

***Computer Networks 6th edition Tanenbaum
Solutions Manual***

SKU: 9780136764052-SOLUTIONS

COMPUTER NETWORKS

SIXTH EDITION

PROBLEM SOLUTIONS

ANDREW S. TANENBAUM

*Vrije Universiteit
Amsterdam, The Netherlands*

NICK FEAMSTER

*University of Chicago
Chicago, Illinois*

DAVID WETHERELL

Google

PRENTICE HALL PTR

UPPER SADDLE RIVER, NJ 07458

PROBLEM SOLUTIONS
SOLUTIONS TO CHAPTER 1 PROBLEMS

1

1. Because the raven flies at an average speed of 40 km/h, it needs $160/40 = 4$ hours for each one-way trip.

(i) The raven makes only one trip because 1.8 terabytes exactly fits on one scroll.

$$\frac{1800 \text{ GB}}{4 \text{ h} \times 3600} = \frac{1}{8} \text{ GB/s} = 1 \text{ Gbps}$$

(ii) To communicate 3.6 TB of data, the raven has to fly back to pick up a second scroll. This means that it needs to fly a total of $3 \times 4 = 12$ hours.

$$\frac{3600 \text{ GB}}{12 \text{ h} \times 3600} = \frac{1}{12} \text{ GB/s} = \frac{2}{3} \text{ Gbps}$$

(iii) The receiving castle receives 1.8 terabytes of data every 8 hours.

$$\frac{1800 \text{ GB}}{8 \text{ h} \times 3600} = \frac{1}{16} \text{ GB/s} = \frac{1}{2} \text{ Gbps}$$

2. There are multiple correct answers. A significant disadvantage is the increased risk of invading people's privacy. The increase in the number of networked devices means a larger attack surface for malicious parties trying to obtain personal information. If the information is not stolen, companies that process and store data from IoT devices could sell it to third parties such as advertising companies.
3. Secondly, wireless networks allow people to move around, instead of tying them to a wall. Secondly, although wireless networks provide lower bandwidth than wired networks, their bandwidth has become large enough to support applications that people find meaningful. Examples include media streaming and video conferencing. Finally, installing wires in (old) buildings can be expensive.
4. An advantage for the company is that they do not have to pay the up-front cost when buying expensive hardware. They lease machines from the data center, paying only for what they use. A disadvantage for the company is that they may not know the underlying infrastructure used by the data center, making it more difficult to obtain high performance from their applications. The large amount of resources available in data centers makes it is easier for the company to scale with user demand, which is an advantage for both. A disadvantage for the users is that it becomes more difficult to track their own data, and what it is used for.

5. The LAN model can be grown incrementally. If the LAN is just a long cable, it cannot be brought down by a single failure (if the servers are replicated). It is probably cheaper. It provides more computing power and better interactive interfaces.
6. A transcontinental fiber link might have many gigabits/sec of bandwidth, but the latency will also be high due to the speed of light propagation over thousands of kilometers. Similarly, a satellite link may run at megabits/sec but have a high latency to send a signal into orbit and back. In contrast, a 56-kbps modem calling a computer in the same building has low bandwidth and low latency. So do low-end local and personal area wireless technologies such as Zigbee.
7. No. The speed of propagation is 200,000 km/sec or 400 meters/ μ sec. In 20 μ sec, the signal travels 4 km. Thus, each switch adds the equivalent of 4 km of extra cable. If the client and server are separated by 5000 km, traversing even 50 switches adds only 200 km to the total path, which is only 4%. Thus, switching delay is not a major factor under these circumstances.
8. The delay is 1% of the total time, which means

$$\frac{100 \mu s \times n}{\frac{29,700 \text{ km}}{300,000 \text{ km/s}} + 100 \mu s \times n} = 0.01$$

, where n is the number of satellites.

$$\frac{29,700 \text{ km}}{300,000 \text{ km/s}} + 100 \mu s \times n = 100 \times 100 \mu s \times n$$

$$\frac{29,700 \text{ km}}{300,000 \text{ km/s}} = 99 \times 100 \mu s \times n$$

$$\frac{29,700 \text{ km}}{300,000 \text{ km/s} \times 99 \times 100 \mu s} = 10 = n$$

This means the signal must pass 10 satellites for the switching delay to be 1% of the total delay.

9. The request has to go up and down, and the response has to go up and down. The total path length traversed is thus 160,000 km. The speed of light in air and vacuum is 300,000 km/sec, so the propagation delay alone is 160,000/300,000 sec or about 533 msec.
10. Traveling at 2/3 the speed of light means 200,000 km/sec. The signal travels for 100 milliseconds, or 0.1 seconds. This means the signal traversed a distance of $200,000 \times 0.1 = 4000$ km.

11. There is obviously no single correct answer here, but the following points seem relevant. The present system has a great deal of inertia (checks and balances) built into it. This inertia may serve to keep the legal, economic, and social systems from being turned upside down every time a different party comes to power. Also, many people hold strong opinions on controversial social issues, without really knowing the facts of the matter. Allowing poorly reasoned opinions be written into law may be undesirable. The potential effects of advertising campaigns by special interest groups of one kind or another also have to be considered. Another major issue is security. A lot of people might worry about some 14-year kid hacking the system and falsifying the results.
12. Call the routers A, B, C, D , and E . There are ten potential lines: $AB, AC, AD, AE, BC, BD, BE, CD, CE$, and DE . Each of these has four possibilities (three speeds or no line), so the total number of topologies is $4^{10} = 1,048,576$. At 50 ms each, it takes 52,428.8 sec, or about 14.6 hours to inspect them all.
13. The mean router-router path is twice the mean router-root path. Number the levels of the tree with the root as 1 and the deepest level as n . The path from the root to level n requires $n - 1$ hops and 0.50 of the routers are at this level. The path from the root to level $n - 1$ has 0.25 of the routers and a length of $n - 2$ hops. Hence, the mean path length, l , is given by

$$l = 0.5 \times (n - 1) + 0.25 \times (n - 2) + 0.125 \times (n - 3) + \cdots$$

or

$$l = \sum_{i=1}^{\text{infinity}} n (0.5)^i - \sum_{i=1}^{\text{infinity}} i(0.5)^i$$

This expression reduces to $l = n - 2$. The mean router-router path is thus $2n - 4$.

14. Distinguish $n + 2$ events. Events 1 through n consist of the corresponding host successfully attempting to use the channel, i.e., without a collision. The probability of each of these events is $p(1 - p)^{n-1}$. Event $n + 1$ is an idle channel, with probability $(1 - p)^n$. Event $n + 2$ is a collision. Since these $n + 2$ events are exhaustive, their probabilities must sum to unity. The probability of a collision, which is equal to the fraction of slots wasted, is then just $1 - np(1 - p)^{n-1} - (1 - p)^n$.
15. Instead of trying to foresee bad things and avoid them from happening, successful networks are fault-tolerant. They allow bad things to happen but isolate or hide them from the rest of the system. Examples include error correction, error detection, and network routing.

16. Because the responsibility of networking is distributed over multiple layers, each layer only has partial knowledge of where the data needs to go. The link layer only knows to which machine the data should be sent next. The network layer knows which machine on the entire network is the correct destination. The transport layer knows to which process on the destination machine to deliver the data.

17.

Guarantee Layer	
Best effort delivery Network	
Reliable Delivery Transport	
In-order Delivery Transport	
Byte-stream abstraction Transport	
Point-to-point link abstraction Link	

18.

Function	Interface
send_bits_over_link(bits)	Physical layer
send_bytes_to_process(dst, src, bytes)	Transport layer
send_bytes_over_link(dst, src, bytes)	Link layer
send_bytes_to_machine(dst, src, bytes)	Network layer

19. $5 \times 1500 = 7,500$ bytes per 100 milliseconds. So, the rate is 75,000 bytes per second.
20. In the OSI protocol model, physical communication between peers takes place only in the lowest layer, not in every layer.
21. Message and byte streams are different. In a message stream, the network keeps track of message boundaries. In a byte stream, it does not. For example, suppose a process writes 1024 bytes to a connection and then a little later writes another 1024 bytes. The receiver then does a read for 2048 bytes. With a message stream, the receiver will get two messages, of 1024 bytes each. With a byte stream, the message boundaries do not count and the receiver will get the full 2048 bytes as a single unit. The fact that there were originally two distinct messages is lost.
22. Negotiation has to do with getting both sides to agree on some parameters or values to be used during the communication. Maximum packet size is one example, but there are many others.

23. The service shown is the service offered by layer k to layer $k + 1$. Another service that must be present is below layer k , namely, the service offered to layer k by the underlying layer $k - 1$.
24. The probability, P_k , of a frame requiring exactly k transmissions is the probability of the first $k - 1$ attempts failing, p^{k-1} , times the probability of the k -th transmission succeeding, $(1 - p)$. The mean number of transmission is then just

$$\sum_{k=1}^{\infty} kP_k = \sum_{k=1}^{\infty} k(1 - p)p^{k-1} = \frac{1}{1 - p}$$

Or, more directly, if the probability of a message getting through is $1 - p$, then the expected number of transmissions per successful message is $1 / (1 - p)$.

25. OSI model: (a) Data link layer. (b) Network layer.
TCP/IP model: (a) Link layer. (b) Internet layer.
26. Frames encapsulate packets. When a packet arrives at the data link layer, the entire thing—header, data, and all—is used as the data field of a frame. The entire packet is put in an envelope (the frame), so to speak (assuming it fits).
27. Each layer considers the data from the layer above it as its payload, and adds their header and/or trailer on the outside of this payload. The resulting order of headers and trailers is: [1][3][5][M][6][4][2].
28. With n layers and h bytes added per layer, the total number of header bytes per message is hn , so the space wasted on headers is hn . The total message size is $M + nh$, so the fraction of bandwidth wasted on headers is $hn/(M + hn)$. This estimate does not take into account fragmentation (one higher layer message is sent as multiple lower layer messages) or aggregation (multiple higher layer messages are carried as one lower layer message) that may be present. If fragmentation is used, it will raise the overhead. If aggregation is used, it will lower the overhead.
29. Many answers are possible. Most have to do with connecting a otherwise isolated devices to the Internet. (a) A wireless modem is the access point for a WiFi network and connects to the cable provider's network. (b) A phone connected to 3G or WiMAX and functioning as the access point for a WiFi network. (c) A phone that is connected to a WiFi network and at the same time acts as the master in a Bluetooth network. (d) A router at an IXP that connects an ISP's network to that from other ISPs. (e) A phone that is connected to two base stations at the same time during a soft handover. This last answer does not connect another device to the Internet, but rather creates the illusion of a continuous connection to the user.

30. Observe that many nodes are connected to three other nodes; the others are connected to more. Three bombs are needed to disconnect one of these nodes. By a quick check there does not appear to be a group of nodes that are connected to the rest of the network by fewer than three other nodes, so we conclude that three bombs are needed to partition the network. For example, the two nodes in the upper-right corner can be disconnected from the rest by three bombs knocking out the three nodes to which they are connected. The system can withstand the loss of any two nodes.
31. Doubling every 18 months means a factor of four gain in 3 years. In 9 years, the gain is then 4^3 or 64, leading to 64 billion hosts. That sounds like a lot, but if every television, cellphone, camera, car, and appliance in the world is online, maybe it is plausible. It would require the average person to have 10 hosts by then given that the estimate is much greater than the expected world population. But if half the world is connected and they average 20 devices per person, it is at least plausible.
32. If the network tends to lose packets, it is better to acknowledge each one separately, so the lost packets can be retransmitted. On the other hand, if the network is highly reliable, sending one acknowledgement at the end of the entire transfer saves bandwidth in the normal case (but requires the entire file to be retransmitted if even a single packet is lost).
33. Having mobile phone operators know the location of users lets the operators learn much personal information about users, such as where they sleep, work, travel, and shop. This information might be sold to others or stolen; it could let the government monitor citizens. On the other hand, knowing the location of the user lets the operator send help to the right place in an emergency. It might also be used to deter fraud, since a person who claims to be you will usually be near your mobile phone.
34. The speed of light in coax is about 200,000 km/sec, which is 200 meters/ μ sec. At 10 Mbps, it takes 0.1 μ sec to transmit a bit. Thus, the bit lasts 0.1 μ sec in time, during which it propagates 20 meters. Thus, a bit is 20 meters long here.
35. The image is $1600 \times 1200 \times 3$ bytes or 5,760,000 bytes. This is 46,080,000 bits. At 56,000 bits/sec, it takes about 822.857 sec. At 1,000,000 bits/sec, it takes 46.080 sec. At 10,000,000 bits/sec, it takes 4.608 sec. At 100,000,000 bits/sec, it takes about 0.461 sec. At 1,000,000,000 bits/sec, it takes about 46 msec.
36. The image is $3840 \times 2160 \times 3$ bytes or 8,294,400 bytes. This is 66,355,200 bits. At 56,000 bits/sec, it takes about 1184.91 sec. At 1,000,000 bits/sec, it takes about 66.36 sec. At 10,000,000 bits/sec, it takes about 6.64 sec. At 100,000,000 bits/sec, it takes about 0.663 sec. At 1,000,000,000 bits/sec it takes about 64.5 msec.

37. Think about the hidden terminal problem. Imagine a wireless network of five stations, *A* through *E*, such that each one is in range of only its immediate neighbors. Then *A* can talk to *B* at the same time *D* is talking to *E*. Wireless networks have potential parallelism, and in this way differ from Ethernet.
38. One disadvantage is security. Every random delivery man who happens to be in the building can listen in on the network. Another disadvantage is reliability. Wireless networks make lots of errors. A third potential problem is battery life, since most wireless devices tend to be mobile.
39. One advantage is that if everyone uses the standard, everyone can talk to everyone. Another advantage is that widespread use of any standard will give it economies of scale, as with VLSI chips. A disadvantage is that the political compromises necessary to achieve standardization frequently lead to poor standards. Another disadvantage is that once a standard has been widely adopted, it is difficult to change,, even if new and better techniques or methods are discovered. Also, by the time it has been accepted, it may be obsolete.
40. There are many examples, of course. Some systems for which there is international standardization include DVD players and their discs, digital cameras and their storage cards, and automated teller machines and bank cards. Areas where such international standardization is lacking include broadcast television (NTSC in the U.S., PAL in parts of Europe, SECAM in other countries), lamps and lightbulbs (different voltages in different countries), electrical sockets and appliance plugs (every country does it differently), photocopiers and paper (8.5 x 11 inches in the U.S., A4 everywhere else), nuts and bolts (English versus metric pitch), etc.
41. Networks are used in a large number of different environments, and each environment imposes different requirements on its protocols. For example, channels with high error rates may need to use error correction codes to achieve reasonable effective throughput, public networks may required additional security not needed in home networks, and commercial networks may need to offer Quality of Service guarantees to paying customers.
42. This has no impact on the operations at layers $k-1$ or $k+1$.
43. There is no impact at layer $k-1$, but operations in $k+1$ have to be reimplemented.
44. Navigating to a webpage will likely trigger multiple GET requests. Each of these requests asks for a certain resource, such as the HTML document, the CSS document, and external resources such as images. Doing these requests separately can improve performance by, for instance, obtaining and rendering the page HTML while downloading an external image.

45. Below are ten possible answers, although many more exist.

1. Buying consumer goods through an online store.
2. Watching video-on-demand.
3. Listening to music using a streaming service.
4. Communicating with others through email or instant messaging.
5. Storing documents and photos on cloud-based storage.
6. Reading the news on the Web.
7. Handing in homework on your school's online learning system
8. Doing your taxes on the government's website.
9. Transferring money via a banking application on your phone.
10. Playing video games with friends.

SOLUTIONS TO CHAPTER 2 PROBLEMS

1. Like a single railroad track, it is half duplex. Oil can flow in either direction, but not both ways at once. A river is an example of a simplex connection while a walkie-talkie is another example of a half-duplex connection.
2. Fiber has many advantages over copper. It can handle much higher bandwidth than copper. It is not affected by power surges, electromagnetic interference, power failures, or corrosive chemicals in the air. It does not leak light and is quite difficult to tap. Finally, it is thin and lightweight, resulting in much lower installation costs. There are some downsides of using fiber over copper. It can be damaged easily by being bent too much. And fiber interfaces cost more than electrical interfaces.
3. Use Eq. (2-1) to convert wavelengths of 1 micron plus/minus 0.05 microns to frequency. We have $f_{low} = 3 \times 10^8 / (1.05 \times 10^{-6}) = 3/1.05 \times 10^{14}$. Similarly $f_{high} = 3/0.95 \times 10^{14}$. Thus $\Delta f = (3/0.95 - 3/1.05) \times 10^{14} = 3 \times 10^{13}$. This is a bandwidth of 30,000 GHz.
4. The data rate is $3840 \times 2160 \times 24 \times 60$ bps, which is about 11,944 Mbps or 11.944 Gbps.
5. In the text it was stated that the bandwidths (i.e., the frequency ranges) of the three bands were approximately equal. From the formula in the text we have $f = c/\lambda$. For a range of wavelengths λ to $\lambda + x$ we have $f_{lower} = c/(\lambda + x)$ and $f_{upper} = c/\lambda$. For a fixed value of x , the range will be larger if λ is smaller. This is why the range of wavelengths on the left is smaller even though the range of frequencies is approximately equal in all bands.
6. This would be a major scientific breakthrough. Compared to electrons, which are subject to electromagnetic forces, it is difficult to make photons interact. Assuming there are computers that work this way, the maximum data rates obtained from information theory would not change. However, it is no longer necessary to translate between optical and electrical signals when using optical fiber communication. This makes it possible to obtain much higher data rates using the same optical fibers.
7. Start with $\lambda f = c$. We know that c is 3×10^8 m/s. For $\lambda = 1$ cm, we get 30 GHz. For $\lambda = 1$ m, we get 300 MHz. Thus, the band covered is 300 MHz to 30 GHz.
8. 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 100 km long is irrelevant.

9. 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.
10. $a_n = \frac{-1}{\pi n}$, $b_n = 0$, $c = 1$.
11. Shannon's theorem gives $5 \times \log_2(1 + 10000) \approx 66.44$ Gbps. Nyquist theorem gives $2 \times 5 \log_2(2) = 10$ Gbps. 10 Gbps is the lowest upper bound. The Nyquist theorem yields a lower upper-bound because the signal is binary. On a channel with a 40 dB signal-to-noise ratio, a higher data rate can be achieved by using more signal levels.
12. 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 3-kHz channel, make 6000 samples/sec. If each sample is 16 bits, the channel can send 96 kbps. If each sample is 1024 bits, the channel can send 8.2 Mbps. The key word here is "noiseless." With a normal 3 kHz channel, the Shannon limit would not allow this. A signal-to-noise ratio of 30 dB means $S/N = 1000$. With $B = 3000$ we get a maximum data rate of about 29.895 kbps.
13. 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 , 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.
14. 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.
15. 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 for a binary signal (with 1 bit per symbol). The bottleneck is therefore the Nyquist limit, giving a maximum channel capacity of 6 kbps.
16. For every four data bits, 5 bits are sent. The channel sends 80 million bits per second, which can be done using a 40 MHz channel.
17. Only the distance from the origin differs, thus Amplitude Shift Keying is used.
18. Yes. QAM-16 uses 16 distinct symbols. Each of these symbols can be assigned to a bit sequence, in which case it can send 4 bits per symbol. Alternatively, some of the symbols can be used as control signals, leaving fewer signals to send bit sequences.
19. In NRZ, the signal completes a cycle at most every 2 bits (alternating 1s and 0s). So, the minimum bandwidth need to achieve B bits/sec data rate is $B/2$ Hz. In MLT-3, the signal completes a cycle at most every 4 bits (a sequence of 1s), thus requiring at least $B/4$ Hz to achieve B bits/sec data rate. Finally, in