

TCP

12강

Transmission Control Protocol

- Most frequently used transport layer protocol
 - The most complex Internet protocol
- Service model and characteristics
 - Connection-oriented
 - Reliable
 - In-order delivery
 - Full-duplex

TCP service model

- Connection-oriented: data exchange possible only after "connection" is set up
- Reliable: retransmission will be done if data packet is not acknowledged
- In-order delivery: in the order the user process provided
- Full-duplex: two data channels are open in one connection

TCP

- TCP sees user data as a stream of bytes
 - No predetermined boundaries
 - Segmentation is performed by TCP
 - Each packet is called "segment"
- Segment size can be adjusted by PMTU discovery
 - ICMP notifies TCP if ICMP error (packet too big but DF set) arrives
- Maximum segment size (MSS) = MTU - 40

TCP

- TCP header format = 20 bytes + options (max 40 bytes)
 - Source port, dest port
 - Sequence number
 - Ack number
 - Header length (In units of 4 bytes)
 - 8 flags
 - Window size
 - Checksum
 - Urgent pointer
 - Options

Sequence number

- Each byte has one 32-bit number
- SYN and FIN segments consume a sequence number
- “Randomly chosen”
 - So that no one can guess
 - Security reason
 - Hijacking possible if the choice is guessed
- The first byte in each segment

Acknowledgement number

- If x is ACK number, it means
 - The receiver has up to $(x-1)$ bytes, so expects x next
 - $x-1$ is the sequence number of the last consecutive bytes the receiver has got
- TCP sender can flush up to $(x-1)$ in the sending socket buffer upon $ACK(x)$

Header length

- Like in IP
- Basic TCP header is 20 bytes
 - "5"
 - Options can come after this
 - TCP options are very useful

Flags

- CWR, ECE: congestion control use
- URG: urgent pointer, means there are urgent bytes came into the byte stream
- ACK: always on except in the first SYN segment
- PSH: "push", today it only means that there is no data byte in the sending socket buffer after this segment (i.e. not functional)

Flags

- RST: "reset", kills the TCP connection
 - Not really a reset (going back to the initial state)
- SYN: "synchronize" 3 pieces of information to establish the connection
 - Later in connection set up
- FIN: "finish"
 - Hangs up the connection in a graceful manner

Window size

- Actually it is the available receive socket buffer space
 - For flow control
- Must be interpreted with respect to the ACK number
 - E.g. 3500 bytes from ACK number 12780

TCP checksum

- The Internet checksum
- Covers the entire TCP segment
 - Header
 - Body
- As in UDP, uses “pseudo-header”
- Mandatory in TCP

Urgent pointer

- Points at the last byte of the urgent data in the byte stream
 - Offset to the position
- In TCP, urgent data cannot preempt preceding normal data
 - They should be served in the same way as other normal bytes
- TCP does not know what kind of urgent data it is; it just knows it's urgent