

Delayed Acknowledge

- Delayed ACKs:
 - 1. Normally TCP does not send an ACK the instant it receives data. Instead, it delays the ACK, hoping to have more data going in the same direction as the ACK, so the ACK can be sent along with the data. This is sometimes called having the ACK piggyback with the data.
 - 2. Most implementations use a 200-ms delay—that is, TCP will delay an ACK up to 200ms to see if there is data to send with the ACK.
 - 3. Every 200ms relative to when the kernel was bootstrapped, NOT relative to when the data received.

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Nagle Algorithm

- Tinygram: 41-bytes packets:
 - 20 : IP header
 - 20 : TCP header
 - 1: data
- The effects of tinygrams:
 - normally not a problem on LANs
 - add to congestion on wide area networks

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Nagle Algorithm

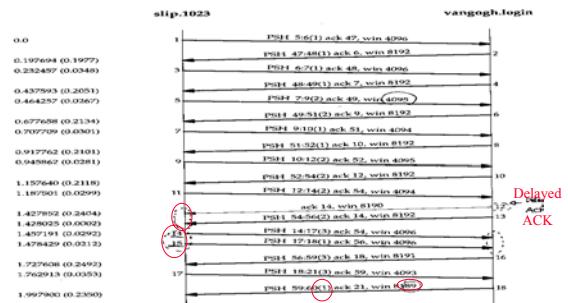
- A simple and elegant solution of tinygrams: Nagle algorithm:
 - When a TCP connection has outstanding data that has not yet been acknowledged, small segments cannot be sent until the outstanding data is acknowledged.
- The beauty of this algorithm: self-clocking
 - the faster the ACKs come back, the faster the data is sent

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Nagle Algorithm

- Examination from slip to vangogh.cs.berkeley:



Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Nagle Algorithm

- Things we notice:
 - the lack of delayed ACKs from slip to vangogh: on this example there is always data ready to send before the delayed ACK timer expires
 - the various amounts of data being sent from the left to the right: by the Nagle algorithm
 - the 14th, 15th segments responded the 12th, 13th segments, that's also obey the Nagle algorithm
 - the 12th segment is a delayed ACK, for without data
 - the 18th segment echoed only 1 byte: this was done by the TCP module in the kernel, and the data not up to the application of the server yet; the advertised window showed 8189 (not 8191) also indicated the same thing

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Disable the Nagle Algorithm

- There are times when the Nagle algorithm needs to be turned off:
 - the X Window System server: small messages (mouse movements) must be delivered without delay to provide real-time feedback for interactive users doing certain operations
 - terminal's special function keys

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

