Department of Computer Science and Engineering

CSE 3213: Communication Networks (Fall 2015)

Instructor: N. Vlajic
Date: Dec 13, 2015

Final Examination

Instructions:

- Examination time: 180 min.
- Print your name and CS student number in the space provided below.
- This examination is closed book and closed notes. Calculator and one-sided cheat-sheet are allowed.
- There are 8 questions. The points for each question are given in square brackets, next to the question title. The overall maximum score is 100.
- Answer each question in the space provided. If you need to continue an answer onto the last page, clearly indicate that and label the continuation with the question number.

FIRST NAME: ______________________________

LAST NAME: ______________________________

STUDENT #: ______________________________

Question	Points
1	/ 15
2	/ 8
3	/ 12
4	/ 17
5	/ 10
6	/ 11
7	/ 14
8	/ 13
Total	/ 100

1. Multiple Choice [15 points]

Circle the letter beside the choice that is the best answer for each question. For each question choose only ONE answer.

(1.1) An IP address is __________ bits long.

(a) 24
(b) 32
(c) 48
(d) 64

(1.2) The port number of the sending machine can generally be found in the header of __________ .

(a) HTTP packets
(b) TCP and UDP packets
(c) IP packets
(d) DNS packets

(1.3) ____________ is considered to be an advantage of full-mesh (direct link) infrastructure over switched infrastructure.

(a) scalability
(b) lower implementation complexity (in case of large networks)
(c) robustness to link failure
(d) none of the above

(1.4) We are transmitting data at 100W, and detect only 90W when receiving, thus the attenuation in decibels is close to:

(a) 2.0
(b) -1.5
(c) -0.5
(d) none of the above

(1.5) __________ is a type of transmission impairment in which the signal loses its form/shape due to the different propagation speeds of each frequency that makes up the signal.

(a) attenuation
(b) delay distortion
(c) thermal noise
(d) cross talk

(1.6) Which of the following line coding methods does not provide for synchronization in a long sequences of 1s?

(a) NRZ-L
(b) NRZ-I
(c) RZ
(d) Bipolar

(1.7) The efficiency of a polling-based system increases as:

(a) overall number of stations (active + inactive) increases
(b) average load per station decreases
(c) average load per station increases
(d) polling time increases

(1.8) Regardless of the actual load in the network, Slotted Aloha improves ______________ over that of regular Aloha.

(a) latency time
(b) propagation time
(c) vulnerable time
(d) frame time

(1.9) The physical length of an Ethernet network is limited by

(a) data rate
(b) packet size
(c) duration of back-off interval
(d) number of back-off intervals

(1.10) In which of the following scenarios Token Ring can be expected to provide better performance than Ethernet? (Assume each active station sends only 1 small-sized packet at a time.)

(a) a mid-size network with only a few active stations at any given time
(b) a mid-size network with a large number of active stations at any given time
(c) both (a) and (b)
(d) none of the above

2. Potpourri [8 points]

Given the following packet

foo header	fing header	yaya header	user data field	foo trailer
5 bytes	10 bytes	20 bytes	Maximum of 150 bytes	4 bytes

2.1 [4 points] Sketch the layered protocol model that applies to this packet by labelling each layer in the figure below with the appropriate layer name. Your choices are foo, fing, yaya, and user data

_______________	Layer
_______________	Layer
_______________	Layer
_______________	Layer

___user data__	**Layer**
_____yaya____	**Layer**
_____fing_____	**Layer**
_____foo______	**Layer**

2.2 [4 points] If the maximum length for the user data field is 150 bytes, what is the overhead (as a percentage) to send a 1600 byte user message?

We need ceil (1600 / 150) = 11 packets.
Each packet has 39 bytes of overhead.
So, 39×11 = 429 bytes of header and trailers (overhead) is required to send 1600 bytes of user data, or 429 / (1600 + 429) = 21.1% overhead

3. Channel Capacity and PCM [12 points]

A researcher wishes to digitally record analog sounds for testing animal hearing with frequencies of up to 100 [kHz].

3.1 [2 points]

What is the minimum sampling rate required to process these sounds?

2 [samples / Hz] * 100 [kHz] = 200 000 [samples / sec]

3.2 [2 points]

If a 16 bit (per sample) PCM (A/D) converter is used, what is the data rate of the resulting digital signal?

C = 16 [bit / sample] * 200 000 [sample / sec] = 3.2 [Mbps]

3.3 [4 points]

Use Shannon's formula to find the minimum signal to noise ratio (in dB!) required to sustain the given data rate over a 500 KHz radio channel.

Shannon capacity theorem is: $C = B \log_2 (1 + SNR)$

$SNR = 2^{C/B} - 1 = 2^{3.2Mbps/0.5MHz} - 1 = 2^{6.4} - 1 = 83.4$

SNR [dB] = 10 log(83.4) = 19.2

3.4 [4 points]

Assume in step 3.2) we decided to deploy Delta modulation instead of PCM. Which data rate would the resultant signal have now? Would you expect better or worst overall system performance? Explain!

In case of Delta modulation, 1 bit / sample is needed.
Hence, C = 1 [bit / sample] * 200 000 [sample / sec] = 0.2 [Mbps]

From Shannon's theorem, lower C implies lower required SNR. This means, the system would be able to tolerate more of the channel noise. Nevertheless, this would not immediately imply a better system performance. Ultimately, the quality of the recorded/reproduced analog system would depend on the selection of DM parameters (e.g. step size δ).

4. Error Control [17 points]

4.1 [4 points]

Two dimensional parity can catch all 3-bit errors, but not all 4-bit errors. By referring to the following diagram illustrate/show 2 different scenarios (i.e., 2 different sets of 4-bit errors) that 2D-parity will not be able to detect.

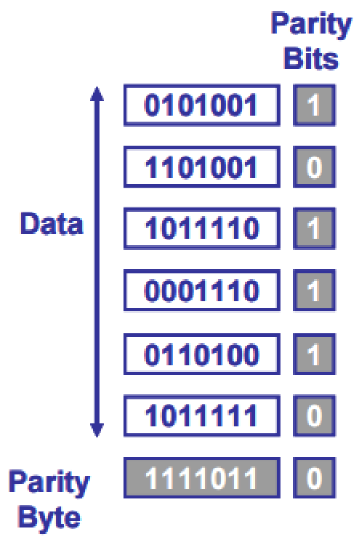

Solution:

Any scenario where the errors form 'corners of a rectangle' will work – two errors in the same row(s), and two in the same column(s).

4.2 [6 points]

The following are the code words from an 8-bit code that can be used either for error detection or error correction.

A 0000000
B 0001111
C 0110011
D 0111100
E 1010101
F 1011010
G 1100110
H 1101001

How many bit error, in particular, can this code correct?

Solution:

The Hamming distance matrix for the given codewords is shown below.

		A	B	C	D	E	F	G	H	
A	0000000 :	0	4	4	4	4	4	4	4	A
B	0001111 :	4	0	4	4	4	4	4	4	B
C	0110011 :	4	4	0	4	4	4	4	4	C
D	0111100 :	4	4	4	0	4	4	4	4	D
E	1010101 :	4	4	4	4	0	4	4	4	E
F	1011010 :	4	4	4	4	4	0	4	4	F
G	1100110 :	4	4	4	4	4	4	0	4	G
H	1101001 :	4	4	4	4	4	4	4	0	H
		A	B	C	D	E	F	G	H	

Given that the minimum Hamming distance of this code is 4, it means that the code can correct t = floor (d – 1)/2 = 1 - one error.

4.3 [7 points]

A message **10110101** is transmitted using the CRC polynomial method. The generator polynomial is $\mathbf{x^3 + 1}$.

What is the actual frame transmitted?

Solution:

5) **a)** The message 10110101 is $x^7 + x^5 + x^4 + x^2 + 1$ in polynomial representation. Degree of generator is 3, so multiply this polynomial by x^3 to obtain

$$x^{10} + x^8 + x^7 + x^5 + x^3$$

Now divide this by $x^3 + 1$. The result is $x^7 + x^5 + 1$, the remainder is 1. Therefore the transmitted message is:

$$x^{10} + x^8 + x^7 + x^5 + x^3 + 1$$

or, equivalently:

10110101001

5. Digital Modulation [10 points]

5.1 [3 points]

Sketch (approximately) BPSK modulated time signal for data stream **01001**, when the bit rate is **1 Mbps** and the carrier frequency is **2 MHz**.

Solution:

There will be 2 periods of carrier per bit duration.

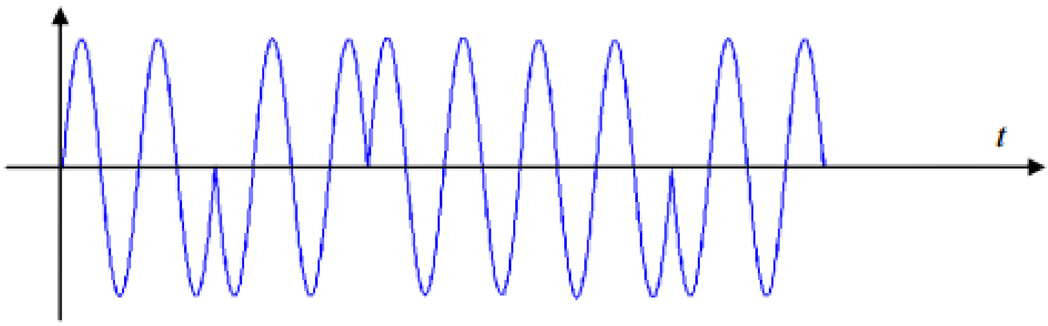

5.2 [2 points]

For which modulation technique the following waveform can be a typical example? (Circle the right answer.)

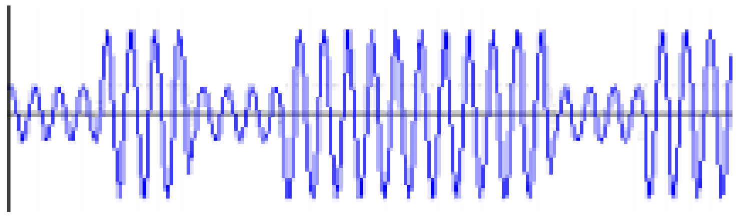

Figure 5.1

ASK

FSK

BSK

QAM

5.2 [5 points]

Which of the frequency spectrum characteristics shown in Figure 5.3 would most likely correspond to the FSK signal shown in Figure 5.2. Recall, in FSK, the carrier(s) are being modulated by a 'proper' digital signal. (Circle the best matching answer.)

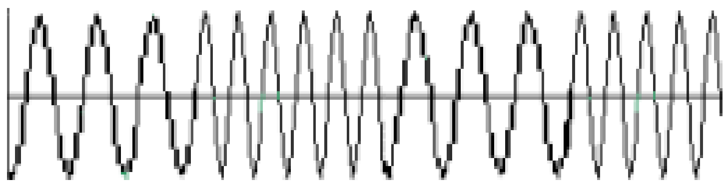

Figure 5.2

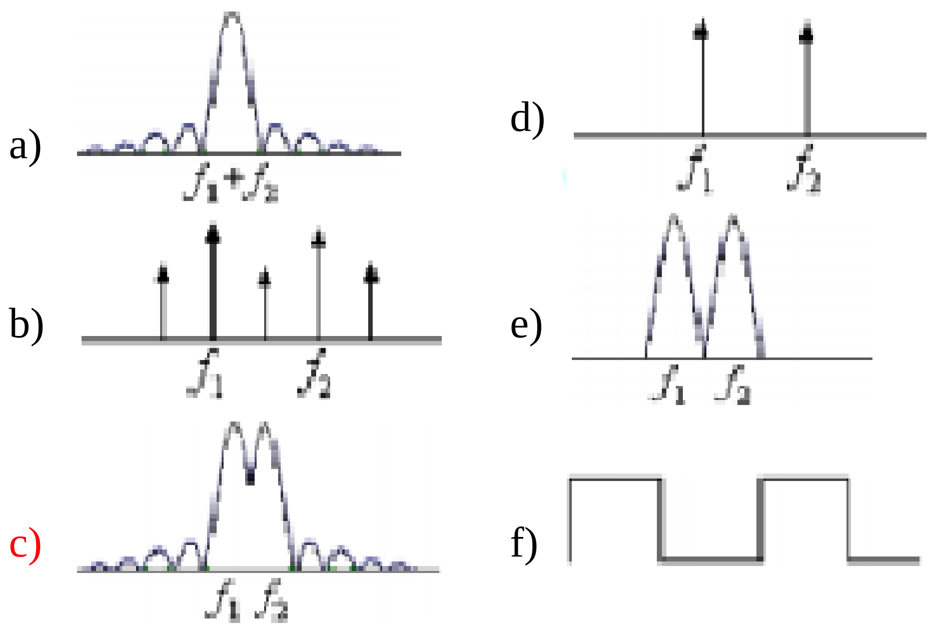

Figure 5.3

6. Random Access [11 points]

6.1 [8 points]

Consider the following multiple access scheme that combines TDMA and slotted-ALOHA. There are 30 users in total, separated in two groups: class-1 group consisting of 4 'prime' users, and class-2 group consisting of the remaining 26 'regular' users. Even time slots (i.e., 0, 2, 4, …) are reserved for the class-1 group, while odd time slots (i.e., 1, 3, 5, …) are reserved for the class-2 group. Contention within each group is resolved by the slotted-ALOHA protocol (e.g., when a user has something to send, it waits for the respective odd/even slot, and then transmits with a probability p).

Now, assume that every user always has something to send, and that each group of users use the optimal probability of transmission p (as discussed in class). How much more likely is a class-1 user to succeed in sending his packet compared to a class-2 user?

(Your answer can be in terms of fractions.)

Solution:

For slotted-ALOHA of N nodes, the probability of a successful transmission of a user in a group of N users is:

$$P_{success} = N*p*(1-p)^{N-1}$$

We have shown in class that the given expression will be maximized if **p=1/N**.

Consequently, the actual probability of success for

a) **Class-1 user is $P_{success}(1) = (1-1/4)^3 = (3/4)^3 = 0.421$**

b) **Class-2 user is $P_{success}(2) = (1-1/26)^{25} = (25/26)^{25} = 0.375$**

Based on these numbers, we conclude that the ratio of these probabilities $P_{success}(1)$ / $P_{success}(2)$ is 1.12.

Or, to put these number in the context of percentages, we could also calculate $P_{success}(1)$ – $P_{success}(1)$ / $P_{success}(1)$, and see that class-1 user is only about 10% more likely to succeed

6.2 [3 points]

Why is it necessary to pad an Ethernet frame to 64 bytes? Explain in no more than 3-4 sentences.

Solution:

To allow stations in the collision domain to transmit long enough in order to detect a collision. Namely, to detect a collision, a station must be in the position to be sending its own and receiving the interfering signal (i.e. compare the two) concurrently.

In the worst case collision scenario (A and B are at two ends of a network. B starts transmitting just a moment before A's packet arrives to B), a packet transmission time must be equal to the network's roundtrip delay.

In real world Ethernet networks, this time translates to 64 byte frame size. Hence, all Ethernet packets have to be padded to (i.e. extended to) this size …

7. Scheduled Access [14 points]

7.1 CDMA [9 points]

7.1.a) [3 points]

Suppose that three CDMA users – A, B and C – are all simultaneously transmitting bit 0, each using the following chip sequence:

A	-1	-1	-1	1	1	-1	1	1
B	-1	-1	1	-1	1	1	1	-1
C	-1	1	-1	1	1	1	-1	-1

What is the resultant chip sequence/signal observe on the transmission channel?

Solution:

Binary 0 is encoded as **-1** signal level. Accordingly, each of the users/stations generates the following resultant chip/signal sequence:

-A	1	1	1	-1	-1	1	-1	-1
-B	1	1	-1	1	-1	-1	-1	1
-C	1	-1	1	-1	-1	-1	1	1
resultant chip	3	1	1	-1	-3	-1	-1	1

7.1.b) [6 points]

A CDMA receiver gets the following chip sequence/signal: **S = -1 +1 -3 +1 -1 -3 +1 +1**. Assuming the chip sequences from 7.1.a), which station(s) transmitted and which bits did each one send?

Solution:

S*A = 1 - 1 +3 +1 -1 +3 +1 +1 = +8 - negative value ⇒ A sent binary 0

S*B = 1 -1 -3 -1 -1 -3 +1 -1 = -8 - positive value ⇒ B set binary 1

S*C = 1 +1 +3 +1-1 -3 -1 -1 = 0 - zero ⇒ C did not send anything

7.2 FDMA [5 points]

Suppose an operator is allocated 2.5 MHz of spectrum. The operator provides service that requires/assumes a 60 kHz half-duplex channel per each served user. A guard band of 30 kHz is needed between any two adjacent channels. (You can assume that the outskirts of the 2.5 MHz spectrum are already protected/isolated from the adjacent spectrums with 100 kHz guard bands, as shown in the below figure.)

What is the maximum number of users that this operator could serve simultaneously?

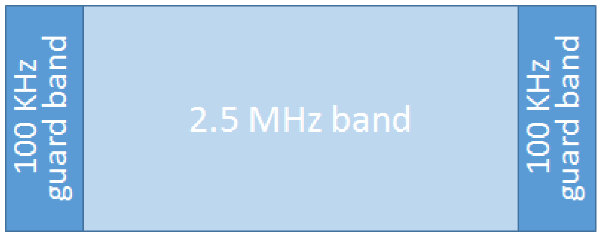

Solution:

For any X simultaneously supported users/channels, there will be (X-1) required guard channels. Hence, the allocated bandwidth has to satisfy the following:

2 500 kHz = X*60kHz + (X-1)*30kHz

2 500 kHz = X*(60+30)*kHz - 30kHz

X = 2530 / 90 = 28.11 => 28 can be supported simultaneously

8. Connecting LANs [13 points]

8.1 [10 points]

Consider the following set of 4 Ethernet segments connected by a central learning switch (not a hub). Each Ethernet segment has two hosts on it. Assume at the beginning that the switch has no state about where each of the hosts reside, and will use self-learning to build its routing table during the sequence of transmissions listed below.

For each step in this sequence of transmissions, list which ports the Ethernet frame is forwarded to by the switch (and, if the frame is NOT forwarded to any ports, just write 'drop').

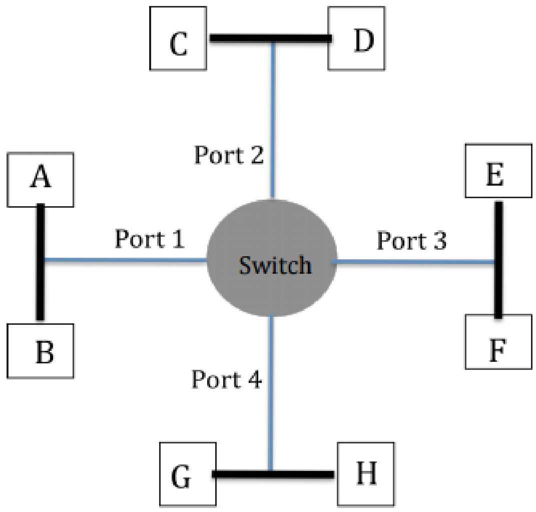

a) A sends to E ______________ **2 3 4**

b) A sends to B ______________ **2 3 4**

c) G sends to E ______________ **1 2 3**

d) E sends to A ______________ **1**

e) E sends to B ______________ **1 2 4**

f) A sends to E ______________ **3**

g) H sends to G ______________ **dropped**

h) H sends to B ______________ **1 2 3**

i) B sends to A ______________ **dropped**

j) F sends to E ______________ **dropped**

8.2 [3 points]

Consider the arrangement of the learning bridges shown in the figure below. Assuming all bridge forwarding tables are initially empty, give the state of the forwarding tables of bridge B1 after the following three transmissions:

D sends to C
C sends to D
A sends to C

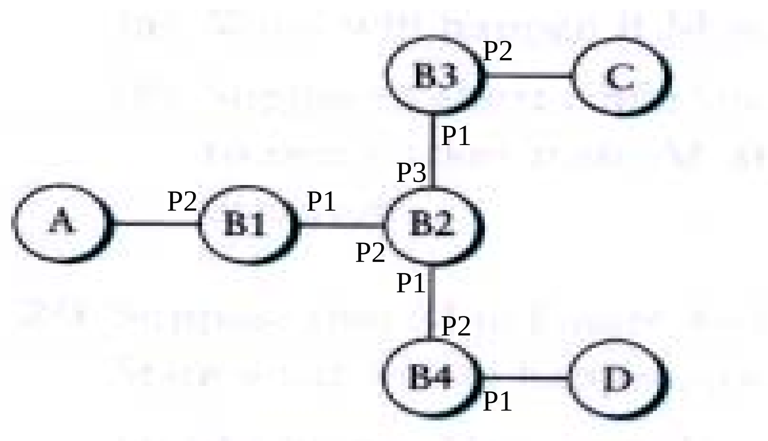

Solution:

Table at B1 after the three transmissions

Station	Port
D	P1
A	P2

Note, the packet sent from C to D will not be forwarded on P2 of B2 (as B2 has already learned the actual location of D, and will only be forwarding this packet on P1). Consequently, this packet will not reach B1, and B1 will not learn about station C.