Reliable Data Transport: Sliding Windows

6.02 Fall 2013 Lecture 23

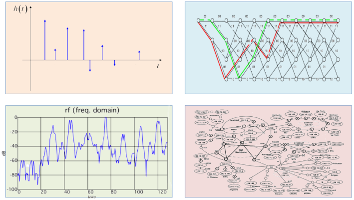

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

Exclusive!

A Brief History of the Internet

guest lecture by Prof. Hari Balakrishnan

Wenesday December 4, 2013, usual 6.02 lecture time

Today's Plan

- Round-Trip Time Statistics
  - average and linear deviation
  - smoothing filters: EWMA
  - timeout for "stop-and-wait"
- Sliding Window Protocol
  - general operation
  - throughput analysis

Round-Trip Time as a Random Variable

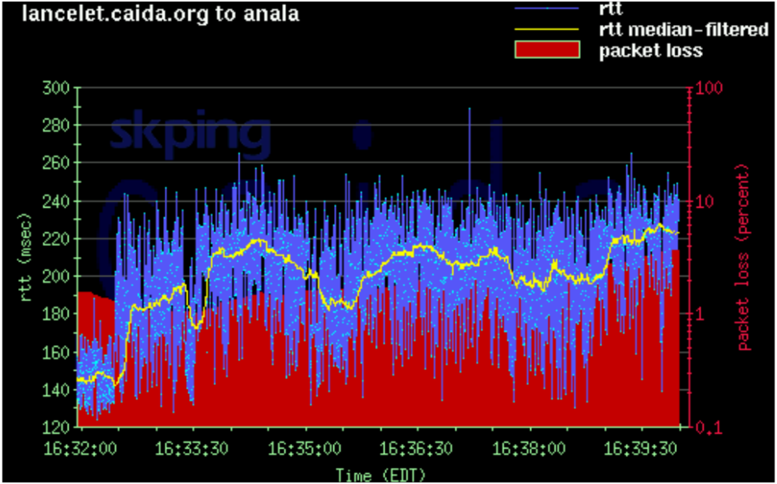

Deciding On The Timeout

Performance analysis of the stop-and-wait protocol, coupled with Markov inequality suggest using

$$T_o = \hat{T}_r + \frac{\hat{D}_r}{q}$$

where

- T_0 is the timeout
- $\hat{T}_r$ is an estimate of $\mathbf{E}[T_r]$ (average RTT)
- $\hat{D}_r$ is an estimate of $\mathbf{E}[|T_r - \hat{T}_r|]$ (linear deviation of RTT)
- q is a bound of spurious retransmission probability

"Static" Estimates

Given RTT samples sequence $T_r[0], T_r[1], T_r[2], \ldots$

$$\begin{aligned}\hat{T}_r[0] &= 0,\\ \hat{T}_r[n] &= \frac{T_r[0] + T_r[1] + \cdots + T_r[n-1]}{n} \quad (n \geq 1).\end{aligned}$$

Equivalently,

$$\begin{aligned}\hat{T}_r[0] &= 0,\\ \hat{T}_r[n] &= \left(1 - \frac{1}{n}\right)\hat{T}_r[n-1] + \frac{1}{n}T_r[n-1] \quad (n \geq 1)\end{aligned}$$

The impact of a single RTT sample diminishes with time!

Replacing $\frac{1}{n}$ with a fixed $\alpha \in (0,1)$ reduces the averaging horizon from "infinite" to "finite" (approximately $n_0 = 1/\alpha$ samples).

Comparing Averages: Reaction to Pure Change

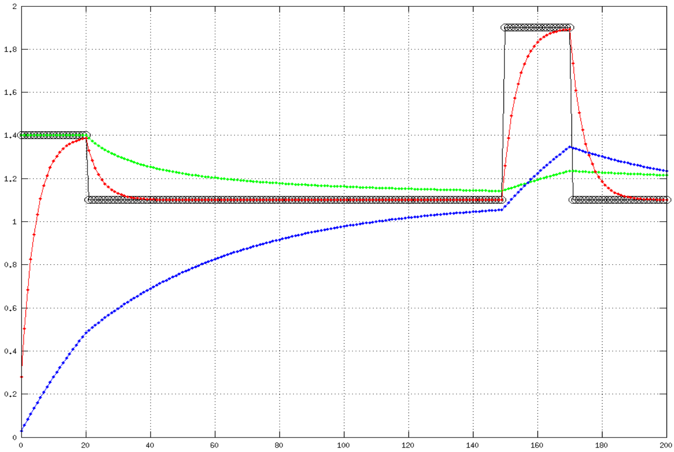

Comparing Averages: Smoothing Power

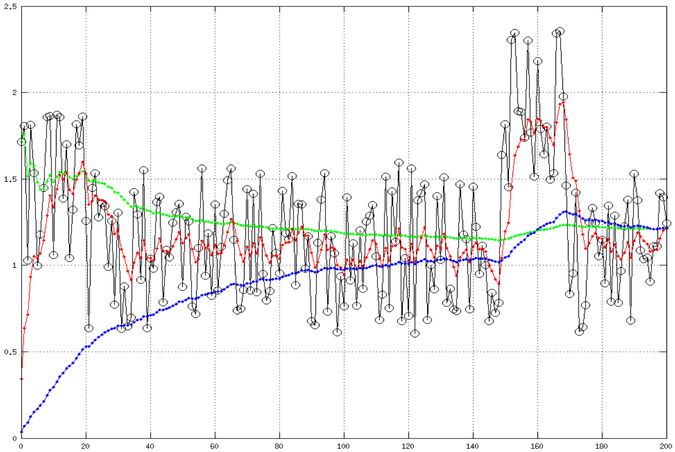

Using Exponentially Weighted Moving Averages (EWMA)

To adapt to changes in the disribution of T_r:

$$
\begin{aligned}
T_o &= \hat{T}_r + q^{-1}\hat{D}_r \\
\hat{T}_r^+ &= (1-\alpha)\hat{T}_r + \alpha T_r \\
\hat{D}_r^+ &= (1-\beta)\hat{D}_r + \beta|T_r - \hat{T}_r|
\end{aligned}
$$

$0 < \alpha, \beta, q < 1$ (TCP uses $\alpha = 1/8$, $\beta = q = 1/4$)

$1/\alpha$, $1/\beta$ are the time constants
q is a bound on the probability of spurious retransmission

EWMA: What Is Going On

LTI model: $y[n+1] = (1-\alpha)y[n] + \alpha x[n]$
Step response: $y_s[n] = 1 - (1-\alpha)^n \ (n \geq 0)$
Frequency response: $H(\Omega) = \alpha/(e^{j\Omega} - 1 + \alpha)$

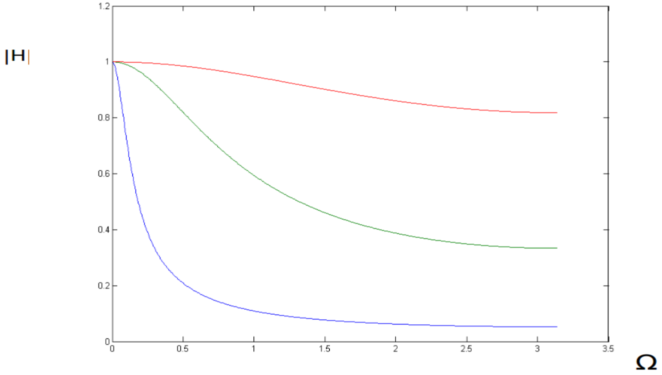

Which color corresponds to the smallest α?

Example: Stop-and-Wait Performance With Constant T_r

Recall: $T = T_r + \frac{pT_o}{1-p}$

If

- bottleneck link can support 100 packets/sec
- $T_r = 100ms$
- $T_o = T_r + \epsilon$

then, using stop-and-wait, the maximum throughput is at most only 10 packets/sec.

Only 10 percent utilization:

we need a better reliable transport protocol!

Sliding Window Protocol

- Allow up to W unacknowledged packets
- Wait T_o to re-transmit long-unacknowledged packets
- Send out next untransmitted packet whenever window permits (and no re-transmission is scheduled)
- Keep a buffer of out-of-order pakets at the receiver

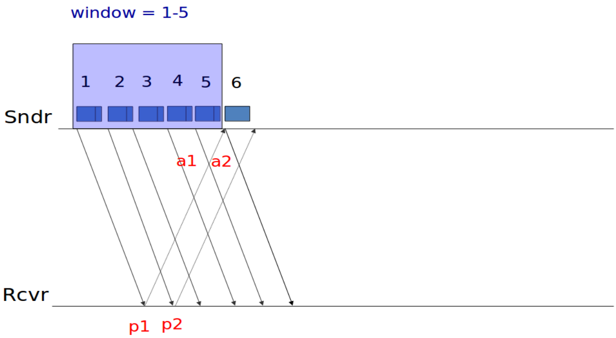

Sliding Window Example ($W = 4,\ T_o = 6$)

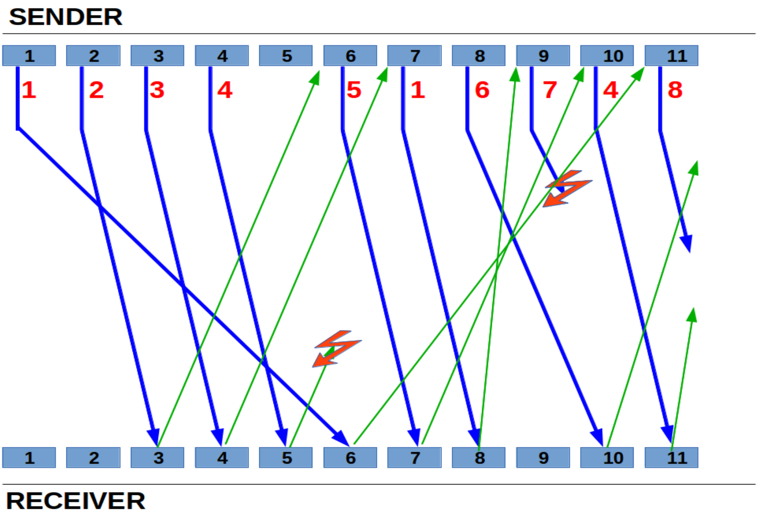

Window Size and Transmission Rate (No Packet Loss)

Size the window for non-stop transmission over single round-trip:

$$\left.\begin{array}{r}\text{bottleneck 100 packets/sec}\\ \text{round-trip time 0.05 sec}\end{array}\right\} \Rightarrow W = 100 \cdot 0.05 = 5$$

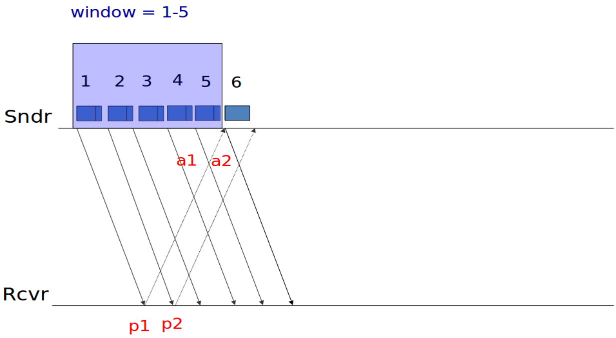

Back To The Little's Law

$W = B \cdot T_r$

(packets in transport) = (throughput rate) × (time in transport)

More accurately: $W = B \cdot T_{r\,\mathrm{min}} + Q$ (Q: no. packets in queues)

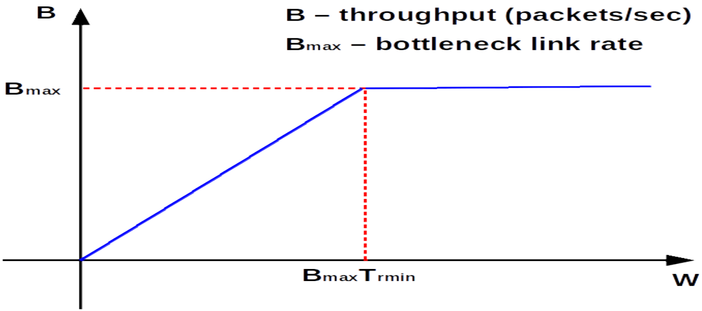

Example

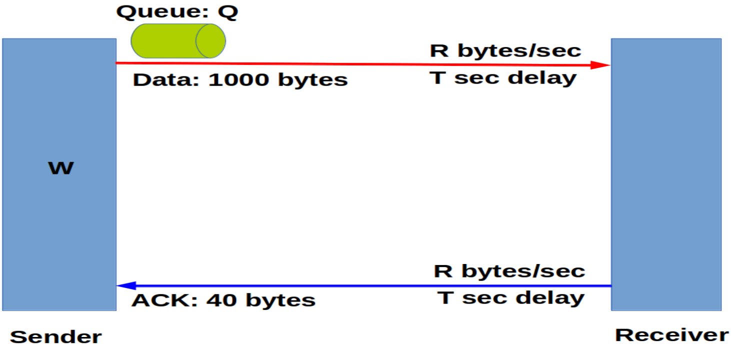

- no packet loss (except in queue overflows)
- throughput (packets/sec) = ?
- $W = B \cdot (2T + \frac{1000}{R} + \frac{40}{R}) + Q, \qquad 1000 \cdot B \leq R, \qquad Q \to \min$

Example: Case 1

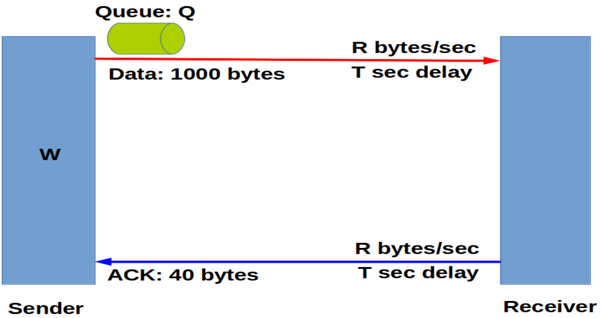

$$W = 10, T = 0.01, R = 10^6$$

$$10 = B \cdot 0.021 + Q, \qquad 1000 \cdot B \leq 10^6 \qquad Q \to \min$$

$$B = 476, \; Q = 0$$

Example: Case 2

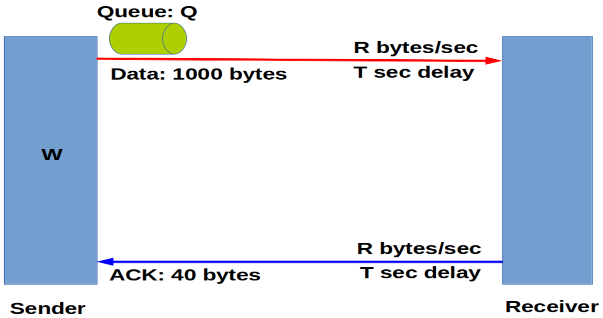

$W = 50, T = 0.01, R = 10^6$

$50 = B \cdot 0.021 + Q, \qquad 1000 \cdot B \leq 10^6 \qquad Q \to \min$

$B = 1000, \; Q = 29$

Summary

- reliability via redundancy (careful retransmissions)
- timeout selection is critical for performance
- round-trip time statistics are essential
- time horizon adjustment in EWMA
- performance improvement with sliding windows
- bandwidth-delay product in throughput analysis