

Chapter 19 – Interactive Dataflow

Objectives of dataflow management

Provide snappy interactive response
Don't let a fast sender overwhelm a slow receiver
Don't let sender inject packets faster than network can deliver them.

Traffic characterization:

Interactive: X, telnet, rlogin
50% of the packets and 10% of the bytes
Bulk: FTP, e-mail, news
50% of the packets and 90% of the bytes
Web: Some characteristics of both

Interactive traffic

Each *keystroke* may require 4 forty-byte packets (160 bytes of total traffic)

- The ascii-code generated by the stroke
- The ack for the ascii-code
- The echo of the character to go on the screen (typically the same)
- The ack for the echo
- Not a big deal for E-net (*min packet size is 64 bytes*).
- A much more significant concern for dial links (*ppp, slip*).

Reducing this overhead

- Withhold acks so that the ack for the forward code goes back with the echo.
- Batch keystrokes and send them in a bundle (works better with local echo as in Web)

TCP has a *delayed ACK timer* that goes off every 200 ms of real time

- All pending acks are sent when the timer goes off
- Thus an ack may be delayed for as little as 1ms or as much as the full 200.

See figure 19.3 in the book for an illustration of:

- The delayed ack phenomenon
- Permits key-stroke acks to ride the echos
- Acks of echos go standalone
- Shows how delayed acks occur at fixed intervals (note t , Δt) columns in the time line at left.
- Multiple byte replies do go as a single message

The Nagle Algorithm

Designed to reduce the number of small segments outstanding in the net.
A connection may have only a single small, un-acked segment on the net.

(Small = smaller than MSS)

Nagle algorithm forces TCP to operate in stop and wait mode.

See figure 19.6 for an illustration of how:

Even though a keystroke generates 3 characters in rapid succession

One byte gets sent followed by the other 2 in the forward direction.

A smarter terminal emulator could avoid these problems and send meta keys and echos as single messages.

Big delays can cause real problems... receipt of meta keys is often timeout based.

Disabling the Nagle Algorithm

For good interactive performance you want the Nagle algorithm disabled

X – servers are a good example

TCP_NODELAY option can be used (in setsockopt) to disable the N.A.

See figure 19.8 for an illustration of:

A dumb terminal emulator sending meta-keys as single byte packets.

Detection and response to a lost packet

Ack immediately with next byte expected.

Repacketization

The three byte unit that was lost gets resent as all four bytes

Window size advertisements (Section 19.5)

TCPs reserves buffer space for a connection.

If the buffers fill with data that has not been read by the application,

TCP *will* ack the data when received (not when passed to the APP),
and may advertise a smaller window reflecting the available buffer space.

This practice is OK and should not be confused with "shrinking" the window which is bad.

Shrinking the window is defined as

expressing willingness to accept a particular byte

and then

expressing unwillingness to accept that same byte

So... if the ACK number increases by at least as much as the window size decreases
you haven't shrunk the window.