Sample Problems for Exam 1

Prof. William A. Maniatty

October 20, 1999

1 Some Hints

Your professor suggests the following preparation strategies:

- The problems in this set should be doable by you in about 100 minutes total, time yourself.
- The number of minutes you are expected to use per problem is posted next to it (to indicate the difficulty level).
- Practice writing neat clean answers, since if the grader cannot understand you on the real exam, it will go badly for you.
- Show your work, if you are guessing the grader will not give much credit (even if you get lucky and guess right).
- Define your notation (you can use tables like the lecture notes if you like).
- Set up the solution symbolically and simplify before plugging numbers in, it is easier to follow for the grader.
- These exercises are meant to be done alone with an open book and notes and a calculator.
- You can solve problems out of order, but keep the work for each problem in one place, and mark it clearly.

2 The Problems

1. CRC (10 Minutes): Given a generator polynomial $P(x) = x^2 + x + 1$ and a message frame 101110, what is the corresponding CRC?

   The bit string corresponding to the coefficients of $P(x)$ is $P =$ 111, and the message frame gets 2 bits appended (because the degree of $P(x)$ is 2), so the augmented message is $M = 10111000$. We compute the remainder of division of P into M using modulo 2 arithmetic as follows:

```
      111110
  ----------
P | 10111000
    111
    ----
     101
     111
     ----
      101
      111
      ----
       100
       111
       ----
        110
        111
        ---
         010 <-Do not divide because of high order 0 bit
         ---
          10 <- Remainder (CRC by definition!)
```

   and we see our remainder is 2 (10 in base 2).

2. Signal Processing (15 Minutes):

   (a) Suppose a signal has a strength of 100 milliwatts and has 10 milliwatts of noise, what is the signal to noise ratio (5 minutes)?

   From our lecture notes we get the formula fir signal to noise ratio in (2) and substitute the given signal strength in for S and the noise strength in for N and solve:

$$\begin{aligned} \text{SNR} &= 10 \log_{10} \frac{S}{N} dB & (1) \\ &= 10 \log_{10} \frac{100\text{mW}}{10\text{mW}} dB & (2) \\ &= 10(\log_{10} 10) dB = (10 \times 1) dB = 10 dB & (3) \end{aligned}$$

(b) Suppose a (complex) signal, $S(t)$ is defined as:

$$S(t) = c_1 \times \sin(t) + c_2 \times \sin(9t) + c_3 \times \sin(11t) \tag{4}$$

where t is time measured in seconds.

i. What is the fundamental frequency of $S(t)$ (5 minutes)?

The fundamental frequency is the frequency of the complex wave form. Recall that the period (by definition) is how long it takes a periodic function to begin repeating itself. Before solving, recall that the function $\sin t$ has a period of 2π. The function $\sin(kt)$, where $k \neq 0$ is a constant, can be shown[1] to have a period of $\frac{2\pi}{k}$. Therefore the periods of the components are:

- $c_1 \sin t$ has period 2π seconds.
- $c_2 \sin(9t)$ has period $\frac{2\pi}{9}$ seconds.
- $c_3 \sin(11t)$ has period $\frac{2\pi}{11}$ seconds.

And we can see that the function will repeat itself after 2π seconds. Since frequency is the reciprocal of the period (by definition), the fundamental frequency is $\frac{1}{2\pi}$.

ii. What is the minimum frequency of sampling needed to accurately reconstruct $S(t)$ (5 minutes)?

The minimum frequency of sampling required for accuracy is defined by the sampling frequency to be twice as fast as the frequency of the highest frequency component. Since frequency is the reciprocal of the period (by definition), the highest frequency component has the smallest period. From the solution to 2(b)i we find that the smallest period of any component is $\frac{2\pi}{11}$ cycles per second, and that component has the highest frequency, $\frac{11}{2\pi}$. The minimum sampling frequency is therefore $\frac{11}{\pi}$ samples per second.

3. Information Theory (15 Minutes): Suppose that an information link transmits a 4 bit messages (3 bits of user data and a 1 bit header). Suppose that the header always contains the value of 1, and each bit of the user data has a probability of 0.5 of being set to 1. What is the entropy of this link?

Let the bit string of the message be $B = b_0, b_1, b_2, b_3$. From our given, assume that $b_0 = 1$ always and define the notation $p_i = \text{Prob}[b_i = 0]$ for $0 \leq i \leq 3$. The entropy of the ith bit in the message is then:

$$H(b_i) = -p_i \log_2 p_i + (1 - p_i) \log_2 (1 - p_i) \tag{5}$$

[1] If you are not convinced, try plotting the following. Consider the case where k is 2. The value of the parameter to the sine function will vary 2 times quicker, so the period would be π. Observe that it works for the case where $k = 1$ too.

From the definition of entropy, the entropy of the message is:

$$
\begin{aligned}
H &= -\sum_{i=0}^{3}[p_i \log_2 p_i + (1-p_i)\log_2(1-p_i)] && (6)\\
&= -[(p_0 \log_2 p_0 + (1-p_0)\log_2(1-p_0)) + (p_1 \log_2 p_1 + (1-p_1)\log_2(1-p_1)) + \\
&\quad (p_2 \log_2 p_2 + (1-p_2)\log_2(1-p_2)) + (p_3 \log_2 p_3 + (1-p_3)\log_2(1-p_3))] && (7)\\
&= -[(0\log_2 0 + 1\log_2 1) + (\frac{1}{2}\log_2\frac{1}{2} + \frac{1}{2}\log_2\frac{1}{2}) + (\frac{1}{2}\log_2\frac{1}{2} + \frac{1}{2}\log_2\frac{1}{2}) + \\
&\quad (\frac{1}{2}\log_2\frac{1}{2} + \frac{1}{2}\log_2\frac{1}{2}) + (\frac{1}{2}\log_2\frac{1}{2} + \frac{1}{2}\log_2\frac{1}{2})] && (8)\\
&= -[(0+0) + (-\frac{1}{2} - \frac{1}{2}) + (-\frac{1}{2} - \frac{1}{2}) + (-\frac{1}{2} - \frac{1}{2})] && (9)\\
&= -[0-1-1-1] = 3 && (10)
\end{aligned}
$$

In case you are curious, by definition (and using substitution) the per bit entropy is then $\frac{3}{4}$.

4. Channel Performance Analysis (15 minutes): Suppose that a satellite has 100 gigabits per second of band width, and is 40000 kilometers from two ground stations. Assume that the signals transmitted to and from the satellite travel at approximately the speed of light (about 300000 kilometers per second). How long will it take to send 1 kilobyte of data from one station to the other?

   Drawing a picture is a powerful design tool, you probably should have drawn one looking like Figure 1. Assuming that the satel-

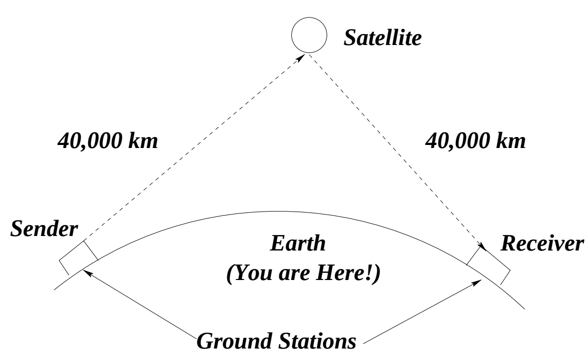

Figure 1: The system in Problem 4

   lite, sender and receiver have negligible overhead, the transmission time from end to end is then:

$$\begin{aligned}
\text{Transmission Time} &= \text{Time of Flight} + \frac{\text{Message Size}}{\text{Bandwidth}} && (11)\\
\text{Time of Flight} &= \frac{\text{Tranmission Distance}}{\text{Signal Velocity}} && (12)\\
&= \frac{2 \times 40000\text{km}}{300000\text{km/sec}} \approx 0.267\text{sec} && (13)\\
\text{Transmission Time} &= 0.267\text{sec} + \frac{1024\text{bytes} \times \frac{8\text{bits}}{\text{byte}}}{\frac{100\times10^9\,\text{bits}}{\text{sec}}} && (14)\\
&= 0.267\text{sec} + 0.00000008192\text{sec} && (15)\\
&= 0.267\text{sec} && (16)
\end{aligned}$$

5. Overhead analysis (10 Minutes):

   (a) Suppose that a packet can hold 1484 bytes of user data and has a 16 byte header. What is the header overhead (5 minutes)?

   The definition of header overhead and substitution of our givens yields:

$$\text{Header Overhead} = \frac{\text{Header Size}}{\text{Header Size} + \text{User Data}} = \frac{16}{16 + 1484} = \frac{16}{1500} \approx 0.0107 \quad (17)$$

   (b) If the ATM forum had adopted 64 byte cells, what would the the header overhead of an ATM packet be (5 minutes)?

   Recall that ATM cells use a 5 byte header to identify their virtual circuit. My original intent when I wrote this problem was to have the amount of user data be 64 bytes, however if you assumed the header data was part of the cell (which it is not according to the ATM standard) it would look like:

5 Byte Header	59 Bytes of User Data

   If you tried it that way, the problem result looks like:

$$\text{Header Overhead} = \frac{\text{Header Size}}{\text{Header Size} + \text{User Data}} = \frac{5}{5 + 59} = \frac{5}{64} \approx 0.078 \quad (18)$$

   However, the wording differs. A strict interpretation of the problem as posed has the structure:

5 Byte Header	64 Bytes of User Data

   Which yields the correct solution:

$$\text{Header Overhead} = \frac{\text{Header Size}}{\text{Header Size} + \text{User Data}} = \frac{5}{5 + 64} = \frac{5}{69} \approx 0.072 \quad (19)$$

6. Short answer (10 minutes):

(a) Why does the US have echo cancellation widely installed, but not Japan or Europe (5 minutes)?

Because echos result by multiplexing 2 lines to carry both the incoming and outbound signals from a telephone, echos occur (just like ringing in electrical circuits). If the lines are too long, the echo is sufficiently delayed that it is distracting to the listener (bad ergonomics). In the U.S. the phone centers are sufficiently far apart that this is a problem, and echo cancellation equipment is installed. Europe and Japan tend to have shorter distances between the offices so they do not require widely installed echo cancellers.

(b) If using layering in protocol design removes complexity at each layer, why not use say 100 layers, or say 2 layers (5 minutes)?

Because each layer requires resources to support it (space and or time), having too many layers limits efficiency. Not having enough layers makes implementation difficult and error prone.

7. Queueing Theory (10 minutes):

(a) Suppose that a router receives (on average) 800 packets per second and takes 1 millisecond to forward a packet, what is the utilization assuming interarrival and service times are exponentially distributed (5 minutes)?

We are given the interarrival rate of $\lambda = \frac{800\text{packets}}{\text{sec}}$ and mean service time of $s = 1\text{msec}$. Using the definition of utilization and substituting our givens we get:

$$\rho = \lambda s = \frac{800\text{packets}}{\text{sec}} \times 1\text{msec} \times \frac{1\text{msec}}{1000\text{sec}} = 0.8 \qquad (20)$$

(b) What is the mean time to route a packet (5 minutes)?

This is an M/M/1 system, so we apply the equations and substitute in our known values (from the problem statement and the results of 7), getting:

$$t_q = \frac{s}{1-\rho} = \frac{1\text{msec}}{1-0.8} = 5\text{msec} \qquad (21)$$

8. Network Systems programming (15 minutes):

(a) Why does Unix sometimes return from a read or write systems call if only part of the data could be transferred? (5 minutes)

Because blocking on a network read or write can cause application problems, particularly if we are using the `select` system call. Furthermore buffer space may be unavailable for sends or data may be unavailable for reads.

(b) What functionality does the gethostbyname system call provide (5 minutes)?

> The gethostbyname call takes the host name as a parameter and returns the ip address. It invokes the DNS protocol.

(c) Describe the functionality of the listen system call (5 minutes).

> In a client server application, only servers execute a listen system call. The listen call blocks on a socket pending an attempted connection by a client and then gives the host the option of accepting the connection (by invoking the accept system call). There is a second parameter specifying the number of pending connection attempts the server will permit.