

TCP Congestion Window Management

14-Nov-07

1

Congestion Window Management

- Slow start / rapid accelerate
- Dynamic window sizing on congestion
- Fast retransmit
- Fast recovery
- Limited transmit

2

Slow Start / Rapid Accelerate

$awnd = \text{MIN}[\text{credit}, cwnd]$

where

$awnd$ = allowed window in segments

$cwnd$ = congestion window in segments

$credit$ = amount of unused credit granted in most recent ack

$cwnd = 1$ for a new connection and increased by 1 for each ack received, up to a maximum; from maximum: increase by 1 per RTT

3




Figure 12.9 Effect of Slow Start

4

Dynamic Window Sizing on Congestion

- A lost segment indicates congestion
- Prudent to reset $cwnd = 1$ and begin slow start process
- May not be conservative enough: “easy to drive a network into saturation but hard for the net to recover” (Jacobson)
- Instead:
 - Set slow-start threshold $ssthresh = cwnd/2$
 - Use slow start from $cwnd=1$ till $cwnd=ssthresh$
 - For $cwnd > ssthresh$, increase only by one each Ack
 - This is called congestion avoidance

5




Figure 12.10 Slow Start and Congestion Avoidance

6

Fast Recovery

- When TCP retransmits a segment using Fast Retransmit, a segment was assumed lost
- Congestion avoidance measures are appropriate at this point
- E.g., slow-start/congestion avoidance procedure
- This may be unnecessarily conservative since multiple acks indicate segments are getting through (and out of the Net)
- Fast Recovery: retransmit lost segment, cut cwnd in half, proceed with linear increase of cwnd
- This avoids initial exponential slow-start

7

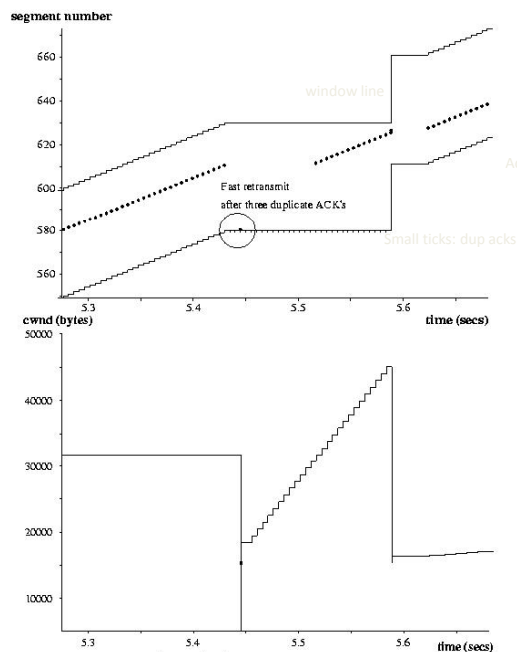


Figure 12.13 Fast Recovery Example

Upon 3rd dup ack:

- Set $ssthresh = cwnd/2$
- Retransmit segment
- Set $cwnd = ssthresh + 3$

Upon additional dup ack:

- $cwnd++$;
- If $pending < cwnd$, send another segment

Upon new ack: set $cwnd = ssthresh$.

8

Limited Transmit

- If congestion window at sender is small, fast retransmit may not get triggered, e.g., $cwnd = 3$
 1. Under what circumstances does sender have small congestion window?
 2. Is the problem common?
 3. If the problem is common, why not reduce number of duplicate acks needed to trigger retransmit?

9

Limited Transmit

- If congestion window at sender is small, fast retransmit may not get triggered, e.g., $cwnd = 3$
 1. Under what circumstances does sender have small congestion window?
 1. Limited amount of data to send
 2. Small limit on receive window (credit)
 3. Small bandwidth*delay (e.g. very low delay)
 2. Is the problem common?
 1. Yes, e.g. about 56% retransmit due to RTO expires, only 44% of them by fast retransmit
 3. If the problem is common, why not reduce number of duplicate acks needed to trigger retransmit?
 1. Packet reordering is not all that rare

10

Limited Transmit Algorithm

RFC 3042

Sender can transmit new segment when 3 conditions are met:

1. Two consecutive duplicate acks are received
2. Destination advertised window allows transmission of segment
3. Amount of outstanding data after sending is less than or equal to $cwnd + 2$

11