ECE-647
TCP: Transmission Control Protocol in the Internet: Part I

Instructor: Xiaojun Lin

Purdue University

TCP/IP Architecture

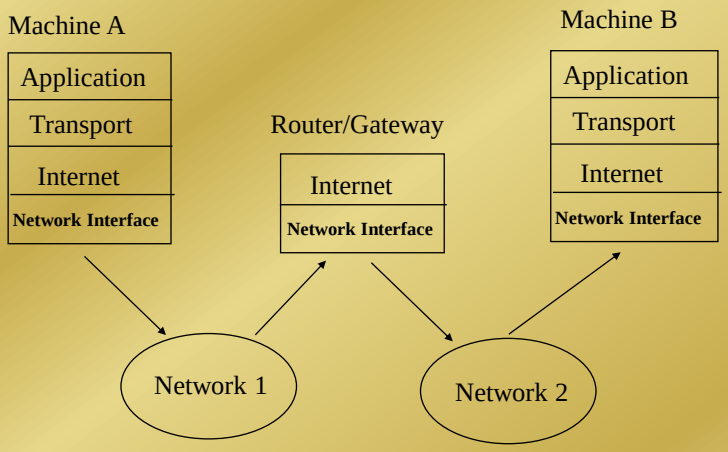

Purdue University

TCP/IP

- *network layer:* IP provides "best-effort" packet delivery between two end-systems
  - Packets could be lost, duplicated, trapped in loops

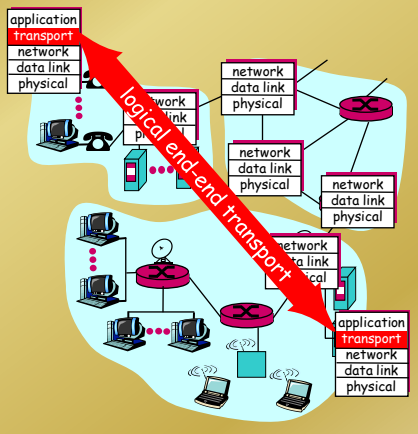

IP Header

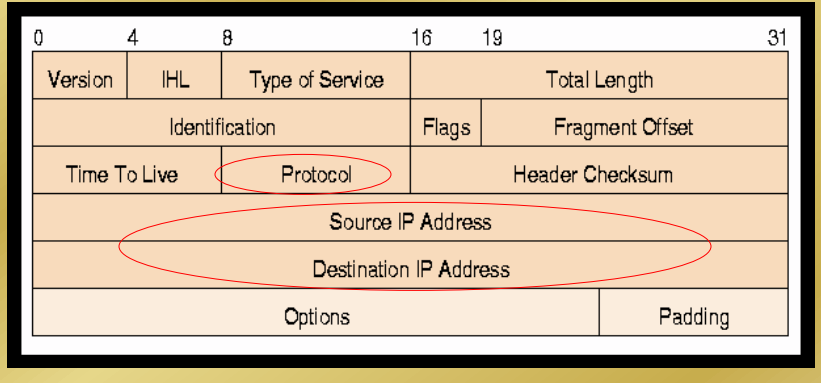

TCP/IP

- *network layer:* IP provides "best-effort" packet delivery between two end-systems
  - Packets could be lost, duplicated, trapped in loops
- *transport layer:* relies on, and enhances, network layer services

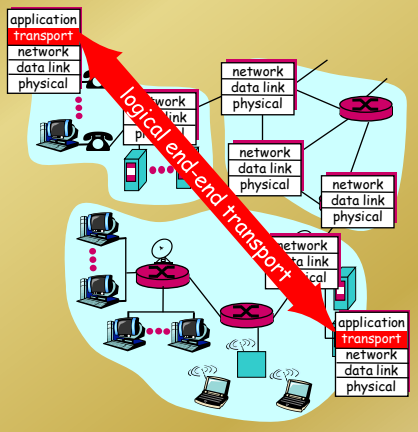

TCP/IP

- The internet offers two basic services that operate on top of IP
- ***UDP*** *(User Datagram Protocol)*
  - unreliable ("best-effort"), unordered unicast or multicast delivery
  - UDP header contains a "port" number that identifies which application process should receive the packet.
- ***TCP*** *(Transmission Control Protocol)*
  - connection oriented, reliable, in-order unicast delivery
  - congestion control
  - flow control

TCP Features

- Connection oriented, duplex, reliable *byte-stream* service with flow-control/congestion-control.
- *Connection oriented*
  - Both end-points maintain the state of the connection (Open/Close/Sequence numbers)
  - However, intermediate routers do NOT maintain the state of the connection (do not reserve resources)
- Reliability:
  - Retransmit a packet when it is not acknowledged.

Purdue University

TCP Header

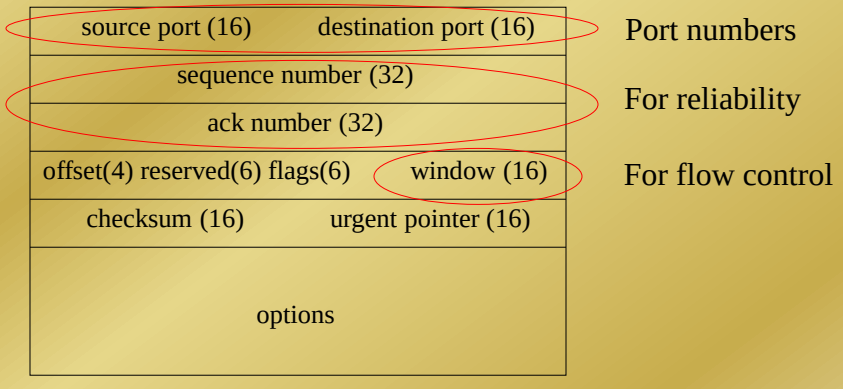

Purdue University

TCP Features

- Every byte of data sent over a TCP connection has a 32-bit sequence number, given by
  - the segment sequence number plus
  - its position in the segment.
- Sequence number of the first segment of a connection is agreed upon by a three way handshake

Purdue University

Three-way Handshake

Client A		**Server B**
Closed		Listen
SYN-SENT	<SEQ=100><CTL=SYN> →	SYN-RECEIVED
Established	← <ACK=101> <CTL=SYN,ACK>	SYN-RECEIVED
Established	<SEQ=101><CTL=ACK> →	Established
Established	<SEQ=101><CTL=ACK> <DATA>→	Established
Established	← <ACK=102>	Established
Established	<SEQ=102><CTL=ACK> <DATA> →	Established

Purdue University

TCP Features

- For each packet sent, the receiver is expected to respond with an acknowledgement
- 32-bit ack is cumulative:
  - Ack of *n* indicates that all bytes up to *n-1* have been received correctly, and *n* is the next expected byte number.
- If the acknowledge is not received with a period of time, the packet is assumed to be lost, and will be retransmitted

Purdue University

TCP Window Control

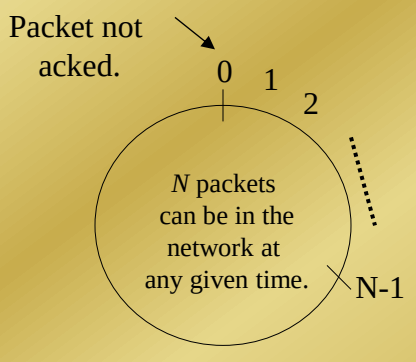

Initial State

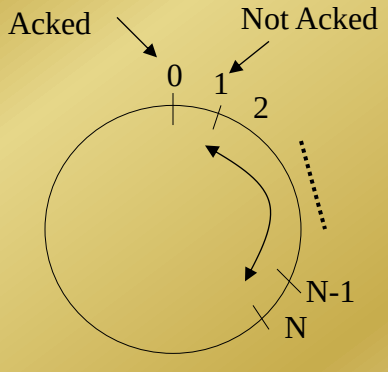

Packet 0 Acked

Purdue University

Window Size

- Window size determines the number of packets that can be sent before waiting for an ack
- Large window size leads to higher rate
- However, the network may become congested.

Purdue University

Congestion Collapse

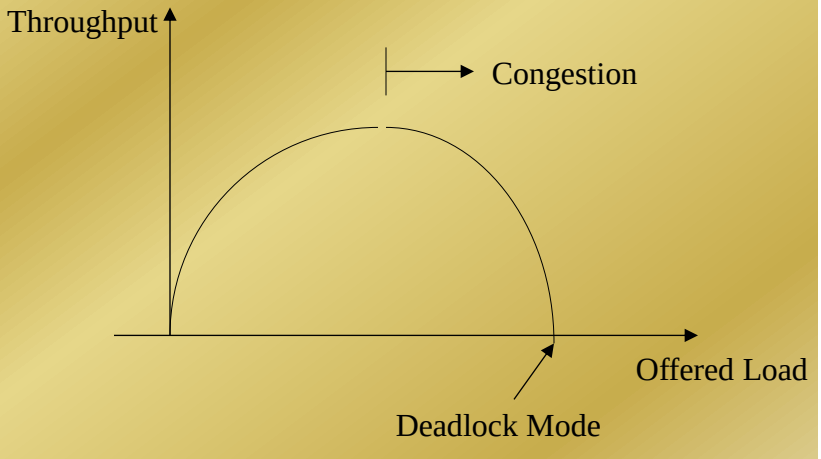

Purdue University