Congestion Control

- The tool available to control the rate of sending is the Effective Window Size.
- What are the indications that congestion is a problem in the network?
  - Timeouts indicating increasing RTT.
  - Source Quench ICMP messages.

Congestion Window $EWS = MIN(CWS, AWS)$

- If there is no sign of congestion, increase the CWS.
- If there appears to be congestion, decrease the CWS.

Additive Increase Multiplicative Decrease

- When an ACK for a full CWS is received, $CWS = CWS + C$.
- When a timeout occurs, let $CWS = CWS \times M$, where $M < 1$.
- Typical values are $C = 1\ MSS$ and $M = 0.5$
- For high speed senders, it would be possible to send an entire CWS and get back an ACK for all. However, to avoid having to keep track of ACK:
- The usual choice is the add $\frac{1}{CWS}\ MSS$ for each ACK.

AIMD Example

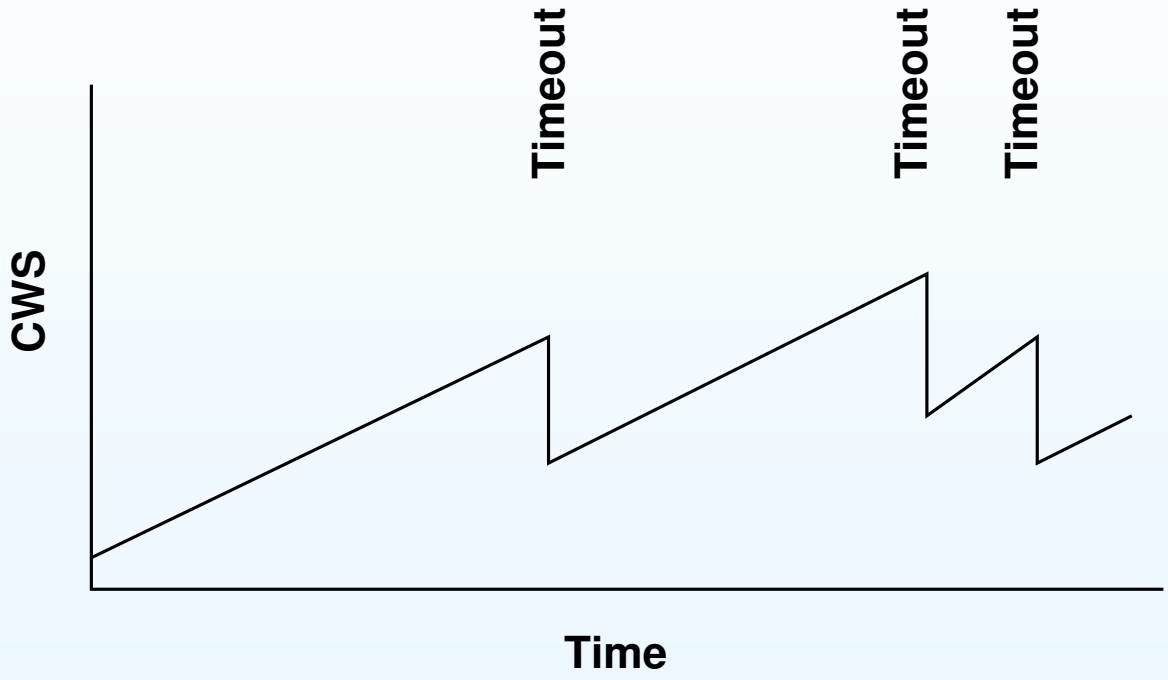

Additive Increase Multiplicative Decrease

Slow Start

AIMD is too slow, so a method to move from an initial CWS of 1 MSS is needed. It is called slow start because it is slower than just starting at AWS.

- For each CWS ack'ed, double CWS until a timeout occurs.
- Then set the cold start threshold at one-half of the CWS. and set the $CWS = 1\ MSS$ and let it grow exponentially to the cold start threshold, and then use AIMD.

Slow Start Example

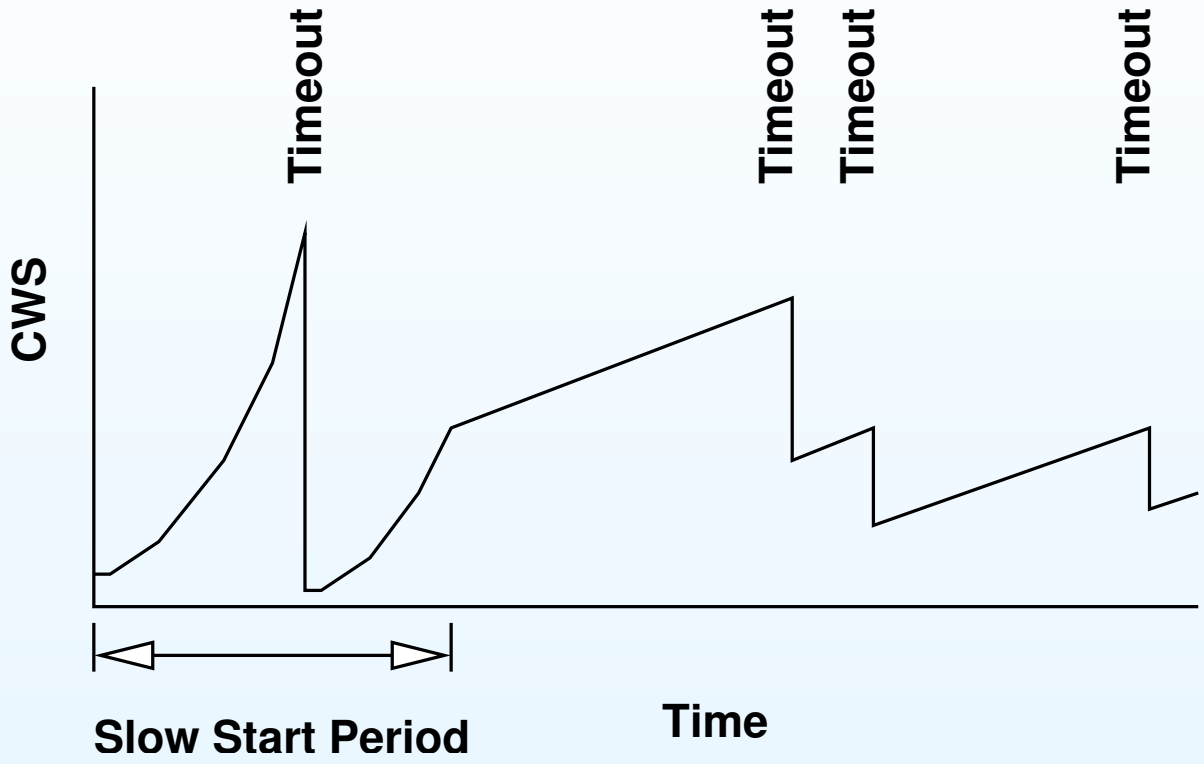

Timeout Types

- Course-grained or dead-channel timeouts are times when the senders EWS goes to zero before a timeout. This usually means major congestion problems. Do a restart with $CWS = 1\,MSS$ and enter slow start with the slow start threshold set to CWS/2 or 2 segments, whichever is greater.
- Fine-grained or live-channel timeouts are times when there is a timeout but the EWS is greater than zero. Reduce the CWS to one-half and stay in AIMD mode.

The TCP Algorithm

```
switch (event)
  ACK Received:
    if (mode == SLOWSTART)
      CWS = CWS + 1 MSS
      if (CWS > threshold) mode = AIMD
    else
      CWS = CWS + MSS/CWS
  Timeout:
     if (EWS == 0)   // Dead Channel??
       threshold = max (CWS/2, 2*MSS)
       CWS = 1; mode = SLOWSTART
     else
       CWS = CWS/2
       if (CWS < threshold)
         mode = SLOWSTART
       else
         mode = AIMD
```

Example

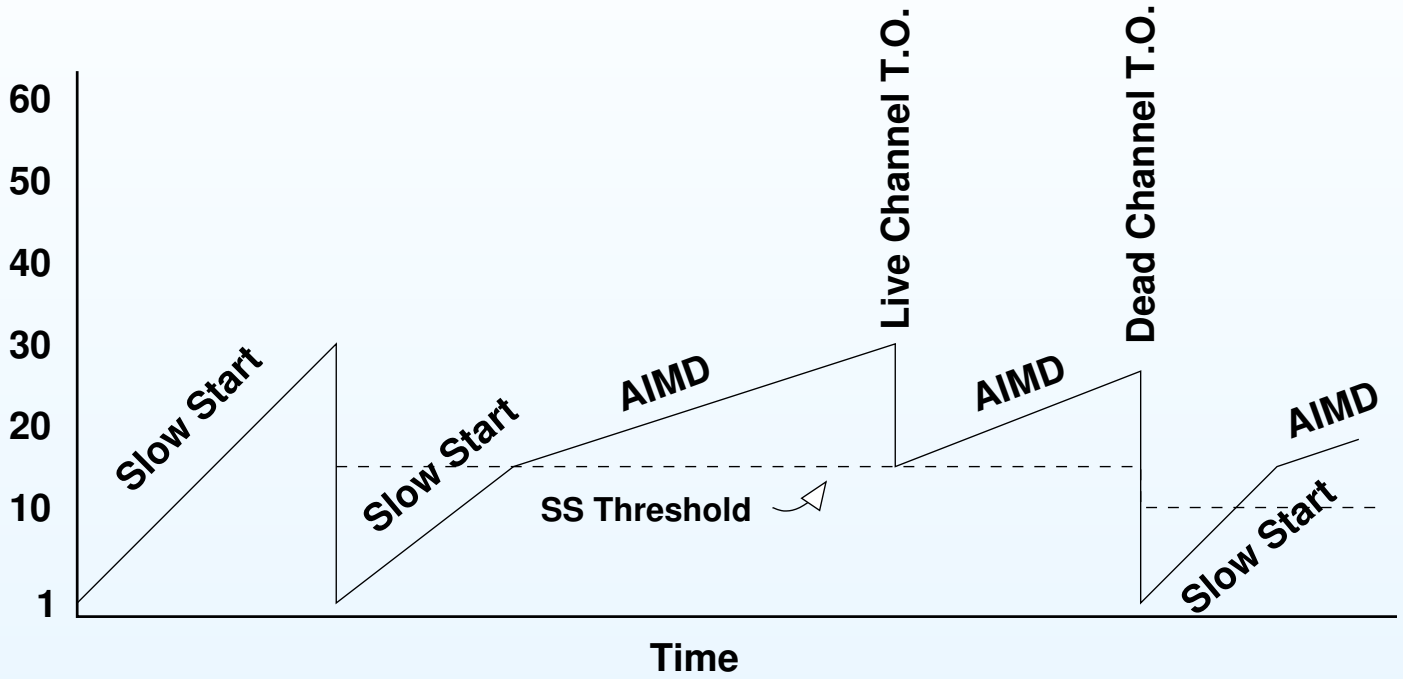
60
50
40
30
20
10
1
Slow Start
Slow Start
AIMD
SS Threshold
Live Channel T.O.
AIMD
Dead Channel T.O.
AIMD
Slow Start
Time

Fast Retransmit

If a receiver notices that it is receiving the same value in the Acknowledgement field, it might guess that segments are arriving, but one has not been received correctly.
Fast Retransmit transmits a duplicate of the missing segment after three duplicate ACK's.

This method eliminates about one-half of dead-channel timeouts.

But a Fast Retransmit could be an indication of a congestion problem, so CWS should be reduced, but by how much?
Fast Recovery reduces the CWS by one-half, but stays in AIMD mode.

Fast Retransmit Example