

TCP/IP 通訊協定及應用

Spring 2002
中央大學 吳曉光博士
<http://wmlab.csie.ncu.edu.tw/course/tcp>



Wireless & Multimedia Network Laboratory™



Chapter 21: TCP Timeout and Retransmission

Wireless & Multimedia Network Laboratory™



Introduction

- Two examples of timeout and retransmission had already seen:
 - 1. In the ICMP port unreachable (Section 6.5)
 - 2. In the ARP example to a nonexistent host (Section 4.5)
- TCP manage four different timers for each connection:
 - 1. A retransmission timer (this chapter)
 - 2. A persist timer (chapter 22)
 - 3. A keepalive timer (chapter 23)
 - 4. A 2MSL timer (section 18.6)
- Simple Timeout and Retransmission Example:
 - Line 4 is the transmission of "hello,world" and line 5 is its ACK.
 - Line 6 shows "and hi". Line 7-18 are 12 retransmissions
 - Line 19 is the TCP finally gives up and sends a reset

Wireless & Multimedia Network Laboratory™



Simple Timeout and Retransmission

```
data, and wait when avr4 discards
bdi: % s:inet avr4 discard
Trying 140.252.13.34...
Connected to avr4.
Escape character is '^['.
hello, world
Connection closed by foreign host.

send this line normally
disconnect cable before sending this line
output when TCP gives up after 9 minutes
```

Figure 21.1 shows the tcpdump output. (We have removed all the type-of-service information that is set by bdi.)

```
1 0.0 bdi:1029 > avr4:discard: S 1747921409:1747921409(0)
2 0.004811 ( 0.0048) avr4:discard > bdi:1029: S 3416685569:3416685569(0)
3 0.006441 ( 0.0016) bdi:1029 > avr4:discard: . ack 1 win 4096
4 6.102296 ( 4.0958) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
5 6.208410 ( 0.1057) avr4:discard > bdi:1029: . ack 15 win 4096
6 24.480158 (18.2207) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
7 25.493733 ( 1.0136) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
8 28.493795 ( 3.0002) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
9 34.493971 ( 6.0002) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
10 46.484427 (11.9905) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
11 70.485105 (24.0007) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
12 118.486408 (48.0023) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
13 182.488144 (64.0018) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
14 246.489921 (64.0018) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
15 310.491678 (64.0018) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
16 374.493431 (64.0018) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
17 438.495196 (64.0018) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
18 502.486941 (63.9917) bdi:1029 > avr4:discard: P 15:23(8) ack 1 win 4096
19 566.488478 (64.0015) bdi:1029 > avr4:discard: R 23:23(0) ack 1 win 4096
```

Figure 21.1 Simple example of TCP's timeout and retransmission.

Wireless & Multimedia Network Laboratory™



Round-Trip Time Measurement

- Two methods of RTO calculate:
 - The original TCP specification method
 - $R \leftarrow \alpha R + (1 - \alpha) M$
 - $RTO = R / \beta$
 - R => smoothed RTT estimator, α is smoothing factor = 0.9
 - β is delay variance factor = 2
 - The Jacobson method
 - $Err = M - A$
 - $A \leftarrow A + g Err$
 - $D \leftarrow D + h(|Err| - D)$
 - $RTO = A + 4D$
 - A is the smoothed RTT average, D is mean deviation
 - The gain g is $1/8(0.125)$, h is 0.25

Wireless & Multimedia Network Laboratory™



Round-Trip Time Measurement

- Karn's Algorithm:
 - A packet is transmitted, a timeout occurs, the packet is retransmitted with the longer RTO, and an acknowledgment is received
 - Is the ACK for the first transmission or the second?
 - This is called the retransmission ambiguity problem
- An RTT Example:
 - sent 32768 bytes of data from slip to vangogh.cs.berkeley.edu
 - `slip% sock -D -i -n32 vangogh.cs.berkeley.edu. discard`
 - slip is connected to the 140.252.1 Ethernet by two SLIP links
 - MTU between slip and bsdi is 296
 - 32 1024-byte => 128 segment with 256 bytes of user data

Wireless & Multimedia Network Laboratory™



An RTT Example




Figure 21.3 Packet exchange and RTT measurement.

Wireless & Multimedia Network Laboratory™



An RTT Example

Round-Trip Time Measurements:

- RTT#1 is 1.061 seconds => 3 clock ticks
- RTT#2 is 0.808 seconds => 1 clock tick
- RTT#3 is 1.015 seconds => 2 clock ticks
- Segment 4,7,9 cannot be timed, since the timer is already being used by segment 3 and 6




Figure 21.3 RTT measurements and clock ticks.

Wireless & Multimedia Network Laboratory™



An RTT Example

- In this complete example 18 RTT samples were collected




Figure 21.4 Measured RTT and TCP's calculated RTO for example.

Wireless & Multimedia Network Laboratory™



An RTT Example

RTT Estimator Calculations:

- The initial $RTO = A + 2D \Rightarrow 0 + 2 \times 3 = 6$ seconds
- After 5.802 seconds $RTO = A + 4D \Rightarrow 0 + 4 \times 3 = 12$ seconds
- The ACK arrives 467 ms after the retransmission. The A and D are not updated because of retransmission ambiguity
- The ACK on line 4 is not timed since it is only an ACK

```

shows the first four lines from the tcpdump output file:
1 0.0 slip-1024 > vanguard-discard: 0 35448001:35448002 (0)
  win 4096, len 2560
2 5.802177 (5.8024) slip-1024 > vanguard-discard: 0 35448001:35448001 (0)
  win 4096, len 2560
3 6.268395 (3.4670) vanguard-discard > slip-1024: 0 1365512703:1365512703 (0)
  seq 35448002, win 3130, len 5120
4 6.270796 (3.0014) slip-1024 > vanguard-discard: 0 seq 1, win 4096
  
```

Figure 21.5 Timeout and retransmission of initial SYN.

Wireless & Multimedia Network Laboratory™



An RTT Example

- RTO calculations
 - first segment arrives => $RTO = 6$ seconds
 - second segment arrives => $RTO = 6.3125$ seconds
 - Fixed-point calculations that are actually used => $RTO = 6$ seconds (not 6.3125)
- Slow Start
 - See the slow start algorithm in Section 20.6
- Congestion Example:
 - Retransmission will appear as motion down and to the right
 - total time was 45 sec, 35 sec for send data segments only, first data segment was sent until 6.3 sec, final took 4.0 sec to receive ACKs

Wireless & Multimedia Network Laboratory™



Congestion Example




Figure 21.6 Sending of 32768 bytes of data from slip to vanguard.

Wireless & Multimedia Network Laboratory™



Congetion Example




Figure 16.7 Ticket exchange for synchronization around the shared resource

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Congetion Example

- ◆ The Jacobson's fast retransmit algorithm:
 - It is followed by his fast recovery algorithm. The third of the duplicate ACKs was received that forces to retransmit
 - Berkeley-derived implementation when the third one is received, assume that a segment has been lost and retransmit only one segment
- ◆ When the missing data arrives(segment 63):
 - The receiving TCP now has data bytes 6657-8960 in its buffer, and passes these 2304 bytes to the user process.
 - All 2304 bytes are acknowledged in segment 72

Wireless & Multimedia Network Laboratory™

**Wireless
Multimedia**

Congestion Avoidance Algorithm

- ◆ What's congestion avoidance?
 - It is a way to deal with lost packets
- ◆ Two indication of packet loss:
 - 1. A timeout occurring
 - 2. The receipt of duplicate ACKs
- ◆ Congestion avoidance algorithm operates:
 - 1. Initialization \Rightarrow cwnd is one segment, ssthresh is 65535 bytes
 - 2. TCP output never sends more than the minimum of cwnd and the receiver's advertised window
 - 3. When congestion occurs, one-half of the current window size is saved in ssthresh. If timeout, cwnd is set to one segment
 - 4. When new data is acknowledged by the other end, increase cwnd If cwnd is less than or equal to ssthresh, doing slow start

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Congestion Avoidance Algorithm




Figure 21.8 Visualization of slow start and congestion avoidance

Wireless & Multimedia Network Laboratory™

**Wireless
Multimedia**

Fast Retransmit and Fast Recovery Algorithms

- ◆ Fast retransmit algorithm:
 - If three or more duplicate ACKs are received in a row, indicate a segment has been lost, then retransmission the missing segment
- ◆ Fast recovery algorithm:
 - Next, congestion avoidance, but not slow start is performed
- ◆ Congestion Example (Continued)
 - In congestion avoidance:
 - ◆ $cwnd \leftarrow cwnd + (egsize \times ssegment) / cwnd + ssegment / 8$
 - By fast retransmit and fast recovery, we can send a new data segment when $cwnd > unacknowledged\ bytes$

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Congestion Example (Continued)

Segment (Figure 21.2)	Action		Variable	
	Sent	Received	Constant	Local (offset)
	95%		initialize	296
	95%		timed interrupt	296
	95%	SYN_ACK		296
	ACK			296
1	125(125)			296
2	207(13226)	ACK 297	slow start	312
3	933(30036)			312
4	768(36826)	ACK 313	slow start	368
5	1025(126126)			368
6	1203(126126)			368
7	1203(126126)	ACK 369	cong. avoid	605
8	1203(126126)			605
9	1203(126126)	ACK 1225	cong. avoid	951
10	1057(276126)			951
11		ACK 1281	cong. avoid	1089
12				1089

Segment (Figure 21.7)	Action		Variable	
	Sent	Received	Constant	Local (offset)
	98	ACK 6657	ACK of new data	2426
	9709(891226)			2426
50		ACK 6657	duplicate ACK #1	2426
61		ACK 6657	duplicate ACK #2	2426
62		ACK 6657	duplicate ACK #3	2426
63		6657(493226)	intertransmission	2426
64		duplicate ACK #4		2426
65		duplicate ACK #5		2426
66		duplicate ACK #6		2426
67	8961(9137226)			2426
68		duplicate ACK #7		2426
69	1127(474226)			2426
70		ACK 6657	duplicate ACK #8	3072
71	1473(474226)			3072
72		ACK 6657	ACK of new data	1206

Figure 21.13 Example of complex interactions (continued).

Figure 21.9 Example of congestion avoidance

Figure 23.31 Example of conversion avoidance (continued)

Wireless & Multimedia Network Laboratory™

Wireless
Multimedia

Congestion Example (Continued)




Figure 21.30 Value of send and sequence number while data is being transmitted.

Wireless & Multimedia Network Laboratory™



ICMP Errors

- Berkeley-based implementations handle ICMP errors as follows:
 - A received source quench cause the cwnd set to one segment to initiate slow start, but the ssthresh is not changed
 - A received host unreachable or network unreachable is effectively ignored, since these two errors are considered transient
- An Example
 - slip% sock aix echo
 - test line
 - test line
 - another line
 -
 - read error: No route to host

SLIP link is brought down here
type this line and retransmissions
SLIP link is reestablished here
after the last line, SLIP link is brought down
TCP finally gives up

Wireless & Multimedia Network Laboratory™



ICMP Errors

```

1 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000
2 0.012273 0.012273 0.012273 0.012273 0.012273 0.012273 0.012273 0.012273
3 0.010485 0.010485 0.010485 0.010485 0.010485 0.010485 0.010485 0.010485
4 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100
5 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100
6 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100
7 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100 174.759100
8 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200
9 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200
10 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200
11 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200 182.150200
12 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783
13 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783 212.149783
14 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774
15 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774
16 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774 252.146774
17 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246
18 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246
19 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246
20 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246 261.977246
21 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330
22 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330 458.155330
23 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604
24 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604
25 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604
26 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604 461.136604
27 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447
28 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447 479.136447
29 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616
30 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616 501.144616
31 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973
32 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973 1000.219973
33 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
34 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
35 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
36 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
37 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
38 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
39 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
40 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
41 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
42 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
43 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
44 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
45 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
46 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
47 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
48 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
49 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
50 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
51 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
52 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
53 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
54 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
55 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
56 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
57 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
58 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
59 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
60 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
61 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
62 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
63 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
64 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
65 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
66 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
67 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
68 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
69 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
70 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
71 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
72 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
73 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
74 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
75 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
76 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
77 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
78 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
79 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
80 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
81 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
82 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
83 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
84 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
85 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
86 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
87 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
88 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
89 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
90 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
91 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
92 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
93 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
94 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
95 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
96 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
97 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
98 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
99 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281
100 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281 1004.201281

```

Figure 21.12 TCP handling of received ICMP host unreachable error.

Wireless & Multimedia Network Laboratory™



Repackitization

- TCP is allowed to perform repackitization which can increase performance
- notice on bytes of line 3 and 6 in following illustration

```

1 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000 0.000000
2 0.140489 0.140489 0.140489 0.140489 0.140489 0.140489 0.140489 0.140489
3 26.407696 26.407696 26.407696 26.407696 26.407696 26.407696 26.407696 26.407696
4 27.639390 27.639390 27.639390 27.639390 27.639390 27.639390 27.639390 27.639390
5 30.639453 30.639453 30.639453 30.639453 30.639453 30.639453 30.639453 30.639453
6 36.639653 36.639653 36.639653 36.639653 36.639653 36.639653 36.639653 36.639653
7 48.640131 48.640131 48.640131 48.640131 48.640131 48.640131 48.640131 48.640131
8 72.640768 72.640768 72.640768 72.640768 72.640768 72.640768 72.640768 72.640768
9 72.719091 72.719091 72.719091 72.719091 72.719091 72.719091 72.719091 72.719091

```

Figure 21.13 Repackitization of data by TCP.

Wireless & Multimedia Network Laboratory™



Summary

- TCP calculates a smoothed RTT estimator and a smoothed mean deviation estimator. Then use these two estimators to calculate the next retransmission timeout value.
- We see many of TCP's algorithm in action:
 - slow start
 - congestion avoidance
 - fast retransmit
 - fast recovery
- We see effect various ICMP errors have on a TCP connection
- We see how TCP is allowed to repackitize its data

Wireless & Multimedia Network Laboratory™

