Internetworking, Examination 2G1305
Date: August 20th 2003 at 9:00 – 13:00

SOLUTIONS

1. General (5p)

a) Place each of the following protocols in the correct TCP/IP layer (Application, Transport, Network or Link): BGP, IGMP, IPsec, PPP, HTTP, RTP (2p)
Application: BGP, HTTP, RTP
Network: IGMP, IPsec
Link: PPP

b) Typically, every layer in the TCP/IP protocol stack adds a header or trailer containing protocol-specific data. What is this called? (1p)
Encapsulation

c) By what mechanism is transport layer de-multiplexing accomplished in IP? (1p)
By the proto field in the IPv4 header

d) How is application layer de-multiplexing accomplished in UDP and TCP? (1p)
By the port field in the TCP and UDP header

2. Link Layer (5p)

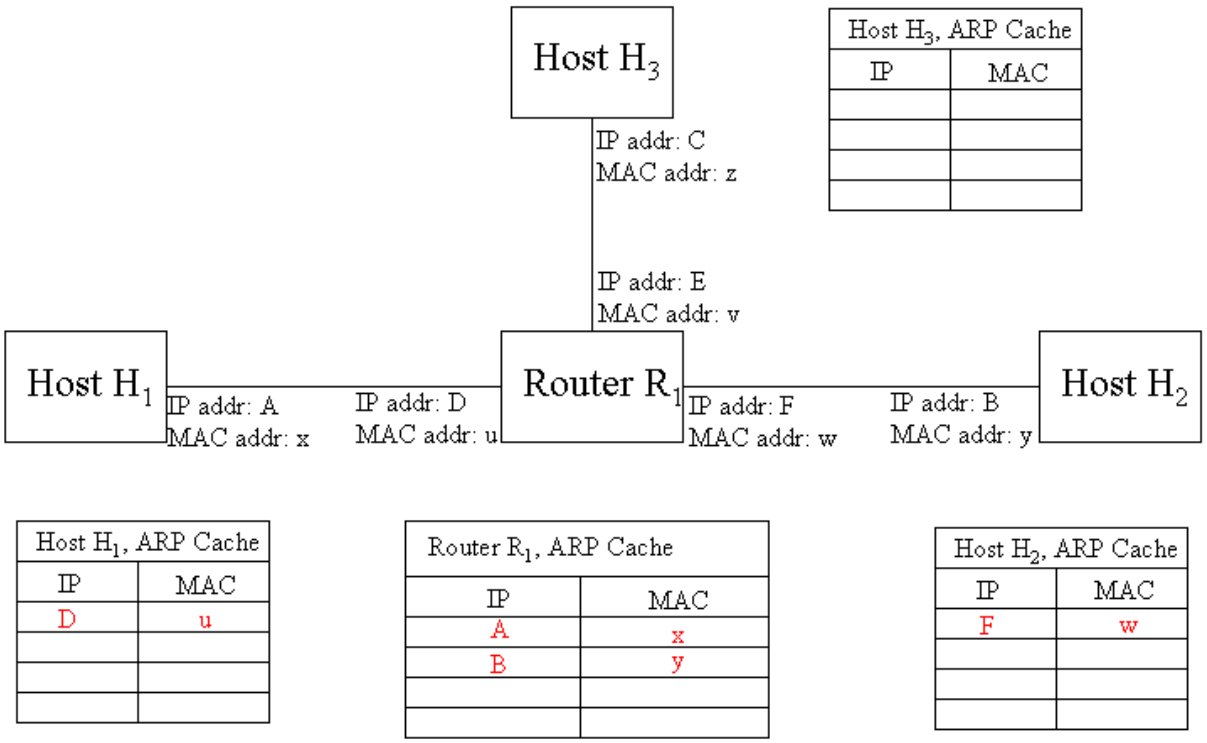

Host H_3, ARP Cache

IP	MAC

Host H_1, ARP Cache

IP	MAC
D	u

Router R_1, ARP Cache

IP	MAC
A	x
B	y

Host H_2, ARP Cache

IP	MAC
F	w

The figure above illustrates three hosts H_1, H_2 and H_3 running IPv4 over a routed network, connected by router R_1. The IP and MAC addresses of the hosts and the router's interfaces are given in the figure. The ARP caches of each host and the router are shown. Assume the ARP caches are initially empty, and that no packets have been sent yet. Now, host H_1 wants to send an IPv4 unicast datagram to host H_2.
Fill in the state of the four ARP caches as they will appear after the IPv4 unicast datagram has been delivered to host H_2, that is, after dynamic ARP resolution has been made. (5p)

3. IPv4 addressing (5p)

Subnet address	Subnet mask	Next Hop / Interface
67.16.120.0	255.255.252.0	Router R1
67.16.124.0	255.255.252.0	Interface A
67.16.123.64	255.255.255.192	Interface B
default		Router R2

Suppose an IPv4 router has built up the routing table shown above. The router can deliver directly over interfaces A or B or it can forward packets to next-hop routers R1 or R2. Describe what the router does with a packet addressed to each of the following destinations:

a) 67.16.123.54 (1p)
Forward to Router R1
b) 67.16.123.105 (1p)
Forward to Interface B
c) 67.16.121.65 (1p)
Forward to Router R1
d) 67.16.127.34 (1p)
Forward to Interface A
e) 67.16.128.89 (1p)
Forward to Router R2

4. IPv4 header (5p)

a) The IPv4 header has two length fields. Which are they and what are the purposes and limitations of each? (2p)
Header length field: 4 bits expressing number of 32-bit words: limit is 60 bytes.
Total length field: 16 bits expressing byte-length of whole datagram. Limit is 64K.
b) What is the purpose of the TTL (time-to-live) field? How do routers handle the TTL field when forwarding a datagram? (2p)
The purpose of the TTL field is to limit the lifetime of a datagram in order to avoid loops.
Routers receiving a datagram with TTL > 1 decrements the TTL and forwards the datagram. If TTL <= 1, the router does not forward the datagram.
c) Is the IPv4 header checksum end-to-end or hop-by-hop? Why? (1p)
Hop-by-hop- the header may change on a hop-by-hop basis. For example, the TTL field decrements, and some options may be changed by a router.

5. TCP and UDP (5p)

a) Compare the performance of TCP's sliding window versus TFTP stop-and-wait protocol as follows. With a sliding window technique we transferred 32768 bytes in about 35 seconds across a link with RTT that averaged around 1.5 seconds. Calculate how long TFTP would take for the same transfer. TFTP sends 512 bytes per packet. (3p)
With 512 bytes per packet, the transfer would require 32768/512 = 64 packets. Using stop-and-wait means that the total time for the transfer would be 64 x 1.5 s = 96 seconds.

b) An application using UDP sends a datagram that gets fragmented into four pieces. Assume that fragments 1, 3, and 4 make it to the destination and fragment 2 is lost. The application times out and retransmits the datagram 10 seconds later with the same fragmentation. Assume that this time 1 and 3 are lost, and 2 and 4 reach the destination. Assume that the reassembly timer on the destination host is 60 seconds (meaning that fragment 1,3, and 4 are not yet discarded). Can the receiver reassemble the IP fragments? Explain your answer. (2p)
No, when the application retransmits, the IP datagram generated by the retransmission has a new identification field. Reassembly is only done for fragments with the same identification field.

6. TCP (5p)

a) What is the difference between *flow control* and *congestion control* in TCP? (2p)
The purpose with flow control is to make sure that the receiver is not overwhelmed with data by the sender. Flow control is a way for the receiver to control the sending TCP window size.
The purpose with congestion control is to avoid buffer overflow and packet loss in the network, i.e. routers between the end hosts. Congestion control is a way for the network to control the sending TCP window size.

b) Name the two phases of TCP congestion control and describe how the window size increases during these phases. (2p)
Congestion control involves 2 phases; Slow Start and Congestion Avoidance. Initially, CWND (congestion window) is set to 1 MSS (Maximum Segment Size). During Slow Start, CWND increases exponentially until it reaches a certain threshold value. Thereafter, Congestion Avoidance takes over. During this phase, CWND increases linearly. The linear growth continues until either the receiver initiated window is reached or congestion is reached.

c) TCP's retransmission time-out (RTO) has to be adaptive. On what parameter does the RTO normally depend? (1p)
The round-trip time (RTT).

7. Routing (5p)

a) BGP uses path-vector routing. Describe how path-vector enhances distance-vector. What is the advantage with the enhancement? (1p)
Path-vector adds the sequence of autonomous systems to pass (a path) in order to reach a destination network. With this information, loops that can occur in distance vector (e.g. count-to-infinity) can be avoided.

b) Which protocol does BGP use as its transport protocol? Name at least one advantage with using this transport protocol. (2p)
BGP uses TCP as transport protocol. With TCP, BGP can assume a reliable delivery, so that the complexity decreases in the protocol, and its implementations.

c) BGP uses four message types. Name these message types and explain briefly what each is used for. (2p)
OPEN – Create neighbourhood relationship
UPDATE – Announce and withdraw routes. Distribute path attributes
NOTIFICATION – Reporting error conditions
KEEPALIVE – Sent regularly between peers to ensure that the peers are alive

8. Applications (5p)

a) Why may delay jitter in an IP network be a problem for real-time applications, such as IP telephony? (2p)
Delay jitter means that packets within a session may arrive with different delays, which means that there is no guarantee that the time relationship between packets is preserved when sending data of an IP network.

b) Real-time applications communicating over IP can use RTP – Real-time Transport Protocol on top of UDP. There are two important features (header fields) of RTP that applications can use to deal with delay jitter. Name the features and explain how they are used? (3p)
The key features are sequence numbers and time stamps. The sequence number is mainly used to detect losses. The time stamps are used to place incoming packets in correct timing order.

9. Advanced IP networking (5p)

a) Describe the difference between traditional IP packet forwarding and packet forwarding in MPLS. (1p)
Traditional IP forwarding is based on the destination IP address and longest prefix match. In MPLS, forwarding is based on a label, which is a small fixed sized field.

b) Describe the difference between independent control and ordered control when it comes to assigning labels to forwarding equivalence classes (FECs) in MPLS. (4p)
Independent control (hop-by-hop routing) means that the LSR (Label Switching Router) assigns a label to every FEC it knows, i.e., at least each address prefix that has been obtained by the routing protocols.
Ordered control (explicit routing) means that label assignment is initiated by either ingress or egress LSR so that the packets of the FEC follow a specific path. This method gives more control of the traffic, but results in a slower LSP (Label Switched Path) establishment.

10. Multicast (5p)

a) IGMP is used for communicating group membership locally between a multicast router and locally connected hosts. In IGMP version 2, name two situations when a host can send an IGMP report to a router? (2p)
1. The host can send an IGMP report when the host joins a group (e.g., a process on the host joins a group that no other process on the host has joined previously).
2. The host can also send an IGMP report as an answer to an IGMP query from the multicast router.

b) With what main functionality does IGMP version 3 extend IGMP version 2? (1p)
In IGMP version 3, hosts can select senders as well as groups. This allows for source-specific multicast.

c) A system uses multicast routing. There are five groups and 100 sources. If a routing algorithm using source-based trees is used, what is the maximum number of different multicast delivery trees that can exist? (1p)
500. Every source can create one tree for every group.

d) Using the same system as in the previous exercise: How many multicast delivery trees may exist if a multicast routing protocol using shared trees is used? (1p)
5. Regardless of the number of sources, one delivery tree can be created per group.

11. IPv6 (5p)

a) Show the shortest form of the following IPv6 address:
2340:0000:0000:000F:7000:119A:A001:0000 (1p)
2340::F:7000:119A:A001:0

b) The IPv6 header includes a Class field and a Flow ID field. Briefly describe the purposes with these fields. (2p)
The Class field is used to distinguish between different priorities of IPv6 packets. The Flow ID field is used by the source to label packets for which it requires special handling by IPv6 routers.

c) Name the 3 strategies devised by IETF, for the transition from IPv4 to IPv6. (2p)
Dual-stack, Tunneling (automatic or configured), and Header Translation.

12. IPsec and Mobile IP (5p)

a) Mobile IP: Suppose a remote host sends a datagram to a mobile host. Describe the roles of the home agent and foreign agent in the delivery of the datagram? (2p)
The datagram goes from the foreign agent to the home agent. The home agent forwards to the foreign agent via tunnelling. The foreign agent delivers it to the mobile host. If co-located addressing is used, the home agent tunnels the datagram directly to the mobile host.

b) Mobile IP suffers from inefficiency in the data delivery phase. Why? (1p)
Datagrams delivered to the mobile must be relayed via the home agent (triangular delivery). In the severe case, datagrams need to pass twice over the same path (double crossing).

c) IPsec uses two modes: tunnel mode and transport mode. Outline in which situations the two different forms are preferably used. Also, describe the encapsulation in the two modes. (2p)
Tunnel mode is used when the two sub-networks are interconnected via a secure tunnel. Tunnel mode encapsulates the complete datagram and adds a new header. Transport mode is used to establish a secure connection between two end-hosts. Transport mode does not add a new header.