SOLUTIONS MANUAL

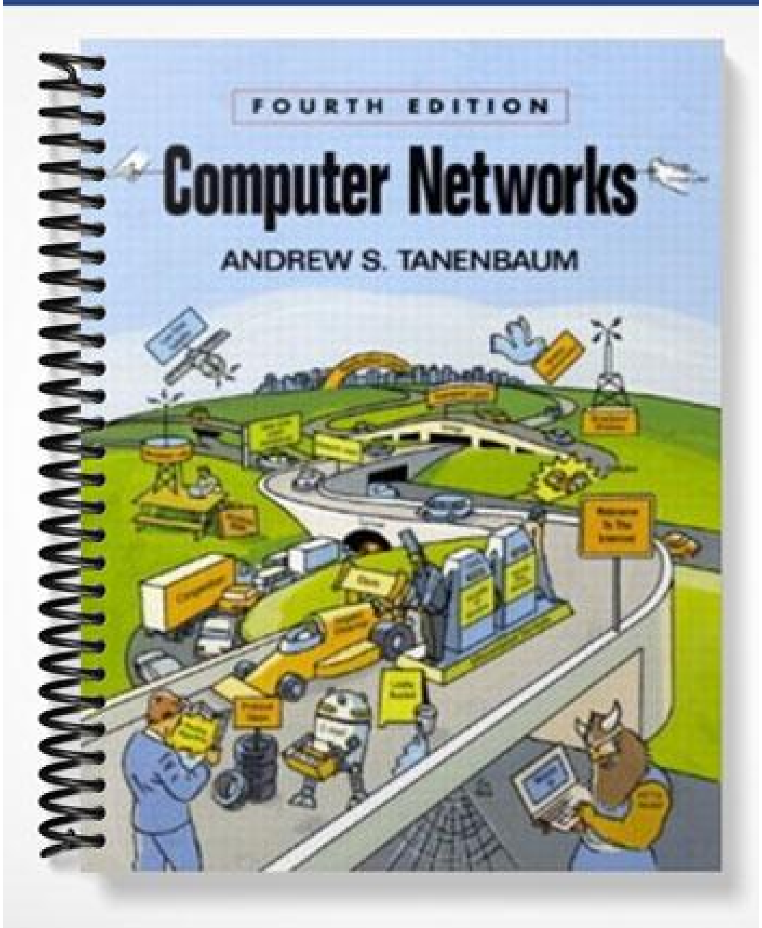

COMPUTER NETWORKS

FOURTH EDITION

PROBLEM SOLUTIONS

ANDREW S. TANENBAUM
Vrije Universiteit
Amsterdam, The Netherlands

PRENTICE HALL PTR

UPPER SADDLE RIVER, NJ 07458

SOLUTIONS TO CHAPTER 1 PROBLEMS

1. The dog can carry 21 gigabytes, or 168 gigabits. A speed of 18 km/hour equals 0.005 km/sec. The time to travel distance x km is $x/0.005 = 200x$ sec, yielding a data rate of $168/200x$ Gbps or $840/x$ Mbps. For $x < 5.6$ km, the dog has a higher rate than the communication line.

2. 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.

3. 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. In contrast, a 56-kbps modem calling a computer in the same building has low bandwidth and low latency.

4. A uniform delivery time is needed for voice, so the amount of jitter in the network is important. This could be expressed as the standard deviation of the delivery time. Having short delay but large variability is actually worse than a somewhat longer delay and low variability.

5. No. The speed of propagation is 200,000 km/sec or 200 meters/μsec. In 10 μsec the signal travels 2 km. Thus, each switch adds the equivalent of 2 km of extra cable. If the client and server are separated by 5000 km, traversing even 50 switches adds only 100 km to the total path, which is only 2%. Thus, switching delay is not a major factor under these circumstances.

6. 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.

7. 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 to 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.

8. 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 100 ms each, it takes 104,857.6 sec, or slightly more than 29 hours to inspect them all.

9. 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

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

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

10. 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 $1 - np(1 - p)^{n-1} - (1 - p)^n$.

11. Among other reasons for using layered protocols, using them leads to breaking up the design problem into smaller, more manageable pieces, and layering means that protocols can be changed without affecting higher or lower ones,

12. No. In the ISO protocol model, physical communication takes place only in the lowest layer, not in every layer.

13. Connection-oriented communication has three phases. In the establishment phase a request is made to set up a connection. Only after this phase has been successfully completed can the data transfer phase be started and data transported. Then comes the release phase. Connectionless communication does not have these phases. It just sends the data.

14. 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.

15. 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.

16. 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 $k - 1$.

17. 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

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

18. (a) Data link layer. (b) Network layer.

19. 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).

20. 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 $hn/(M + hn)$.

21. Both models are based on layered protocols. Both have a network, transport, and application layer. In both models, the transport service can provide a reliable end-to-end byte stream. On the other hand, they differ in several ways. The number of layers is different, the TCP/IP does not have session or presentation layers, OSI does not support internetworking, and OSI has both connection-oriented and connectionless service in the network layer.

22. TCP is connection oriented, whereas UDP is a connectionless service.

23. 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.

24. 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 6.4 billion hosts. My intuition says that is much too conservative, since by then probably every television in the world and possibly billions of other appliances will be on home LANs connected to the Internet. The average person in the developed world may have dozens of Internet hosts by then.

25. 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).

26. Small, fixed-length cells can be routed through switches quickly, and completely in hardware. Small, fixed-size cells also make it easier to build hardware that handles many cells in parallel. Also, they do not block transmission lines for very long, making it easier to provide quality-of-service guarantees.

27. The speed of light in coax is about 200,000 km/sec, which is 200 meters/μsec. At 10 Mbps, it takes 0.1 μsec to transmit a bit. Thus, the bit lasts 0.1 μsec in time, during which it propagates 20 meters. Thus, a bit is 20 meters long here.

28. The image is $1024 \times 768 \times 3$ bytes or 2,359,296 bytes. This is 18,874,368 bits. At 56,000 bits/sec, it takes about 337.042 sec. At 1,000,000 bits/sec, it takes about 18.874 sec. At 10,000,000 bits/sec, it takes about 1.887 sec. At 100,000,000 bits/sec, it takes about 0.189 sec.

29. 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.

30. 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.

31. 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.

32. There are many examples, of course. Some systems for which there is international standardization include compact disc players and their discs, Walkman tape players and audio cassettes, cameras and 35mm film, and automated teller machines and bank cards. Areas where such international standardization is lacking include VCRs and videotapes (NTSC VHS in the U.S., PAL VHS in parts of Europe, SECAM VHS in other countries), portable telephones, 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.

SOLUTIONS TO CHAPTER 2 PROBLEMS

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

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