm you want to set p in these functions, including the ones discussed in lecture. Good schemes achieve high utilization, but ensure also that fairness is high (as close to 1 as possible -- lower than 0.9 is a sign that there is significant unfairness when the number of nodes is between 8 and 16) and avoid the capture effect. There is no absolute correct answer (though there are bad methods!), so feel free to be creative if you think you have a good idea. Note that you are not allowed to use the number of backlogged nodes in your scheme, because that information would not be available in practice.

Run your code as follows for a few different settings of p_{\min} and p_{\max} and observe the utilization and inter-node fairness values.

```
python PS6_3.py -r -n 8 --pmin=value --pmax=value
```

(Note the two dashes in front of the p_{\max} and p_{\min} options.)

When you're ready, please submit the file with your code using the field below.

File to upload for Task 3:

Browse...

(points: 1)

Questions:

- A. Run `python PS6_3.py -r -n 8 --pmin=0 --pmax=1`. Would you recommend running a real network with these parameters for p_{\min} and p_{\max} ? Briefly explain your answer.

No - some hit 1 + dominate

(points: 0.5)

p change?
have to find

lecture notes
multiplicative

but still $v \approx 0.4$
which is what I got

- B. Set p_{\min} to 0.01 and pick p_{\max} so that the fairness is as large as possible when the number of nodes is 8 and there is no load skew. What is the utilization of your protocol when the packet size is 10 slots? How does it compare to the utilization when the packet size is 1?

For the first case (packet size of 10), run

```
python PS6_3.py -s 10 -t 70000 -r -n 8 --pmin=0.01 --pmax=value
```

For the second case (packet size of 1), run

```
python PS6_3.py -r -n 8 --pmin=0.01 --pmax=value
```

Additive

$$\begin{aligned} \max &= 2 \rightarrow U = .39 \quad D = .01 \\ &12 \quad .39 \quad \Delta = .01 \end{aligned}$$

Same - wrong!

(points: 0.5)

Python Task #4: CSMA

Useful download link:

PS6_4.py -- template file for this task

Multi

18 max	$U = .55$	$f = .75$
18 max	$U = .56$	$U = .77$

In this task, we will try to get the best utilization and fairness we can for a MAC protocol that uses CSMA. To check if the channel is idle, you can use the

`self.network.channel_idle()` from inside `channel_access()`. This function returns True if there is no on-going transmission in the **current time slot**, and False otherwise.

Every carrier sensing mechanism has a *detection time*, defined as the time interval between the ending of a previous transmission and the detection of the channel as "idle" by a node.

For the purposes of this lab, we will assume that the detection time is 0. Hence, a node can sense that the carrier is idle in the immediate next slot after the termination of the previous transmission, when it does the check for whether the channel is busy. However, it is still possible for collisions to occur, for multiple nodes could simultaneously sense the channel at the beginning of a slot, and conclude that the channel is "idle", and possibly attempt a transmission in that time slot (or in the same future time slot).

Write your code PS6_4.py for the `channel_access()` method assuming that the node has the ability to sense the carrier. Obviously, you should fill in the steps for the `on_collision()` and `on_xmit_success()` methods as well.

Test your code as follows. The `-s` option is important because it causes the packets to be longer than 1 slot, allowing a node to sense whether another transmission is in progress during a time slot.

065, 18 lecture

```
python PS6_4.py -r -n 8 -s 10 -t 100000 --pmin=value --pmax=value
```

Note that you should run the above for 100000 time slots, because we have scaled up the packet size to 10 (from 1).

When you're ready, please submit the file with your code using the field below.

File to upload for Task 4:

Browse...

(points: 1)

Questions:

A. What is the utilization and fairness of your protocol when $p_{\min} = 0$ and $p_{\max} = 1$?

*$U = .79$
 $f = .92$*

(points: 0.5)

B. How would you set p_{\min} and p_{\max} in your protocol to make the fairness number be over 0.95 consistently while still maintaining good utilization?

used in lecture

(points: 0.5)

Python Task #5: Practical CSMA, as in WiFi (802.11) and Ethernet (802.3)

Useful download link:

PS6_5.py -- template file for this task

The ALOHA and CSMA schemes in the previous tasks pick a probability p for transmitting a packet, and adapt p to stabilize the protocol. The advantage of this method is that it is easy to analyze. In practice, however, real-world CSMA protocols like the popular 802.11 WiFi standard and the 802.3 Ethernet standard implement something a bit different, as explained below. Your task will be to write the code for the key parts of this scheme.

Rather than decide whether any given slot should have a transmission with probability p , each node maintains a contention window, which we denote by cw . cw is initially set to

c_{wmin} , which is a small positive integer (say, 1). Denote the current time slot by C . If the sender is backlogged, it picks a random integer t in $[1, c_w]$ and decides to send a packet in time slot $C + t$. *we should do token*

Of course, with carrier sense in place, the sender should only send a packet if the channel is idle in time slot $C + t$. So, the sender senses the carrier in that time slot, and then sends a packet only if the channel is idle then. If the channel is not idle, it waits until the channel is idle, and then sends the packet. *wait some random time $[1, c_w]$*

Whenever a collision occurs, the node doubles c_w , but makes sure c_w never exceeds c_{wmax} (say, 512). Whenever a transmission is successful, the node might reset c_w to c_{wmin} , or might halve its current value of c_w . You will note that this scheme is similar in spirit to the probabilistic transmission scheme from Task #4, but a crucial difference is that here each backlogged node is **guaranteed** to send a packet within a finite time, unlike in the probabilistic case where there is always a small probability that the node cannot send within any given number of time slots. In mathematical terms, the probability distribution that governs whether a node transmits a packet in a given time slot is uniform in this scheme (Task #5), while it is geometrically distributed in the previous case (Task #4).

You will first implement the scheme described thus far. It will turn out not to do as well as we would like, and in Task 5.2, you will fix an important weakness.

Implement this scheme by writing the appropriate code for `channel_access()`, `on_collision()`, and `on_xmit_success()` in `PS6_5.py`. How well does it work? To answer this question, measure the utilization and fairness by running

```
python PS6_5.py -r -s 10 -n 16 -t 100000 -w 256
```

The `-w` option sets the maximum contention window size. The minimum contention window is 1 (you can change it using the `-w` option if you like (lower case "w")). Running the above, you will note that even with carrier sense being used, the utilization is quite a bit lower than in Task 4.

- A. Report the utilization and fairness. Briefly explain why the utilization is low. Hint: Think about what happens if more than one node is backlogged and waiting for an on-going transmission to complete; what happens when the on-going one finishes?

$$U = .78$$

$$f = .93$$

(points: 0.5)

To fix this problem, each node needs to ignore the time slots when other nodes are transmitting data. That is, if a node picks a time slot t in $[1, c_w]$ to transmit, it should wait for that many **idle** slots before attempting its own transmission. Of course, before transmitting data, it should ensure that the channel is idle.

Modify your code to include the above suggestion and run the same command as before:

```
python PS6_5.py -r -s 10 -n 8 -t 100000 -W 256
```

- B. Report the utilization and fairness for your scheme after including the suggestion and running the same command as before.

$U = .83$
 $f = .92$

(points: 0.5)

Please upload the final version of your Task #5 code:

File to upload for Task 5:

Browse...

(points: 1)

You can save your work at any time by clicking the Save button below. You can revisit this page, revise your answers and SAVE as often as you like.

Save

✓ 8:41 PM

To submit the assignment, click on the Submit button below. YOU CAN SUBMIT ONLY ONCE after which you will not be able to make any further changes to your answers. Once an assignment is submitted, solutions will be visible after the due date and the graders will have access to your answers. When the grading is complete, points and grader comments will be shown on this page.

Submit

6 cont

$$\frac{N \cdot T_d}{T_d + [N - s + 1] T_h}$$

for each to transmit

- but only $\frac{s}{N}$ want to transmit

$$\frac{s}{N} \cdot N = s$$

$$\frac{s \cdot T_d}{T_d + [N - s + 1] T_h}$$

$$\frac{2 \cdot 10}{19} \text{ No!}$$

s in denom

$$\frac{20}{29} = \frac{2}{3}!$$

$$\frac{s T_d}{s T_d + [N - s + 1] T_h}$$

Good to figure it out!

PS 6 - Lipy TDMA

- fill in channel_access()
- node gets # self. get Id()
- So $t_{mod} Id == 1$ - transmit
when
- Have functions > 1 time slot
 - so is that given by s_i
 - I guess
- What is skew?
 - unequal offered loads
 - so not backlogged
- Oh simulator runs for each node
- Instructions says use numnodes
- do $time \% numnodes == id$
b/c above way had made by 0
- better
- do you have to implement on_collision
on_xmit_access?

②

Look at error list

Everything is failing

Oh did not do == Id

Succeeded!

Now. multy thing lenght

- will always be thing
- So ~~and~~ take time % ptime == 1 - Then try to trans
- works for 7
- But fails for 1

Do time to ptime == 0

Finished 25 min

- slow, careful, relaxing in lab

③ #2 Slotted Aloha

- transmits probabilistic
- just generates p r#
- how to generate random #?
- decimal
- say draw decimal $[0,1]$
if $< p$ - fire
- Worked!
 - 15 min
 - 2 lines of code!
- Sim GUI time undefined
- I should look Piazza
- Just forget GUI
- done - whole thing 30 min

(4)

#3 Need to fix 2 for multi length
- seems to work pretty good

I wish it showed ending p

My system is always hitting wall - wrong way²

- no util = .48

#4 CSMA

what is this again?

Sense

self, network, channel_idle()

- True if current time slot

detection time = () here

- but multiple can still sense and hit

So send if idle

or 1 after idle

Just send if idle

Oh no still have to do p and p_{min}, max

9

No need to check time % ptime

- \uparrow U from .4 to .6
- if collision restart next slot trying, not waiting time

do max

- put at cap - don't skip

do multiplicative

- fairness when up much faster .6 \rightarrow .8
- and util .4 \rightarrow .55

Since who is backlogged changes over time

#5 still sense, stable

'where is cw'