

Computer Networks

Problem Set 1A

Instructor: Shashi Prabh

Peterson and Davie, Chapter 1. Foundation

1.1

Use anonymous FTP to connect to ftp.rfc-editor.org (directory in-notes), and retrieve the RFC index. Also, retrieve the protocol specifications for TCP, IP, and UDP.

1.2 Calculate the total time required to transfer a 1000-KB file in the following cases, assuming an RTT of 50 ms, a packet size of 1 KB data, and an initial $2 \times$ RTT of “handshaking” before data is sent:

- a. The bandwidth is 1.5 Mbps, and data packets can be sent continuously.
- b. The bandwidth is 1.5 Mbps, but after we finish sending each data packet we must wait one RTT before sending the next.
- c. The bandwidth is “infinite,” meaning that we take transmit time to be zero, and up to 20 packets can be sent per RTT.
- d. The bandwidth is infinite, and during the first RTT we can send one packet ($2^1 - 1$), during the second RTT we can send two packets ($2^2 - 1$), during the third we can send four ($2^3 - 1$), and so on. (A justification for such an exponential increase will be given in Chapter 6.)

1.3 Consider a point-to-point link 4 km in length. At what bandwidth would propagation delay (at a speed of 2×10^8 m/s) equal the transmit delay for 100-byte packets? What about 512-byte packets?

1.4 One property of addresses is that they are unique; if two nodes had the same address, it would be impossible to distinguish between them. What other properties might be useful for network addresses to have? Can you think of any situations in which network (or postal or telephone) addresses might not be unique?

1.5 Give an example of a situation in which multicast addresses might be beneficial.

1.6 What differences in traffic patterns account for the fact that STDM is a cost-effective form of multiplexing for a voice telephone network and FDM is a cost-effective form of multiplexing for television and radio networks, yet we reject both as not being cost effective for a general-purpose computer network?

1.7 How “wide” is a bit on a 10-Gbps link? How long is a bit in copper wire, where the speed of propagation is 2.3×10^8 m/s?

1.8 How long does it take to transmit x KB over a y -Mbps link? Give your answer as a ratio of x and y .

1.9 Suppose a 1-Gbps point-to-point link is being set up between the Earth and a new lunar colony. The distance from the moon to the Earth is approximately 385,000 km, and data travels over the link at the speed of light— 3×10^8 m/s.

- a. Calculate the minimum RTT for the link.
- b. Using the RTT as the delay, calculate the delay \times bandwidth product for the link.
- c. What is the significance of the delay \times bandwidth product computed in (b)?
- d. A camera on the lunar base takes pictures of the Earth and saves them in digital format to disk. Suppose Mission Control on Earth wishes to download the most current image, which is 25 MB. What is the minimum amount of time that will elapse between when the request for the data goes out and the transfer is finished?

1.10 Calculate the latency (from first bit sent to last bit received) for the following:

- a. 100-Mbps Ethernet with a single store-and-forward switch in the path and a packet size of 12,000 bits. Assume that each link introduces a propagation delay of $10 \mu\text{s}$ and that the switch begins retransmitting immediately after it has finished receiving the packet.
- b. Same as (a), but with three switches.
- c. Same as (a), but assume the switch implements “cut-through” switching; it is able to begin retransmitting the packet after the first 200 bits have been received.

1.11 Hosts A and B are each connected to a switch S via 100-Mbps links as in Figure 1. The propagation delay on each link is $20 \mu\text{s}$. S is a store-and-forward device; it begins retransmitting a received packet $35 \mu\text{s}$ after it has finished receiving it. Calculate the total time required to transmit 10,000 bits from A to B:

- a. As a single packet.
- b. As two 5000-bit packets sent one right after the other.



■ **FIGURE 1.21** Diagram for Exercise 20.

Figure 1: Exercise 1.12: Network setup with switch S

1.12 Suppose that a certain communications protocol involves a per-packet overhead of 50 bytes for headers and framing. We send 1 million bytes of data using this protocol; however, one data byte is corrupted and the entire packet containing it is thus lost. Give the total number of overhead + loss bytes for packet data sizes of 1000, 10,000, and 20,000 bytes. Which size is optimal?

1.13 Assume you wish to transfer an n -byte file along a path composed of the source, destination, 7 point-to-point links, and 5 switches. Suppose each link has a propagation delay of 2 ms and a bandwidth of 4 Mbps, and that the switches support both circuit and packet switching. Thus, you can either break the file up into 1-KB packets or set up a circuit through the switches and send the file as one contiguous bitstream. Suppose that packets have 24 B of packet header information and 1000 B of payload, store-and-forward packet processing at each switch incurs a 1-ms delay after the packet has been completely received, packets may be sent continuously without waiting for acknowledgments, and circuit setup requires a 1-KB message to make one round trip on the path, incurring a 1-ms delay at each switch after the message has been completely received. Assume switches introduce no delay to data traversing a circuit. You may also assume that file size is a multiple of 1000 B.

- a. For what file size n is the total number of bytes sent across the network less for circuits than for packets?
- b. For what file size n is the total latency incurred before the entire file arrives at the destination less for circuits than for packets?
- c. How sensitive are these results to the number of switches along the path? To the bandwidth of the links? To the ratio of packet size to packet header size?
- d. How accurate do you think this model of the relative merits of circuits and packets is? Does it ignore important considerations that discredit one or the other approach? If so, what are they?

1.14 Obtain and build the `simplex-talk` sample socket program shown in the text. Start one server and one client, in separate windows. While the first client is running, start 10 other clients that connect to the same server; these other clients should most likely be started in the background with their input redirected from a file.

- What happens to these 10 clients?
- Do their `connect()` calls fail, or time out, or succeed?
- Do any other calls block?
- Now let the first client exit. What happens?
- Try this with the server value `MAX_PENDING` set to 1 as well.

1.15 Modify the `simplex-talk` socket program so that it uses UDP as the transport protocol, rather than TCP. You will have to change `SOCK_STREAM` to `SOCK_DGRAM` in both the client and the server.

- In the server, remove the calls to `listen()` and `accept()`.
- Replace the two nested loops at the end with a single loop that calls `recv()` with socket `s`.
- Run two UDP clients simultaneously and observe what happens when both connect to the same UDP server.

- Compare the behavior with that of the TCP version.

1.16 The Unix utility `ping` can be used to find the round-trip time (RTT) to various Internet hosts.

- Read the man page for `ping`.
- Use it to find the RTT to `www.cs.princeton.edu` in New Jersey and `www.cisco.com` in California.
- Measure the RTT values at different times of the day and compare the results.
- Discuss what factors might account for the differences in RTT.

1.17 Use the `traceroute` command to map out some of the routers within your organization or network.

- Run `traceroute` to several external destinations.
- Analyze the number of hops and latency per hop.
- Determine if internal routers are revealed or if the network directly connects to the internet.