Summary of Post-1987 TCP Specification[1] (5 March 1998)

Let

RTT_i	=	estimated round trip time just after arrival of acknowledgement of the i-th segment sent without retransmission (initially 0 [p. 213 in Comer Vol. II])
DEV_i	=	estimated variance in RTT_i (initially 0 in Comer, 1.5 in Stevens)
RTO_i	=	retransmission time value just after arrival of acknowledgement of the i-th segment sent without retransmission; TCP uses one timer per segment sent (Comer Vol. II 12.11); initially 0.5 sec. in Comer Vol. II
SSTHRESH	=	*current* value of threshold to switch from slow start to additive increase (initially 2^{16} bytes in Comer Vol. II 11.20.)
CWND	=	*current* value of congestion window (initally 1)
MSS	=	max segment size; either MTU−40 if connection endpoints are directly connected; otherwise 536 (from RFC 1122). MTU is max transmission unit, or data size that a layer 2 frame can carry. For example, for a TCP connection between two computers attached to the same Ethernet network, MSS=1460 because MTU=1500 for Ethernet and 40 is size of TCP and IP headers in bytes (20 bytes for each).

1. *Round trip variance estimation* (Jacobson Appendix A, Comer Vol. I, 12.1, Comer Vol. II 14.9, and Stevens Section 8.4)[2]:

1a.	RTT_sample	=	(time when ack for S arrived) - (time when S was sent)	
1b.	DIFF	=	RTT_sample - RTT_i	
1c.	RTT_{i+1}	=	$RTT_i + \delta * DIFF$	(Comer Vol. 1 wrongly uses "-" not "+")
1d.	DEV_{i+1}	=	$DEV_i + \delta * (\|DIFF\| - DEV_i)$	(Typically $\delta = 1/8$)
1e.	RTO_{i+1}	=	$RTT_{i+1} + 2DEV_{i+1}$	(Comer Vol II p. 276 code is incorrect)

Notes:

- Above formulas are used only upon reception of an ack for segment that was not retransmitted
- Acks that are received out of order are ignored, so they are not used in the above formulas.

2. *Exponential retransmit timer backoff* (Comer Vol. I 12.18 and Vol. II 14.4) -- When retransmitting a segment, schedule another timer expiration at this time:

timer expiration = $2^{count} * RTO$, where *RTO* is *current* retransmission timer value and *count* is number of times segment has been retransmitted.

Note: In Comer's implementation, when the timer expires for a segment, the segment is retransmitted *regardless* of the value of CWND. Even if CWND=1, all segments whose timers pop are retransmitted.

3. *Slow start* (Comer, p. 194) coupled with *multiplicative decrease/additive increase congestion control* (see pp. 320-321 and Appendix B of Jacobson):

3a. Applied when retransmitting a segment:

SSTHRESH	=	max(CWND/2,1)	
CWND	=	1	(See Comer II 14.4.2.)

---

[1] References are to

- Douglas Comer, *Internetworking with TCP/IP*. Vol I is 2nd ed. (1991); Vol II is 1991 edition.
- Richard Stephens, *UNIX Network Programming*, 1st. edition.

[2] The entire computation is scaled by 1/d to perform integer arithmetic only; see Jacobson App. A.2 and Comer Vol. II pp. 276-278.

3b. Upon arrival of first acknowledgement of a segment sent without retransmission (Comer II, 14.8):

```
if      (CWND < SSTHRESH)
then    CWND = CWND + 1              /* do slow start */
else    CWND = CWND + 1/CWND         /* do additive increase */
endif
```

Note: Sometimes there is not a one-to-one correspondence between segments sent and received. For example, some TCP implementations acknowledge once for every two segments received. Or an acknowledgement could be lost. In these cases, letting $n \geq 0$ denote the number of previously unacknowledged segments that are acknowledged by an arriving segment, replace "1" by "n" in formula 3b above.

4. Always holds true (See Comer Vol. II 12.11 and 14.8.):

Sender window size = min(receiver advertisement, CWND*MSS)