Solutions to exercise 5 in
EITF25 Internet - Techniques and Applications - 2013

October 26, 2013

1

a

If a bit occupies 1 millisecond = 10^{-3} seconds, then a second is occupied by $\frac{1}{10^{-3}} = 10^3$ bits = 1 kbps.

Answer: 1 kbps

b

If a bit occupies 2 microseconds = 2×10^{-6} seconds, then a second is occupied by $\frac{1}{2\times 10^{-6}} = 0.5 * 10^6$ bits = 500 kbps.

Answer: 500 kbps

2

a

If bits are produced/received at a rate of $100 * 10^3 = 10^5$ per second, then one bit occupies $\frac{1}{10^5} = 10^{-5}$ seconds = 10 microseconds.

Answer: 10 microseconds

b

If bits are produced/received at a rate of 2×10^6 per second, then one bit occupies $\frac{1}{2\times 10^6} = 0.5 \times 10^{-6}$ seconds = 0.5 microseconds.

Answer: 0.5 microseconds

3

a

In a NRZ (Non-Return-to-Zero) modulated transmission ones and zeros are a represented by a specific output level, constrained to a specified duration, see figure 1.

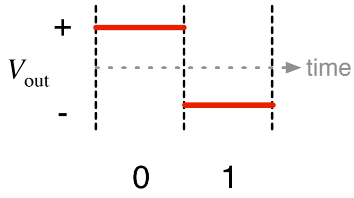

Figure 1: NRZ bit vs. voltage level assignment

As a consequence, unless you know when the transmission started and the duration a bit occupies, you will be unable to distinguish between multiple consecutive bits of the same sign. In a communication system, both receiver and transmitter thus need to be synchronized.

Answer:

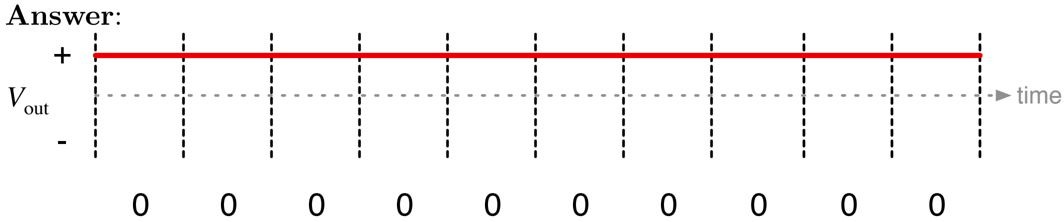

b

In a Manchester modulated signal, each bit i demarked by an output transition during the allocated bit duration (symbol) as opposed to a constant level seen in NRZ modulation. In Figure 2 a zero is represented by a falling edge, while a one is represented by a rising edge.

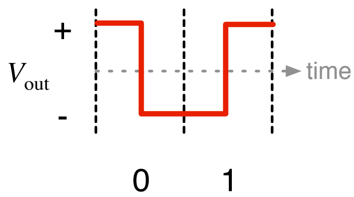

Figure 2: Machester bit representation as according to IEEE 802.3

As each bit infers a output level transition, with Manchester coding, each bit can be identified without prior clock synchronization. Nevertheless, both receiver and transmitter needs to agree on which transition to represent which bit. Moreover, In this case, we are dealing with a sequence of all zeros, the signal is thus modulated with all falling edges for each bit. Manchester coding has a relative low throughput compared to more complex modulation schemes and is today mainly used in systems such as 10BASE-T ethernet (IEEE 802.3) and NFC.

Answer:

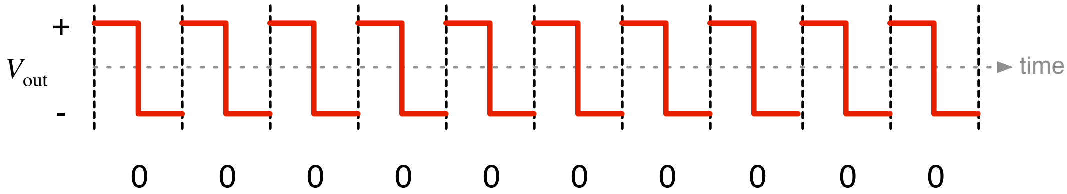

c

As you will see in problem 5, Differential Manchester coding deals with transitions between symbols rather than output levels within a symbol. A sequence with all zeros is represented by the absence of symbol transitions, and the waveform is therefore identical to the Machester coded one in problem 3b.

Answer:

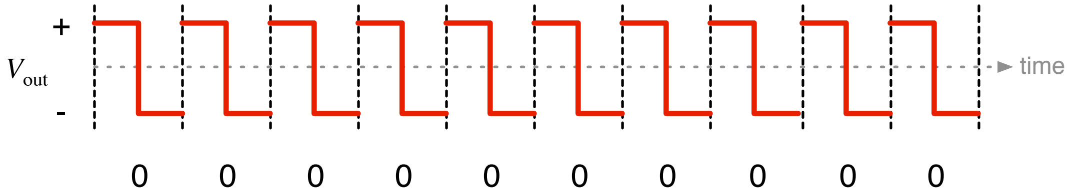

4

When applying the Manchester modulation scheme presented in Figure 2 to the observed sequence, we arrive at the bit sequence presented in Figures 3, 4, and 5.

a

Following the above convention yields.

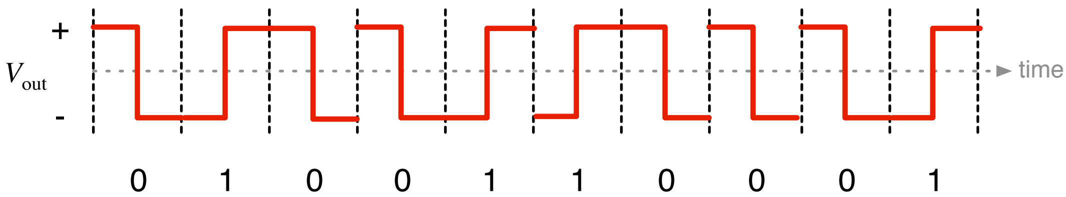

Figure 3: Decoded 4a Manchester waveform

Answer: 0100110001

b

Following the above convention yields.

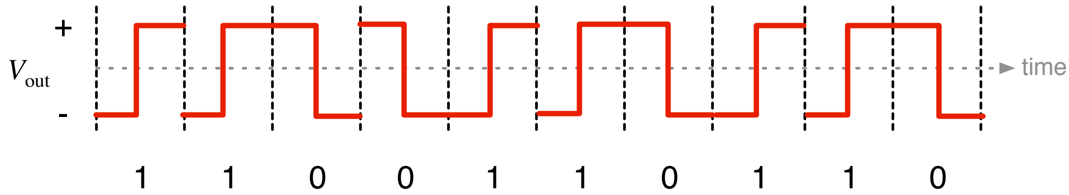

Figure 4: Decoded 4b Manchester waveform

Answer: 1100110110

c

Following the above convention yields.

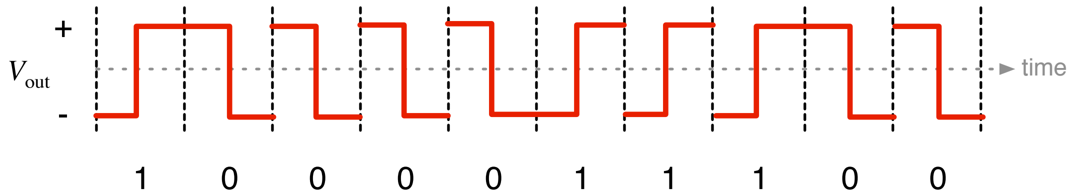

Figure 5: Decoded 4c Manchester waveform

Answer: 1000011100

5

As opposed to Manchester coding, Differential Manchester coding only represents transitions between $0 \rightarrow 1$. See Figure 6.

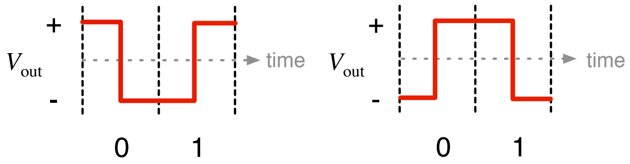

(a) Without guard interval (b) With guard interval, dark

Figure 6: Transitions

As such, several transitions represent just as many ones. A single transition followed by a consecutively repeated symbol represents a one followed by zeros, until the next transition. As such, the absence of a transition signifies a zero. See Figure 7

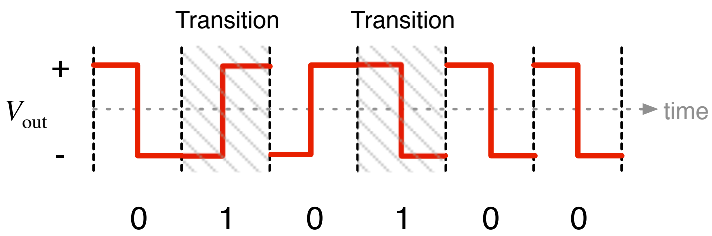

Figure 7: Sequence with transitions

When observing an arbitrary waveform, as we do not know whether it is a continuation of the previous symbol or a transition, the first bit is inherently unknown. See Figure 8.

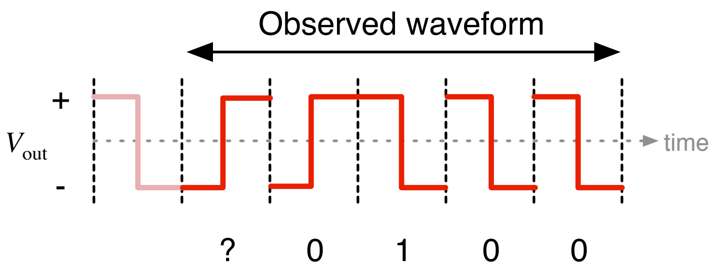

Figure 8: Decoding the first bit

As you can see in Figure 8, we do not know if first symbol is zero that transitions into a one, or the tail of a sequence of transitions representing several ones. Differential Manchester is used because a transition is less likely to be misinterpreted than a individual symbols by the receiver, given a noisy channel. Differential Manchester encoding is predominantly used in magnetic and optical storage.

a

Following the above convention yields.

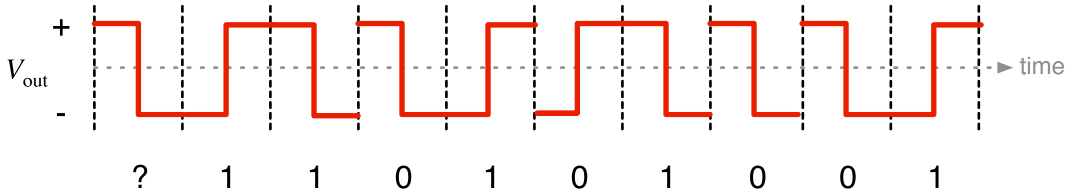

Figure 9: Decoded 5a Differential Manchester waveform.

Answer: 110101001

b

Following the above convention yields.

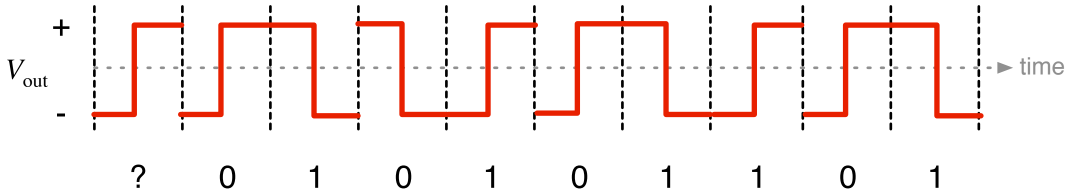

Figure 10: Decoded 5b Differential Manchester waveform.

Answer: 010101101

c

Following the above convention yields.

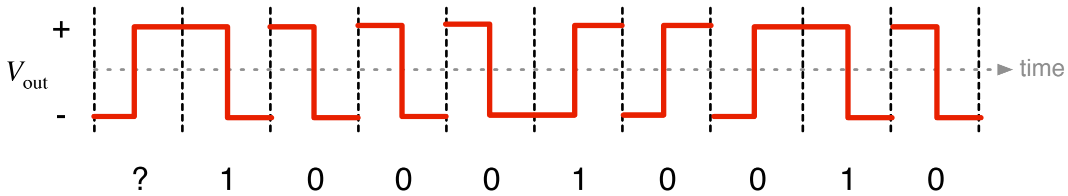

Figure 11: Decoded 5c Differential Manchester waveform.

Answer: 100010010

6

a

Following the convention established in problem 4.

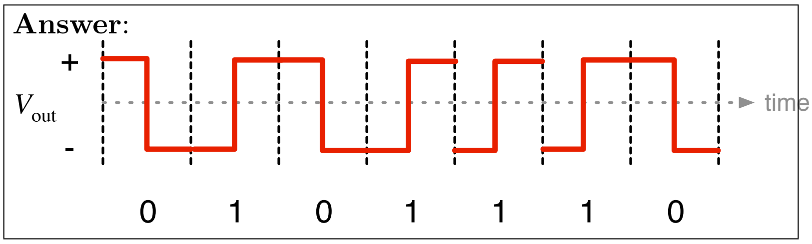

b

Following the convention established in problem 5. The first symbol is to illustrate the state of the signal.

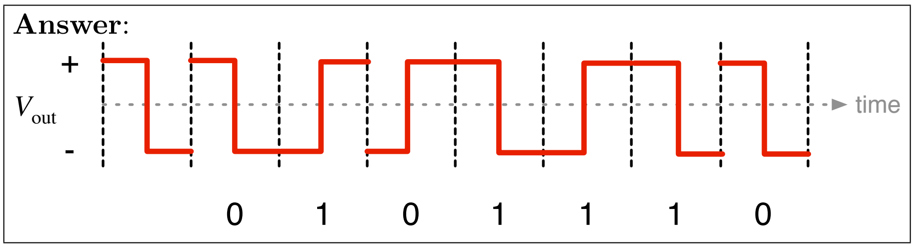

7

Being it a wireless, an optical, or an electrical link, in FDM (Frequency Division Multiple access) the entire or a portion of that mediums frequency spectrum bandwidth is divided into channels, see Figure 12. Moreover, this is one approach for accommodating multiple connections in a common link, where the transmitter (Tx) and receiver (Rx) are transmitting and receiving in the same allocated frequency range, channel. In most FDM based communication systems, a frequency guard interval is introduced to separate the channels, see Figure 12b. This is to ensure that any leakage from one channel is not interfering with its adjacent channels. The guard interval does not carry any intentional information. Note that guard intervals are only needed to separate the channels in the medium, and not the channels from beyond the frequency range of the medium. Unless there is another communication link at a frequency range above or below, in which case, a guard interval is introduced between the links. However, this is not taken into account when referring the bandwidth of the observed link.

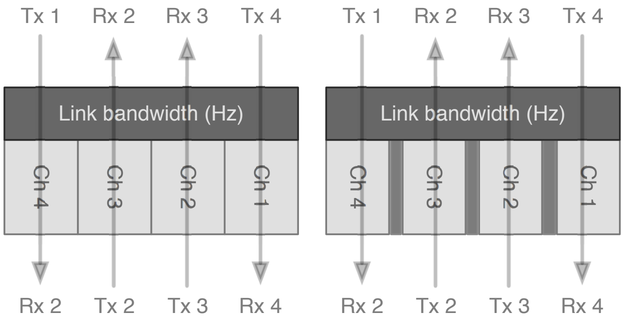

(a) Without guard interval (b) With guard interval, dark

Figure 12: Link multiplexed across 4 frequency channels

In this instance, it is specified that each channel needs a bandwidth of 4000 Hz = 4 KHz with a guard interval of 200 Hz = 0.2 kHz. SInce there are five channels, we need four guard intervals to separate them. Moreover, this leaves us with a total link bandwidth of 5×4 kHz + 4×0.2 kHz = 20.8 kHz. See Figure 13

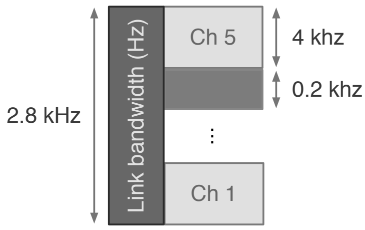

Figure 13: Link bandwidth breakdown, guard intervals are shaded (Not to proportion)

Answer: 20.8 kHz

8

In this example we are allocated a total link bandwidth of 7900 Hz = 7.9 kHz. It is specified that the link will carry 3 channels and that the channels shall be separated by unused frequency space, or guard intervals of 200 Hz = 0.2 kHz each. Since there are three channels, they will be separated by a total of 2 guard intervals. See Figure 14.

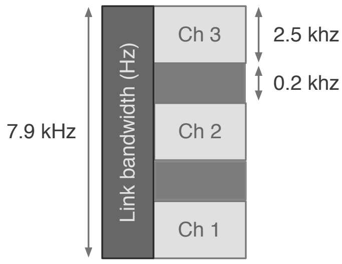

Figure 14: Link bandwidth breakdown, guard intervals are shaded (Not to proportion)

Subtracting the sum of the two guard intervals of 0.2 kHz each from the total allocated link bandwidth leaves us with the remaining bandwidth for the sum of the three channels, 7.9 kHz - 2 $\times$ 0.2 kHz = 7.5 kHz. Furthermore, as there are three channels, each channel can be allowed a bandwidth of $\frac{7.5}{3}$ kHz = 2.5 kHz.

Answer: 2.5 kHz

9

a

In a Synchronous TDM (Time Division Multiplexing) system, each sub-link is assigned a reoccurring time slot at a constant frequency on the multiplexed link. The capacity of the multiplexed link is shared proportionally amongst sub-links. For example, as illustrated in Figure 15, if the main link has a capacity of 3 Mbit and is shared indiscriminately amongst the three sub-links, they will each a access to a throughput of 1 Mbps.

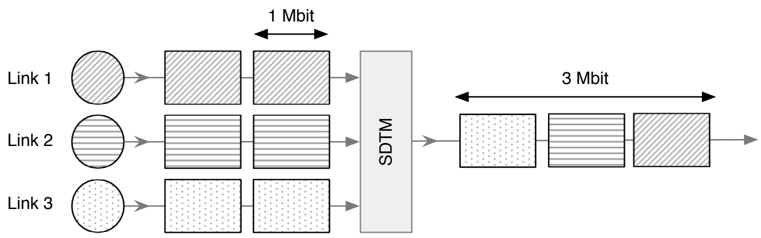

Figure 15: STDM bitrate

Furthermore, it is specified that each each of the 100 sub-links have a capacity of 14,4 kbps. In order to ensure that the each sub-links capacity is atleast maintained, the multiplexed link needs to have minimum capacity of 100 $\times$ 14.4 kbps = 1.44 Mbps.

Answer: 1.44 Mbps

b

In a Synchronous TDM system, unused time slots are left unused if a sub-link has nothing to transmit. See Figure 16.

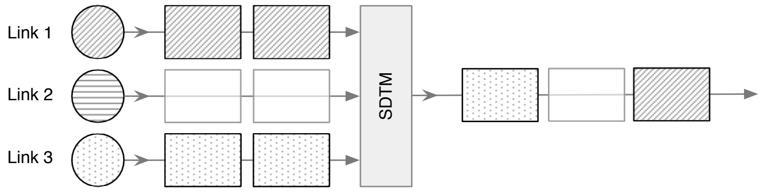

Figure 16: STDM utilization

As such, if only 70 out of 100 sub-links have information to transmit, then 30 % of the congruent time slots are left unused.

Answer: 30%

10

It is given that each character occupies a time frame. Since the links carry different quantities of data, there will be vacant time slots. See Figure 17.

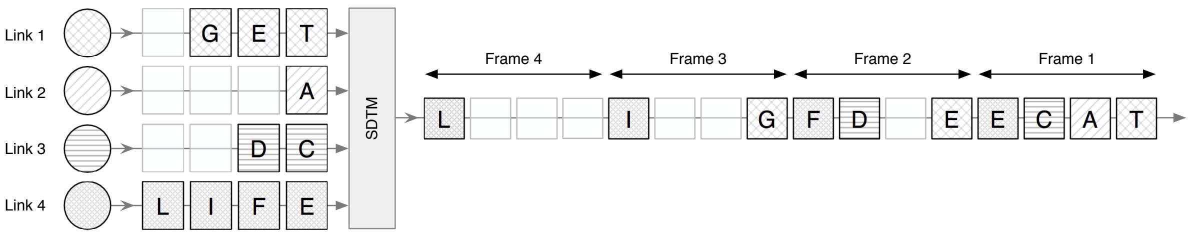

Figure 17: STDM example

Assuming that the sender on each link starts to send their letters simultaneously, the first frame will contain the first letter transmitted on each link. The subsequent frame will contain the second letter, if any. Moreover, in this instance 6 out of 16 frames will go unused, resulting in a link utilization of $\frac{10}{16} = 62.5\%$.

Answer:

11

In a Statistical TDM system, the multiplex fills each subsequent and squall sized time slot with the next available frame from the buffer. In other words, if a sub-link has nothing to send in this time frame the multiplexer will try to fill that time slot with data from the next sub-link that has data to send in the multiplexers buffer. Consequently, this ensures that all slots are utilized, given that there is data to send. See Figure 18.

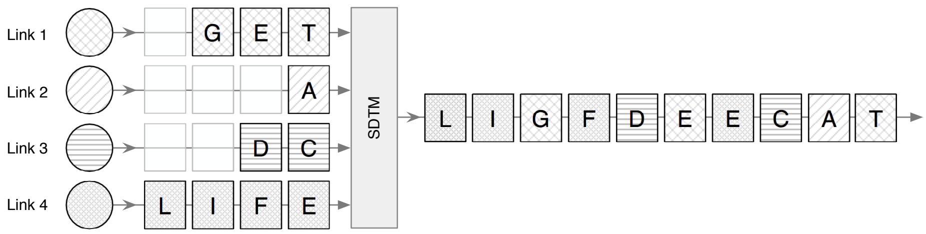

Figure 18: Statistical TDM.

As such, compared to Synchronous TDM, Statistical TDM will not transmit the intermediate vacant time frames.

Answer:
L I G F D E E C A T

12

The described Synchronous TDM system combines three sub-links of 300 kbps each. A frame starts with a synchronization bit, followed by three time slots of 3 bits each, resulting in a frame size of 10 bits. See Figure 19.

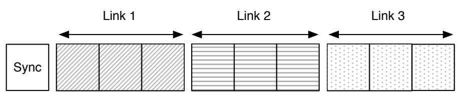

Figure 19: Frame structure

a

The three streams are congruent, as a result there are no vacant time slots in the resulting Synchronous TDM transmission. When applying the frame structure in Figure 19 to the specified bit streams, assuming that the first sync bit is 0, the resulting first frame will have the following content:

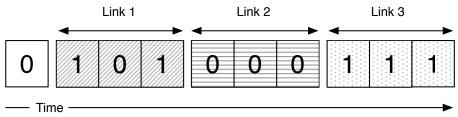

Figure 20: First frame, with sync bit = 0

As you can see the time slot is occupied by the first 3 bits from Link 1, followed by three consecutive bits each from Link 2 and Link 3 respectively. The following frame will start with the sync bit set to 1, prefixed to the following three consecutive bits from each Link.

Answer:
1 0 1 | 1 1 1 | 1 0 1 | 1 | 0 0 0 | 1 1 1 | 1 1 0 | 0 | 0 0 1 | 1 0 0 | 1 1 1 | 1 | 1 1 1 | 0 0 0 | 1 0 1 | 0

b

It is specified that each sub-link has a capacity of 300 kbps, as such the multiplexed link should have a capacity of at least 3×300 kbps $= 900$ kbps. However, the sync bit in each multiplexed frame of 9 data bits add one bit of overhead. As a result instead of sending 9 bits we are now sending 10 bits per frame. The link thus needs $\frac{10}{9} \approx 11\%$ extra capacity in addition to what is expect of the sub-links to maintaing the oncoming traffic from the sub-links and to accommodate the sync bit. Moreover, applying the additional capacity needed to the minimum multiplex link capacity yields, $\frac{10}{9} \times 900$ kbps $= 1000$ kbps $= 1$ Mbps.

Answer: 1 Mbps

c

In 12b we concluded that the multiplexed link needs to have a minimum bandwidth of 1Mbps. Furthermore, each one of the bits transmitted on the multiplex link would have to occupy at least $\frac{1}{10*6} = 10^{-6}$ seconds = 1 microsecond.

Answer: 1 microsecond

d

Using what we deducted in 12c, if one bit occupies 1 microsecond, then a frame of 10 bits will occupy 10 microseconds. Consequently, one second is occupied by $\frac{1}{10 \times 10^{-6}} = 10^5 = 100000$ frames.

Answer: 100 000

e

As deducted in 12d, if one bit occupies 1 microsecond, then a frame of 10 bits occupy 10 microseconds.

Answer: 10 microseconds

13

a

The observed bit stream consists of 24 bits and can thus be broken down into 2 frames, each consisting of 3 time slots carrying 4 bits each. See Figure 21.

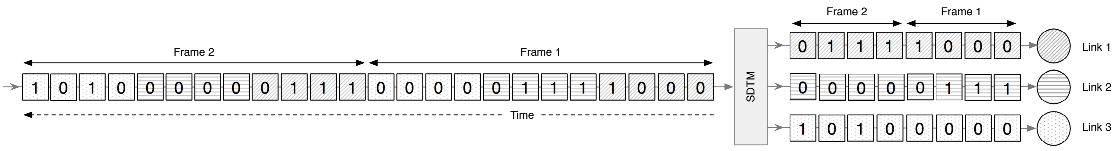

Figure 21: Stream demultiplexed into 3 sub-streams

As there are no synchronizations bits in the stream, each frame can trivially be broken down and separated into each sub-links time slots and then reassembled, as illustrated in Figure 21.

Answer:

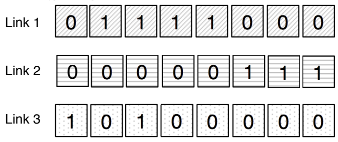

b

The main line is demultiplexed into three congruent sub streams. The specified frame only carries data, we can thus conclude that each resulting sub-link is is feed data at a rate of $\frac{9}{3} = 3$ Mbps.

Answer: 3 Mbps