Computer Networks
Problem Set 2A

Instructor: Shashi Prabh

Peterson and Davie. Chapter 2 - Getting Connected

2.1 Show the NRZ, Manchester, and NRZI encodings for the bit pattern shown in Figure 1. Assume that the NRZI signal starts out low.

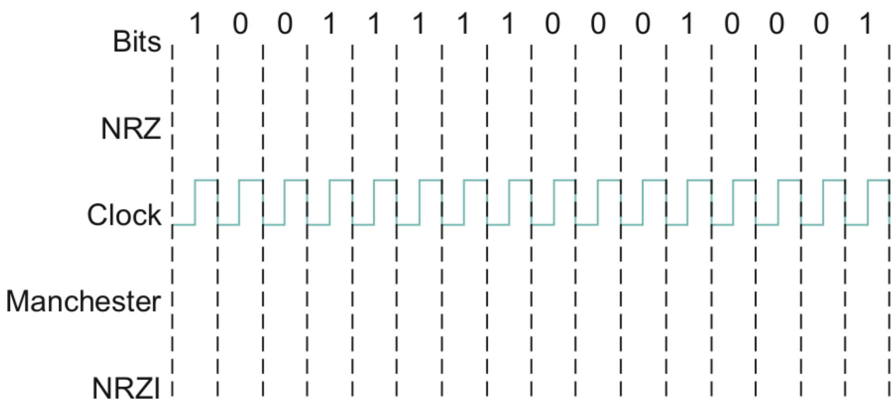

Figure 1: Exercise 2.1

2.2 Show the 4B/5B encoding,and the resulting NRZI signal, for the following bit sequence: `1101 1110 1010 1101 1011 1110 1110 1111`

2.3 Assuming a framing protocol that uses bit stuffing, show the bit sequence transmitted over the link when the frame contains the following bit sequence: `110101111101011111101011111110`. Mark the stuffed bits.

2.4 Suppose the following sequence of bits arrives over a link:
`0110 1011 1110 1010 0111 1111 0110 0111 1110`. Show the resulting frame after any stuffed bits have been removed. Indicate any errors that might have been introduced into the frame.

2.5 Show that two-dimensional parity allows detection of all 3-bit errors.

2.6 Show that the Internet checksum will never be 0xFFFF (that is, the final value of sum will not be 0x0000) unless every byte in the buffer is 0. (Internet specifications in fact require that a checksum of 0x0000 be transmitted as 0xFFFF; the value 0x0000 is then reserved for an omitted checksum. Note that, in ones complement arithmetic, 0x0000 and 0xFFFF are both representations of the number 0.)

2.7 Suppose we want to transmit the message : `1011001001001011` and protect it from errors using the CRC8 polynomial

$$x^8 + x^2 + x^1 + 1$$

.

- **a.** Use polynomial long division to determine the message that should be transmitted.
- **b.** Suppose the left most bit of the message is inverted due to noise on the transmission link. What is the result of the receiver's CRC calculation? How does the receiver know that an error has occurred?

2.8 With 1 parity bit we can detect all 1-bit errors. Show that at least one generalization fails, as follows:

- **a.** Show that if messages m are 8 bits long, then there is no error detection code $e = e(m)$ of size 2 bits that can detect all 2-bit errors. Hint: Consider the set M of all 8-bit messages with a single 1 bit; note that any message from M can be transmuted into any other with a 2-bit error, and show that some pair of messages m_1 and m_2 in M must have the same error code e.
- **b.** Find an N (not necessarily minimal) such that no 32-bit error detection code applied to N-bit blocks can detect all errors altering up to 8 bits.

2.9 Consider an ARQ protocol that uses only negative acknowledgments (NAKs), but no positive acknowledgments (ACKs). Describe what timeouts would have to be scheduled. Explain why an ACK-based protocol is usually preferred to a NAK-based protocol.

2.10 Suppose you are designing a sliding window protocol for a 1-Mbps point-to-point link to the moon, which has a one-way latency of 1.25 seconds. Assuming that each frame carries 1 KB of data, what is the minimum number of bits you need for the sequence number?

2.11 Suppose you are designing a sliding window protocol for a 1-Mbps point-to-point link to the stationary satellite revolving around the Earth at an altitude of $3 * 10^4$ km. Assuming that each frame carries 1 KB of data, what is the minimum number of bits you need for the sequence number in the following cases? Assume the speed of light is $3 * 10^8$ m/s.

- **a.** RWS = 1
- **b.** RWS = SWS

2.12 In stop-and-wait transmission, suppose that both sender and receiver retransmit their last frame immediately on receipt of a duplicate ACK or data frame; such a strategy is superficially reasonable because receipt of such a duplicate is most likely to mean the other side has experienced a timeout.

- **a.** Draw a timeline showing what will happen if the first data frame is somehow duplicated, but no frame is lost. How long will the duplications continue? This situation is known as the Sorcerer's Apprentice bug.
- **b.** Suppose that, like data, ACKs are retransmitted if there is no response within the timeout period. Suppose also that both sides use the same timeout interval. Identify a reasonably likely scenario for triggering the Sorcerer's Apprentice bug.

2.13 Describe a protocol combining the sliding window algorithm with selective ACKs. Your protocol should retransmit promptly, but not if a frame simply arrives one or two positions out of order. Your protocol should also make explicit what happens if several consecutive frames are lost.

2.14 Suppose that we attempt to run the sliding window algorithm with SWS = RWS = 3 and with MaxSeqNum = 5. The N th packet DATA[N] thus actually contains N mod 5 in its sequence number field. Give an example in which the algorithm becomes confused; that is, a scenario in which the receiver expects DATA[5] and accepts DATA[0]—which has the same transmitted sequence number—in its stead. No packets may arrive out of order. Note that this implies MaxSeqNum $\geq$ 6 is necessary as well as sufficient.

2.15 What kind of problems can arise when two hosts on the same Ethernet share the same hardware address? Describe what happens and why that behavior is a problem.

2.16 Draw a timeline diagram for the sliding window algorithm with SWS = RWS = 4 frames in the following two situations. Assume the receiver sends a duplicate acknowledgment if it does not receive the expected frame. For example, it sends DUPACK[2] when it expects to see Frame[2] but receives Frame[3] instead. Also, the receiver sends a cumulative acknowledgment after it receives all the outstanding frames. For example, it sends ACK[5] when it receives the lost frame Frame[2] after it already received Frame[3], Frame[4], and Frame[5]. Use a timeout interval of about 2×RTT.

a. Frame 2 is lost. Retransmission takes place upon timeout (as usual).

b. Frame 2 is lost. Retransmission takes place either upon receipt of the first DUPACK or upon timeout. Does this scheme reduce the transaction time? (Note that some end-to-end protocols, such as variants of TCP, use similar schemes for fast retransmission.)

2.17 Coaxial cable Ethernet was limited to a maximum of 500 m between repeaters, which regenerate the signal to 100% of its original amplitude. Along one 500-m segment, the signal could decay to no less than 14% of its original value (8.5 dB). Along 1500 m,then, the decay might be $0.14^3 = 0.3\%$. Such a signal, even along 2500 m,is still strong enough to be read; why then are repeaters required every 500 m?

2.18 Let A and B be two stations attempting to transmit on an Ethernet. Each has a steady queue of frames ready to send; A's frames will be numbered A1, A2,and soon,and B's similarly. Let T =51.2 μ s be the exponential back-off base unit.

Suppose A and B simultaneously attempt to send frame 1, collide, and happen to choose back-off times of 0×T and 1×T, respectively, meaning A wins the race and transmits A1 while B waits. At the end of this transmission, B will attempt to retransmit B1 while A will attempt to transmit A2. These first attempts will collide, but now A backs off for either 0×T or 1×T,while B backs off for time equal to one of 0×T,...,3×T.

a. Give the probability that A wins this second back-off race immediately after this first collision; that is, A's first choice of back-off time k x 51.2 is less than B's.

b. Suppose A wins this second back-off race. A transmits A3, and when it is finished, A and B collide again as A tries to transmit A4 and B tries once more to transmit B1.Give the probability that A wins this third back-off race immediately after the first collision.

c. Give a reasonable lower bound for the probability that A wins all the remaining back-off races.

d. What then happens to the frame B1? This scenario is known as the Ethernet capture effect.

2.19 Suppose Ethernet physical addresses are chosen at random (using true random bits).

a. What is the probability that on a 1024-host network, two addresses will be the same?

b. What is the probability that the above event will occur on one or more of 220 networks?

c. What is the probability that, of the 230 hosts in all the networks of (b), some pair has the same address?

Hint: The calculation for (a) and (c) is a variant of that used in solving the so-called Birthday Problem: Given N people, what is the probability that two of their birthdays (addresses) will be the same? The second person has probability $1 - \frac{1}{365}$ of having a different birthday from the first, the third has probability $1 - \frac{2}{365}$ of having a different birthday from the first two, and so on. The probability that all birthdays are different is thus

$$\prod_{n=1}^{N-1} (1 - \frac{n}{365})$$

which for smallish N is about:

$$1 - \frac{1 + 2 + ... + (N-1)}{365}$$

2.20 How can a wireless node interfere with the communications of another node when the two nodes are separated by a distance greater than the transmission range of either node?

2.21 Why is collision detection more complex in wireless networks than in wired networks such as Ethernet?

2.22 How can hidden terminals be detected in 802.11 networks?

2.23 Why might a wireless mesh topology be superior to a base station topology for communications in a natural disaster?

2.24 Why isn't it practical for each node in a sensor net to learn its location by using GPS? Describe a practical alternative.