

**University of Toronto
Faculty of Applied Science and Engineering**

Final Exam, December 2007

ECE 461: Internetworking
Examiner: J. Liebeherr

- Exam Type: A
- Calculator: Type 2

- There are a total of 10 problems.
- Note the binary-decimal conversion table on the last page.
- Write your solutions into an answer book. Make sure your name is on the answer book.

Problem 1. (15 Points)

Consider the network shown in Figure 1 with three hosts (HostA, HostB, HostC), one router (Router1), and two Ethernet segments. The figure includes the network configuration, IP addresses, netmasks, and MAC addresses. For simplicity, the MAC address has a length of two bytes.




Figure 1.

- a. (5 Points) In addition to the addresses and netmasks, describe the minimally required configuration on HostA, HostB, HostC, and Router1, so that all three hosts can successfully send IP datagrams to each other. (Provide your description in words. Do not provide configuration commands.)

Use the configurations of your answer in (a) for your answer to the following two questions.


- b. (5 Points) Describe in detail the ARP and ICMP packets which are transmitted on the Ethernet segments when HostA executes the command “ping 10.1.2.188”. (Assume that the ARP tables of all hosts and the router are initially empty.)
- c. (5 Points) Describe in detail the ARP and ICMP packets which are transmitted on the Ethernet segments when HostB executes the command “ping 10.1.3.100”. (Assume that the ARP tables of all hosts and the router are initially empty.)

Note:

- For each packet, you need to specify the source and destination MAC addresses, the source and destination IP addresses (if applicable), and a description of the packet content.
- Recall that a successful ping involves two ICMP packets: an ICMP Echo Request from the host that issues the ping, and an ICMP Echo Reply from the host which is queried.

Problem 2. (10 Points)

In the network topology below with seven routers, the objective is that traffic from A to E traverses the path A-C-D-E and traffic from B to E traverses the path B-C-F-G-E.



- (5 Points) Describe a key problem in achieving this objective with Internet routing protocols.
- (5 Points) Propose a method to overcome the problem.

Problem 3. (10 Points)

- (5 Points) Explain the concept of “soft state” as a design principle encountered in numerous protocols for the Internet. Describe the rationale for using soft state in protocols. Provide three different examples where the concept of soft state is realized in a protocol.
- (5 Points) Explain the concept of “exponential backoff” found in Internet protocols that need to determine a timer value. Describe the purpose of exponential backoff. Provide three different examples where exponential backoff is realized in a protocol.

Problem 4. (10 Points)

Suppose you have a program available that permits you to set the header fields and payload of IP packets and ARP packets, and transmit these packets. This program can be used to stage the following malicious attacks:

Attack 1: Modify routing table entries at a remote host. (Without running a routing protocol such as RIP or OSPF).

Attack 2: Redirect traffic that is destined to a host with a given IP address to your host.

Attack 3: Stage a denial of service attack, where a large number of hosts send traffic to a target host.


For each of the attacks:

- Describe how the attack is staged. Describe which packets need to be transmitted and how the headers need to be set to stage the attack.
- Describe the limitations of the attack.

c. Describe how to protect a host against the attack.

Problem 5. (10 Points)

Consider the statically routed IP network in the figure below with three routers. The interfaces of the routers are indicated as I1, I2, and I3.



The routing tables are as follows:

Router 1:	Destination	Interface
	128.100.11.32/27	I2
	128.100.11.0/24	I3
	0.0.0.0/0	I1

Router 2:	Destination	Interface
	128.100.11.32/27	I2
	0.0.0.0/0	I1


Router 3:	Destination	Interface
	128.100.11.16/28	I1
	128.100.11.0/24	I3
	0.0.0.0/0	I2

The routing tables are misconfigured and will result that packets in a certain range of IP addresses will loop in a circle.

- (5 Points) Specify the full range of IP addresses that will loop. Specify which routers the looping packets will traverse.
- (5 Points) You are allowed to delete one line of the routing tables in one of the routers. Your objective is to eliminate the loop, while making sure that hosts on the IP network 128.100.11.0/24 can send traffic to and from the rest of the Internet.

Problem 6. (10 Points)

Consider the interconnected IP networks as shown in the Figure.




Assume that the routers run a routing protocol that minimizes the hop count to a destination network (similar as in RIP). If a router has multiple routes with the same hop count to a destination, the router breaks the tie by picking the route which maximizes its opportunities for route aggregation in its routing table. If that leaves unresolved ties, the remaining ties are broken arbitrarily.

- a. (5 Points) Provide the routing table for Router 1 that follows the stated routing objectives. Provide an explanation for each successful route aggregation. Provide an explanation why the remaining routing table entries cannot be further aggregated.

- b. (5 Points) Consider a routing strategy which selects routes solely on the basis of maximizing opportunities for route aggregation. Such a routing strategy should result in the smallest possible routing tables. However, there are problems with realizing this routing strategy. Provide two distinct potential problems of such an approach to routing.

Problem 7. (10 Points) Consider the network topology in the figure below. There are four hosts (H1, H2, H3, and H4) and six routers (R1, R2, R3, R4, R5, and R6). Links have a cost as indicated in the figure. The routers run a shortest path unicast routing protocol that considers link costs (e.g., OSPF). The routers also run a multicast routing protocol that uses Reverse Path forwarding with Flood-and-Prune (e.g., PIM-DM).



Suppose that H1 is the sender in the multicast group, and hosts H2, H3, and H4 are members of the multicast group.

- (4 Points) Indicate the interfaces where each router receives and forwards packets in the flooding stage. Indicate which of the forwarded packets get discarded at the sender.
- (3 Points) Indicate the Prune messages that are transmitted in the above topology and draw the distribution tree that is a result of the flood-and-prune algorithm.
- (3 Points) When H4 leaves the multicast group, describe what will happen in the following round of flood-and-prune.

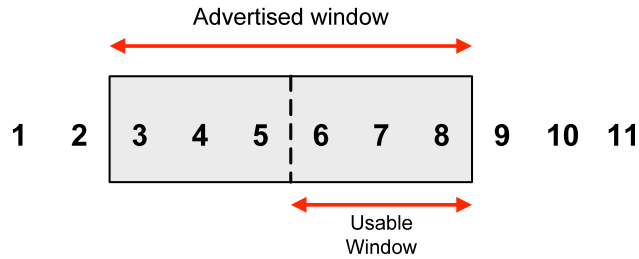
Problem 8. (10 Points) The spanning tree protocol (STP) for bridges is known to adapt only slowly after a failure of a link or a bridge. The following are two suggested modifications to the spanning tree protocol which improve the convergence after a link or bridge failure:

- **Uplink Fast:** Each bridge keeps track of all received BPDUs that provide an alternate path to the root bridge. If a bridge detects a failure of the link to the bridge reachable via its root port, the root port is switched to the port on which the BPDU with the next lowest cost was received.
- **Backbone Fast:** Consider a Bridge A that has Bridge R as root bridge. Suppose that Bridge A receives BPDUs from a designated bridge (say, Bridge X), advertising Bridge R as root bridge. If Bridge X all of a sudden starts to advertise an inferior Bridge S as root bridge ($S > R$), then Bridge A initiates the Backbone Fast procedure as follows. Bridge A sends Root Link Query (RLQ) request messages to determine if there still exists a path from Bridge A to Bridge R. Bridge A sends an RLQ request on all its ports, except its designated ports and except the port leading to Bridge X. Bridge A includes in each RLQ request the ID of its root, R. Any bridge that receives an RLQ request immediately answers with an *RLQ response* if (1) it knows it has lost connection to root R (i.e., it has a different root bridge), or (2) if it is the root (i.e., it is Bridge R). If a bridge that receives an RLQ request does not send an RLQ response, it forwards the RLQ request to its root port. The sender of an RLQ response puts the ID of its root bridge in the PDU. RLQ responses are flooded on designated ports. Once Bridge A receives RLQ responses for each transmitted RLQ request, and no response specifies Bridge R as root bridge, Bridge A initiates a re-computation of the spanning tree.
 - a. (4 Points) Describe a scenario where Uplink Fast improves the convergence of the STP protocol after a link or bridge failure. Explain the improvement over the standard STP protocol.


The following questions relate to the Backbone Fast improvement:

- b. (2 Points) Describe the valid conclusions that Bridge A can draw if it receives a BPDU from the designated Bridge X, which starts to advertise an inferior root bridge.
- c. (2 Points) Explain why Bridge A does not send RLQ requests to its designated ports and the port leading to Bridge X.
- d. (2 Points) Describe the valid conclusions that Bridge A can draw if none of the RLQ responses specifies Bridge R as root bridge. Describe how the re-computation of the spanning tree by Bridge A results in an improvement over the standard STP algorithm.


Problem 9. (5 Points) Consider the state of a sliding window at the sending side of a TCP connections as shown in the figure. (Each number corresponds to one byte).



(a) (3 Points) From the initial figure, which actions by the sender and/or the receiver bring the sliding window into the following state:



(b) (2 Points) From the initial figure, which actions by the sender and/or the receiver bring the sliding window into the following state:



Problem 10. (10 Points)

Explain how the following types of DNS queries work. Give an example for each type:

- a. Inverse queries
- b. Iterative queries
- c. Recursive queries