




15-441 Computer Networking

Lecture 18 – More TCP & Congestion Control

Good Ideas So Far...



- Flow control
 - Stop & wait
 - Parallel stop & wait
 - Sliding window (e.g., **advertised windows**)
- Loss recovery
 - Timeouts
 - Acknowledgment-driven recovery (selective repeat or cumulative acknowledgment)
- Congestion control
 - AIMD → fairness and efficiency
- **How does TCP actually implement these?**

10-21-2004

Lecture 18: TCP Details

3

Outline




- **THE SPOOKY PARTS** of TCP
 - If it doesn't scare you now... it will in the future!
- TCP connection setup/data transfer
 - **The Candy-exchange Protocol (TCP)**
- TCP reliability
 - **How to recover your DEAD packets**
- TCP congestion avoidance
 - **Avoiding the death-trap of overloaded routers**

10-21-2004


Lecture 18: TCP Details

3

Sequence Number Space



- Each byte in byte stream is numbered.
 - 32 bit value
 - Wraps around
 - Initial values selected at start up time
- TCP breaks up the byte stream into packets.
 - Packet size is limited to the Maximum Segment Size
- Each packet has a sequence number.
 - Indicates where it fits in the byte stream




10-21-2004

Lecture 18: TCP Details

4

TCP Connection Teardown Example



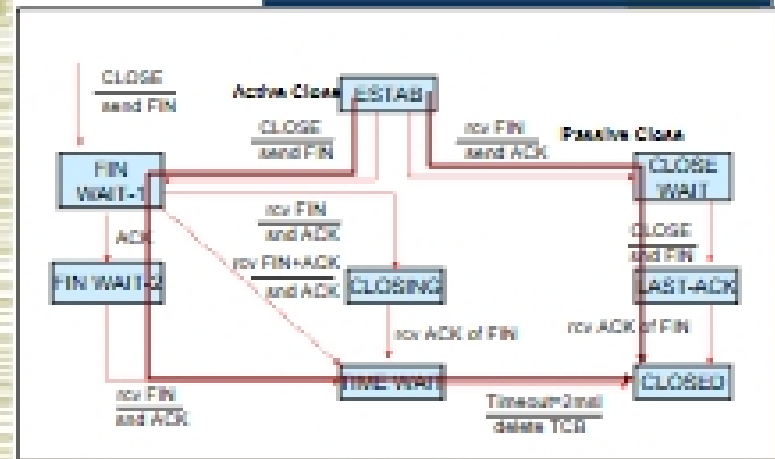
```
08:24:17.000000 IP 128.2.222.194 > 128.2.210.194: P
1480294582.1480294582 (0) seq 1480294582 win 0 (0)

08:24:17.001000 IP 128.2.210.194 > 128.2.222.194: P
1480294582.1480294582 (0) seq 1480294582 win 0 (0)


08:24:17.001800 IP 128.2.222.194 > 128.2.210.194: . seq
1480294582 win 0 (0)
```

- Session
 - Echo client on 128.2.222.194, server on 128.2.210.194
- Client FIN
 - SeqS: 1480294581
- Server ACK + FIN
 - Ack: 1480294582 (= SeqC+1)
 - SeqS: 1480294580
- Client ACK
 - Ack: 1480294580 (= SeqS+1)

State Diagram: Connection Tear-down



Outline



- TCP connection setup/data transfer
- **TCP reliability**

Reliability Challenges



- Congestion related losses
- Variable packet delays
 - What should the timeout be?
- Reordering of packets
 - How to tell the difference between a delayed packet and a lost one?