

### **NWEN 243**

**Networked Applications** 

Layer 4 – TCP and UDP



### About the second lecturer

- Aaron Chen
- Office: AM405
- Phone: 463 5114
- Email: <u>aaron.chen@ecs.vuw.ac.nz</u>
- Transport layer and application layer protocols
- Building mobile applications and Web Services
- Multimedia streaming and P2P (optional)



#### Transport layer protocols

• Entire network seen as a pipe



### Transport layer port numbers

- How can ensure it is delivered to the right application on that machine?
- We need to be able to address the applications on a machine, just like we had to address the machines themselves.
- We use port numbers.



IP: 192.168.0. Port: 3000

### Layer 4 – the Transport Layer.

- The network layer performs machine to machine delivery of datagrams
- The transport layer performs application to application delivery.
- Ports are (16 bit) numbers (like house numbers) that form the address space of a protocol.
  - i.e. you can have tcp/53 and udp/53.
- A socket is a software structure associated with a port.
- An application must 'bind' (associate) a socket to a port before it can be used.
- The Socket API is the interface for applications to gain access to the network.



### Socket API

- Client/server paradigm
- Two levels of service in the socket API:
  - User Datagram Protocol (UDP)
    - Best effort protocol, transmits datagrams.
  - Transmission Control Protocol (TCP)
    - reliable, byte stream-oriented, with capacity control, transmits segments.



### Simple Example

- Client is assigned some random unused port x by its host.
- The SMTP server runs on well known port 25.
- Reply is to destination port x, with src port 25



port used: smtp

### **Delivery Without Guarantees**

- Remember, IP Datagrams may be:
  - Duplicated,
  - Delivered out of order,
  - Lost/Discarded
  - Corrupted
- By the network layer.
- The TCP protocol will need to resolve these, while the UDP protocol ignores them.



### Stop And Wait



### Stop and Wait Performance







### **Pipelined Protocols**

Pipelining: sender allows multiple, "in-flight", yet-to-beacknowledged packets.



(a) a stop-and-wait protocol in operation

(b) a pipelined protocol in operation

Two generic forms of pipelined protocols: go-Back-N, selective repeat

# How do we manage all these inflight packets?

- We can use a software construct called a sliding window.
  - We have a fixed size window (i.e. n packets)
    - This means we can have up to n packets on the 'wire'.
    - When we successfully send a packet (i.e. its been ack'd by the receiver), we can move the window on by I.
    - We need buffering at sender and/or receiver





### Quick exercise

- When using the sliding-window protocol, will the utilization be improved because of the following?
  - A. Reduce propagation delay
  - B. Reduce sliding window size
  - C. Increase sliding window size
  - D. Reduce packet size



### Protocol summary

- The important features of a sliding window protocol are:
  - Sequence number assignment
  - The resend protocol
  - Window management
- We're going to look at GoBackN followed by Selective repeat.



Timeout on first packet, then all UNACKNOWLEDGED packets resent

8

7

5

sent

not ACK'd

9

available

10

window (size = 8)

12 13 14 15 16 17

Not available

18 19



- Should a later packet be ACKed, consider all prior packets in the window to also be ACKed, so in other words:
  - ACK is cumulative, so an ACK for a later sequence number effectively ACKs all preceding packets.
- Lost ACK(6) and ACK(7)
- Get ACK(8)



### Go Back N in action

#### Sender

Receiver



Sender's sliding window size is 4



### Quick exercise

- In a Go-Back-N protocol, if the window size is 63, what is the minimum range of sequence numbers?
  - A. 0-62
  - B. 0-63
  - C. 0-64
  - D. I-63



### An example

sequence number range is 0-2 and the window size is 3.

A -> 0 -> B

A -> I -> B

**B** is expecting to receive a new packet **0** 

A -> 2 -> B

(ACKs from B to A all lost)



### Selective Repeat



- Can send fewer packets at the cost of making the protocol more complicated.
- Packets are individually acknowledged.
- Only one packet resent on timer expiry.

### Selective Repeat in Action



22



### Quick exercise

- In Selective Repeat, if 5 is the number of bits for the sequence number, then the maximum size of the sliding window must be \_\_\_\_\_
  - A. 15
    B. 16
    C. 31
    D. 1



### An example

sequence number range is 0-3 and the window size is 3.

A -> 0 -> B

A -> I -> B

B's sliding window={3,0, I}

A -> 2 -> B

(ACKs from B to A all lost)



#### User datagram protocol

- Connection-Less
  - (no handshaking)



- UDP packets (Datagrams)
  - Each application interacts with UDP transport sw to produce EXACTLY ONE UDP datagram!



This is why, improperly, we use the term UDP packets



#### UDP datagram format

8 bytes header + variable payload

0	7	15	23	3
	source port		destination port	
	length (bytes)		Checksum	
		Da	ta	

- <sup>1</sup> UDP length field
  - all UDP datagram
  - (header + payload)
  - payload sizes allowed:
    - Empty (0)
    - 65527 bytes (65535-8)

- UDP functions limited to:
  - Error checking
    - which may even be disabled for performance
  - addressing
    - which is the only strictly necessary role of a transport protocol



#### IPv4 pseudo header

• The pseudo header is only used for the checksum computation





#### Quick exercise

- Which of the following is not true about disabling checksum?
  - A. Checksum in UDP is not always necessary since IP packet checksum includes packet payload.
  - B. UDP checksum is often disabled to speed up implementation.
  - C. Lack of UDP checksum is tolerable for communication in LANs.
  - D. Lack of UDP checksum is definitely dangerous in the Internet.

### UDP: a lightweight protocol

- No connection establishment
  - no initial overhead due to handshaking
- No connection state
  - greater number of supported connections by a server!
- Small packet header overhead
  - 8 bytes only vs 20 in TCP
- originally intended for simple applications, oriented to short information exchange
  - DNS
  - management (e.g. SNMP)
  - Distributed file system support (e.g. NFS)
  - etc



#### Unregulated send rate in UDP

- No rate limitations
- No re-transmission
- Extremely important features for today multimedia applications
- Be careful: UDP ok for multimedia because it does not provide anything at all (no features = no limits!).



### Quick exercise

- Which of the following about UDP is correct?
  - A. UDP supports a self-regulating "throttle" feature that prevents network saturation
  - B. UDP consumes fewer computer resources by not maintaining connection state
  - C. UDP guarantees that individual packets of a transmission will arrive "in order"
  - D. None of the above



### Question to ponder

 Why would a TCP and UDP "phone call" likely have equivalent performance characteristics in practice?



#### RTP: sublayer of transport



#### **Connection Oriented Transport**





### TCP

- point-to-point:
  - one sender, one receiver
  - no multicast
- reliable, in-order byte steam:
  - no "message boundaries"

- full duplex data:
  - bi-directional data flow in same connection
- connection-oriented:
  - handshaking (exchange of control msgs) init's sender, receiver state before data exchange





### Quick exercise

- TCP does not support virtual circuit, why?
  - A.TCP relies on IP and IP does not support virtual circuit
  - B.TCP maintains connection information at two communicating end systems only
  - C.Virtual circuit can only be established at the network layer
  - D. None of the above



### The TCP Header

Byte stream sent as sequence of segments

• Segment may be 0 to 64k bytes

Offsets Octet						0				1									2								3									
Octet	Bit	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31			
0	0	Source port Destination port																																		
4	32	Sequence number																																		
8	64	Acknowledgment number (if ACK set)																																		
12	96			Be	sen	ved	N	c	B	U	A	P	R	s	F																					
		Data offset		set	0 0 0	0	s	W	C	R	С	S	S	Y	I		Window					w Size														
																REG		H	т	N	N															
16	128	Checksum Urgent pointer (if URG set)																																		
20	160	Options (if Data Offset > 5, padded at the end with "0" bytes if necessary)																																		





## PiggyBacking

- There is one important way to increase efficiency.
  - Assume that any connection is duplex,
  - with data flowing evenly back and forth.
- Now, we don't need to send a separate acknowledgment, just attach to a returning data segment.



### Quick exercise

- - A. no
  - B. one
  - C. two
  - D. three

### **TCP Connection Establishment**

TCP uses a **three-way handshake** to open a connection:

- I. ACTIVE OPEN: Client sends a segment with
  - SYN bit set
  - port number of client and server
  - initial sequence number (X) of client
- 2. **PASSIVE OPEN:** Server responds with a segment with
  - SYN bit set
  - initial sequence number (Y) of server
  - ACK for ISN of client (X+I)
- 3. Client acknowledges by sending a segment with:
  - ACK ISN of server (Y+I)



### Establishing a Connection

- Normal connection
- A initiates and chooses its initial sequence number.
- B replies with its own initial sequence number and acknowledges host I.
- A begins data transmission.





### **Duplicate Request**

- A previous connection request appears from nowhere.
- B replies with its own initial sequence number and acknowledges host
   I.
- A rejects the connection request.





### **Duplicate Connection**

- A previous connection request appears from nowhere.
- B replies.
- The first data packet of the previous connection also appears.
- A rejects the connection request.





### Quick exercise

- The connection establishment procedure in TCP is susceptible to a serious security problem called the \_\_\_\_\_\_ attack.
  - A.ACK flooding
  - B. FIN flooding
  - C. SYN flooding
  - D. None of the above



 Each end of the data flow must be shut down independently ("half-close")

Four steps involved:

 (I) A sends a FIN to B (active close)
 (2) B ACKs the FIN,

(at this time: B can still send data to A)

(3) and B sends a FIN to A (passive close)(4) A ACKs the FIN.



### Quick exercise

- Which process can be used to terminate a TCP connection?
  - A. three-way handshake
  - B. four-way handshake
  - C. two-termination sequence
  - D. None of the above



### **TCP** Sequencing

 Retransmission due to lost acknowledgment.





### TCP Sequencing cont.

 Cumulative acknowledgment avoids retransmission.





### TCP Sequencing cont.

 Delayed ACK causes retransmission





### 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.

- If sender receives 3 ACKs for the same data, it supposes that segment after ACKed data was lost:
  - <u>fast retransmit</u>: resend segment before timer expires



### Fast retransmit algorithm:



#### TCP ACK Generation @ Receiver

Event	TCP Receiver action
in-order segment arrival, no gaps, everything else already ACKed	delayed ACK.Wait up to 500ms for next segment. If no next segment, send ACK
in-order segment arrival, no gaps, one delayed ACK pending	immediately send single cumulative ACK
out-of-order segment arrival 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 ACK if segment starts at lower end of gap