DATA COMMUNICATOIN NETWORKING

Instructor: Ouldooz Baghban Karimi

Course Book: Computer Networking, A Top-Down Approach
By: Kurose, Ross

Course Overview

- **Basics of Computer Networks**
  - Internet & Protocol Stack
  - Application Layer
  - Transport Layer
  - Network Layer
  - Data Link Layer

- **Advanced Topics**
  - Case Studies of Computer Networks
  - Internet Applications
  - Network Management
  - Network Security

TCP Sender Events

Data received from App

- Create segment with sequence number
- Sequence number is byte-stream number of first data byte in segment
- Start timer if not already running
  - Think of timer as for oldest unACKed segment
  - Expiration interval: `TimeOutInterval`

Timeout

- Retransmit segment that caused timeout
- Restart timer

ACK received

- If ACK acknowledges previously unACKed segments
  - Update what is known to be ACKed
  - Start timer if there are still unACKed segments

TCP Sender Events

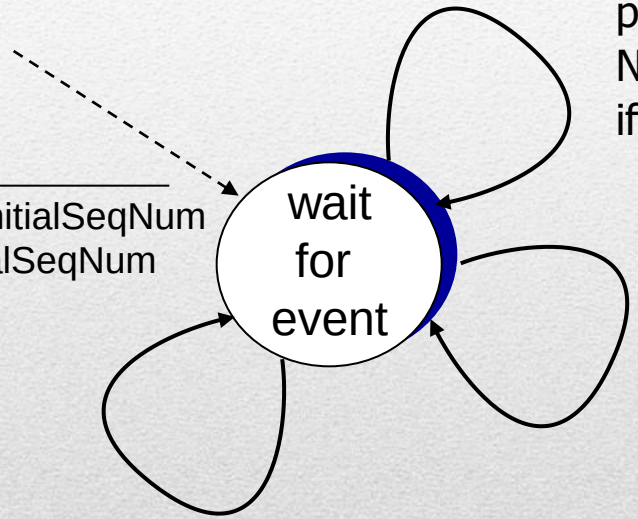

Λ

NextSeqNum = InitialSeqNum
SendBase = InitialSeqNum

data received from application above

create segment, seq. #: NextSeqNum
pass segment to IP (i.e., "send")
NextSeqNum = NextSeqNum + length(data)
if (timer currently not running)
    start timer

timeout

retransmit not-yet-acked segment
        with smallest seq. #
start timer

ACK received, with ACK field value y

```
if (y > SendBase) {
   SendBase = y
   /* SendBase–1: last cumulatively ACKed byte */
   if (there are currently not-yet-acked segments)
      start timer
    else stop timer
   }
```

TCP: Retransmission Scenarios

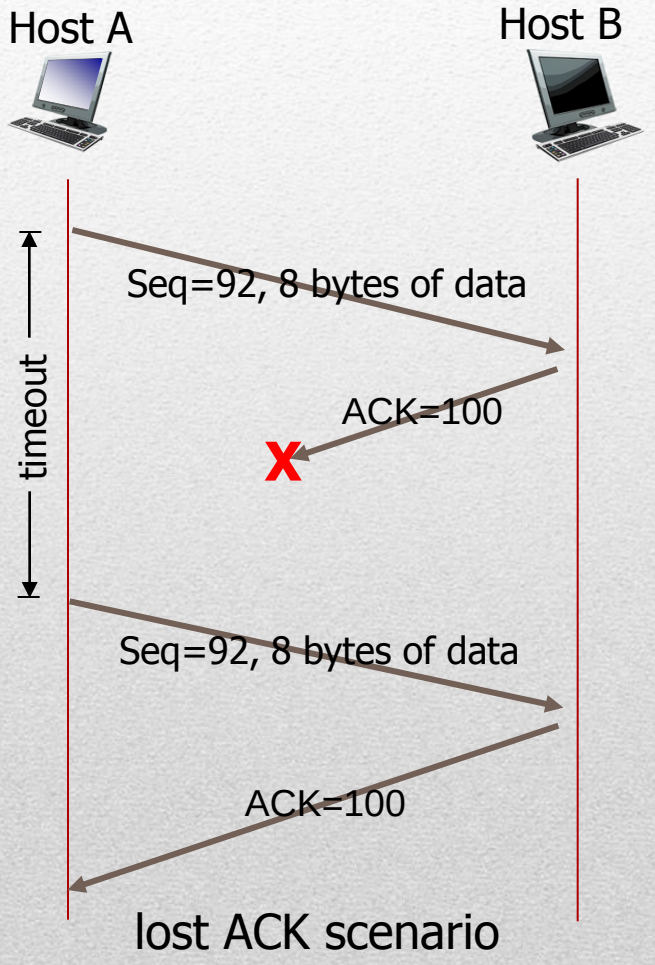

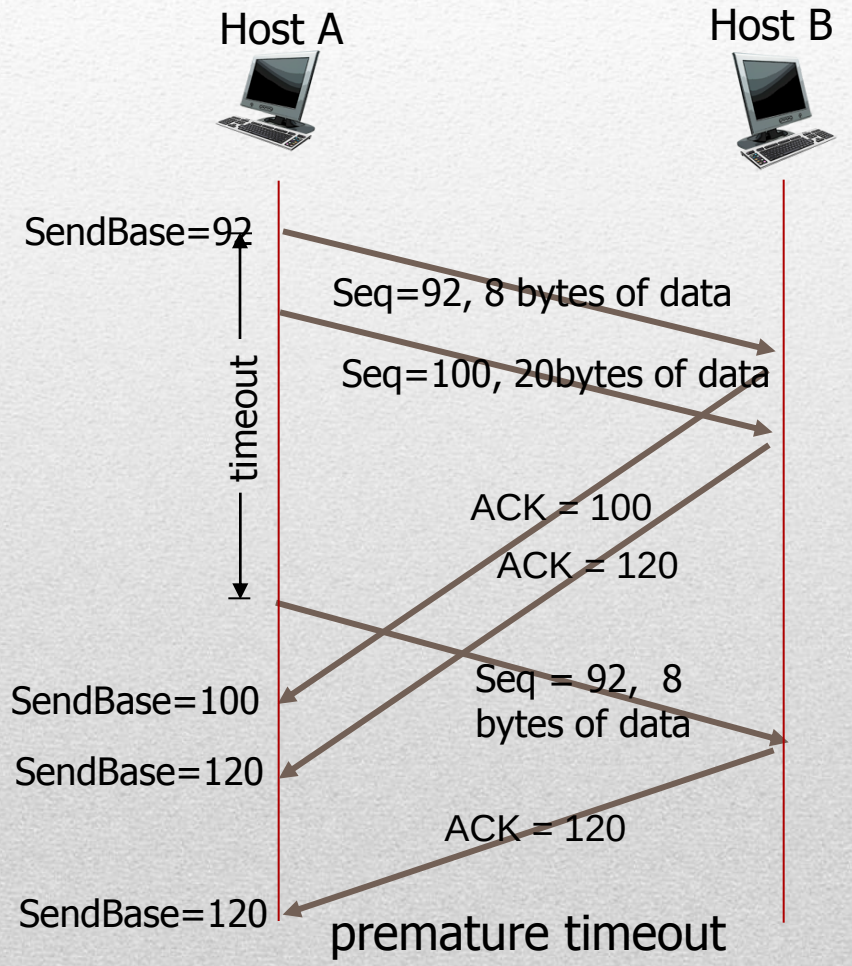

TCP: Retransmission Scenarios

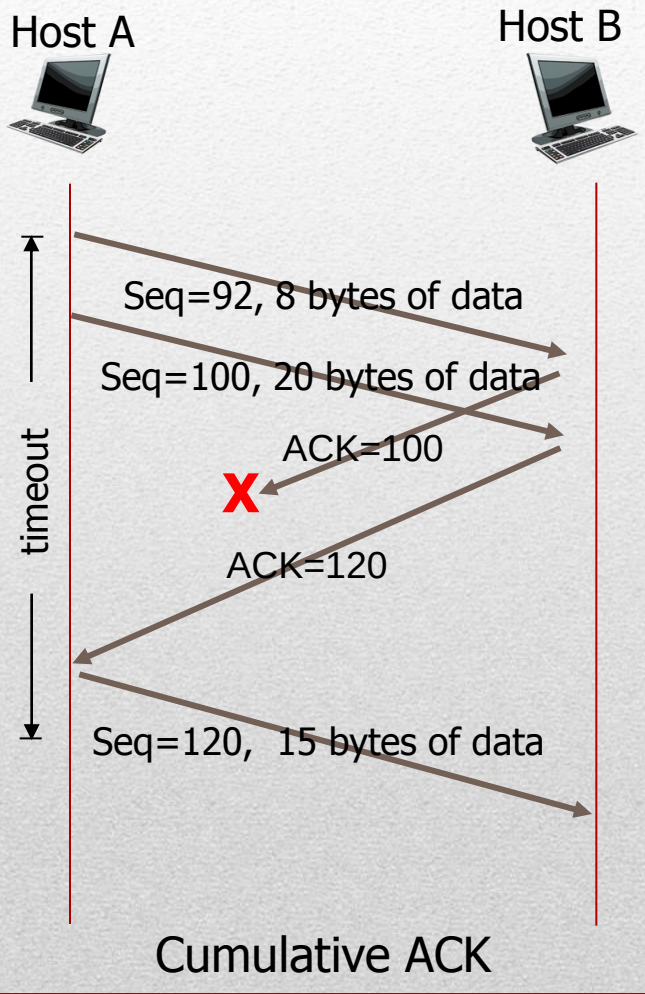

TCP ACK Generation

Event at receiver	TCP receiver action
Arrival of in-order segment with expected seq #. All data up to expected seq # already ACKed	Delayed ACK. Wait up to 500ms for next segment. If no next segment, send ACK
Arrival of in-order segment with expected seq #. One other segment has ACK pending	Immediately send single cumulative ACK, ACKing both in-order segments
Arrival of out-of-order segment higher-than-expect seq. # . Gap detected	Immediately send *duplicate ACK*, indicating seq. # of next expected byte
Arrival of segment that partially or completely fills gap	Immediate send ACK, provided that segment starts at lower end of gap

TCP FAST Retransmit

- **Time-out period often relatively long**
  - Long delay before resending lost packet
- **Detect lost segments via duplicate ACKs**
  - Sender often sends many segments back-to-back
  - If segment is lost, there will likely be many duplicate ACKs.

TCP fast retransmit

If sender receives 3 ACKs for same data ("triple duplicate ACKs"), resend unACKed segment with smallest sequence number

- Likely that unACKed segment lost, so do not wait for timeout

TCP FAST Retransmit

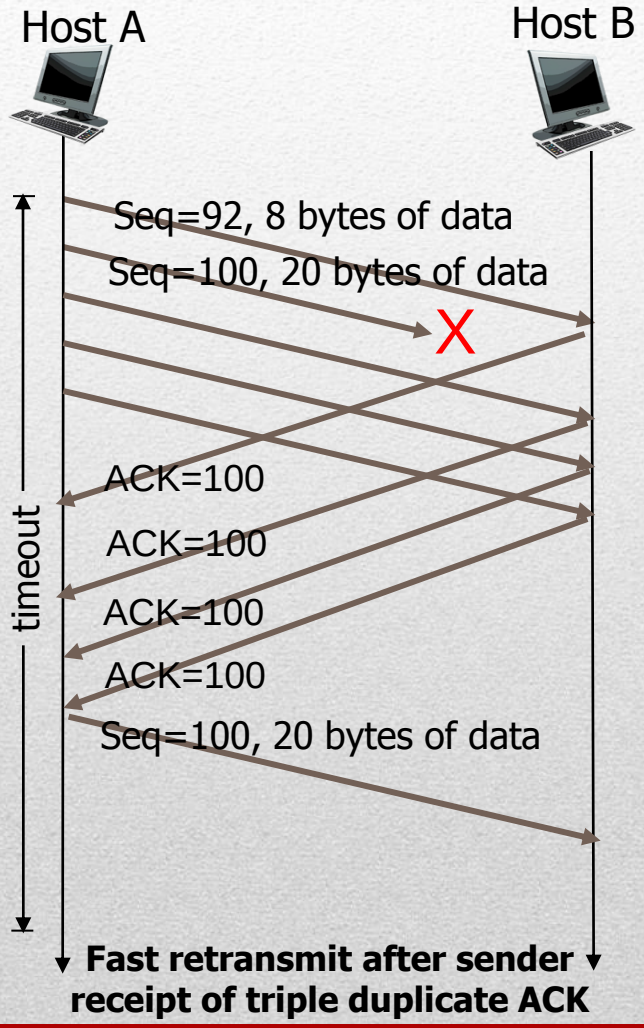

TCP Flow Control

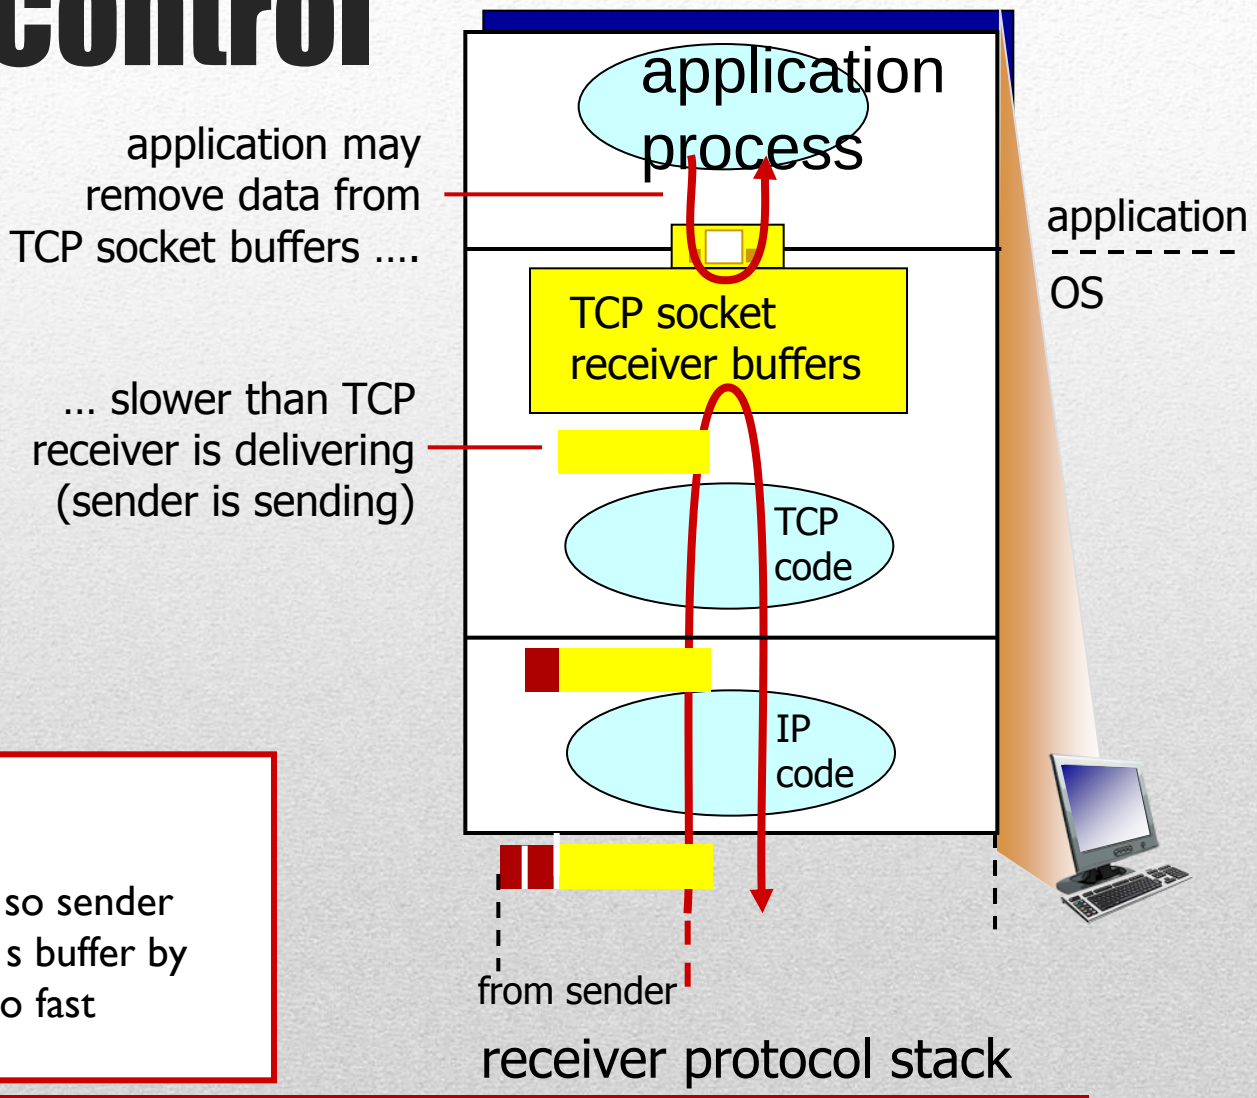

Flow Control

Receiver controls sender, so sender won't overflow receiver's buffer by transmitting too much, too fast

TCP Flow Control

- Receiver "advertises" free buffer space by including **`rwnd`** value in TCP header of receiver-to-sender segments
  - **`RcvBuffer`** size set via socket options (typical default is 4096 bytes)
  - Many operating systems autoadjust **`RcvBuffer`**
- Sender limits amount of unACKed ("in-flight") data to receiver's **`rwnd`** value
- Guarantees receive buffer will not overflow

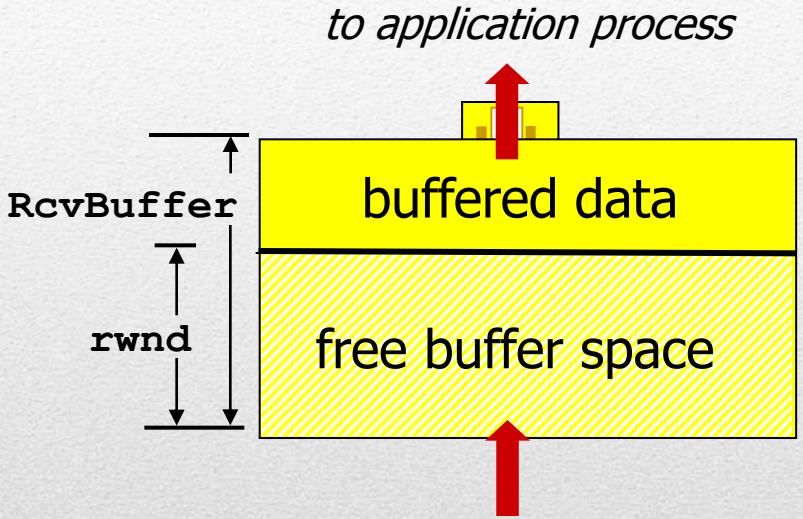

receiver-side buffering

Connection Management

Before exchanging data, sender/receiver "handshake":

- Agree to establish connection (each knowing the other willing to establish connection)
- Agree on connection parameters

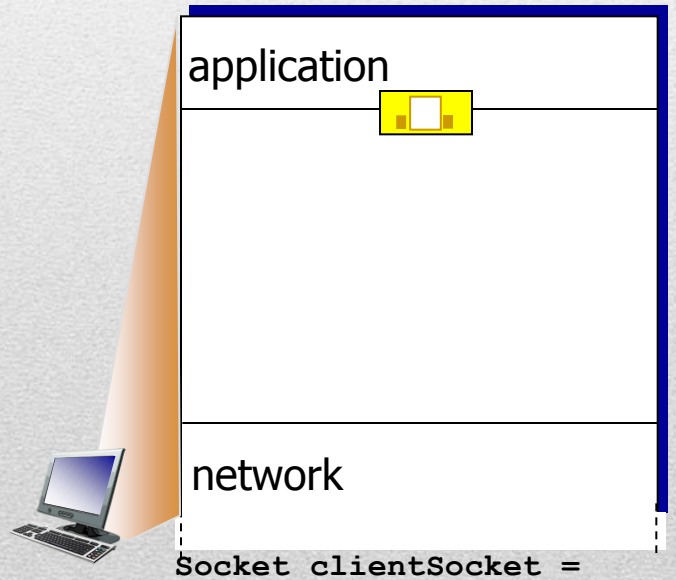


```
Socket clientSocket =
  newSocket("hostname","port
   number");
```

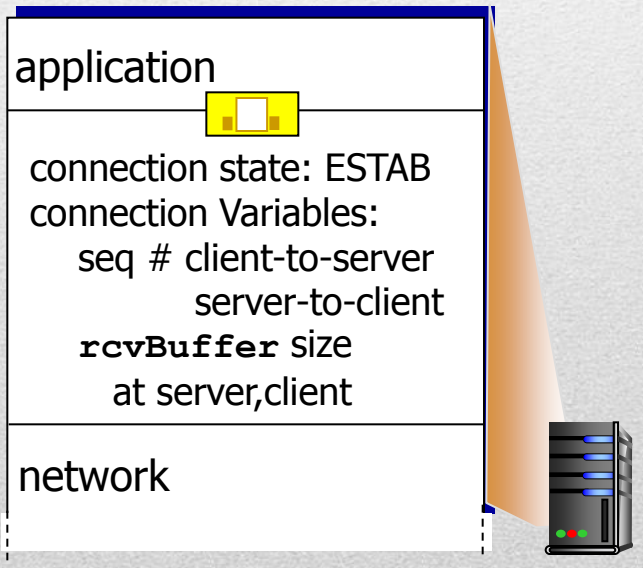


```
Socket connectionSocket =
  welcomeSocket.accept();
```

Agreeing to Establish a Connection

Two-way handshake

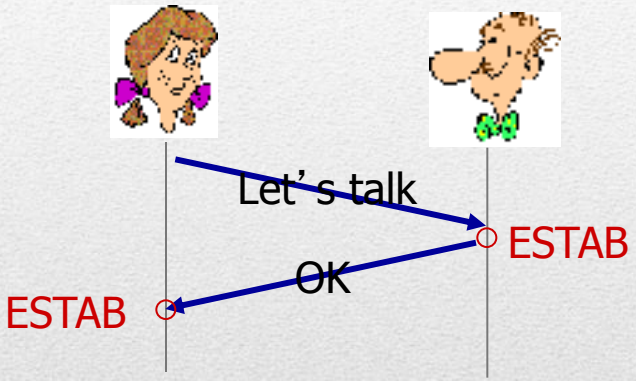

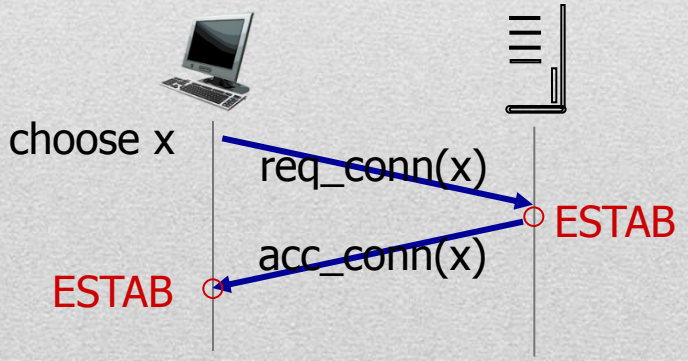

Q: Will 2-way handshake always work in network?

- Variable delays
- Retransmitted messages (e.g. req_conn(x)) due to message loss
- Message reordering
- Cannot "see" other side

Two-way Handshake Failure

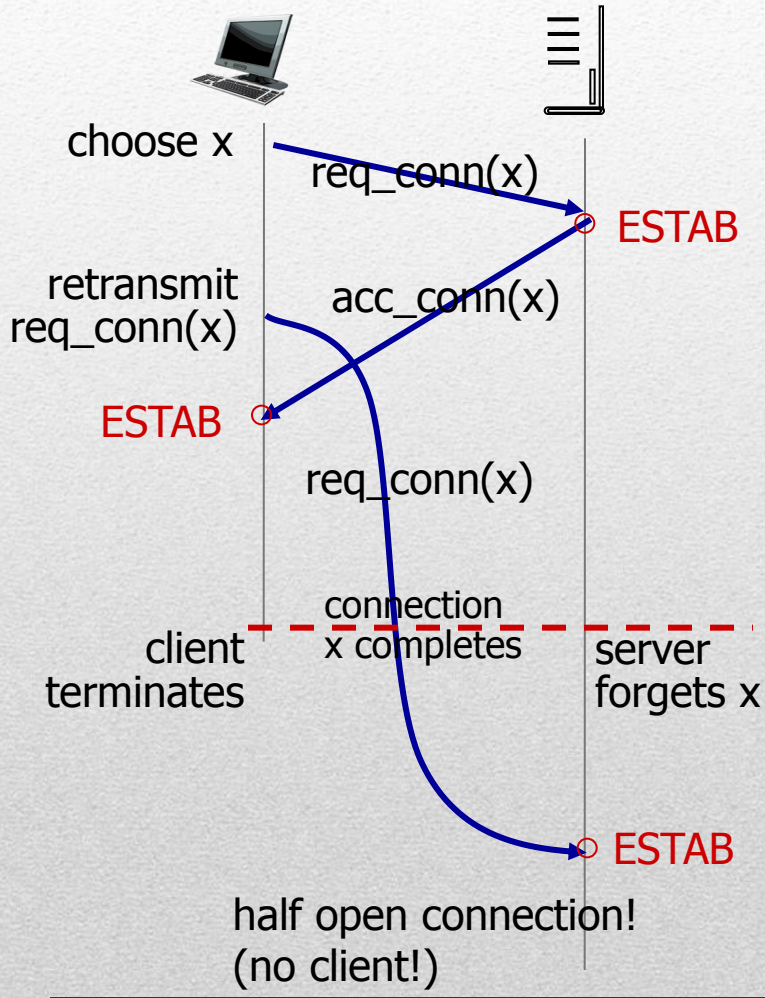

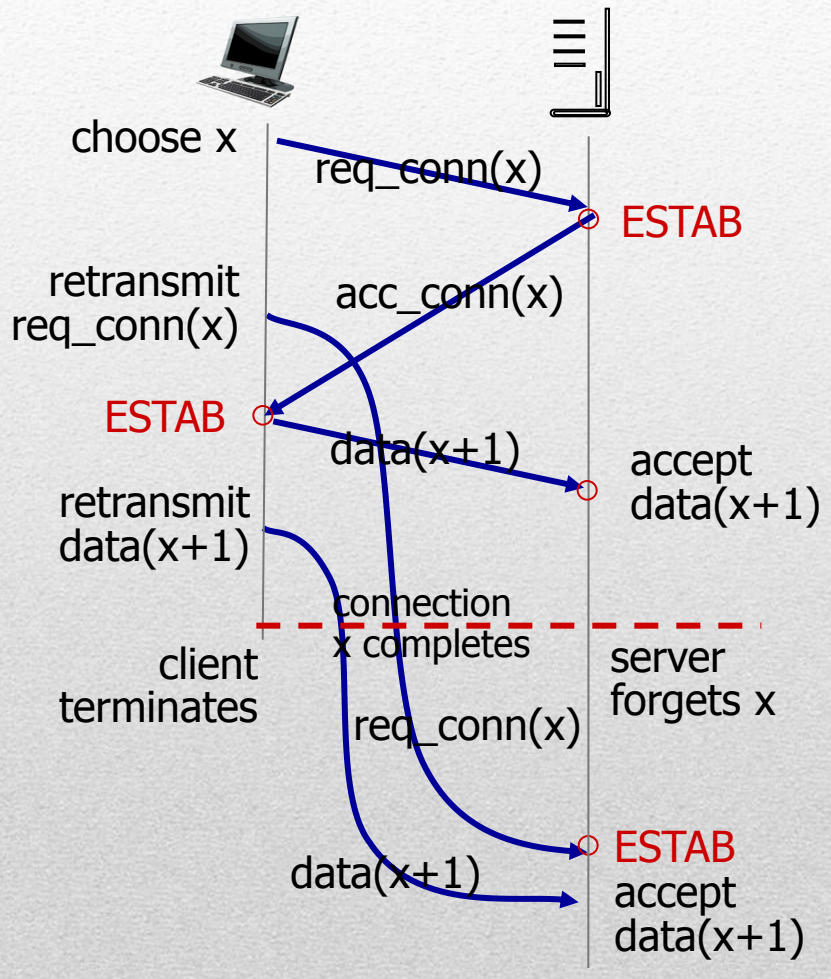

TCP Three-way Handshake

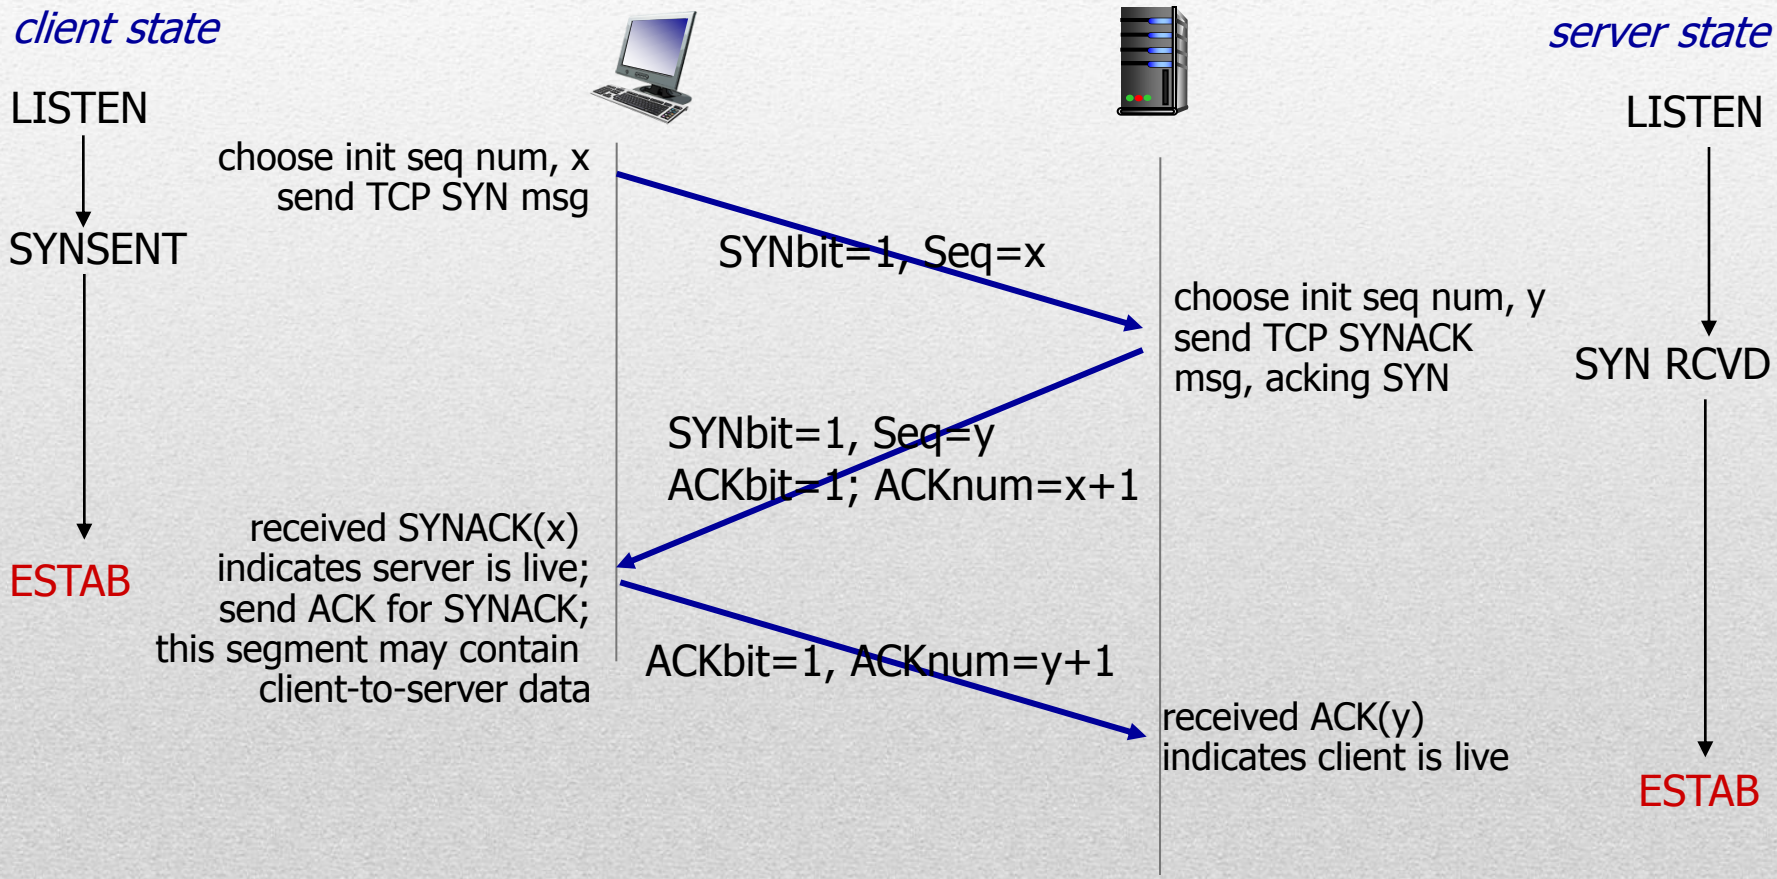

TCP Three-way Handshake

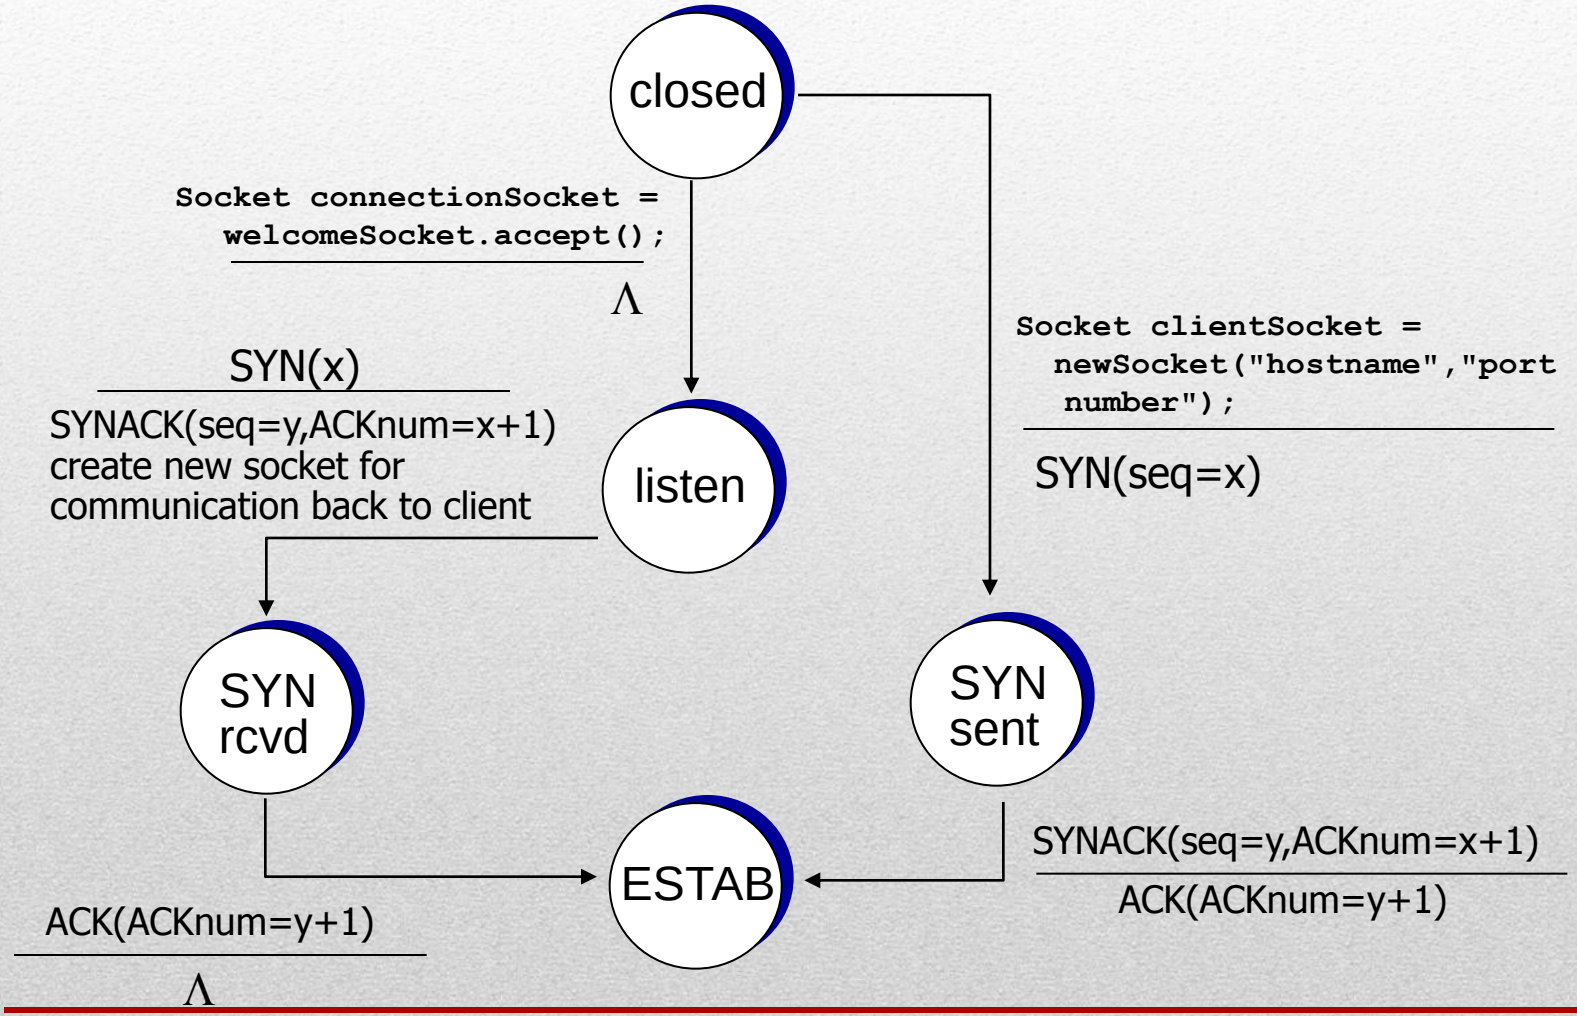

Closing a TCP Connection

- Client, server each close their side of connection
  - Send TCP segment with FIN bit = 1
- Respond to received FIN with ACK
  - On receiving FIN, ACK can be combined with own FIN
- Simultaneous FIN exchanges can be handled

Closing a TCP Connection