SOLUTIONS TO CHAPTER 2 PROBLEMS

1. $a_n = \frac{1}{\pi n}$, $b_n = 0$, $c = 1$.

2. A noiseless channel can carry an arbitrarily large amount of information, no matter how often it is sampled. Just send a lot of data per sample. For the 4 kHz channel, make 8000 samples/sec. If each sample is 16 bits, the channel can send 128 kbps. If each sample is 1024 bits, the channel can send 8.2 Mbps. The key word here is ‘‘noiseless.’’ With a normal 4 kHz channel, the Shannon limit would not allow this.

3. Using the Nyquist theorem, we can sample 12 million times/sec. Four-level signals provide 2 bits per sample, for a total data rate of 24 Mbps.

4. A signal-to-noise ratio of 20 dB means $S/N = 100$. Since $\log_2 101$ is about 6.658, the Shannon limit is about 19.975 kbps. The Nyquist limit is 6 kbps. The bottleneck is therefore the Nyquist limit, giving a maximum channel capacity of 6 kbps.

5. To send a T1 signal we need $H \log_2 (1 + S/N) = 1.544 \times 10^6$ with $H = 50,000$. This yields $S/N = 2^{30} - 1$, which is about 93 dB.

6. A passive star has no electronics. The light from one fiber illuminates a number of others. An active repeater converts the optical signal to an electrical one for further processing.

7. Use $\Delta f = c \Delta\lambda / \lambda^2$ with $\Delta\lambda = 10^{-7}$ meters and $\lambda = 10^{-6}$ meters. This gives a bandwidth (Δf) of 30,000 GHz.

8. The data rate is $480 \times 640 \times 24 \times 60$ bps, which is 442 Mbps. For simplicity, let us assume 1 bps per Hz. From Eq. (2-3) we get $\Delta\lambda = \lambda^2 \Delta f / c$. We have $\Delta f = 4.42 \times 10^8$, so $\Delta\lambda = 2.5 \times 10^{-6}$ microns. The range of wavelengths used is very short.

9. The Nyquist theorem is a property of mathematics and has nothing to do with technology. It says that if you have a function whose Fourier spectrum does not contain any sines or cosines above f, then by sampling the function at a frequency of $2f$ you capture all the information there is. Thus, the Nyquist theorem is true for all media.

10. In the text it was stated that the bandwidths (i.e., the frequency ranges) of the three bands were approximately equal. From the formula $\Delta f = c\Delta\lambda/\lambda^2$, it is clear that to get a constant Δf, the higher the frequency, the larger $\Delta\lambda$ has to be. The x-axis in the figure is λ, so the higher the frequency, the more $\Delta\lambda$ you need. In fact, $\Delta\lambda$ is quadratic in λ. The fact that the bands are approximately equal is an accidental property of the kind of silicon used.

11. Start with $\lambda f = c$. We know that c is 3×10^8 m/s. For $\lambda = 1$ cm, we get 30 GHz. For $\lambda = 5$ m, we get 60 MHz. Thus, the band covered is 60 MHz to 30 GHz.

12. At 1 GHz, the waves are 30 cm long. If one wave travels 15 cm more than the other, they will arrive out of phase. The fact that the link is 50 km long is irrelevant.

13. If the beam is off by 1 mm at the end, it misses the detector. This amounts to a triangle with base 100 m and height 0.001 m. The angle is one whose tangent is thus 0.00001. This angle is about 0.00057 degrees.

14. With 66/6 or 11 satellites per necklace, every 90 minutes 11 satellites pass overhead. This means there is a transit every 491 seconds. Thus, there will be a handoff about every 8 minutes and 11 seconds.

15. The satellite moves from being directly overhead toward the southern horizon, with a maximum excursion from the vertical of 2ϕ. It takes 24 hours to go from directly overhead to maximum excursion and then back.

16. The number of area codes was $8 \times 2 \times 10$, which is 160. The number of prefixes was $8 \times 8 \times 10$, or 640. Thus, the number of end offices was limited to 102,400. This limit is not a problem.

17. With a 10-digit telephone number, there could be 10^{10} numbers, although many of the area codes are illegal, such as 000. However, a much tighter limit is given by the number of end offices. There are 22,000 end offices, each with a maximum of 10,000 lines. This gives a maximum of 220 million telephones. There is simply no place to connect more of them. This could never be achieved in practice because some end offices are not full. An end office in a small town in Wyoming may not have 10,000 customers near it, so those lines are wasted.

18. Each telephone makes 0.5 calls/hour at 6 minutes each. Thus, a telephone occupies a circuit for 3 minutes/hour. Twenty telephones can share a circuit, although having the load be close to 100% ($\rho = 1$ in queueing terms) implies very long wait times). Since 10% of the calls are long distance, it takes 200 telephones to occupy a long-distance circuit full time. The interoffice trunk has 1,000,000/4000 = 250 circuits multiplexed onto it. With 200 telephones per circuit, an end office can support $200 \times 250 = 50,000$ telephones.

19. The cross-section of each strand of a twisted pair is $\pi/4$ square mm. A 10-km length of this material, with two strands per pair has a volume of $2\pi/4 \times 10^{-2}$ m^3. This volume is about 15,708 cm^3. With a specific gravity of 9.0, each local loop has a mass of 141 kg. The phone company thus owns 1.4×10^9 kg of copper. At 3 dollars each, the copper is worth about 4.2 billion dollars.

20. Like a single railroad track, it is half duplex. Oil can flow in either direction, but not both ways at once.

21. Traditionally, bits have been sent over the line without any error correcting scheme in the physical layer. The presence of a CPU in each modem makes it possible to include an error correcting code in layer 1 to greatly reduce the effective error rate seen by layer 2. The error handling by the modems can be done totally transparently to layer 2. Many modems now have built in error correction.

22. There are four legal values per baud, so the bit rate is twice the baud rate. At 1200 baud, the data rate is 2400 bps.

23. The phase shift is always 0, but two amplitudes are used, so this is straight amplitude modulation.

24. If all the points are equidistant from the origin, they all have the same amplitude, so amplitude modulation is not being used. Frequency modulation is never used in constellation diagrams, so the encoding is pure phase shift keying.

25. Two, one for upstream and one for downstream. The modulation scheme itself just uses amplitude and phase. The frequency is not modulated.

26. There are 256 channels in all, minus 6 for POTS and 2 for control, leaving 248 for data. If 3/4 of these are for downstream, that gives 186 channels for downstream. ADSL modulation is at 4000 baud, so with QAM-64 (6 bits/baud) we have 24,000 bps in each of the 186 channels. The total bandwidth is then 4.464 Mbps downstream.

27. A 5-KB Web page has 40,000 bits. The download time over a 36 Mbps channel is 1.1 msec. If the queueing delay is also 1.1 msec, the total time is 2.2 msec. Over ADSL there is no queueing delay, so the download time at 1 Mbps is 40 msec. At 56 kbps it is 714 msec.

28. There are ten 4000 Hz signals. We need nine guard bands to avoid any interference. The minimum bandwidth required is $4000 \times 10 + 400 \times 9 = 43,600$ Hz.

29. A sampling time of 125 μsec corresponds to 8000 samples per second. According to the Nyquist theorem, this is the sampling frequency needed to capture all the information in a 4 kHz channel, such as a telephone channel. (Actually the nominal bandwidth is somewhat less, but the cutoff is not sharp.)

30. The end users get $7 \times 24 = 168$ of the 193 bits in a frame. The overhead is therefore 25/193 = 13%.

31. In both cases 8000 samples/sec are possible. With dibit encoding, two bits are sent per sample. With T1, 7 bits are sent per period. The respective data rates are 16 kbps and 56 kbps.

32. Ten frames. The probability of some random pattern being 0101010101 (on a digital channel) is 1/1024.

33. A coder accepts an arbitrary analog signal and generates a digital signal from it. A demodulator accepts a modulated sine wave only and generates a digital signal.

34. (a) 64 kbps. (b) 32 kbps. (c) 8 kbps.

35. The signal must go from 0 to A in one quarter of a wave—that is, in a time $T/4$. In order to track the signal, 8 steps must fit into the quarter wave, or 32 samples per full wave. The time per sample is $1/x$ so the full period must be long enough to contain 32 samples—that is, $T > 32/x$ or $f_{max} = x/32$.

36. A drift rate of 10^{-9} means 1 second in 10^9 seconds or 1 nsec per second. At OC-1 speed, say, 50 Mbps, for simplicity, a bit lasts for 20 nsec. This means it takes only 20 seconds for the clock to drift off by one bit. Consequently, the clocks must be continuously synchronized to keep them from getting too far apart. Certainly every 10 sec, preferably much more often.

37. Of the 90 columns, 86 are available for user data in OC-1. Thus, the user capacity is $86 \times 9 = 774$ bytes/frame. With 8 bits/byte, 8000 frames/sec, and 3 OC-1 carriers multiplexed together, the total user capacity is $3 \times 774 \times 8 \times 8000$, or 148.608 Mbps.

38. VT1.5 can accommodate 8000 frames/sec × 3 columns × 9 rows × 8 bits = 1.728 Mbps. It can be used to accommodate DS-1. VT2 can accommodate 8000 frames/sec × 4 columns × 9 rows × 8 bits = 2.304 Mbps. It can be used to accommodate European CEPT-1 service. VT6 can accommodate 8000 frames/sec × 12 columns × 9 rows × 8 bits = 6.912 Mbps. It can be used to accommodate DS-2 service.

39. Message switching sends data units that can be arbitrarily long. Packet switching has a maximum packet size. Any message longer than that is split up into multiple packets.

40. The OC-12c frames are $12 \times 90 = 1080$ columns of 9 rows. Of these, $12 \times 3 = 36$ columns are taken up by line and section overhead. This leaves an SPE of 1044 columns. One SPE column is taken up by path overhead, leaving 1043 columns for user data. Since each column holds 9 bytes of 8 bits, an OC-12c frame holds 75,096 user data bits. With 8000 frames/sec, the user data rate is 600.768 Mbps.

41. The three networks have the following properties:
star: best case = 2, average case = 2, worst case = 2
ring: best case = 1, average case = $n/4$, worst case = $n/2$
full interconnect: best case = 1, average case = 1, worst case = 1

42. With circuit switching, at $t = s$ the circuit is set up; at $t = s + x/b$ the last bit is sent; at $t = s + x/b + kd$ the message arrives. With packet switching, the last bit is sent at $t = x/b$. To get to the final destination, the last packet must be retransmitted $k - 1$ times by intermediate routers, each retransmission taking p/b sec, so the total delay is $x/b + (k - 1)p/b + kd$. Packet switching is faster if $s > (k - 1)p/b$.

43. The total number of packets needed is x/p, so the total data + header traffic is $(p + h)x/p$ bits. The source requires $(p + h)x/pb$ sec to transmit these bits. The retransmissions of the last packet by the intermediate routers take up a total of $(k - 1)(p + h)/b$ sec. Adding up the time for the source to send all the bits, plus the time for the routers to carry the last packet to the destination, thus clearing the pipeline, we get a total time of $(p + h)x/pb + (p + h)(k - 1)/b$ sec. Minimizing this quantity with respect to p, we find $p = \sqrt{hx/(k - 1)}$.

44. Each cell has six neighbors. If the central cell uses frequency group A, its six neighbors can use B, C, B, C, B, and C respectively. In other words, only 3 unique cells are needed. Consequently, each cell can have 280 frequencies.

45. First, initial deployment simply placed cells in regions where there was high density of human or vehicle population. Once they were there, the operator often did not want to go to the trouble of moving them. Second, antennas are typically placed on tall buildings or mountains. Depending on the exact location of such structures, the area covered by a cell may be irregular due to obstacles near the transmitter. Third, some communities or property owners do not allow building a tower at a location where the center of a cell falls. In such cases, directional antennas are placed at a location not at the cell center.

46. If we assume that each microcell is a circle 100 m in diameter, then each cell has an area of 2500π. If we take the area of San Francisco, 1.2×10^8 m^2 and divide it by the area of 1 microcell, we get 15,279 microcells. Of course, it is impossible to tile the plane with circles (and San Francisco is decidedly three-dimensional), but with 20,000 microcells we could probably do the job.

47. Frequencies cannot be reused in adjacent cells, so when a user moves from one cell to another, a new frequency must be allocated for the call. If a user moves into a cell, all of whose frequencies are currently in use, the user's call must be terminated.

48. It is not caused directly by the need for backward compatibility. The 30 kHz channel was indeed a requirement, but the designers of D-AMPS did not have to stuff three users into it. They could have put two users in each channel, increasing the payload before error correction from $260 \times 50 = 13$ kbps to $260 \times 75 = 19.5$ kbps. Thus, the quality loss was an intentional trade-off to put more users per cell and thus get away with bigger cells.

49. D-AMPS uses 832 channels (in each direction) with three users sharing a single channel. This allows D-AMPS to support up to 2496 users simultaneously per cell. GSM uses 124 channels with eight users sharing a single channel. This allows GSM to support up to 992 users simultaneously. Both systems use about the same amount of spectrum (25 MHz in each direction). D-AMPS uses 30 KHz × 892 = 26.76 MHz. GSM uses 200 KHz × 124 = 24.80 MHz. The difference can be mainly attributed to the better speech quality provided by GSM (13 Kbps per user) over D-AMPS (8 Kbps per user).

50. The result is obtained by negating each of *A*, *B*, and *C* and then adding the three chip sequences. Alternatively the three can be added and then negated. The result is (+3 +1 +1 −1 −3 −1 −1 +1).

51. By definition

$$\mathbf{S} \bullet \mathbf{T} \equiv \frac{1}{m} \sum_{i=1}^{m} S_i T_i$$

If T sends a 0 bit instead of 1 bit, its chip sequence is negated, with the i-th element becoming $-T_i$. Thus,

$$\mathbf{S} \bullet \mathbf{T} \equiv \frac{1}{m} \sum_{i=1}^{m} S_i(-T_i) = -\frac{1}{m} \sum_{i=1}^{m} S_i T_i = 0$$

52. When two elements match, their product is +1. When they do not match, their product is −1. To make the sum 0, there must be as many matches as mismatches. Thus, two chip sequences are orthogonal if exactly half of the corresponding elements match and exactly half do not match.

53. Just compute the four normalized inner products:

(−1 +1 −3 +1 −1 −3 +1 +1) • (−1 −1 −1 +1 +1 −1 +1 +1)/8 = 1
(−1 +1 −3 +1 −1 −3 +1 +1) • (−1 −1 +1 −1 +1 +1 +1 −1)/8 = −1

$$(-1\ +1\ -3\ +1\ -1\ -3\ +1\ +1) \bullet (-1\ +1\ -1\ +1\ +1\ +1\ -1\ -1)/8 = 0$$
$$(-1\ +1\ -3\ +1\ -1\ -3\ +1\ +1) \bullet (-1\ +1\ -1\ -1\ -1\ -1\ +1\ -1)/8 = 1$$

The result is that *A* and *D* sent 1 bits, *B* sent a 0 bit, and *C* was silent.

54. Ignoring speech compression, a digital PCM telephone needs 64 kbps. If we divide 10 Gbps by 64 kbps we get 156,250 houses per cable. Current systems have hundreds of houses per cable.

55. It is both. Each of the 100 channels is assigned its own frequency band (FDM), and on each channel the two logical streams are intermixed by TDM. This example is the same as the AM radio example given in the text, but neither is a fantastic example of TDM because the alternation is irregular.

56. A 2-Mbps downstream bandwidth guarantee to each house implies at most 50 houses per coaxial cable. Thus, the cable company will need to split up the existing cable into 100 coaxial cables and connect each of them directly to a fiber node.

57. The upstream bandwidth is 37 MHz. Using QPSK with 2 bits/Hz, we get 74 Mbps upstream. Downstream we have 200 MHz. Using QAM-64, this is 1200 Mbps. Using QAM-256, this is 1600 Mbps.

58. Even if the downstream channel works at 27 Mbps, the user interface is nearly always 10-Mbps Ethernet. There is no way to get bits to the computer any faster than 10-Mbps under these circumstances. If the connection between the PC and cable modem is fast Ethernet, then the full 27 Mbps may be available. Usually, cable operators specify 10 Mbps Ethernet because they do not want one user sucking up the entire bandwidth.