

Communication Networks (0368-3030) / Spring 2011

The Blavatnik School of Computer Science,
Tel-Aviv University

Allon Wagner

TCP Connection Management

Kurose & Ross, Chapter 3 (5th ed.)

Many slides adapted from:
J. Kurose & K. Ross \ Computer Networking: A Top Down Approach (5th ed.) Addison-Wesley, April 2009.
Copyright 1996-2010, J.F Kurose and K.W. Ross, All Rights Reserved.

TCP: Overview

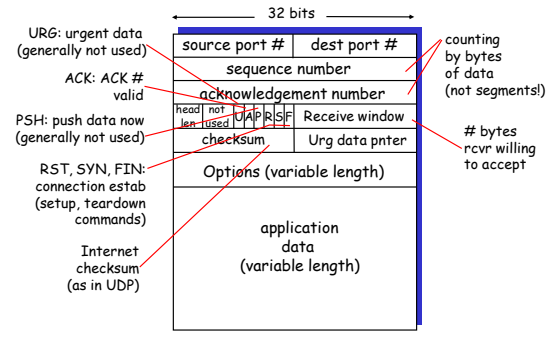
RFCs: 793, 1122, 1323, 2018, 2581

- ❖ **point-to-point:**
 - one sender, one receiver
- ❖ **reliable, in-order byte stream:**
 - no "message boundaries"
- ❖ **pipelined:**
 - TCP congestion and flow control set window size
- ❖ **send & receive buffers**
- ❖ **full duplex data:**
 - bi-directional data flow in same connection
 - MSS: maximum segment size
- ❖ **connection-oriented:**
 - handshaking (exchange of control msgs) inits sender, receiver state before data exchange
- ❖ **flow controlled:**
 - sender will not overwhelm receiver



Transport Layer 3-3

TCP segment structure



Transport Layer 3-4

TCP Connection Management

Recall: TCP sender, receiver establish "connection" before exchanging data segments

- ❖ initialize TCP variables:
 - seq. #s
 - buffers, flow control info (e.g. RevWindow)
- ❖ **client:** connection initiator
- ❖ **server:** contacted by client

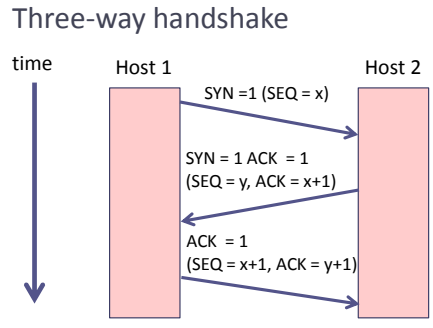
```
Socket clientSocket = new Socket("hostname", "port number");
Socket connectionSocket = welcomeSocket.accept();
```

Three way handshake:

- Step 1:** client host sends TCP SYN segment to server
 - specifies initial seq #
 - no data
- Step 2:** server host receives SYN, replies with SYNACK segment
 - server allocates buffers
 - specifies server initial seq. #
- Step 3:** client receives SYNACK, replies with ACK segment, which may contain data

Transport Layer 3-5

Three-way handshake



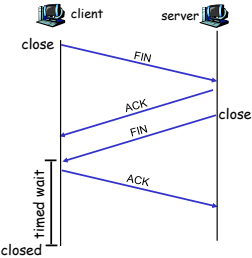
TCP Connection Management (cont.)

Closing a connection:

client closes socket:
`clientSocket.close();`

Step 1: client end system sends TCP FIN control segment to server

Step 2: server receives FIN, replies with ACK. Closes connection, sends FIN.



Transport Layer 3-7

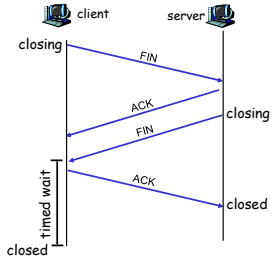
TCP Connection Management (cont.)

Step 3: client receives FIN, replies with ACK.

- Enters "timed wait" - will respond with ACK to received FINs

Step 4: server, receives ACK. Connection closed.

Note: with small modification, can handle simultaneous FINs.



Transport Layer 3-8



Transport Layer

TCP's statechart

- On board
 - Statechart appears in RFC 793
- Discussion of:
 - TIME_WAIT state
 - Syn flood attacks