# TCP connection establishment (3-way handshake): 



- $X, Y$ are chosen randomly
$\rightarrow$ sequence number prediction
- piggybacking

2-person consensus problem: are $A$ and $B$ in agreement about the state of affairs after 3 -way handshake?
$\longrightarrow$ in general: impossible
$\longrightarrow$ can be proven
$\longrightarrow$ "acknowledging the ACK problem"
$\longrightarrow$ also TCP session ending
$\longrightarrow$ lunch date problem

## Call Collision:


$\longrightarrow$ only single TCB gets allocated
$\longrightarrow$ unique full association

## TCP connection termination:



- full duplex
- half duplex


# More generally, finite state machine representation of TCP's control mechanism: 

$\longrightarrow$ state transition diagram


Features to notice:

- Connection set-up:
- client's transition to ESTABLISHED state without ACK
- how is server to reach ESTABLISHED if client ACK is lost?
- ESTABLISHED is macrostate (partial diagram)
- Connection tear-down:
- three normal cases
- special issue with TIME WAIT state
- employs hack


## Basic TCP data transfer:



## TCP's sliding window protocol



- sender, receiver maintain buffers MaxSendBuffer, MaxRcvBuffer

Note asynchrony between TCP module and application.

Sender side: maintain invariants

- LastByteAcked $\leq$ LastByteSent $\leq$ LastByteWritten
- LastByteWritten-LastByteAcked < MaxSendBuffer
$\longrightarrow$ buffer flushing (advance window)
$\longrightarrow$ application blocking
- LastByteSent-LastByteAcked $\leq$ AdvertisedWindow

Thus,
EffectiveWindow = AdvertisedWindow-
(LastByteSent - LastByteAcked)
$\longrightarrow$ upper bound on new send volume

Actually, one additional refinement:
$\longrightarrow$ CongestionWindow

EffectiveWindow update procedure:
EffectiveWindow $=$ MaxWindow -
(LastByteSent - LastByteAcked)
where
MaxWindow $=$
min\{ AdvertisedWindow, CongestionWindow \}

How to set CongestionWindow.
$\longrightarrow$ domain of TCP congestion control

Receiver side: maintain invariants

- LastByteRead $<$ NextByteExpected $\leq$ LastByteRcvd +1
- LastByteRcvd - NextByteRead < MaxRcvBuffer
$\longrightarrow$ buffer flushing (advance window)
$\longrightarrow$ application blocking

Thus,
AdvertisedWindow $=$ MaxRcvBuffer-
(LastByteRcvd - LastByteRead)

Issues:

How to let sender know of change in receiver window size after AdvertisedWindow becomes 0 ?

- trigger ACK event on receiver side when

AdvertisedWindow becomes positive

- sender periodically sends 1 -byte probing packet
$\longrightarrow$ design choice: smart sender/dumb receiver
$\longrightarrow$ same situation for congestion control

Silly window syndrome: Assuming receiver buffer is full, what if application reads one byte at a time with long pauses?

- can cause excessive 1-byte traffic
- if AdvertisedWindow $<$ MSS then set AdvertisedWindow $\leftarrow 0$

Do not want to send too many 1 B payload packets.

Nagle's algorithm:

- rule: connection can have only one such unacknowledged packet outstanding
- while waiting for ACK, incoming bytes are accumulated (i.e., buffered)
... compromise between real-time constraints and efficiency.
$\longrightarrow$ useful for telnet-type applications

Sequence number wrap-around problem: recall sufficient condition
SenderWindowSize < (MaxSeqNum + 1)/2
$\longrightarrow$ 32-bit sequence space/16-bit window space

However, more importantly, time until wrap-around important due to possibility of roaming packets.

| bandwidth | time until wrap-around $\dagger$ |
| :--- | :---: |
| T1 $(1.5 \mathrm{Mbps})$ | 6.4 hrs |
| Ethernet $(10 \mathrm{Mbps})$ | 57 min |
| T3 $(45 \mathrm{Mbps})$ | 13 min |
| F/E $(100 \mathrm{Mbps})$ | 6 min |
| OC-3 $(155 \mathrm{Mbps})$ | 4 min |
| OC-12 $(622 \mathrm{Mbps})$ | 55 sec |
| OC-24 $(1.2 \mathrm{Gbps})$ | 28 sec |

Even more importantly, "keeping-the-pipe-full" consideration.

| bandwidth | delay-bandwidth product $\dagger$ |
| :--- | :---: |
| T1 (1.5 Mbps) | 18 kB |
| Ethernet (10 Mbps) | 122 kB |
| T3 (45 Mbps) | 549 kB |
| FDDI (100 Mbps) | 1.2 MB |
| OC-3 (155 Mbps) | 1.8 MB |
| OC-12 (622 Mbps) | 7.4 MB |
| OC-24 (1.2 Gbps) | 14.8 MB |
| $\longrightarrow 100 \mathrm{~ms}$ latency |  |

Also, throughput limitation imposed by TCP receiver window size.
$\longrightarrow$ e.g., high-performance grid apps

## RTT estimation

... important to not underestimate nor overestimate.

Karn/Partridge: Maintain running average with precautions

EstimateRTT $\leftarrow \alpha \cdot$ EstimateRTT $+\beta \cdot$ SampleRTT

- SampleRTT computed by sender using timer
- $\alpha+\beta=1 ; \quad 0.8 \leq \alpha \leq 0.9,0.1 \leq \beta \leq 0.2$
- TimeOut $\leftarrow 2 \cdot$ EstimateRTT or

TimeOut $\leftarrow 2 \cdot$ TimeOut (if retransmit)
$\longrightarrow$ need to be careful when taking SampleRTT
$\longrightarrow$ infusion of complexity
$\longrightarrow$ still remaining problems

## Hypothetical RTT distribution:



Jacobson/Karels:

- Difference $=$ SampleRTT - EstimatedRTT
- EstimatedRTT $=$ EstimatedRTT $+\delta \cdot$ Difference
- Deviation $=$ Deviation $+\delta(\mid$ Difference $\mid-$ Deviation $)$

Here $0<\delta<1$.

Finally,

- TimeOut $=\mu \cdot$ EstimatedRTT $+\phi \cdot$ Deviation
where $\mu=1, \phi=4$.
$\longrightarrow$ persistence timer
$\longrightarrow$ how to keep multiple timers in UNIX

