

**SACK TCP: The sender's congestion control  
algorithms for the implementation "sack1" in the**

**LBL's "ns" simulator.**

Sally Floyd

Lawrence Berkeley National Laboratory

Contributors to code: Kevin Fall, Sally Floyd, Steve McCanne  
(LBL),

and Janshid Mahdavi (PSC)

LA IETF

March 1996

## **The internet draft:**

“The congestion control algorithms present in the de facto standard TCP implementations **MUST** be preserved.”

- Accommodates out-of-order delivery.
- Congestion window algorithms.
- Use of time-outs.


## **The implementation in the “ns” simulator:**

- Three dup acks required to trigger Fast Recovery.
- Reduce congestion window by half; don't Slow-start.
- Response to further dup acks.

Main difference from Reno: When multiple packets are lost from a single window of data.

# Two states: Regular and Fast Recovery

3 duplicate acknowledgements



acknowledgement for everything that was sent  
before Fast Recovery was entered  
(`"recover_"`)

## **On entering Fast Recovery:**

- Retransmit one packet.
- Cut the congestion window in half (“cwnd\_”).
- Estimate the number of packets in the pipe (“pipe\_”).

## **Behavior in Fast Recovery:**

- **When and how much to send:** whenever the number of packets in the pipe is less than the congestion window.
- **What to send:** Fill “holes”, one packet at a time, in sequence number order. If there are no holes, send new packets.
- **If a retransmitted packet is itself dropped, then slow-start.** (The current implementation in ns waits for a retransmit timer to detect the dropped packet.)

## **Behavior in Fast Recovery: receiving ack packets**

- Duplicate ACKs: Decrement “pipe\_”, call “send”.
- An ACK that ends Fast Recovery: Call “send”.
- An ACK that does not end Fast Recovery:  
Decrement “pipe\_” by two packets, once for the retransmitted packet, and once for the original packet (now presumed to have been dropped). Call “send”.

## **Behavior in Fast Recovery: details of sending data packets**

- Send if the number of packets in the pipe (“pipe\_”) is less than the congestion window (“cwnd\_”).
- Use the SACK scoreboard to determine which packet to send.
- Increment “pipe\_”.



## **Details:**

- MaxBurst parameter
- Overhead parameter - just for the simulator.

**Ns simulator available from:**

<http://www-nrg.ee.lbl.gov/ns>

**These viewgraphs available from:**

<ftp://ftp.ee.lbl.gov/talks/sacks.ps>

**Papers available from:**

<http://www-nrg.ee.lbl.gov>