

Reliable Transport Services “Lab”

Tom Kelliher, CS 325

Mar. 12, 2008

1 Administrivia

Announcements

Assignment due NOW.

Assignment

Read 3.5–3.6.

From Last Time

Intro to transport-layer protocols; UDP.

Outline

1. Class design of a reliable transport protocol.

Coming Up

TCP; congestion control principles.

2 Class Design of a Reliable Transport Protocol

2.1 Lab Setup

1. $\frac{1}{2}$ of class is source host; $\frac{1}{2}$ of class is destination host; I'm the link.

We'll just do one-way communication; two-way communication is analogous.

We'll do this quietly — passing sheets of paper back and forth so neither side has global information.

2. Streams will be sentences containing a few words. Segments will be a single word.
3. In all cases, we'll assume the link doesn't reorder messages (How key is that assumption?), but will degrade the link in other ways as we go along:
 - (a) Perfect link.
 - (b) Garbles segments. (First data, then ACK).
 - (c) Drops segments.
4. We'll start with a stop & wait protocol. Then look at how wasteful that is and move on to two pipelined protocols:
 - (a) Go Back N.
 - (b) Selective repeat.

2.2 Experiments

Stop & Wait:

1. Perfect link — no check info needed.
2. Link garbles data segments — checksum, ACK/NACK needed.

3. Link garbles ACK/NACK — resend segment; one bit sequence number needed.

(An ACK for the ACK? No — resend segment.)

Drop NACK.

4. Link drops segments — timer needed to timeout and resend segment.

Stop & wait kills performance: consider a 1 Gbps link between the coasts with a RTT of 30 ms (speed of light) and a segment size of 1 KB. Link utilization is (transmit time)/(transmit time + RTT):

$$\frac{0.008}{30 + 0.008} = 0.00027$$

where transmit time is $\frac{8 \times 1000}{10^9}$, or 8 microseconds.

Pipelined protocol semantics:

1. Source can only have as many segments in flight as it has segments in the window.
2. Source slides window past received ACKs.
3. Destination slides window past sent ACKs.
4. Destination only accepts segments in order. (Simplification.)

Go Back N:

1. Sequence number space size vs. window size — Use two bit sequence numbers and a three segment window. Drop the first ACK.

Selective repeat.

Discussion question: What are the limitations on window size? (At most half sequence number space size, buffer space at hosts, congestion control, total buffer space should be less than bandwidth delay product, etc.)

What about out of order delivery? Can it even happen? (Yes — multiple routes between hosts.)