COMPUTER NETWORKS

FIFTH EDITION

PROBLEM SOLUTIONS

ANDREW S. TANENBAUM
Vrije Universiteit
Amsterdam, The Netherlands

and

DAVID J. WETHERALL
University of Washington
Seattle, WA

Revision of 1 Feb 2012

PRENTICE HALL
Upper Saddle River, NJ

Instructors:

Correct problems and solutions are important to us. If you find any errors, please notify us (with what you believe to be the correct answer!) by sending mail to David Wetherall (djw@uw.edu).

Please also check the online errata (at http://computernetworks5e.org) for revisions to the text that include a small number of end of chapter problems.

This date of last revision of these solutions is given on the cover page to help you check that you have the latest version.

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.

(i) If dog's speed is doubled, maximum value of x is also doubled.

(ii) If tape capacity is doubled, value of x is also doubled.

(iii) If data rate of the transmission line is doubled, value of x is halved.

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. 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 techologies such as Zigbee.

4. A uniform delivery time is needed for voice as well as video, 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. For financial transaction traffic, reliability and security are very important.

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

10. 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. One possible disadvantage is the performance of a layered system is likely to be worse than the performance of a monolithic system, although it is extremely difficult to implement and manage a monolithic system.

11. In the OSI protocol model, physical communication between peers takes place only in the lowest layer, not in every layer.

12. 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.

13. 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.

14. 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 $k - 1$.

15. 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

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

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

16. 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.

17. TCP is connection oriented, whereas UDP is a connectionless service. Alternatively, TCP provides a reliable service, whereas UDP provides an unreliable service.

18. 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.

19. 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 38.4 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 dozens of hosts by then given that the estimate is much greater than the expected world population.

20. 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).

21. 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.

22. 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.

23. 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.

24. 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.

25. 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.

26. 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.

27. This has no impact on the operations at layers k-1 or k+1.

28. There is no impact at layer k-1, but operations in k+1 have to be reimplemented.

29. One reason is request or response messages may get corrupted or lost during transmission, which will necessitate a retransmission delay. Another reason is the processing unit in the client may get overloaded processing several requests at once. Yet another reason is that the request or response messages may be queued in the network along with other messages.

30. Small-sized cells result in large header-to-payload overhead. Fixed-size cells result in wastage of unused bytes in the payload.