La Sapienza
Università degli Studi di Roma
Dipartimento di Informatica e Sistemistica

Computer Networks II
Exercise collection 1 – TCP/IP

Luca Becchetti/Ugo Colesanti /Emiliano Trevisani

colesanti@dis.uniroma1.it

A.A. 2008/2009

Exercise 1:

- Consider the following scenario:

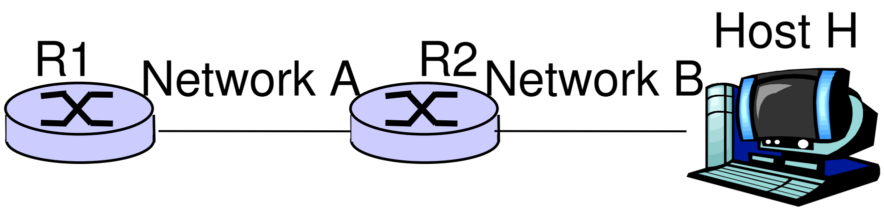

- Assume an MTU OF 1500 Bytes for network A and 532 bytes for network B respectively.
- Assume R1 receives an IP packet [packet P] directed at H and containing a 2000 bytes TCP segment [TCP header + data].
- Assume 20 bytes for IP header.

- Quiz 1: describe the M bit [More Fragment] and the Offset field for IP fragments of packet P sent from R1 on network A
- Quiz 2: answer the same question packets sent from R2 over network B .

Quiz 1 - solution 1:

- Q1: MTU = 1500 · 20 header + 1480 payload + 20 header + 520 payload \ 2 frammenti
  - First fragment: M=1 , Offset= 0
  - Second fragment: M=0 , Offset= 1480/8 = 185
- Q2:
  - First packet originating from R1: 20+512+20+512+20+456 ▪ 3 fragments
    - M = 1 , Offset= 0
    - M = 1, Offset = 512/8 = 64
    - M = 1, Offset = 1024/8 = 128
  - Second packet originating from R1: 20+512+20+8 ▪ 2 fragments
    - M=1, Offset = 1480/8 = 185
    - M=0, Offset = 1992/8 = 249
- Remark:
  - All fragments contain same value as packet P **for identifier field**

Quiz 2:

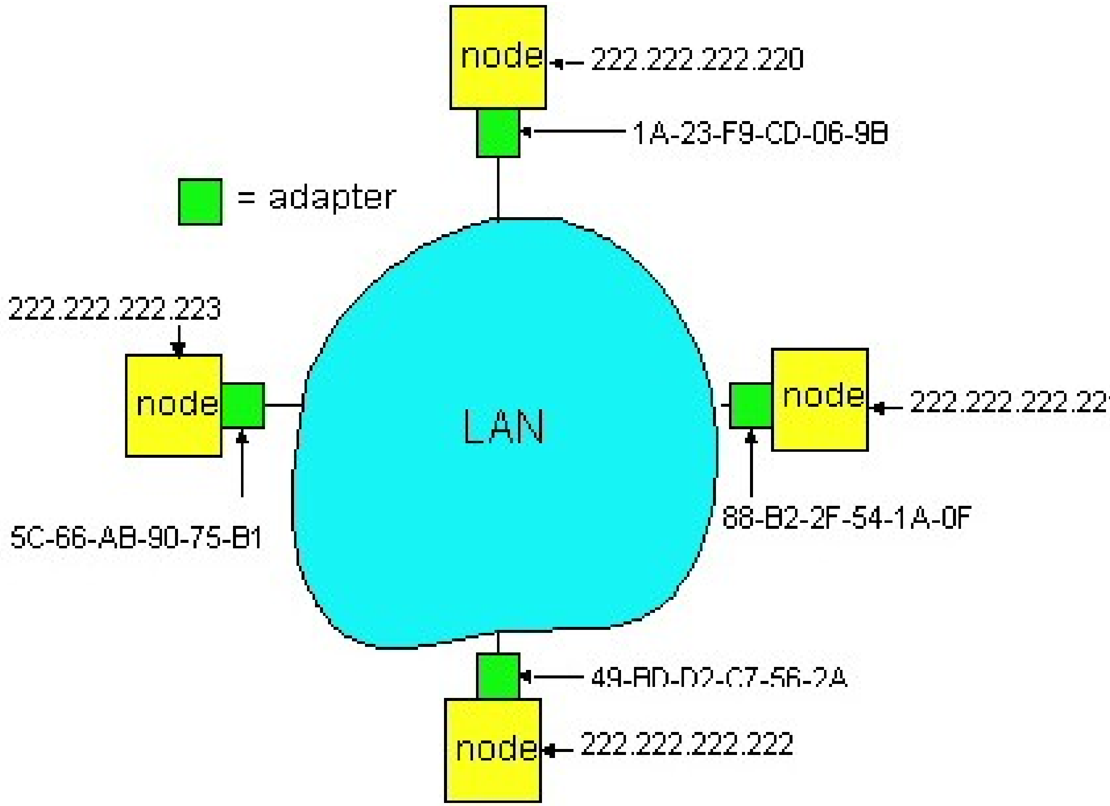

- Describe the process of sending a datagram from IP node 222.222.222.223 to node 222.222.222.220. Describe tasks performed by IP and ARP protocols. Assume the ARP table of 222.222.222.223 initially contains no entry for address 222.222.222.220.

Quiz 2 - solution

- IP hands datagram to link layer
- ARP looks up IP address in local ARP table ▪ it does not find it
- ARP sends broadcast request over local network
- ARP receives reply from node 222.222.222.220
- ARP updates local table
- ARP encapsulates IP datagram within a frame and sends it over the local network

- Which protocol layer does ARP logically belong to?
- In which protocol layer's data unit is an ARP PDU encapsulated

Appendix:
Window based flow control and Bandwidth – Delay product 1/3

- let W [expressed as number of equal size segments] be current size of Advertised Window:
  - At most W segments can be transmitted before ACK to first segment [or to any segment transmitted in case of cumulative ACKs] is received
  - Let L be size [in bits] of every transmitted segment [assumed constant for simplicity]
  - In window based flow control over a link we have:

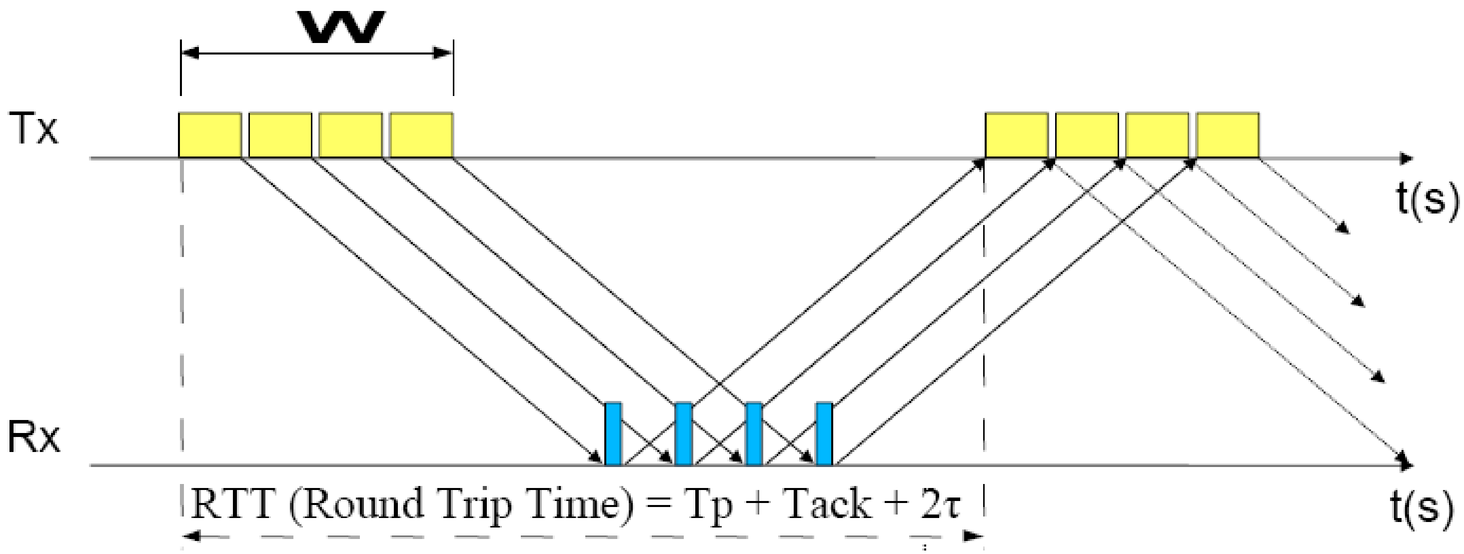

Appendix: Window based flow control and Bandwidth – Delay product 2/3

- Transmission speed is therefore bounded as follows:

$$B_{eff} = \min(B, \frac{W \cdot L}{RTT})$$

  - B denotes bandwidth available over communication channel

- So, in order to exploit all available bandwidth, we must have:

$$B \cdot RTT \geq W \cdot L$$

  - B·RTT known as *delay-bandwidth* product is an important index of efficiency
  - **Remark**: effective window size varies dynamically according to flow and congestion control mechanisms – its value is typically lowe than ideal one

Appendix: Window based flow control and Bandwidth – Delay product 3/3

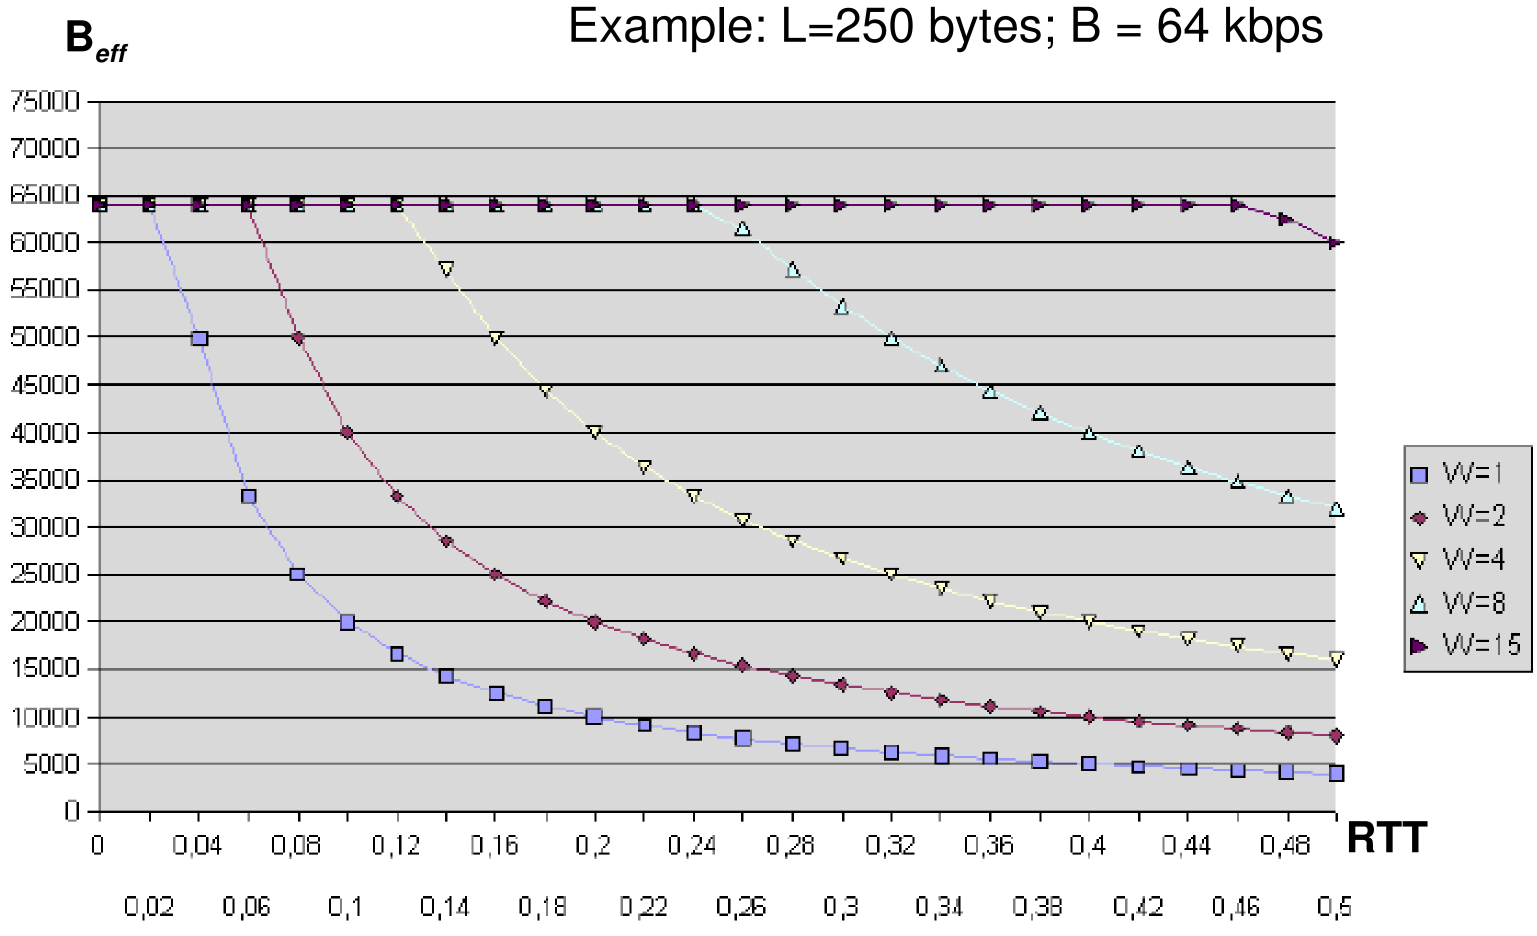

Appendix:
Maximum Segment Lifetime [MSL] 1/2

- Can sequence numbers be reused?I
  - Only if we are sure that no more old segments using these numbers are still in the network
  - Finite sequence number space requires a limit on segments' lifetime
  - Maximum Segment Lifetime o MSL must be publicly known
    - MSL = 2 mins (RFC 793)
    - On a 2 Megabits/sec link 4.5 hours needed to exhaust 2^{32} byte numbering
    - On a 100 Megabits/sec link 5.4 minutes needed to exhaust 2^{32} byte numbering
    - What happens with bandwidth values 1 Gbit/s or higher?
  - MSL denotes maximum time a ind TCP segment may persist in the network before being dropped

Appendix:
Maximum Segment Lifetime [MSL] 2/2

- At the end of the active shut down of a connection a TCP peer entity:
  - Remains in TIME-WAIT state for a time of 2 MSLs
  - At the end of this interval the connection is finally CLOSED
- Permanence in TIME-WAIT state guarantees all segments pertaining to the current connection are "out" of the network
  - This way, a TCP receiving entity can safely assume that a segment with the same sequence number as a previous one is duplicate and does not belong to a previous connection with the same remote TCP entity
- Permanence in TIME_WAIT state guarantees that a client cannot open a connection towards the same address as the one just closed using the same local port

Quiz 3:

- Assume you must design a byte-based reliable protocol using a TCP-like sliding window mechanism. Assume the protocol must operate on a 1 Gbps network with the following characteristics:
  - Round Trip Time (RTT) = 140ms
  - Maximum Segment Lifetime (MSL) = 60s

Which are the minimum sizes for the AdvertisedWindow and SeqNum fields in order to avoid wraparound and use available bandwidth efficiently?

Quiz 3 - solution

- Delay – bandwith product:

  $B \cdot RTT = (1/8)x10^9 \; x \; 140x10^{-3} = 17,5 \cdot 10^6$ B   16,7 MB

- MinWindowSize = $17,5 \cdot 10^6$ B ¿25bit (2^{25} = 33.554.432)
  - Segments headers neglected
  - MSL = 60s ▪ $(60x1x10^9)/8$ = 7,5 GB ▪ 33 bits
  - From the time a segment is transmitted 7,5 Gbytes can be transmitted until it expires; we must ensure that sequence number not reset before this amount of bytes are transmitted.

Quiz 4:

- TCP over 40 Gbps links
- A single TCP connection is allowed to use all available bandwidth

Quiz: how long before wraparound occurs?

Quiz 4 - solution:

- 32 bits NumSeq
- Time to wraparound: $2^{32}/((40/8)x10^{9})$ = 0,859s

Quiz 5:

- TCP entity A sends 10KB message to TCP entity B
- Connection already established
- MSS used by A = 500 bytes
- Initial NumSeq of A = 50
- All packets reach destination

Quiz: which are the sequence numbers of segments sent from A to B ?

Quiz 5 - solution:

- 20 500 bytes segments (header not included in MSS)
- -> NumSeq = 50 + (i-1)*500 , i= 1,2,....20

- Remark from RFC 879: “TCP maximum segment size and related topics”
  - The MSS counts only data octets in the segment, it does not count the TCP header or the IP header.

Quiz 6:

- Given the following routing table of a router R:

Destination	Gateway	Owner	Netif
10.0.0.0/8	directly connected	-	eth1
11.0.0.0/8	directly connected	-	eth2
132.44.0.0/16	11.0.0.2	Static	eth2
193.204.161.0/24	10.0.0.1	RIP	eth1

Propose a network topology that is compatible with the routing table above

Quiz 6 -possible solution