CSE 461 Homework 3
Due Friday November 12, 1999

1. (From Peterson and Davie, p 342) The sequence number field in the TCP header is 32 bits long, which is big enough to cover over 4 billion bytes (4Gb) of data. Even if this many bytes were never transferred over a single connection, why might this sequence number still wrap from $2^{32} - 1$ to 0?

2. (Based on Peterson and Davie, p 342) You have been hired to design a reliable byte stream protocol that uses a sliding window, like TCP. This protocol will run over a 100Mbps network. The round trip time is 100 ms, and the maximum segment lifetime is 60 seconds. Your supervisor foolishly believes that by saving a few bits here and there in the transport packet header, that the delivered performance of this protocol will improve. They're paying you an obscene amount of money, so you go along with it.

   (a) What is the minimum number of bits you would have to include in the Advertised Window field?

   (b) What is the minimum number of bits you would have to include in the Sequence Number field?

3. Exercise 12.14 on page 393.

4. (From Tom Anderson's class) If two stations on an Ethernet communicate via a sliding window protocol, why might we see a collision on nearly every packet?

5. A TCP receiver may send an acknowledgement that advertises a zero-length window. This tells the sender that the window is "closed."

   (a) Why would a receiver send such an acknowledgement?

   (b) If the application reads a byte from the buffer, should the kernel send a new acknowledgement advertising a 1-byte window? Why or why not?

6. Suppose an application wants to send 100 bytes reliably using TCP. The Maximum Segment Size is 1460 bytes, so these 100 bytes clearly fit within a single packet. In the absence of loss, and over the entire duration of the connection,

   (a) How many packets will be sent by the sender?

   (b) How many packets will be sent by the receiver?

   (c) In multiples of the round trip time of the network between sender and receiver, how long will the connection last?

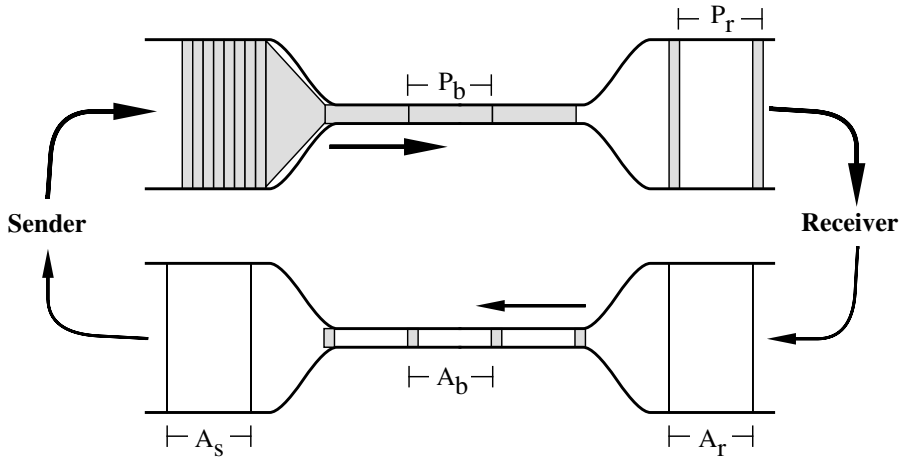

7. One mechanism for estimating the bottleneck link bandwidth sends a pair of back to back packets into the network, and measures the time between the receipt of acknowledgements. As these packets traverse the bottleneck link, they stretch: taking more time because of the reduced bandwidth. When those packets find a high bandwidth link, they contract in time, but preserve the spacing they acquired because of the bottleneck link. Acknowledgements for these packets also preserve the same spacing.

Because this scheme uses a pair of packets, it is referred to as "packet pair." Keshav describes this in the context of flow control in Section 13.4.8, but, aside from figure 13.10, you shouldn't need to refer to content there to answer this question.

Assume that the source sends back to back 1500 byte packets, and that the space between acknowledgements is 1ms.

What is the bandwidth of the bottleneck link?

As an editorial aside, there are two problems to know about this scheme. One is called "ack compression," where the returning acks might have to wait behind queued packets on the return path, which reduces the inter-ack spacing, making the bandwidth of the bottleneck appear larger.

Another is that PC clocks generally tick once every hundredth of a second, and it is somewhat difficult to measure time on a finer granularity. While packets traversing a modem can span hundreds of milliseconds, packets traversing a 10 Mbit Ethernet don't even show up. Since neither of these problems is presented in the book, it seemed worthwhile to let you in on the secret.