TCP Congestion Control [Jacobson88a]

(started Feb. 9; finished Feb. 16)

CSci551: Computer Networks SP2006 Thursday Section John Heidemann

6b_Jacobson88a: CSci551 SP2006 © John Heideman

Key ideas • [Jacboson88a]: - new algo to calculate RTO, RTT variance - congestion control for TCP - fast retransmission • problem: network was not stable -why? congestion

7

13

18

• oscillations in congetion

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

Agenda

- · connection setup and teardown
- flow control
- congestion control theory
- congestion control practice (in TCP) - slow start
 - congestion avoidance
- loss recovery
- putting it together
- security
- .Jperformance6 © John Heidemann

TCP Congestion Control

- · three mechanisms:
 - slow start: goal: come up to speed quickly and determine what bitrate the network can support
 - congestion avoidance: goal: TCP at "equilibria", goal stability (don't be aggressive), but keep adding a little bit of tfc to see if there's more room in the network, also to eventually converge on a "fair" allocation
 - congestion recovery: goal: backoff after congestion (indicated by loss or ECN bit) and let network recover; go to slow start (cwnd=1) and halve ssthresh

· interacts very closely with loss repair:

- good retransmit timeout (RTO) estimation
- fast retransmit and recovery – why?
 - main signal from the net about congestion or overutilization is loss
 - · we must exepct loss as part of regular TCP operation

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

TCP Congestion Principals

- underlying principle: packet conservation
- at equilibrium, inject packet into network only when one is removed
 - basis for stability of physical systems
- components:
 - how to get there: slow start
 - how to stay there: congestion avoidance
 - if we overshoot: congestion recovery

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

17

1

TCP Congestion Control Mechanisms

- new congestion window cwnd
 - what the *network* can handle
 - vs. flow control window (wnd): what the other end can handle
- sender limits tx
 - min (wnd, cwnd)
 - (and also considers outstanding data)

6b_Jacobson88a: CSci551 SP2006 © John Heidemann















- upon receiving ACK
 - Increase cwnd by 1/cwnd
 - This is *additive* increase (over 1 RTT it adds up to increasing by 1 segment)
- why not multiplicative increase?
 - don't want to overly congest the network
 assume ssthresh is good estimate of target rate, so increase slowly after that
 - rate, so mercase slowly after
- 6b_Jacobson88a: CSci551 SP2006 © John Heidemann





33

TCP Loss Recovery

- timeout and retransmit
- fast retransmit
- fast recovery
- New-Reno partial ACKs
- SACK

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

37





6b_Jacobson88a: CSci551 SP2006 © John Heidemann

Initial Round-trip Estimator

Round trip times exponentially averaged:

- New RTT = α (old RTT) + (1 α) (new sample)
- Recommended value for α: 0.8 0.9
- Retransmit timer set to β RTT, where $\beta = 2$

43

45

48

· Every RTO expiration, increase it multiplicatively

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

Retransmission Ambiguity what does the ACK indicate? A в A в Original transmission Original transmission RTO RTO ACK Sample retransmission Sample cretransmission RTT RTT 6b_Jacobson88a: CSci551 SP2006 © John Heidemann



- Keep backed off time-out for next packet
- Reuse RTT estimate only after one successful transmission



• Key observation:

- Using β RTT for timeout doesn't work (not adaptive enough with fixed β : at high loads, variance is high)
- Solution:
 - If D denotes mean variation (measured)
 - Timeout = RTT + 4D
 - $-\beta$ is now *adaptive*

(and can do it with integer math in few l.o.c.)

```
6b_Jacobson88a: CSci551 SP2006 © John Heidemann
```



42



• timeout and retransmit

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

- fast retransmit
- fast recovery
- New-Reno partial ACKs

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

• SACK

Why is recovery important?

- goal is high utilization
- want recover quickly because slow start is a big penalty
 - to recover quickly, need to try and keep the ACK clock going
 - so need to keep some packets in flight
- and have to recover the packet
 want to quickly determine what was lost and resend it
- 6b_Jacobson88a: CSci551 SP2006 © John Heidemann

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

Fallback Recovery: Timeout and Retransmit

- timeout after RTO
- pros

• how

- simple (no help from routers)
- doesn't require extra net traffic (like NACKs)
- can depend on this
- cons
 - need good RTO estimation
 - need timers (can be expensive)
 - can be slow (at most one pkt retx per RTT)

54

6b_Jacobson88a: CSci551 SP2006 © John Heidemann





Fast Recovery [see Fall96a]

- Problem: fast retx still forces slow-start, breaking the ACK clock
- Fast Recovery Solution: artificially *inflate* the *cwnd* as more dup ACKs come in
 - cut *cwnd*, but instead of slow start, do additive increase for each ACK
 - justification: each dup ACK represents a packet leaving the network, so we can increase *cwnd*
 - exit when out of dup ACKs

6b_Jacobson88a: CSci551 SP2006 © John Heidemann

57

51

55

























