Internetworking Examination 2G1305
Date April 23rd, at 10:00 – 13:00

LCN IMIT KTH

- ?? ***No help material is allowed.***
- ?? *You may answer questions in English or Swedish.*
- ?? *Please answer each question on a separate page.*
- ?? *The grading of the exam will completed no later than May 14 2003.*
- ?? *After grading, the exams will be availablefor inspection at KTH STEX at Campus.*
- ?? *Deadline for written complaints is May 28 2003.*
- ?? *Course responsible is Olof Hagsand, phone 08-790 42 61.*

Your name:……………………………………………………………………………………

Your social security number (personnummer): ………………………………………………

Your major (utbildningslinje):…………………………………………………………..

Total Points: .……	Grade: ………..

Question	Answered	Potential points	Received points
1		5	
2		5	
3		5	
4		5	
5		5	
6		5	
7		5	
8		5	
9		5	
10		5	
11		5	
12		5	

Total	60	

1. General (5p)

a) Layering – Which are the four layers defined in the TCP/IP protocol stack? (1p)
Application, Transport, Network, Link layer

b) Place each of these protocols in the correct layer: IP, ICMP, OSPF, UDP, Ethernet, HTTP (2p)
OSPF, HTTP – Application layer
UDP – Transport layer
IP, ICMP – Network layer
Ethernet – Link layer

c) Define the term *encapsulation* (1p)
The technique of placing a data unit of a protocol in the data field portion of another protocol's data unit.

d) Define the term *multiplexing* (1p)
The technique of running several protocols on top of one lower level protocol.
E.g., running TCP and UDP on top of IP.

2. Link Layer / ARP (5p)

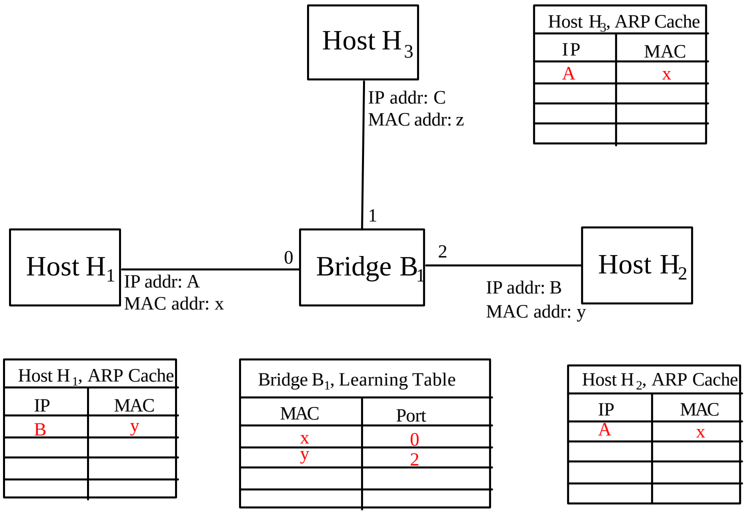

Host H_3, ARP Cache

IP	MAC
A	x

Host H_1, ARP Cache

IP	MAC
B	y

Bridge B_1, Learning Table

MAC	Port
x	0
y	2

Host H_2, ARP Cache

IP	MAC
A	x

The figure above illustrates three hosts H_1, H_2 and H_3 running IPv4 over an Ethernet bridged by bridge B_1. The IP and MAC addresses of the hosts, and the bridge port numbers are given in the figure. The ARP cache of each host is shown, as well as the learning table (station cache) of bridge B_1. Assume the ARP caches and the learning table is initially empty, and that no packets have been sent by either host. Now, host H_1 wants to send an IPv4 unicast datagram to host H_2.
Fill in the state of the three ARP caches and the learning table (station cache) of bridge B_1 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. IP Addressing (5p)

Subnet address	Subnet mask	Next Hop / Interface
189.139.29.0	255.255.255.128	Interface 0
189.139.29.128	255.255.255.128	Interface 1
189.139.30.0	255.255.255.128	R2
167.0.213.0	255.255.255.192	R3
default		R4

Suppose an IPv4 router has built up the routing table shown above. The router can deliver directly over interfaces 0 and 1, or it can forward packets to next-hop routers R2, R3 and R4. Describe what the router does with a packet addressed to each of the following destinations:

a) 189.139.29.32 (1p)
Forward to Interface 0

b) 167.0.213.91 (1p)
Forward to R4

c) 189.139.30.16 (1p)
Forward to R2

d) 167.0.213.16 (1p)
Forward to R3

e) 189.139.30.163 (1p)
Forward to R4

4. IP Fragmentation (5p)

A message from the application consists of 3500 bytes. The message is encapsulated as a UDP datagram over IPv4 (no options). The message is then transferred from a sender to a receiver over (exactly) two local area networks that both use 14 bytes of headers (trailers omitted in this exercise). The first local area network has an MTU of 1500 and the second has a MTU of 1000 bytes. A router is situated between the two local area networks. How many bytes (headers included) will arrive to the link level interface at the receiver? Assume that no PATH MTU discovery is performed: the sender fragments according to the MTU of the first link. An IP header is 20 bytes (without options), and a UDP header is 8 bytes. (5p)

3 frames on the first link:

1) 14 + 20 + 8 + 1472 (Link header, IP header, and UDP header)
2) 14 + 20 + 1480 (Link header and IP header)
3) 14 + 20 + 548 (Link header and IP header)

5 frames on the second link:

1) 14 + 20 + 8 + 972
2) 14 + 20 + 500
3) 14 + 20 + 980
4) 14 + 20 + 500
5) 14 + 20 + 548

No of bytes received at link layer: 5*14 + 5*20 + 8 + 3500 = 3678 bytes

5. TCP Flow Control (5p)

Suppose that a 100Mbps link is being set up between earth and a communication satellite at an altitude of 36,000 km. An image file of 20 MB should be transferred from station A to station B on earth. Assume the speed of light is 300,000,000 m/s.

a) Calculate the minimum RTT for the link. (2p)

Min RTT = 2*up-link delay + 2*down-link delay =
2*36,000,000/300,000,000 + 2*36,000,000/300,000,000 = 0.48s

b) Calculate the delay * bandwidth product for the link. (1p)

bw*RTT = 100 Mbps * 0,48s = 48,000,0000 bits

c) Explain the meaning of the bandwidth-delay product. (2p)

The bw-delay product is the TCP window size that should be kept for optimal throughput.

6. TCP (5p)

Important algorithms included in TCP.

a) Retransmission timeout (RTO) in TCP is based on the round trip time (RTT). Karn's algorithm is used to solve a certain problem regarding the calculation of RTO. What problem is solved by Karn's algorithm and how does the algorithm work? (3p)
Suppose a segment is not ACKed in time and gets retransmitted. When the ACK finally arrives, there is no way to tell if it is the original segment or the retransmitted segment that is ACKed. Thus, we cannot trust the measured value of RTT.
Karn's algorithm:
1) Do exponential back-off by doubling the RTO
2) Don't consider RTT of the retransmitted segment
3) Reuse the backed off RTO for following transmissions
4) Don't update RTT until a segment gets ACKed without need for retransmission

b) An application that generates data too slowly can cause a silly window syndrome. Briefly describe how Nagle's algorithm deals with this problem. (2p)
Nagle's algorithm states that a sending TCP can have only one outstanding (unacknowledged) tinygram at a time. While waiting, sender data is accumulated and gets sent when ACK is received or when MSS can be filled.

7. DHCP (5p)

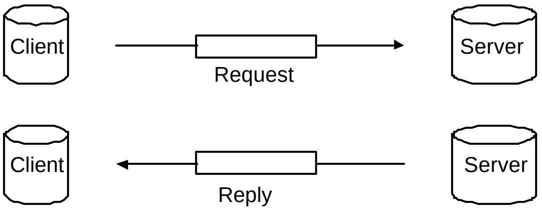

The figure above illustrates a DHCP client sending an initial request (to find out its IP address) to a DHCP server and getting a reply back.

a) Which transport protocol is used for the DHCP request and reply communication? (1p)
UDP

b) Which IP addresses are used as source and destination for the request? (2p)
Src IP address: 0.0.0.0, Dst IP address: 255.255.255.255

c) Which IP addresses are used as source and destination for the reply? (2p)
Src IP address: server's IP address, Dst IP address: 255.255.255.255 or client's new IP address.

8. IP routing: Count to Infinity (5p)

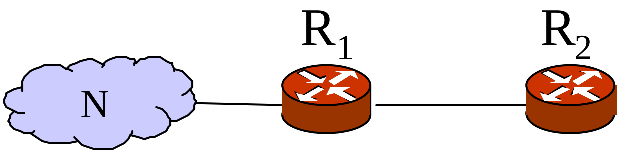

In the figure above, network N is reachable via routers R_1 and R_2. Suppose R_1 and R_2 runs RIP. If the link between R_1 and N is broken, the count-to-infinity problem may occur.

a) What is the count-to-infinity problem and why does it occur? (2p)
When the link is broken, R1 will set the cost to N to infinity. The count-to-infinity problem occurs if R2 sends an update to R1 before R1 sends it to R2. Then, R1 will believe that R2 has a valid route to N and will update its route to N to go via R2, and a loop will occur. The problem will resolve eventually after R1 and R2 have sent periodical updates, each time incrementing the cost to N, until it reaches infinity.

Suppose we implemented the following functionalities in RIP:

b) How can the problem be avoided by *triggered updates*? (1p)
When the link goes down, and N becomes unreachable, R1 immediately sends sets the cost to infinity in its routing table and sends this to R2, which in turn changes its route to infinity.

c) How can the problem be avoided by *hold-downs*? (1p)
Ignore information about a route for a fixed period of time following the message that a network is unreachable. In our case, R1 would ignore R2:s message for the hold-down time, instead, R1's information will eventually spread to R2.

d) How can the problem be avoided by *split horizon*? (1p)
If a router receives route update information from an interface, then it may not send back updated information on that interface. This means that R2 will not send an update to R1 on L2 and R1 will not be fooled into updating its entry to 3.

9. IP routing (5p)

a) What is an autonomous system? (1p)
An autonomous system is a group of networks and routers controlled by a single administrative authority.

b) Explain the difference between interior and exterior routing protocols. (1p)
An interior routing protocol is a dynamic routing protocol that handles routing within an autonomous system. An exterior routing protocol handles routing between autonomous systems.

c) What is a static route and a dynamic route? (1p)
Static routes are configured manually. Dynamic routes are inserted by a dynamic routing protocol, that is, they are derived, computed, or imported from an external node, by an automatic mechanism.

d) Explain at least three differences between RIP and OSPF (2p)
1) RIP uses distance vector (Bellman-Ford), OSPF uses Link-state routing (Dijkstra).
2) RIP can be only used on a smaller scale - one area and max 15 hop-count. OSPF has an hierarchical area structure and does not have the hop limitation.
3) RIP uses UDP, OSPF is implemented directly on IP
4) OSPF supports load balancing (if forwarding supports it) RIP computes single

routes only.
5) OSPF supports authentication, RIP does not.

10. IP Multicasting (5p)

a) What protocol handles the signalling of group memberships between a multicast router and hosts? (1p)
IGMP

b) Multicast delivery trees: Explain the difference between a shared tree and a source-based tree. (2p)
A shared tree uses the same delivery tree within a domain, for all senders. Senders delivers the multicast IP datagram to the root, or a branch, of the tree. A source-based tree builds a tree for each specific sender.

c) In multicast routing protocols based on distance-vector the multicast routers perform a *reverse lookup*. Explain how this works. (2p)
The router uses the *source* of an IP multicast datagram to make a lookup in the unicast routing table. The result should match the interface the datagram was received on.

11. IP QoS (5p)

a) Briefly describe the procedure used by RSVP to set up a reservation in the network. What types of messages are sent between sender/receiver and what information is communicated in these messages? (3p)
Senders send RSVP PATH messages, carrying classification info and TSpec (Traffic Specification). Receivers send RSVP RESV messages carrying RSpec (Resource requirements) back along the path.

b) RSVP maintains soft state in the routers along the path between sender and receiver(s). Explain what this means. (2p)
States in the routers are created by PATH and RESV messages, and the states will automatically time out unless they are periodically refreshed by PATH and RESV messages.

12. Mobile IP (5p)

a) Describe the triangular delivery and two crossing problem that appears in Mobile IP. (2p)

Triangular delivery: When a mobile node M communicates with a destination node A, M sends datagrams directly to the A, but A sends to the home agent which relies them to M.

Two crossing: if M and A are close to each other, every datagram will essentially cross the same network twice, ie, it is a special case of triangular delivery.

b) Discovery, Registration and Tunneling are three basic capabilities of Mobile IP. Describe them briefly. (3p)

Discovery: A mobile node uses a discovery procedure to identify home agents and foreign agents.

Registration: A mobile node uses an authenticated registration procedure to inform its home agent of its care-of address.

Tunneling: Used by the home agent to forward IP-datagrams to a care-of address.