

**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, the IP addresses, the netmasks, and the 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 type.
- 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.

a)

- Default gateway on HostA must be set to 10.1.1.1
- Note: HostB and HostC do not need a default gateway.
- Proxy ARP must be installed on both interfaces of Router 1

- IP forwarding must be enabled
- There must be a routing table entry:
 10.1.0.0/16 Link #1
 10.1.3.0/24 Link #2

b)

ICMP Echo Request:

1. Ethernet Header(Src=HostA, Dest=broadcast)
 ARP who-has 10.1.1.1 tell 10.1.1.88
2. Ethernet Header (Src= Router1@Link#1, Dest = HostA),
 ARP reply 10.1.1.1 is-at fd:e3
3. Ethernet Header (Src= HostA, Dest = Router1@Link#1),
 IP Header (Src= HostA, Dest = HostB),
 ICMP Echo Request
4. Ethernet Header(Src=Router1@Link#1, Dest=broadcast)
 ARP who-has 10.1.2.188 tell 10.1.1.1
5. Ethernet Header (Src= HostB, Dest = Router1@Link#1),
 ARP reply 10.1.2.188 is-at b6:da
6. Ethernet Header (Src= Router1@Link#1, Dest = HostB),
 IP Header (Src= HostA, Dest = HostB),
 ICMP Echo Request

ICMP Echo Reply:

1. Ethernet Header(Src=HostB, Dest=broadcast)
 ARP who-has 10.1.1.88 tell 10.1.2.188
 /* Note that HostB believes that Host A is on the same subnetwork.
 Therefore, it will try to send the packet directly to HostA*/
2. Ethernet Header (Src= HostA, Dest = HostB),
 ARP reply 10.1.1.88 is-at f4:0a
3. Ethernet Header (Src= HostB, Dest = HostA),
 IP Header (Src= HostB, Dest = HostA),
 ICMP Echo Reply

c)

ICMP Echo Request:

1. Ethernet Header(Src=HostB, Dest=broadcast)
 ARP who-has 10.1.3.100 tell 10.1.2.188

2. Ethernet Header (Src= Router1@Link#1, Dest = HostB),
ARP reply 10.1.3.100 is-at fd:e3 /* **This is a Proxy ARP** */
3. Ethernet Header (Src= HostB, Dest = Router1@Link#1),
IP Header (Src= HostB, Dest = HostC),
ICMP Echo Request
4. Ethernet Header(Src=Router1@Link#2, Dest=broadcast)
ARP who-has 10.1.3.100 tell 10.1.3.1
5. Ethernet Header (Src= HostC, Dest = Router1@Link#2),
ARP reply 10.1.3.100 is-at c8:f1
6. Ethernet Header (Src= Router1@Link#1, Dest = HostC),
IP Header (Src= HostB, Dest = HostC),
ICMP Echo Request

ICMP Echo Reply:


7. Ethernet Header(Src=HostC, Dest=broadcast)
ARP who-has 10.1.3.188 tell 10.1.3.100
8. Ethernet Header (Src= Router1@Link#2, Dest = HostC),
ARP reply 10.1.3.188 is-at5a:ea /* **This is another Proxy ARP** */
9. Ethernet Header (Src= HostC, Dest = Router1@Link#2),
IP Header (Src= HostC, Dest = HostB)
ICMP Echo Reply

/* The next two are not needed if part (b) was run previously*/

7. Ethernet Header(Src=Router1@Link#1, Dest=broadcast)
ARP who-has 10.1.2.188 tell 10.1.1.1
 8. Ethernet Header (Src= HostB, Dest = Router1@Link#1),
ARP reply 10.1.2.188 is-at b6:da
-
10. Ethernet Header (Src= Router1@Link#1, Dest = HostB),
IP Header (Src= HostC, Dest = HostB),
ICMP Echo Reply

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 should traverse the path B-C-F-G-E.



- d. (5 points) Describe a key problem in achieving this objective in an IP network.
- e. (5 points) Propose a method to overcome the problem.

Solution:

The arrangement of routers in Figure 3 is referred to as "the fish tail topology" (the fish's tail is to the left). It is the classical example to describe the limitations of the destination-based routing approach in IP networks.

The problem can only be solved by fundamentally changing the behavior or IP routing:

- source routing (but then the route has to be known at A and B)
- consider source and destination address when performing routing
- rely on a connection-oriented approach.

Problem 3. (10 points)

- a. (5 Points) Explain the concept of “soft state” as a design principle encountered in numerous protocols for the Internet. Describe the purpose of soft state protocols. Provide three different examples where the concept of soft state is realized in a protocol.

- b. (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 the concept of soft state 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:

- a. 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.
- b. Describe the limitations of the attack.
- c. Describe how to protect a host against the attack.

(4 Points for Attack 1, 3 Points each for Attack 2 and Attack 3)

Multiple answers are possible. Here are the most obvious answers:

Attack 1:

- a. Attack: Send ICMP Route Redirect packets
- b. Limitations: Must be on the same IP network or only effects routing caches (in Linux)
- c. Prevention: Do not react to ICMP Route Redirect or clear route cache frequently

Attack 2:


- a. Attack: Send ARP Replies
- b. Limitations: Must be on the same broadcast network.
- c. Prevention: Use static ARP entries, match ARP replies with previously sent ARP requests

Attack 3:

- a. Attack: Send ICMP Echo Request with broadcast IP address as destination and target host as source.
- b. Limitations: Needs to have access to a large number of hosts.
- c. Prevention: Do not respond to ICMP Echo Requests with broadcast address

Problem 5. (10 points)

Consider the statically routed IP network in the figure below with three routers. Router 3 is connected to an Ethernet network. 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	I2
128.100.11.0/24	I3
0.0.0.0/0	I1

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

- a. (3 points) Specify the full range of IP addresses that will loop. Specify which routers the looping packets will traverse.

128.100.11.16 to 128.100.11.31


The loop is R1- R3-R2-R1

- b. (2 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.

Delete the line 128.100.11.16/28 from Router 3

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 the destination (similar as RIP). If there are multiple routers with the same hop count, a router breaks the tie by picking the route which maximizes its opportunities for route aggregation. If that leaves unresolved ties, the remaining ties are broken arbitrarily.

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

Network Address	Interface

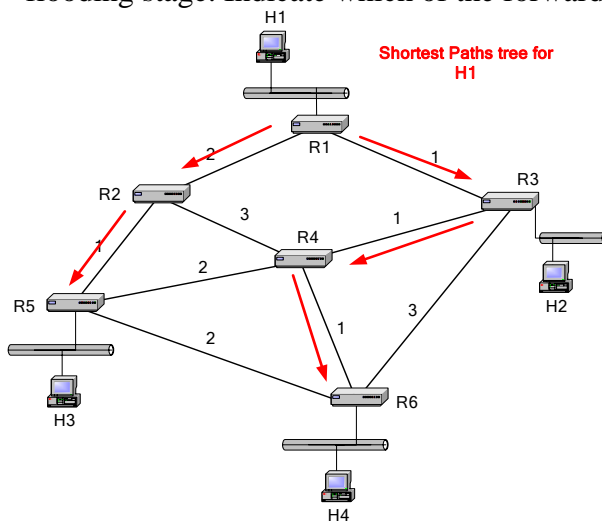
- b. (5 points) Consider a routing strategy which selects routes solely on the basis of maximizing opportunities for route aggregation. Such a routing strategies should result in the smallest possible routing tables. However, there are problems with realizing this routing strategy. Provide at least two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

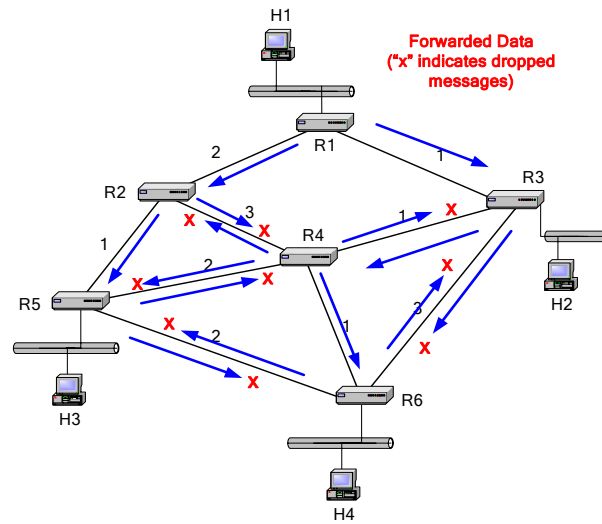
Problem 7. (10 points) Consider the network topology in the figure below. There are four hosts (H1, H2, H3, and H4) and six routers (R1, R2, R3, R4, R5, and R6). Links have a cost as indicated in the figure. The routers run a shortest path unicast routing protocol that considers the link costs (e.g.,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



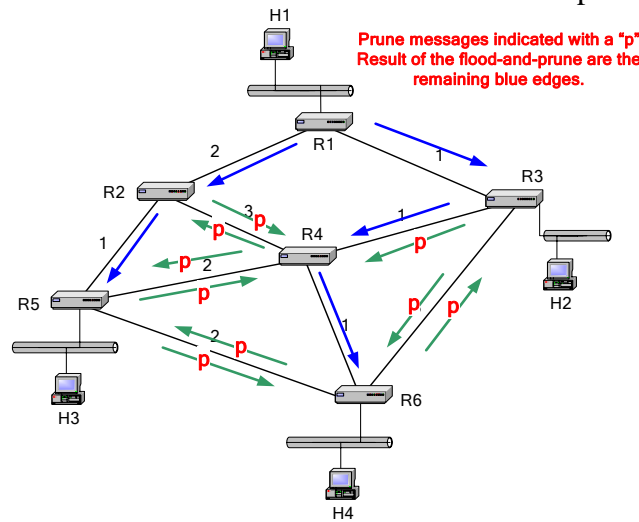
Suppose that H1 is the sender in the multicast group, and hosts H2, H3, and H4 are members of the multicast group.

- a. (3 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.





- b. (2 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.



- c. (3 points) When H4 leaves the multicast group, describe what will happen in the following round of flood-and-prune.

In the next round R6 will send a prune message to R4. As a result, R4 will not forward packets from H1 to R6.

Problem 8. (10 points) The spanning tree algorithm for bridges (LAN switches) is known to adapt only slowly after a failure of a link or a bridge. The following are two improvements to the spanning tree protocol to improve the convergence of the STP protocol after a link or bridge failure:

- **UplinkFast:** Each bridge keeps track of all received BPDUs that provide an alternate connection to the root bridge. If the root port fails (that is, the node detects a failure of the link to the bridge reachable via the 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 with Bridge R as root bridge. When Bridge A receives a BPDU from a designated bridge (say, Bridge X), which specifies Bridge S as root bridge, and $S > R$, 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 to all ports, except the port leading to Bridge X and except all ports where Bridge A is the designated bridge. The purpose of sending RLQ requests is to determine if there still exists a path from Bridge A to Bridge R. The RLQ request includes the ID of root R. Any bridge that receives a RLQ request immediately answers with a *RLQ response* if (1) it knows it has lost connection to root R (i.e., it has a different root bridge), or if it is the root (i.e., it is Bridge R). If a bridge does not send a RLQ response, it forwards the RLQ request to its root port. The sender of the 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. (5 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. (3 Points) Describe the valid conclusions that Bridge A can draw if it receives a BPDU from Bridge X, which specifies an inferior root bridge.
- c. (3 Points) Explain why Bridge A does not send RLQ requests to ports leading to Bridge X and to ports where Bridge A is the designated bridge.
- d. (3 Points) Describe the valid conclusions that Bridge A can draw if all RLQ responses specify a different root bridge. Describe how the re-computation of the spanning tree by Bridge A results in an improvement over the standard STP algorithm.

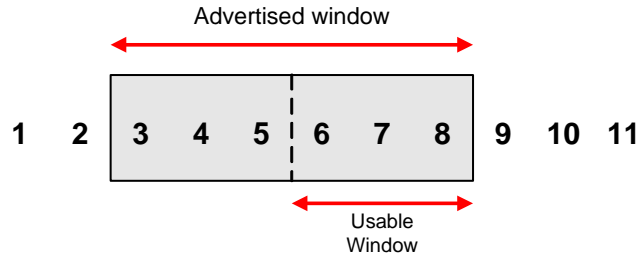
Description of Uplink Fast:

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a0080094641.shtml

Description of Backbone Fast:

http://www.cisco.com/en/US/tech/tk389/tk621/technologies_tech_note09186a00800c2548.shtml

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 or the receiver bring the sliding window into the following state:



Sender transmits a 1-byte long segment.
 A segment with (AckNo=4, Window size = 6) is received.

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



This is not feasible.

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

Binary-Decimal Conversion

Last 4 bits→	0000	0001	0010	0011	0100	0101	0110	0111	1000	1001	1010	1011	1100	1101	1110	1111
First 4 bits↓																
0000	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
0001	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	31
0010	32	33	34	35	36	37	38	39	40	41	42	43	44	45	46	47
0011	48	49	50	51	52	53	54	55	56	57	58	59	60	61	62	63
0100	64	65	66	67	68	69	70	71	72	73	74	75	76	77	78	79
0101	80	81	82	83	84	85	86	87	88	89	90	91	92	93	94	95
0110	96	97	98	99	100	101	102	103	104	105	106	107	108	109	110	111
0111	112	113	114	115	116	117	118	119	120	121	122	123	124	125	126	127
1000	128	129	130	131	132	133	134	135	136	137	138	139	140	141	142	143
1001	144	145	146	147	148	149	150	151	152	153	154	155	156	157	158	159
1010	160	161	162	163	164	165	166	167	168	169	170	171	172	173	174	175
1011	176	177	178	179	180	181	182	183	184	185	186	187	188	189	190	191
1100	192	193	194	195	196	197	198	199	200	201	202	203	204	205	206	207
1101	208	209	210	211	212	213	214	215	216	217	218	219	220	221	222	223
1110	224	225	226	227	228	229	230	231	232	233	234	235	236	237	238	239
1111	240	241	242	243	244	245	246	247	248	249	250	251	252	253	254	255