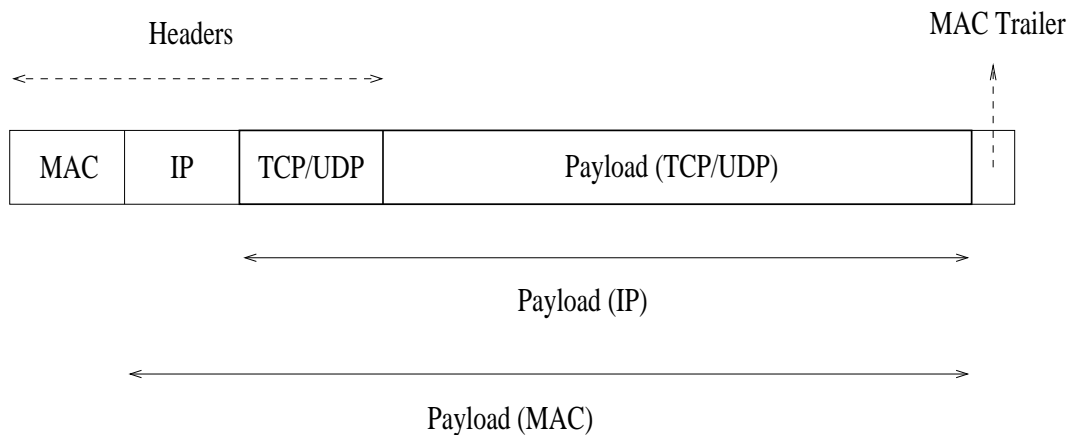


Transport Protocols: TCP/UDP Structure

- end-to-end mechanism
- runs on top of link-based mechanism
- treat network layer as black box

Three-level encapsulation:



Network layer assumptions:

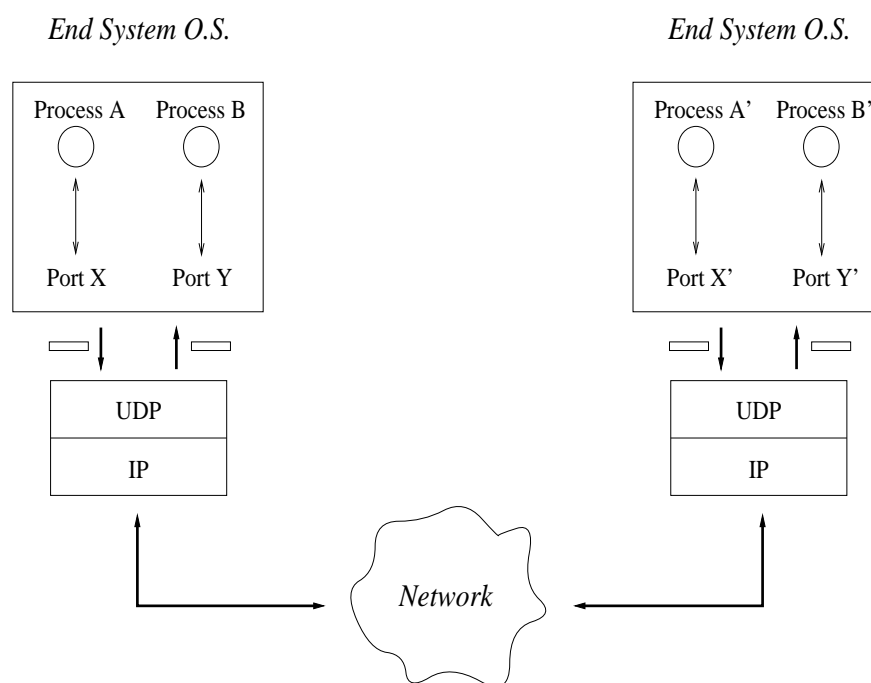
- unreliable
- out-of-order delivery (in general)
- absence of QoS guarantees (delay, throughput etc.)
- insecure (IPv4)

Additional (informal) performance properties:

- works “fine” under low load conditions
- can break down under high load conditions
- behavior range predictable (to certain extent)

Goal of UDP: Process identification (“multiplexing”).

→ port number as process demux key



- form of end host processing (O.S.)
- generally: end system support (e.g., scheduling)

UDP packet format:

2	2
Source Port	Destination Port
Length	Checksum
Payload	

Checksum calculation (pseudo header):

4		
Source Address		
Destination Address		
00 ... 0	Protocol	UDP Length

Goals of TCP:


- process identification
 - reliable communication (ARQ)
 - speedy communication (congestion/flow control)
 - segmentation
- connection-oriented (i.e., stateful)
- complex mixture of functionalities

Segmentation task: Provide “stream” interface to higher level protocols

—→ view: contiguous stream of bytes

- segment stream of bytes into blocks or *segments* of fixed size
- segment size determined by TCP MTU (Maximum Transmission Unit)
- use also for reliability mechanism

TCP packet format:

Source Port		Destination Port						
Sequence Number								
Acknowledgement Number								
Header Length		U R G	A C K	P S H	R S T	S S Y	F I N	Window Size
Checksum				Urgent Pointer				
Options (if any)								
DATA (if any)								

- Sequence Number: position of first byte of payload
- Acknowledgement: next byte of data expected (receiver)
- Header Length (4 bits): 4 B units
- URG: urgent pointer flag
- ACK: ACK packet flag
- PSH: override TCP buffering
- RST: reset connection
- SYN: establish connection
- FIN: close connection
- Window Size: receiver's advertised window size
- Checksum: prepend pseudo-header
- Urgent Pointer: byte offset in current payload where urgent data begins
- Options: MTU; take min of sender & receiver (default 556 B)

Checksum calculation (pseudo header):

4

Source Address		
Destination Address		
00 ... 0	Protocol	UDP Length

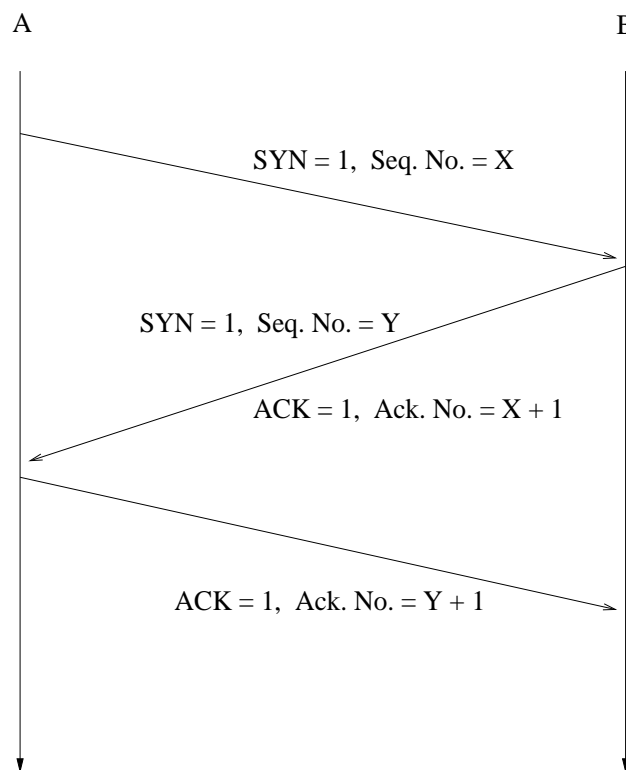
Nagle's algorithm:

- do not want to send too many 1 B payload packets
- 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.

→ useful for `telnet`-type applications

TCP connection establishment (3-way handshake):




- X, Y are chosen randomly
- piggybacking
- sequence number prediction
- lingering packet problem

2-person consensus problem: Are A and B in agreement about the state of affairs after 3-way handshake?

→ impossibility, in general


→ lunch date problem

Call Collision:



- only single TCB gets allocated
- unique full association

TCP connection termination:



- full duplex
- half duplex

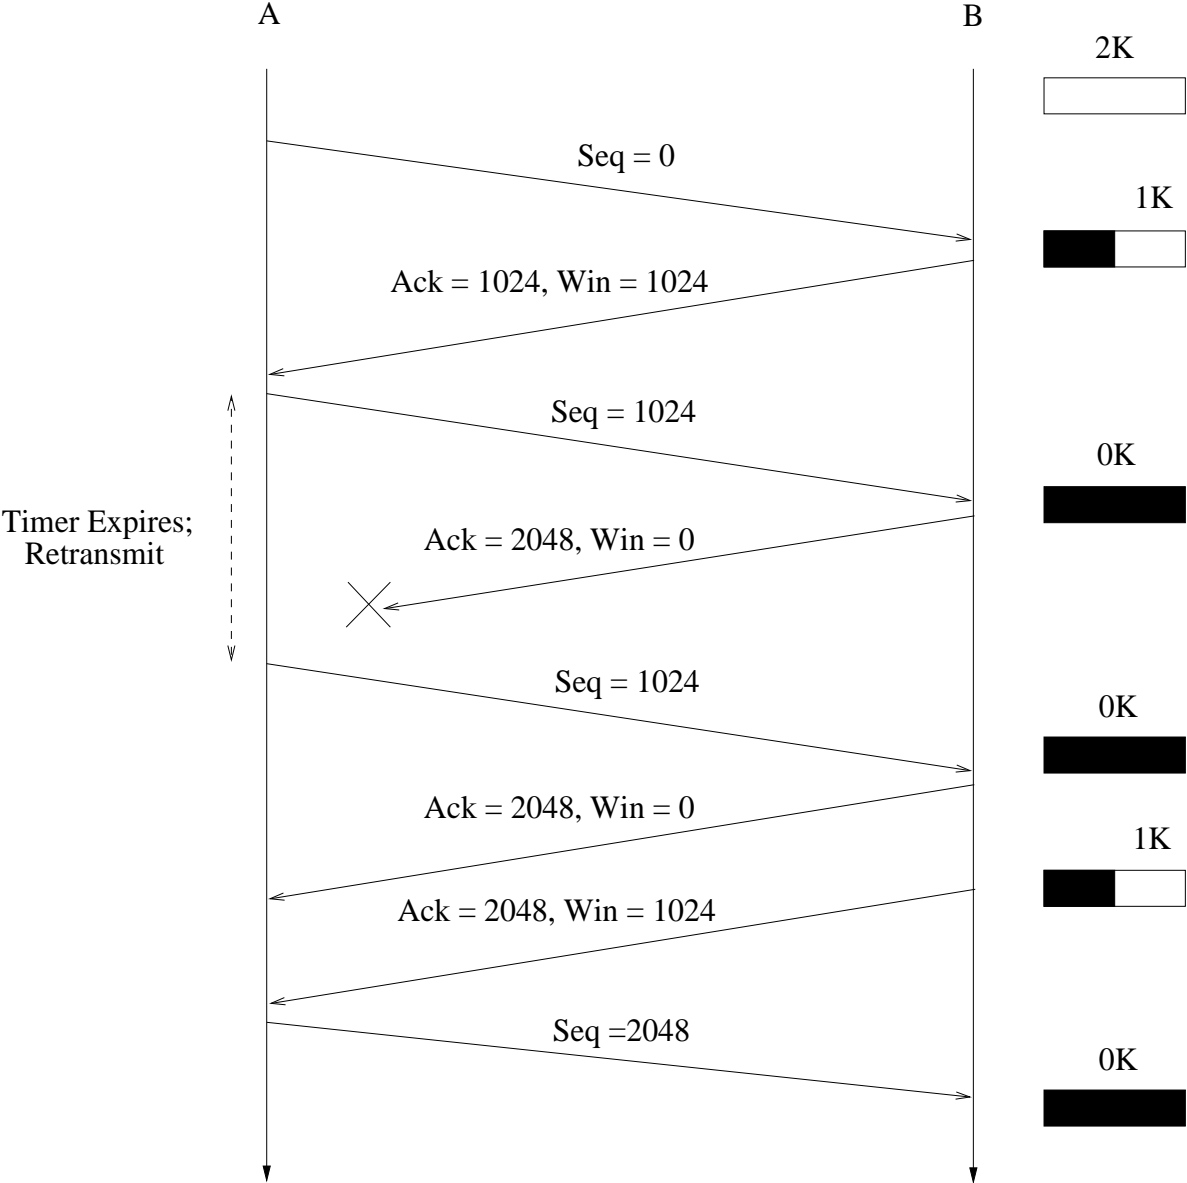
More generally, finite state machine representation of TCP's control mechanism:

TCP's State-transition Diagram comes here

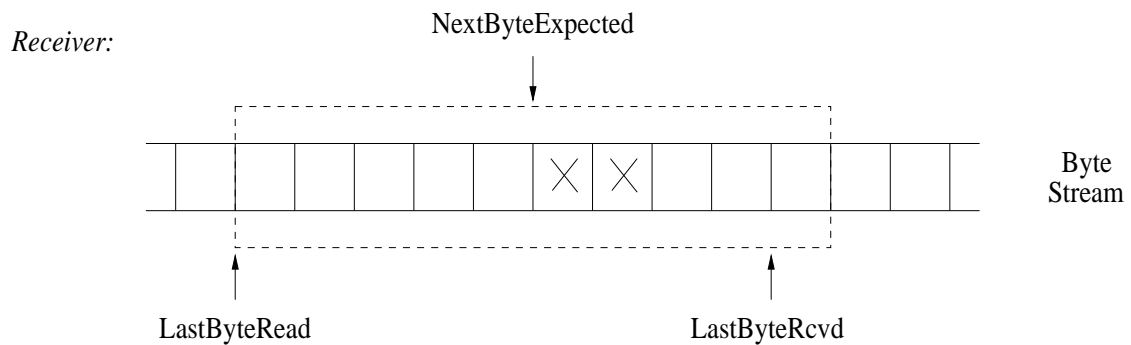
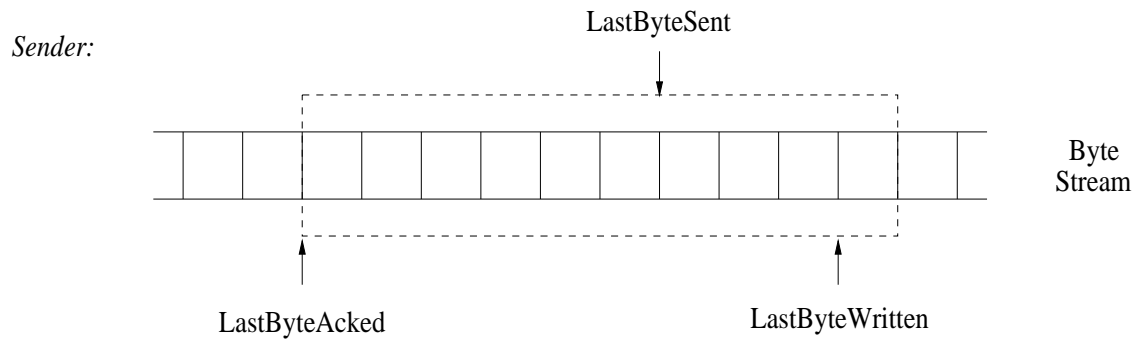
Features to notice:

- Connection set-up:
 - client's transition to **ESTABLISHED** state without **ACK**
 - how is server to reach **ESTABLISHED** if client **ACK** is lost?
 - TCP: default **ACKing** executed by all data packets; no extra overhead incurred
 - note: **ESTABLISHED** is macrostate
 - not a complete transition diagram
- Connection tear-down:
 - three normal cases
 - special issue with **TIME WAIT** state

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

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

Thus,

$$\text{EffectiveWindow} = \text{AdvertisedWindow} - (\text{LastByteSent} - \text{LastByteAcked})$$

→ upper bound on new send volume

Receiver side: maintain invariants

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

Thus,

$$\text{AdvertisedWindow} = \text{MaxRcvBuffer} - (\text{LastByteRcvd} - \text{LastByteRead})$$

Three problems:

How to let sender know of changed 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
→ design choice: smart sender/dumb receiver

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 ← 0`

Sequence number wrap-around problem: recall sufficient condition

$$\text{SenderWindowSize} < (\text{MaxSeqNum} + 1)/2$$

→ 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 †
T1 (1.5 Mbps)	6.4 hrs
Ethernet (10 Mbps)	57 min
T3 (45 Mbps)	13 min
FDDI (100 Mbps)	6 min
OC-3 (155 Mbps)	4 min
OC-12 (622 Mbps)	55 sec
OC-24 (1.2 Gbps)	28 sec

† From P & D for 32-bit sequence space

Even more importantly, “keeping-the-pipe-full” consideration.

bandwidth	delay-bandwidth product †
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

† From P & D for 100 ms latency

RTT estimation

... important to not underestimate nor overestimate.

Karn/Partridge: Maintain running average with precautions

$$\text{EstimateRTT} \leftarrow \alpha \cdot \text{EstimateRTT} + \beta \cdot \text{SampleRTT}$$


- **SampleRTT** computed by sender using timer
- $\alpha + \beta = 1$; $0.8 \leq \alpha \leq 0.9$, $0.1 \leq \beta \leq 0.2$
- **TimeOut** $\leftarrow 2 \cdot \text{EstimateRTT}$ or
TimeOut $\leftarrow 2 \cdot \text{TimeOut}$ (if retransmit)

→ need to be careful when taking **SampleRTT**

→ infusion of complexity

→ still remaining problems

Hypothetical RTT distribution:



→ need to account for variance

Jacobson/Karels:

- $\text{Difference} = \text{SampleRTT} - \text{EstimatedRTT}$
- $\text{EstimatedRTT} = \text{EstimatedRTT} + \delta \cdot \text{Difference}$
- $\text{Deviation} = \text{Deviation} + \delta (|\text{Difference}| - \text{Deviation})$

Here $0 < \delta < 1$.

Finally,

- $\text{TimeOut} = \mu \cdot \text{EstimatedRTT} + \phi \cdot \text{Deviation}$

where $\mu = 1$, $\phi = 4$.

→ persistence timer

→ how to keep multiple timers in UNIX