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

SOLUTIONS TO CHAPTER 1 PROBLEMS

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.

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

(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 × 4 = 12 hours.

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

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

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

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 underlaying 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/μsec. In 20 μsec, 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

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

, where n is the number of satellites.

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

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

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

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 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.

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

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

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.