Reliable Data Transport

6.02 Fall 2013 Lecture 22

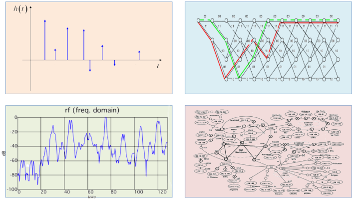

INTRODUCTION TO EECS II

DIGITAL COMMUNICATION SYSTEMS

Today's Plan

- Tools:
  - sequence number
  - time stamp
  - acknowledgements
  - retransmission
- Basic protocols:
  - stop-and-wait
  - sliding window
- Throughput analysis
- Round-trip time statistics

Reliable Data Transport: Objectives and Challenges

- sender sends a sequence of packets (DATA) to receiver
- receiver sends acknowledgements (ACK) of receipt
- packets/acknowledgements may be lost arbitrarily
- packets/acknowledgements may be reordered arbitrarily
- packet/acknowledgements delays are variable (queueing)
- packets may even be duplicated
- sender wants to know which data was received
- receiver wants to assemble a correct sequence of all packets

Basic Plan

- Transmitter
  - receive packets in sequence for transmission
  - attach sequence numbers and time stamps before transmission
  - keep "transmission time, packet sequence number" record
  - keep "packet acknowledged" and "round-trip time" records
  - periodically re-transmit long-unacknowledged packets
- Receiver:
  - send acknowledgements for received packet sequence numbers
  - delier packets in sequence

The "Stop-and-Wait" Protocol

- transmit next packet only after receiving ACK (for the previous one)
- re-transmit unacknowledged packet in time T_o (timeout) (after sending out its last copy)
- adjust T_o using round-trip time and packet loss statistics

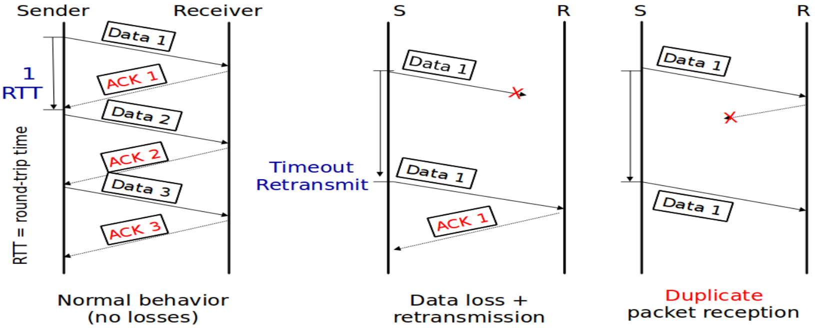

Choosing timeout T_o

- Too small:
  unnecessary re-transmissions of acknowledged packets
- Too large:
  poor throughput waiting for lost packets' acknowledgements

- Analysis depends on network properties: too tricky
- Trying to get insight based on simplistic models

Round-Trip Time as a Random Variable

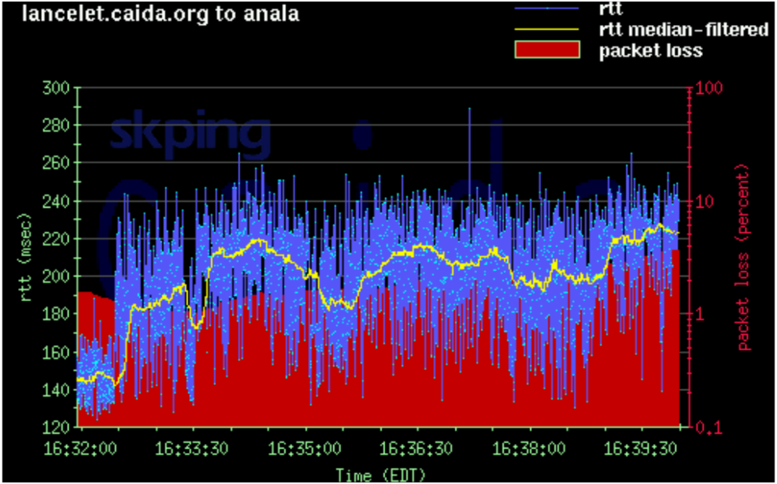

Round-Trip Time as a Random Variable

Figure 1: Round-trip time during a TCP download on the Verizon LTE network in Cambridge, Mass., Oct. 14, 2011 at 3 p.m.

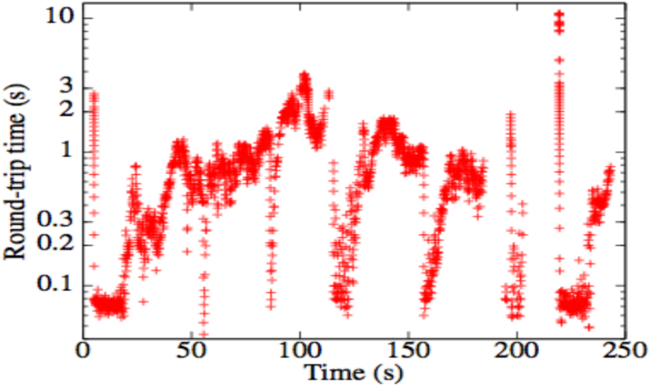

http://nms.csail.mit.edu/papers/index.php?detail=208

Two Simplistic Performance Numbers

(a) Assume round-trip time T_r is a random variable, but nothing is lost. Then $\mathbf{P}(\text{re-transmission}) = \mathbf{P}(T_r > T_o)$ should be small enough. Since $\mathbf{P}(T_r > T_o)$ decreases as T_o increases, T_o must be large enough

(b) Assume round-trip time $T_r < T_o$ is fixed, but packets are lost with probability p, causing time T_a to receiving acknowledgement to be random. Then average time per packet $T = \mathbf{E}[T_s]$ should be small enough. Since $T = (1-p)T_r + p(T + T_o)$ implies $T = T_r + \frac{pT_o}{1-p}$, T_o must be small enough

- **Use (a) as a guide to get T_o from the statistics of T_r**
- **Otherwise, use (b) to keep T_o as small as possible**

Markov Inequality as a Guide

How do you find a threshold γ such that $\mathbf{P}(X \geq \gamma) \leq q$ when we do not know much about the distribution of X?

If $X \geq 0$ and $\gamma \geq 0$ then $\gamma \mathbf{P}(X \geq \gamma) \leq \mathbf{E}[X]$

- Apply with $X = |T_r - \mathbf{E}[T_r]|^2$: **Chebyshev inequality**
- Apply with $X = |T_r - \mathbf{E}[T_r]|$: common in traffic control

$$T_o = \mathbf{E}[T_r] + \frac{\mathbf{E}[|T_r - \mathbf{E}[T_r]|]}{q}.$$

How do we get $\mathbf{E}[T_r]$ and $\mathbf{E}[|T_r - \mathbf{E}[T_r]|]$?

Estimating $\hat{T}_r = \mathbf{E}[T_r]$ and $\hat{V}_r = \mathbf{E}[|T_r - \mathbf{E}[T_r]|]$

The "time-invariant" case:

$$\hat{T}_r = \frac{T_{r1} + T_{r2} + \cdots + T_{rm}}{m}$$

$$\hat{V}_r = \frac{|T_{r1} - \hat{T}_r| + |T_{r2} - \hat{T}_r| + \cdots + |T_{rm} - \hat{T}_r|}{m}$$

with m round-trip measurements $T_{r1}, \ldots, T_{rm}$

When m is large, this becomes very slow to adapt to changes in $\hat{T}_r = \mathbf{E}[T_r]$ and $\hat{V}_r = \mathbf{E}[|T_r - \mathbf{E}[T_r]|]$

Exponentially Weighted Moving Average

To adapt to changes in T_r distribution:

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

$0 < \alpha, \beta < 1$ (TCP uses $\alpha = 1/8$, $\beta = 1/4$)
previous slide is the "limit as $\alpha, \beta \to 0$

Meaning: $1/\alpha$, $1/\beta$ are the time constants of the network

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: $H(\Omega) = \dfrac{\alpha}{e^{j\Omega} - 1 + \alpha}$

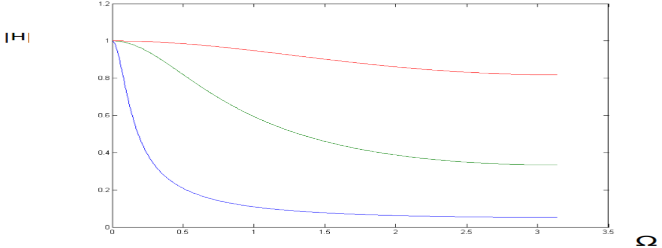

Which color corresponds to the smallest α?

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, $T_r = 100ms$, and $T_o = T_r + \epsilon$ 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
- Packet k acknowledged – send out packet $k + W$
- Keep a buffer of out-of-order pakets at the receiver