# University of Toronto
# Faculty of Applied Science and Engineering

## Final Exam, December 2010

ECE 461: Internetworking
Examiner: J. Liebeherr

- Exam Type: B
- Calculator: Type 2

- There are a total of 10 problems.
- Note the information about header formats and a binary-decimal conversion table on the last pages.
- Write your solutions into an answer book. Make sure your name is on the answer book.
- Do not write answers in this handout.

**Problem 1.** *(10 points)*

Below is the traffic capture of a packet in hexadecimal notation. The capture consists of an Ethernet II header, followed by an IP header, followed by a TCP header. (Hint: Each digit corresponds to 4 bits.)

```
00 13 72 30 84 8d 00 1f f3 56 8a 9f 08 00 45 00
00 34 2c 9e 40 00 40 06 1a bb 8e 96 eb 21 8e 96
eb 1c c4 e4 00 50 7a 4f 7f 03 ef 7c db 51 80 11
ff ff f3 91 00 00 01 01 08 0a 35 b9 51 3b c1 67
d1 28
```

Note: The format of headers is provided on attached sheets.

a. *(2 points)* Provide the value of the Source IP Address and the Destination IP address (Use dotted decimal notation !)

b. *(2 points)* How can you tell that the Ethernet header is followed by an IP header? How can you tell that the IP header is followed by a TCP header? Include the relevant information from the captured data in your answer.

c. *(2 points)* Provide the values of the TCP source and destination port numbers (as decimal numbers). Which application protocol has sent the above TCP segment?

d. *(2 points)* The above segment has the ACK bit set. After receiving the segment, describe the range of sequence numbers that the receiver of the segment is allowed to transmit. (Sequence numbers can be given as hexadecimals).

e. *(2 points)* How can you tell that there is no payload following the TCP header? Since there is no payload in the above packet, what is the purpose of this packet?

```
(a)  Src: 142.150.235.33,
     Dst: 142.150.235.28

(b)  Ethernet Header: Type: IP (0x0800)
     IP Header: Protocol: TCP (0x06)

(c)  Src Port: 50404, Dst Port: 80,
     HTTP Protocol (Port 80 is the well-known port for HTTP servers)

    Answer of "Web traffic" is correct, but must be justified with port
    number.
```

(d) Allowed range of sequence numbers that sender is allowed to send is AckNo, … AckNo+Window size -1

```
AckNo:          ef 7c db 51
Window size: ff ff

Range of sequence number:
Hexadecimal:     from ef7cdb51 to ef7ddb50
Decimals:    from 4017937233 to 4018002768
```

(e)

```
IP datagram total length = 52  bytes
IP header size  = 5 x 4 = 20  bytes
TCP header size = 8 x 4 = 32 bytes
```

→ IP and TCP header lengths equal to IP datagram size, thus no header

Purpose of packet is that of a control packet, indicated by a flag being set. Correct solution does need not identify which flags are set, but wrong guesses should lead to point deduction:

```
Flags: 0x11 (FIN, ACK)
    0... .... = Congestion Window Reduced (CWR): Not set
    .0.. .... = ECN-Echo: Not set
    ..0. .... = Urgent: Not set
    ...1 .... = Acknowledgement: Set
    .... 0... = Push: Not set
    .... .0.. = Reset: Not set
    .... ..0. = Syn: Not set
    .... ...1 = Fin: Set
```

**Problem 2.** *(10 points)*
Consider the following figure with two Ethernet segments where

- H1 and H2 are hosts;
- R1 is a router;
- X1 is a bridge.

The IP addresses of the interfaces are indicated in the figure.
Proxy ARP is not available, i.e., cannot be used.

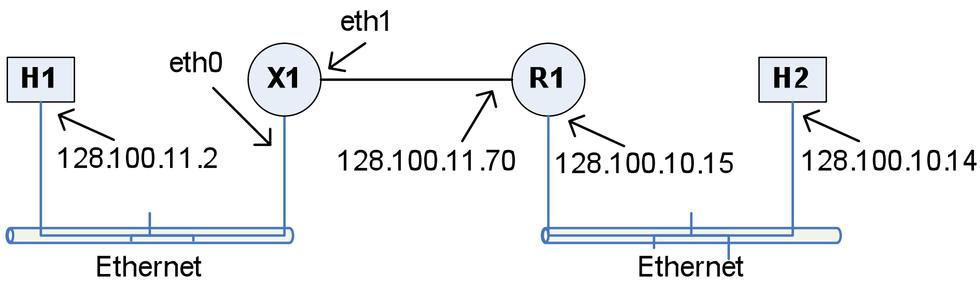


a. *(5 points)* Select subnetmasks for H1, H2, and R1. The subnetmasks should be selected so that the (extended) network prefix has maximal length.

b. *(5 points)* Suppose the configuration X1 is set to an IP router. Provide the required changes to the IP configuration in the network, so that H1 and H2 can communicate. The IP addresses of H1, H2, and R1 cannot be changed. The configuration includes:
   - IP addresses and subnetmasks for interfaces eth0 and eth1 at X1.
   - Subnetmasks at H1, H2, and R1 (If you choose the same as in part (a), state so).
   - Routing table entries at X1 and R1.

| | IP address | Binary (bits 1-16) | Binary (bits 17-32) |
|---|---|---|---|
| H1 | 128.100.11.2 | 10000000.01100100 | **00001011.0**0000010 |
| R1 left | 128.100.11.70 | 10000000.01100100 | **00001011.0**1000110 |
| R1 right | 128.100.10.15 | 10000000.01100100 | **00001010.000**01111 |
| H2 | 128.100.10.14 | 10000000.01100100 | **00001010.000**01110 |

```
(a)
For H1 and R1 (left):  max. prefix length is 25 bits  "255.255.255.128"
For R1 (right) and H2: max. prefix length is 27 bits  "255.255.255.252"

For the /27 network, a longer prefix is invalid since R1 (right) would
collide with the broadcast address of that longer prefix. With /27 the host
number address of R1 (right) is not "all  1"s.

The IP addresses of the networks are:
128.100.11.0/25
128.100.10.0/27
```

(b) There are multiple valid solutions.
It is crucial to create a new subnet.

- there is no need to change the subnetmasks of R1 (right) and H2.
- There is no requirement that the subnetmasks have maximal length
- The following solution makes an even split of IP addresses in 128.100.11.128/25

Network of H1 and X1 (eth0): 128.100.11.0/26 (subnetmask:255.255.255.192)
Network of X1(eth1) and R1(left): 128.100.11.64/26 (subnetmask:255.255.255.192)

| | IP address | Binary (bits 1-16) | Binary (bits 17-32) | Subnetmask |
|---|---|---|---|---|
| H1 | 128.100.11.2 | 10000000.01100100 | **00001011.00**000010 | **255.255.255.192** |
| X1 (eth0) | **128.100.11.3** | 10000000.01100100 | **00001011.00**000011 | **255.255.255.192** |
| X1 (eth1) | **128.100.11.71** | 10000000.01100100 | **00001011.01**000111 | **255.255.255.192** |
| R1 left | 128.100.11.70 | 10000000.01100100 | **00001011.01**000110 | **255.255.255.192** |

**Routing table of X1**

| Network with Prefix | Next hop |
|---|---|
| 128.100.11.0/26 | Direct (eth0) |
| 128.100.11.64/26 | Direct (eth1)) |
| 128.100.10.0/27 | 128.100.11.70 |

→ it is important that X1 has an entry for the network of R1 (right) and H2.

**Routing table of R1**

| Network with Prefix | Next hop |
|---|---|
| 128.100.11.0/26 | **128.100.11.71 (address of X1(eth1))** |
| 128.100.11.64/26 | direct |
| 128.100.10.0/27 | direct |

→ it is important that R1 has an entry for the network of X1 (eth0) and H1.

**Problem 3.** *(10 Points)*
Consider the interconnected IP networks as shown in the figure.

Note there is  typo, with two routers marked as Router 2. (Both choices are accepted).
Here, the routers are labeled Router2a (left) and Router2b (corner)

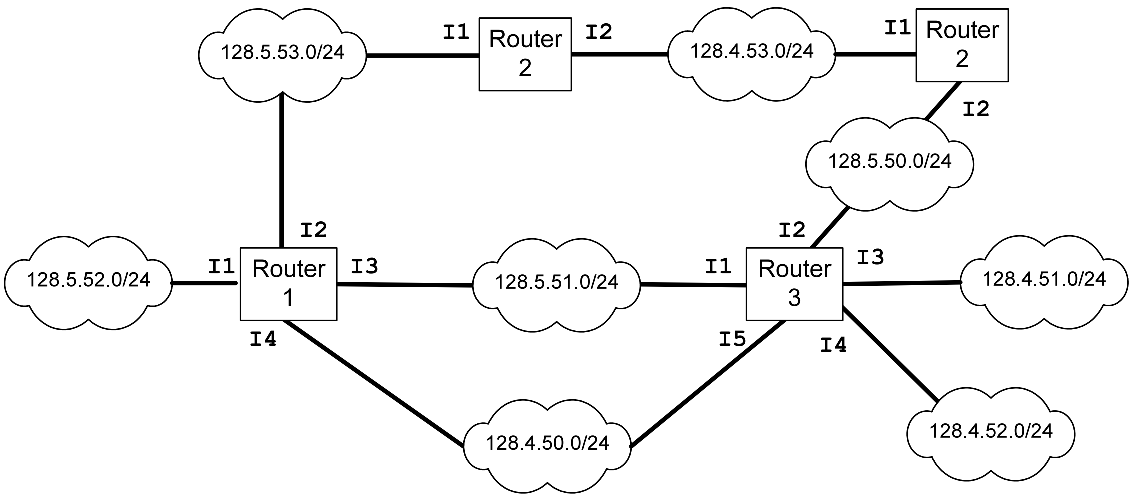


Assume that the routers run a routing protocol that minimizes the hop count to the destination (similar as RIP). If there are multiple routers with the same hop count, a router breaks the tie by picking  the route which maximizes its opportunities for route aggregation. If that leaves unresolved ties,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

a. *(5 points)* Provide the routing table for Router 2. The routing table should have two columns: (1) Network Address (including prefix length), and (2) Interface. The routing table must take advantage of all opportunities for route aggregation.

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.

(a)

**For Router 2a:**

First without aggregation: when 2 interfaces are listed then the tie must be broken by maximizing route aggregation.

| Network Address | Binary | Interface |
|---|---|---|
| 128.5.53.0/24 | 10000000.00000101.00110101.0 | I1 |
| 128.5.52.0/24 | 10000000.00000101.00110100.0 | I1 |
| 128.5.51.0/24 | 10000000.00000101.00110011.0 | I1 |
| 128.5.50.0/24 | 10000000.00000101.00110010.0 | I2 |
| | | |
| 128.4.53.0/24 | 10000000.00000100.00110101.0 | I2 |
| 128.4.52.0/24 | 10000000.00000100.00110100.0 | I2/I1 |
| 128.4.51.0/24 | 10000000.00000100.00110011.0 | I2/I1 |
| 128.4.50.0/24 | 10000000.00000100.00110010.0 | I1 |

→ Note: for 128.4.52.0/24 we should pick I2 since it can then be aggregated. For 128.4.51.0/24 we should pick I2.

Routing table:

| Network Address | Interface |
|---|---|
| 128.5.52.0/23 | I1 |
| 128.5.51.0/24 | I1 |
| 128.5.50.0/24 | I2 |
| | |
| 128.4.52.0/23 | I2 |
| 128.4.50.0/23 | I1 |

**For Router 2b:**

First without aggregation: when 2 interfaces are listed then the tie must be broken by maximizing route aggregation.

| Network Address | Binary | Interface |
|---|---|---|
| 128.5.53.0/24 | 10000000.00000101.00110101.0 | I1 |
| 128.5.52.0/24 | 10000000.00000101.00110100.0 | I1/I2 |
| 128.5.51.0/24 | 10000000.00000101.00110011.0 | I2 |
| 128.5.50.0/24 | 10000000.00000101.00110010.0 | I2 |
| | | |
| 128.4.53.0/24 | 10000000.00000100.00110101.0 | I1 |
| 128.4.52.0/24 | 10000000.00000100.00110100.0 | I2 |
| 128.4.51.0/24 | 10000000.00000100.00110011.0 | I2 |
| 128.4.50.0/24 | 10000000.00000100.00110010.0 | I2 |

→ Note: for 128.5.52.0/24 we should pick I1 since it can then be aggregated.

Routing table:

| Network Address | Interface |
|---|---|
| 128.5.52.0/23 | I1 |
| 128.5.50.0/23 | I2 |
| 128.4.53.0/24 | I1 |
| 128.4.52.0/24 | I2 |
| 128.4.50.0/23 | I2 |

(b) Several solutions are possible:

1. Solutions can be suboptimal (when a very long route is selected if it leads to better aggregation)
2. Addition of new networks can drastically change routing tables in all routers.
3. Convergence of the routing protocol is slower, and there may not be convergence.

**Problem 4.** *(10 Points)*
Pretend for a moment that IP addresses are 12 bits instead of 32, and suppose you are allocated the block of addresses 1011xxxxxxxx (i.e., all addresses that start with `1011' in binary notation). Your task is to create address prefixes for four networks A, B, C, D, with the requirement that

- Network A has 100 hosts;
- Network B has 50 hosts;
- Network C has 25 hosts;
- Network D has 10 hosts.

a. *(8 points)* For each network prefix, provide the IP network address in binary notation and the length of the network prefix.

b. *(2 points)* Create a network (Network E) that uses the entire unused portion of the address block. Provide the IP address and prefix length of the network, as well as the maximum number of hosts.

a)

- There are multiple correct solutions.
- The subnets must have enough IP addresses for each network, and the address blocks must not overlap.
- The size of the subnet is a power of 2 larger than the number of hosts

A: needs 128 = 2^7 addresses: `10110 /5`
Here "10110 / 5" means a 5-bit prefix with leading bits 10110.

B: needs 64 = 2^6 addresses: `101110 /6`
C: needs 32 = 2^5 addresses: `1011110 /7`
D: needs 16 = 2^4 addresses: `10111110 /8`

b) There are needs 16 = 2^4 addresses left. This leaves the subnet
E: 2^4 addresses: `10111111 /8`

Subtracting the network address and the broadcast address, network E can have 14 hosts.

**Problem 5.** *(10 Points)*
Consider the network of autonomous systems in the figure, where AS A is a customer of AS X and AS U, and AS B is a customer of AS X and AS V. AS X, in turn, has two provider ASes U and V, while AS C is a customer AS of AS W and AS U. The three ASes U, V and W have peering relationships among themselves.

- AS C owns the prefix 128.1.0.0/16.
- AS A owns two network prefixes 64.1.10.0/24, and 128.101.34.0/24,

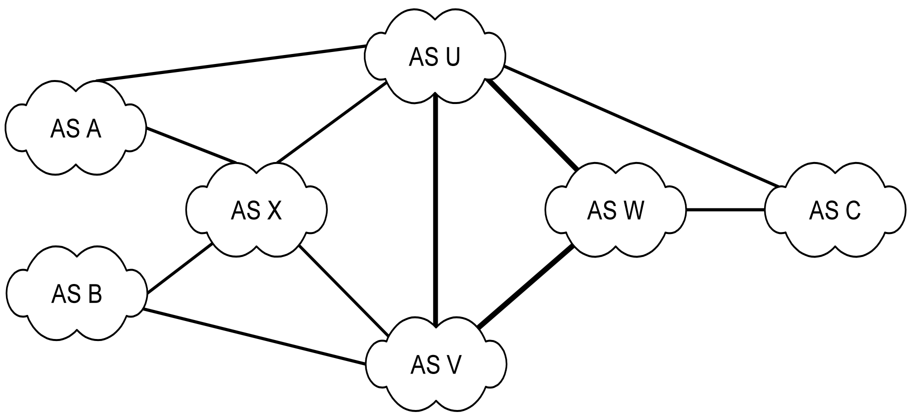


The inter-domain routing protocol BGP is used among the ASes to exchange routing information.

a. *(5 points)* For a packet from a host in AS A with the destination IP address 128.1.34.35, which ASes would this packet most likely traverse to reach its final destination? Explain your answer.

b. *(5 points)* Suppose AS A wants the traffic to its prefix 128.101.34.0/24 to come from AS U and the traffic to its prefix 64.1.10.0/24 to come from AS X. Which routes should AS A announce to AS U, and what routes should AS A announce to AS X? For each announcement, provide the prefix and the AS-PATH attribute. Explain your answers.

a) The IP address is in AS C.
Based on the customer/provider relationships both ASX and ASU will advertise the prefix to 128.1.0.0/16. Since the path A → U → C is shorter than A → X → U → C, the shorter path is selected.

b) The announcements should be:
A → U: 128.101.34.0/24, AS-PATH{AS A}
A → X: 64.1.10.0/24, AS-PATH{AS A}

**Problem 6.** *(10 points)*

Consider a TCP connections between A and B. Assume that A has sent to B the sequence numbers 100 through 500 (measured in bytes).

(a) *(3 points)* Describe the difference between the advertised window and the usable window in TCP.

(b) *(3 points)* Show the advertised and usable windows at A, after A receives from B a segment with *(AckNo=500, Window size = 500)*. Which sequence numbers can A send to B?

(c) *(4 points)* After sending the segment with *(AckNo=500, Window size = 500)* in part (b), B wants to further reduce the advertised window by setting *Window size = 200.* Describe how B should proceed.

a) Advertised window is the range of sequence numbers that the receiver is willing to accept. The usable window is the part of the advertised window which has not been transmitted ("used up") by the sender.

b) The ACK does not acknowledge the last byte (SeqNo=500!).
The advertised window is [500,999].
So, A can send the numbers from 501 to 999.

c) B should not send an *(AckNo=500, Window size = 200), since A may have already transmitted numbers in the range 701…999.*
The larger principle is that the receiver should not reduce the upper bound of the advertised window.
Instead, the receiver should wait for more data from A and not increase the upper bound of the advertised window, while still acknowledging all data received.
E.g., (assuming that A sends 100 byte segments)
with *(AckNo=600, Window size = 400),*
*(AckNo=700, Window size = 300)*
*(AckNo=800, Window size = 200)*

In this way, the advertised window shrinks by increasing the lower bound of the window.

**Problem 7.** *(10 points)*
Assume you modify the TCP implementation on your computer as follows:

- After a retransmission timeout and subsequent retransmission, do not modify the values of the congestion window (*cwnd*) or the slow-start threshold value (*ssthresh*).
- After a retransmission timeout and subsequent retransmission, do not apply Karn's algorithm.
- Disable the round-trip time (RTT) measurements, i.e., always use the initial values of RTO.
- Do not use Fast Retransmit after a third duplicate acknowledgement (instead only retransmit using timeouts).
- Never leave the Slow-Start phase.

(a) *(5 points)* Describe the impact of each of the modifications on the maximum data rate at which your computer can transmit on a TCP connection. (Consider each modification individually).

(b) *(5 points)* Describe the impact of each of the modification when all computers on the Internet are modified. (Again, consider each modification individually).

**Problem 8.** *(10 points)*
Consider Bridge 12 with 5 ports as shown in the figure below. The bridge is receiving configuration messages (BPDUs) as shown in the figure. Specifically, a message (R, C, B, P) where R is the value of the root ID, C is the value of the root path cost, B is the bridge ID, and P is the port ID, is interpreted as follows: "I am bridge B and I am sending from my port P. I believe R to be the root bridge, and the cost of my path to the root bridge is C."

(a) What is the root bridge of Bridge 12, and what is its Root Path Cost?
(b) What is Bridge 12's root port?
(c) What is Bridge 12's configuration message?
(d) On which ports does Bridge 12 send its configuration message?
(e) Which ports of Bridge 12 are marked as "blocked"?

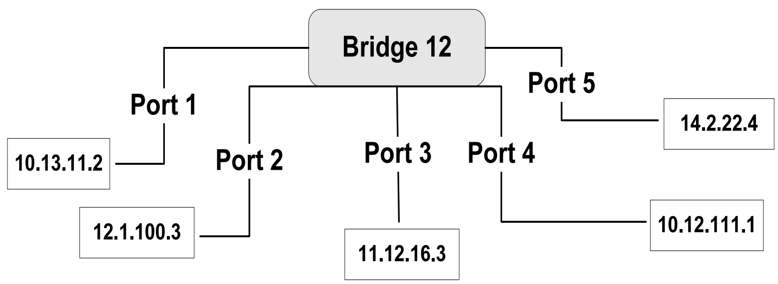


a. Root bridge: 10, root path cost is 13.
b. 4
c. 10.13.12.x (where x is port number)
d. Ports 2, 3, 5.
e. Port 1

f.

**Problem 9.** *(10 points)*
Consider the network in the figure with two private IP networks, which are each connected to the Internet by a NAT router. Both routers R1 and R2 perform IP masquerading (Port Address Translation). R1 and R2 have public addresses, 142.150.235.16 and 128.143.71.21, respectively, which have been acquired using DHCP.

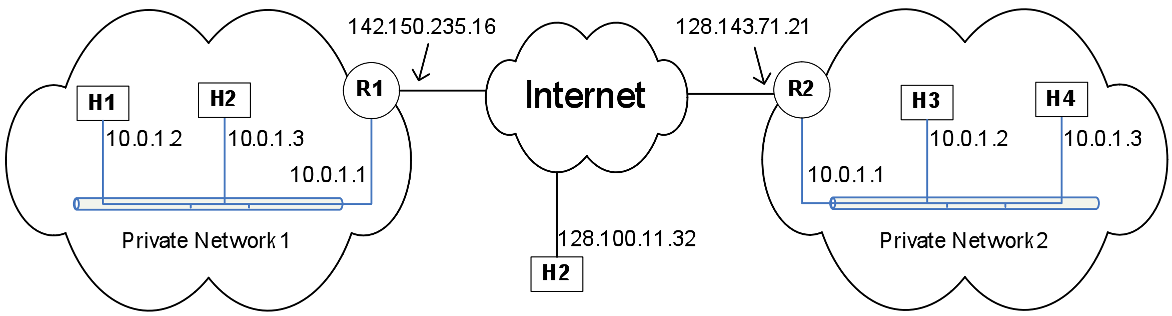


(a) *(3 points)* Suppose both H1 and H2 issue the command "ping 128.100.11.32". Describe the difficulty for the NAT router of handling these ping commands. Explain what NAT router R1 must do to handle the traffic resulting from the ping commands.

(b) *(3 points)* Suppose users on H1 and H3 want to set up a Voice-over-IP application (e.g., Skype) with each other. Describe the difficulty with setting up this application, and describe how it can be solved.

(c) *(4 points)* Suppose that H1 wants to run a web server that can be accessed from any computer on the Internet. The restriction is that the IP addresses in Private Network 1 cannot be changed. Also, NAT router R1 must continue to perform IP masquerading with a single public IP address. Describe changes to the configuration of R1, so that H1 can run a web server which can be accessed by H3, H4, and H5.

**Problem 10.** *(10 points)* Consider the following set of prefixes

A. 010*
B. 1100*
C. 01111*
D. 00*
E. 0111*
F. 01*
G. 1*
H. 11*

a) *(4 points)* Construct a disjoint-prefix binary trie for these prefixes.

b) *(3 points)* Construct the disjoint-prefix binary trie after adding the prefixes "X. 1111*" and "Y. 1010*".

c) *(3 points)* Construct the disjoint-prefix binary trie after removing A and C.

Binary trie for a and b:

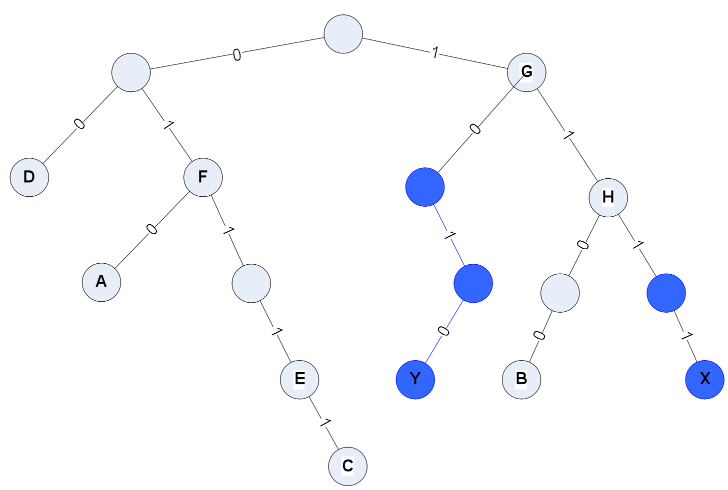


a) Disjoint binary trie

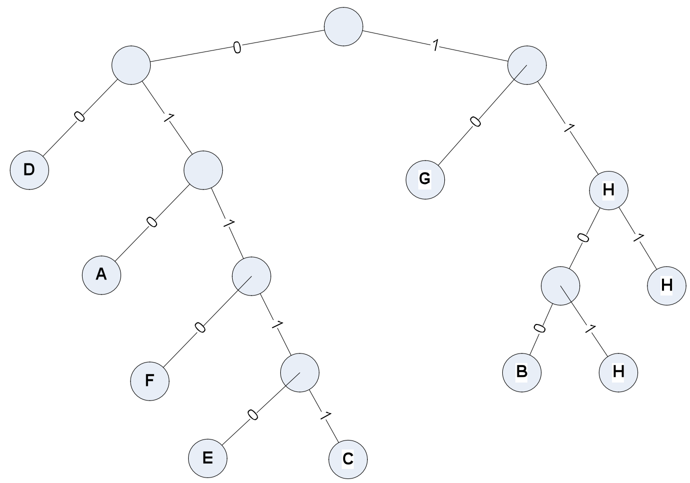


b) Disjoing binary trie for after additions

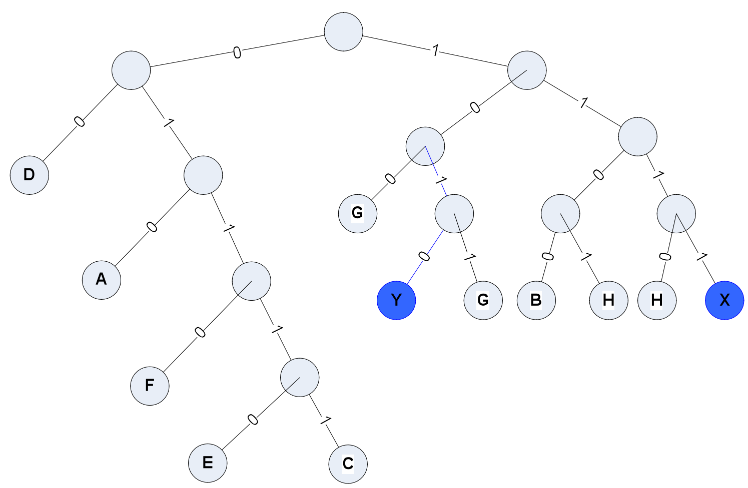


c) Disjoing binary trie for after deletions

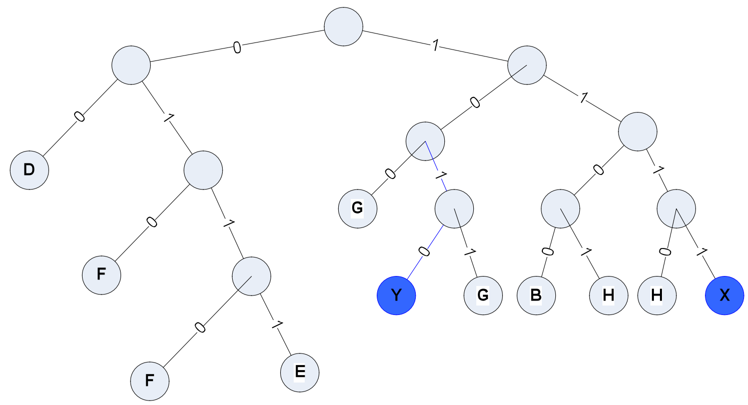
0
1
0
1
0
1
D
0
1
0
1
0
1
0
1
0
1
F
G
Y
G
B
H
H
X
0
1
F
E

**Traffic capture from Problem 1:**

```
 00 13 72 30 84 8d 00 1f  f3 56 8a 9f 08 00 45 00
 00 34 2c 9e 40 00 40 06  1a bb 8e 96 eb 21 8e 96
 eb 1c c4 e4 00 50 7a 4f  7f 03 ef 7c db 51 80 11
 ff ff f3 91 00 00 01 01  08 0a 35 b9 51 3b c1 67
 d1 28
```

```
No.     Time        Source                Destination           Protocol Info
     69 14.772114   142.150.235.33        142.150.235.28        TCP      50404
> 80 [FIN, ACK] Seq=912 Ack=11497 Win=524280 Len=0 TSV=901337403
TSER=3244806440

F   Frame Length: 66 bytes
    Capture Length: 66 bytes
    Ethernet II, Src: 00:1f:f3:56:8a:9f (00:1f:f3:56:8a:9f), Dst:
00:13:72:30:84:8d (00:13:72:30:84:8d)
    Destination: 00:13:72:30:84:8d (00:13:72:30:84:8d)
    Source: 00:1f:f3:56:8a:9f (00:1f:f3:56:8a:9f)
    Type: IP (0x0800)
Internet Protocol, Src: 142.150.235.33 (142.150.235.33), Dst: 142.150.235.28
(142.150.235.28)
    Version: 4
    Header length: 20 bytes
    Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
        0000 00.. = Differentiated Services Codepoint: Default (0x00)
        .... ..0. = ECN-Capable Transport (ECT): 0
        .... ...0 = ECN-CE: 0
    Total Length: 52
    Identification: 0x2c9e (11422)
    Flags: 0x02 (Don't Fragment)
        0.. = Reserved bit: Not Set
        .1. = Don't fragment: Set
        ..0 = More fragments: Not Set
    Fragment offset: 0
    Time to live: 64
    Protocol: TCP (0x06)
    Header checksum: 0x1abb [correct]
    Source: 142.150.235.33 (142.150.235.33)
    Destination: 142.150.235.28 (142.150.235.28)
Transmission Control Protocol, Src Port: 50404 (50404), Dst Port: 80 (80), Seq:
912, Ack: 11497, Len: 0
    Source port: 50404 (50404)
    Destination port: 80 (80)
    Sequence number: 912    (relative sequence number)
    Acknowledgement number: 11497    (relative ack number)
    Header length: 32 bytes
    Flags: 0x11 (FIN, ACK)
        0... .... = Congestion Window Reduced (CWR): Not set
        .0.. .... = ECN-Echo: Not set
        ..0. .... = Urgent: Not set
        ...1 .... = Acknowledgement: Set
        .... 0... = Push: Not set
        .... .0.. = Reset: Not set
```

```
    .... ..0. = Syn: Not set
    .... ...1 = Fin: Set
Window size: 524280 (scaled)  <-- actual value is ff ff
Checksum: 0xf391 [validation disabled]
Options: (12 bytes)
    NOP
    NOP
    Timestamps: TSval 901337403, TSecr 3244806440
```