





### Solution: nonce



- Label request and ack with unique identifier that is never re-used

### Engineering a nonce

- Use sequence numbers
- Challenges:
  - Wrap around?
  - Failures?




### Timer value

- Fixed is bad. RTT changes depending on congestion
  - Pick a value that's too big, wait too long to retransmit a packet
  - Pick a value too small, generates a duplicate (retransmitted packet).
- Adapt the estimate of RTT to adaptive timeout

### RTT Measurements

(collected by Caida)


lancelet.caida.org to anala



### Adaptive Timeout: Exponential weighted moving averages


- Samples  $S_1, S_2, S_3, \dots$
- Algorithm
  - EstimatedRTT =  $T_0$
  - EstimatedRTT =  $\alpha S + (1 - \alpha)$  EstimatedRTT
  - where  $0 \leq \alpha \leq 1$
- What values should one pick for  $\alpha$  and  $T_0$ ?
  - Adaptive timeout is also hard

### At Most Once Challenges




- Server shouldn't process req 1
- Server should send result preferably

## Idea: remember sequence number




- Server remembers also last few responses

## Problem: failures




- Performed request 1 twice!
- How to maintain the last nonce per sender (tombstone)?
  - Write to non-volatile storage?
  - Move the problem? (e.g., different port number)
  - Make probability of mistake small?
- How about exactly once? (Need transactions)

## How fast should the sender sends?




- Waiting for acks is too slow
- Throughput is one packet/RTT
  - Say packet is 500 bytes
  - RTT 100ms
  - à Throughput = 40Kb/s, **Awful!**
- Overlap pkt transmission

## Send a window of packets



- Assume the receiver is the bottleneck
  - Maybe because the receiver is a slow machine
- Receiver needs to tell the sender when and how much it can send
- The window advances once all previous packets are acked à **too slow**

## Sliding Window




- Senders advances the window whenever it receives an ack à **sliding window**
- But what is the right value for the window?

## The Right Window Size


- Assume server is bottleneck
  - Goal: make idle time on server zero
  - Assume: server rate is B bytes/s
  - Window size = B x RTT
  - Danger: sequence number wrap around
- What if network is bottleneck?
  - Many senders?
  - Sharing?
  - Next lecture

## “Negative” ACK




- Minimize reliance on timer
- Add sequence numbers to packets
- Send a Nack when the receiver finds a hole in the sequence numbers
- Difficulties
  - Reordering
  - Cannot eliminate acks, because we need to ack the last packet

## E2E layer in Internet




The 4-layer Internet model

## UDP

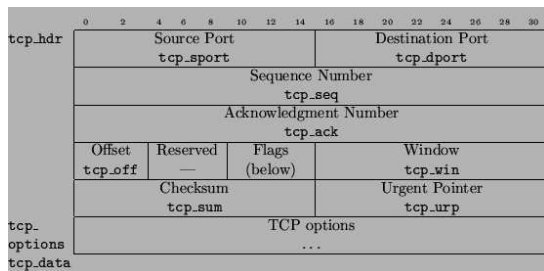


## Transmission Control Protocol (TCP)

- Connection-oriented
- Delivers bytes at-most-once
- Bidirectional
  - ACKs are piggybacked



## TCP header



## Closing a TCP connection

