Solutions to Final Exam

Question 1

A communication system transmits three different messages with the following probabilities:

message	probability
A	50%
B	x %
C	6.25%

(a) What is the probability of message B?

(b) How much information, in bits, is communicated by each message?

(c) What is entropy of this source in bits per message?

Solution

(a) The probability of message B is 1 minus the sum of the other message probabilities: $1 - 0.5 - 0.0625 = \boxed{43.75\%}$.

(b) The amount of information, in bits, communicated by message i is given by $-\log_2(P_i)$. These are $-\log_2(0.5) = \boxed{1 \text{ bits/message}}$, $-\log_2(0.4375) = \boxed{1.19 \text{ bits/message}}$, and $-\log_2(0.0625) = \boxed{4 \text{ bits/message}}$.

(c) The entropy is the average information per message or $0.5 \times 1 + 0.4375 \times 1.19 + 0.0625 \times 4 = 0.5 + 0.52 + 0.25 = \boxed{1.27 \text{ bits/message}}$.

Question 2

The Arabic letter Hah (ح) has the Unicode code point U+062D.

(a) How many bytes would be required to transmit this character using the UTF-8 encoding?

(b) What are the values of those bytes? Give your answer in hexadecimal.

Solution

(a) The bit pattern for 0x062D is `0110 0010 1101` and has 11 non-zero bits. From the encoding table this can be represented as $\boxed{\text{two (2) bytes}}$.

(b) The values of the first five non-zero bits are y =`11000` and the second six bits are x =`101101`. When the prefixes `110` and `10` are added the bytes are `110 1 1000` and `10 10 1101` which in hexadecimal are $\boxed{\text{D8 AD}}$.

Question 3

A communication system transmits data using MLT-3 with levels of +1, 0 and −1 volts. The decision threshold are at ±0.5 V. The channel adds Gaussian noise with an RMS voltage of 0.16 V to the signal. If a signal level of 0 V is transmitted, what is the probability that it will not be received correctly (i.e. that the receiver will decide that either a +1 or a −1 was transmitted instead)?

Solution

An errors will be made if the noise is greater than 0.5 of less than −0.5. The normalized thresholds are $t = 0.5/0.16 \approx 3$ and $t = -0.5/0.16 \approx -3$. The probability of normalized Gaussian noise being less than −3 is (from the graph or using a calculator) 0.00135 which by symmetry is also the probability of being greater than 3. The probability of either event is the sum of the two areas or $2 \times 0.00135 \approx \boxed{0.27\%}$.

Question 4

A differential transmitter uses two conductors (D+ and D−) to transmit data.

- When a '1' is being transmitted D+ is at 5V relative to ground and D− is at 0V (also relative to ground).

- When a '0' is being transmitted D− is at 5V relative to ground and D+ is at 0V (also relative to ground).

Assume the differential voltage is positive when D+ is more positive than D−.

(a) What are the common-mode and differential voltages when a '1' is being transmited?

(b) What are the common-mode and differential voltages when a '0' is being transmited?

(c) If the channel adds an offset voltage of 5 V to both conductors is the differential voltage affected? Is the common-mode voltage affected?

Solution

(a) The common-mode voltage when a '1' is being transmited is $\frac{5+0}{2} = 2.5$ V. The differential voltage when a '1' is being transmitted is $5 - 0 = 5$ V.

(b) The common-mode voltage when a '0' is being transmited is $\frac{0+5}{2} = 2.5$ V. The differential voltage when a '1' is being transmitted is $0 - 5 = -5$ V.

(c) If the channel adds an offset voltage of 5 V to both conductors the voltages are 10 and 5 for a '1' and 5 and 10 for a '0'. Repeating the above calculation shows the differential voltages are not affected because the 5 V offsets cancels out when the voltages are subtracted. The common-mode voltage becomes $\frac{10+5}{2} = 7.5$ V and so the common-mode voltages are affected.

Question 5

The following waveform transmits a single value (a "character") over an asynchronous serial interface:

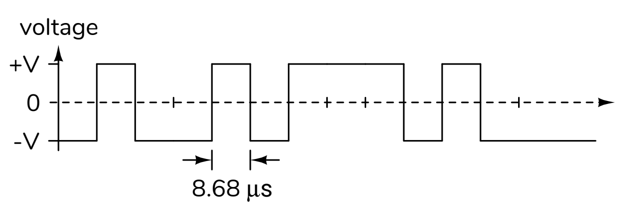

(a) What was the baud rate?

(b) Was parity used? If so, was it even or odd parity?

(c) What is the value of the character that was transmitted. Give your answer as a hexadecimal value.

Solution

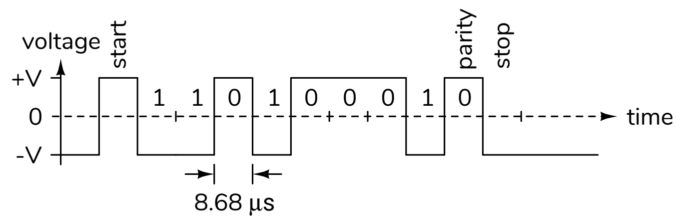

(a) The baud rate is $\frac{1}{8.68 \times 10^{-6}} =$ 115.2 kb/s.

(b) Since there are 8 bits plus a high (mark) after the start bit then parity must have been used since serial interfaces typically don't transmit more than 8 bits per character. The values of the other bits were as shown above (in order transmitted: `1101 0001` followed by a `0` parity bit. The number of ones is four, an even number so even parity was used.

(c) What is the value of the character that was transmitted. Give your answer as a hexadecimal value. In order from most- to least-signifcant bit the value as `1000 1011` which in hexadecimal is 8B.

Question 6

The following waveform shows data that has been encoded using a differential NRZ line code using the convention that a change is a '1' and no change is a '0':

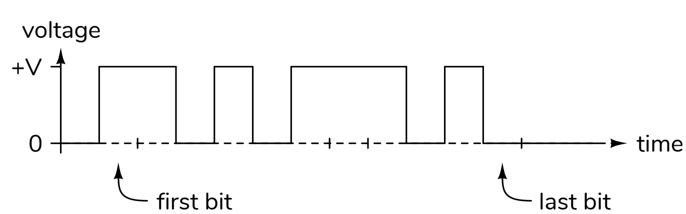

What sequence of bits was transmitted between the indicated bits? Give your answer as a sequence of 1's and 0's in the order the bits were transmitted.

Solution

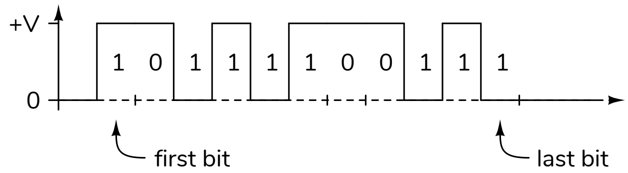

Following the rules above, the bits can be decoded as shown above. The sequence of bits was `1011 1100 111`.

Question 7

You measure the following transition between two levels on a communication system waveform:

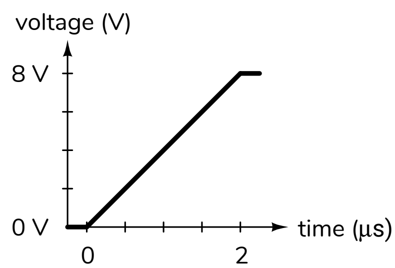

(a) What is the slew rate?

(b) If the decision threshold is half-way between the two levels, what is the noise margin?

Solution

(a) The slew rate is the slope of the voltage versus time curve. In this case it's $\frac{8}{2} = \boxed{4\text{ V}/\mu s}$.

(b) A decision threshold half-way between the two levels would be at $\frac{8}{2} = 4$ V. The noise margin is the magnitude of the voltage required to cause an error which is the same for both levels: $|8 - 4| = 4$ V or $|4 - 8| = \boxed{4\text{ V}}$.

Question 8

The following sequence of bits includes an HDLC frame:

```
1 0 1 1
  0 1 1 1 1 1 1 0
1 0 1 1 1 1 1 0 1 1
  0 1 1 1 1 1 1 0
  0 1 0
```

What are the values of the data bits within the HDLC frame? Give your answer a sequence of 0's and 1's in the order they appear in the frame.

Solution

The bits in the question show the bit sequence divided up by finding the 8-bit flag sequences. After removing the stuffed zero bit after five consecutive ones the sequence of bits in the frame is: `1 0 1 1 1 1 1 1 1`.

Question 9

A communication system uses the generator polynomial $G(x) = x^4 + x + 1$ to protect a message consisting of the bits `10110`.

(a) What is the length of the CRC in bits?

(b) Use the simple CRC algorithm described in the lectures to compute the CRC that would be appended to this message. Give your answer as a binary number.

Solution

(a) The length of the CRC in on less than the number of bits in the generator polynomial or the order of the highest term: $\boxed{4}$.

(b) The CRC that would be appended to this message is calculated as shown below. In binary it's `1111`.

```
101100000
10011
-----
 01010
 -----
  10100
  10011
  -----
   01110
   -----
    11100
```

```
10011
-----
 1111 = CRC
```

check:

```
101101111
10011
-----
 01011
 -----
  10111
  10011
  -----
   01001
   -----
    10011
    10011
    -----
        0 => OK
```

Question 10

A code has four valid codewords:

- `0000000`
- `1110000`
- `0001111`
- `1000011`

(a) What is the minimum distance of this code?

(b) How many errors can it correct?

(c) How many errors can it detect?

(d) If the codeword **1111111** was received, was there an error?

(e) What codeword was most likely transmitted?

Solution

(a) The distances between the codewords are:

	0000000	1110000	0001111	1000011
0000000		3	4	3
1110000			7	3
0001111				3
1000011				

and so the minimum distance is $\boxed{d_{min} = 3}$.

(b) It can correct $\lfloor \frac{d_{min}-1}{2} \rfloor = \boxed{\text{1 error}}$?

(c) It can it detect $d_{min} - 1 = \boxed{\text{2 errors}}$.

(d) Since the codeword **1111111** is not a valid codeword $\boxed{\text{yes, there was an error}}$.

(e) The codeword at the smallest distance to the received codeword is the most likely to have been transmitted. This is $\boxed{\texttt{0001111}}$ at a distance of 3.

Question 11

A channel has the following frequency response:

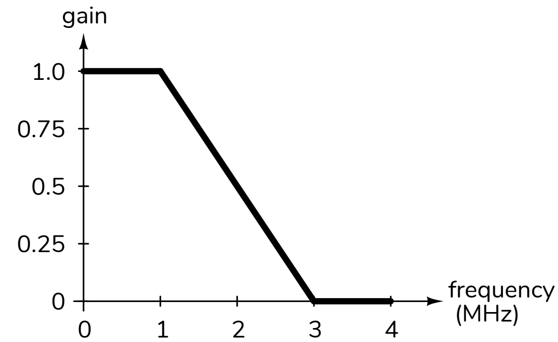

Is it possible to transmit data over this channel without ISI? If so, what is the maximum symbol rate for which this is possible?

Solution

The channel has the symmetry required for the Nyquist no-ISI criteria in frequency so $\boxed{\text{yes}}$ it is possible to transmit data over this channel without ISI. The channel is symmetrical about 2 MHz so the maximum symbol rate is twice this or $\boxed{\text{4 MHz}}$.

Question 12

What is the minimum bandwidth required for any communication system to be able to achieve a very low error rate while transmitting 1 Mb/s of information over an AWGN channel that has an SNR of 0 dB?

Solution

The capacity of an AWGN channel is $C = B\log_2(1 + S/N)$. In this case $S/N = 10^{0/10} = 1$ and $C = 1$ Mb/s so $B = \frac{C}{\log_2(1+1)} = \boxed{\text{B=1 MHz}}$.

Question 13

The following bytes (in hexadecimal) follow the preamble in an Ethernet frame:

```
00 e0 4c 6b e4 9b 50 3e aa 0c 77 7e 08 00 50 3e aa 0c 77 7e
```

(a) What is the destination address?

(b) What is the value of the type field?

(c) What is the value of the first byte in the (Ethernet frame's) payload?

Solution

(a) The destination address is the first six bytes of the header or `00 e0 4c 6b e4 9b`.

(b) The type field follows the source address and in this case it's `08 00`.

(c) The payload follows the type field and value of the first byte is `50`.

Question 14

1. Would TDMA, FDMA or CSMA be most appropriate for a communication system primarily used by customers viewing web pages? *Briefly* explain.

2. Would (i) stop-and-wait ARQ, (ii) go-back-N ARQ or (iii) selective repeat ARQ be most appropriate for use on a link between North America and Asia operating at 10 Gb/s with very low error rates (e.g. one frame error per month)? *Briefly* explain.

Solution

1. TDMA and FDMA reserve time (slots) or frequency (channels) even when no data is being sent. Customers viewing web pages generate bursty traffic for which CSMA would make the most efficient use of the channel.

2. stop-and-wait ARQ would be inefficient due to the long delay relative to the frame duration. Go-back-N ARQ or selective repeat ARQ have good performance but the additional complexity of selective repeat ARQ would not be justified given the low error rate. Thus go-back-N ARQ would be most appropriate.