CS-202 Exercises on Network Layer: IP & Routing (L15-L16)

Before we start

A few basic rules on how to allocate an IP prefix to a subnet:

- An IP prefix A/M is the range of IP addresses whose M most significant bits are the same as A's M most significant bits. E.g., 10.0.0.16 belongs to IP prefix 10.0.0.0/24 , because the 24 most significant bits of 10.0.0.16 are the same as the 24 most significant bits of 10.0.0.0.
- This implies that a /M IP prefix contains 2^{32-M} IP addresses. E.g., a /24 IP prefix contains $2^{32-24} = 2^{8} = 256$ IP addresses. In other words, we don't have the freedom to create an IP prefix that contains an arbitrary number of IP addresses, it must always contain a number that is a power of 2.
- /M is called the "subnet mask", representing the bitmask consisting of M consecutive 1's followed by enough 0's to reach 32 bits. The subnet mask is applied to any IP address using bitwise AND to obtain the IP range the address belongs to. E.g. 100.52.12.18 belongs to the prefix 100.52.12.16/28 because: (1) the /28 mask represents the mask 1111 1111.1111 1111.1111 1111.1111 0000, (2) applying the mask to 100.52.12.18 (0110 0100.0011 0100.0000 1100.0001 0010 in binary) results in 100.52.12.16 (0110 0100.0011 0100.0000 1100.0001 0000 in binary).
- Each IP subnet must have its own IP prefix. Hence, IP prefixes allocated to different IP subnets must not overlap.

In addition to the network interfaces, each subnet has two IP addresses that are sometimes reserved for special use:

- The first IP address in the subnet's IP prefix (called "network address"). E.g., the first IP address in 10.0.0.0/24 is 10.0.0.0. This address is sometimes reserved for special uses, e.g., a discovery service provided by the subnet.
- The last IP address in the subnet's prefix (called "broadcast address"). E.g., the last IP address in 10.0.0.0/24 is 10.0.0.255. This address is sometimes reserved to be used as the subnet's broadcast address, i.e., as the destination IP address for packets that should be received by all end-systems in a subnet.

In practice, operators do not assign these addresses to any network interface. However, in this course you are free to assign them to a network interface or not except when it is clearly stated in the problem description.

Exercise 1: IP prefix allocation

IP subnets A, B and C contain 10, 5, and 3 network interfaces, respectively. Allocate an IP prefix to each subnet, and assign an IP address to each network interface, from IP prefix 1.2.3.0/27.

Consider two cases for allocating prefixes to subnets. In each case, follow the given order:

(a) A, B, C

(b) B, A, C

In the context of this exercise, when we say that allocation "follows a given order," we mean that, if Subnet X comes before Subnet Y in that order, the IP addresses for Subnet X should be arithmetically smaller than the IP addresses for Subnet Y (in the sense that IP address 1.2.3.4 is arithmetically smaller than IP address 1.2.3.5).

Note: Allocating addresses might be infeasible in some cases.

Exercise 2: network configuration

Consider the topology shown in Figure 1. There are three IP subnets (A, B and C) that contain some end-systems, and two IP subnets (D and E) that contain no end-systems. The green boxes (a, b, c, . . . g) denote network interfaces for routers R1, R2 and R3.

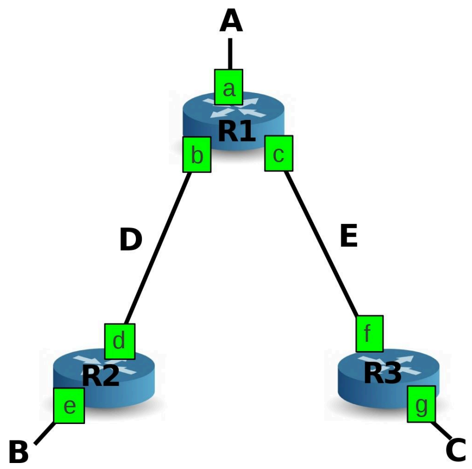

Figure 1: Network Topology.

1. Allocate an IP prefix to each subnet. Your allocation must respect the following constraints:
   - All prefixes must be allocated from 214.97.254.0/23.
   - Subnet A should have enough addresses to support 250 interfaces.
   - Subnet B should have enough addresses to support 120 interfaces.
   - Subnet C should have enough addresses to support 60 interfaces.
   - Each of subnets D and E should have enough addresses to support 2 interfaces.

2. Using your previous answer, provide the forwarding tables for each of the three routers (R1, R2, R3). Each table should contain two columns which show (i) the destination IP prefix, and (ii) the corresponding output link.

3. Can you reduce the number of entries of each forwarding table, i.e., for each table create an equivalent one, which has the same outcome but consists of fewer entries?

Exercise 3: practice prefix allocation

Consider the graph shown in Figure 2, which consists of four network subnets, *A*, *B*, *C* and *D*, all connected to the Internet through Router 3.

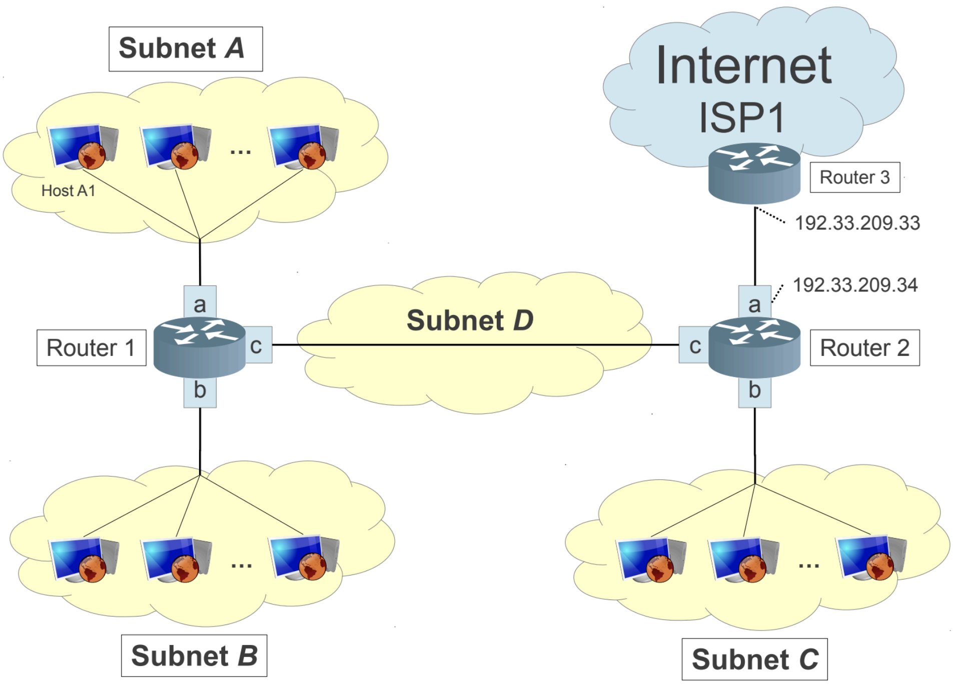

Figure 2: Network topology.

- Subnet *A* is attached to port *a* of Router 1.
- Subnet *B* is attached to port *b* of Router 1.
- Subnet *C* is attached to port *b* of Router 2.
- Subnet *D* has no hosts.
- All hosts have access to the Internet, through an Internet Service Provider "ISP1".

Assign network addresses to each of these four subnets, with the following constraints:

- All addresses must be allocated from 1.0.2.0/23 (in other words, they should have binary format 00000001.00000000.0000001x.xxxxxxxx);
- subnets *A*, *B* and *C* should have enough addresses to support 200, 100 and 50 interfaces, respectively;
- you should allocate the smallest possible range of IP addresses to each subnet;
- only in this exercise, assume subnets do not have a broadcast address (i.e., you should not allocate it);
- for each subnet, the assignment should take the form a.b.c.d/x.

Exercise 4: identify subnets

Consider the network in Figure 3, consisting of hosts A, B, C, D, G and H, DNS server E, web server F, router R1 and switches S1, S2, S3 and S4.

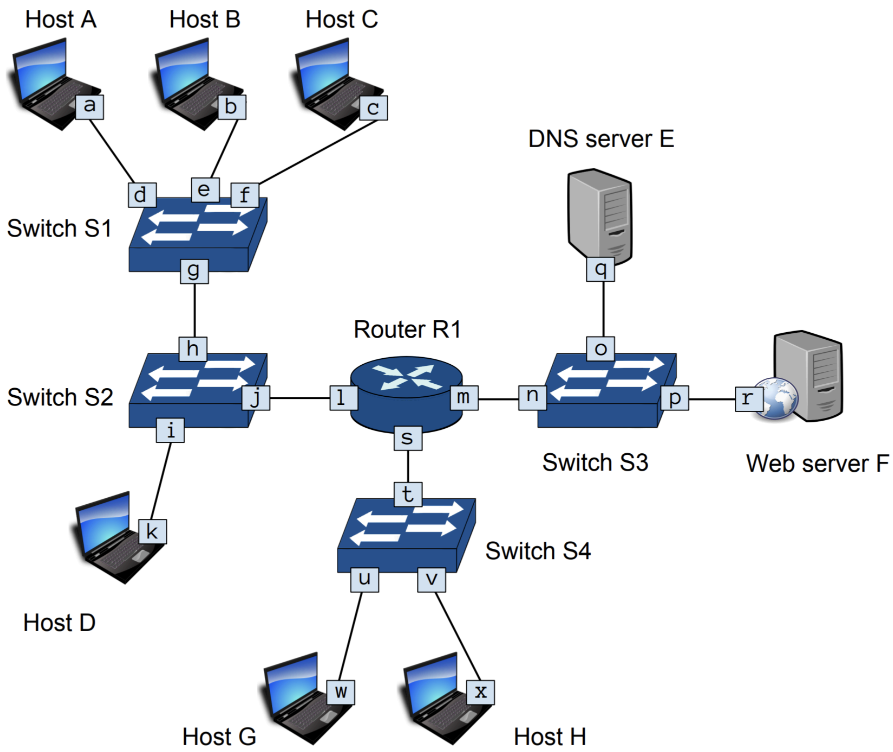

Figure 3: Network topology.

1. Allocate an IP prefix to each IP subnet that contains end-systems, following these rules:
   - All addresses must be allocated from 10.0.0.0/24 (they should have the binary format 00001010.00000000.00000000.xxxxxxxx), following the basic rules for allocating IP addresses.
   - You should allocate the smallest possible range of IP addresses to each subnet (*Note: be careful to consider* ***only*** *the network interfaces that should be assigned an IP address*)
2. Fill Table 1 with the IP address for the network interfaces that should be assigned an IP address. write "−" for interfaces for which an IP address is not needed.

Network interface	IP address
Example: y	1.2.3.4
Example: z	–
a	
b	
c	
d	
e	
f	
g	
h	
i	
j	
k	
l	
m	
n	
o	
p	
q	
r	
s	
t	
u	
v	
w	
x	

Table 1: IP address allocations for the interfaces from Figure 3

Exercise 5: link-state routing

Consider the network in Figure 4. Execute the link-state (Dijkstra's) algorithm we saw in class to compute the least-cost path from each of *x*, *v*, and *t* to all the other routers.

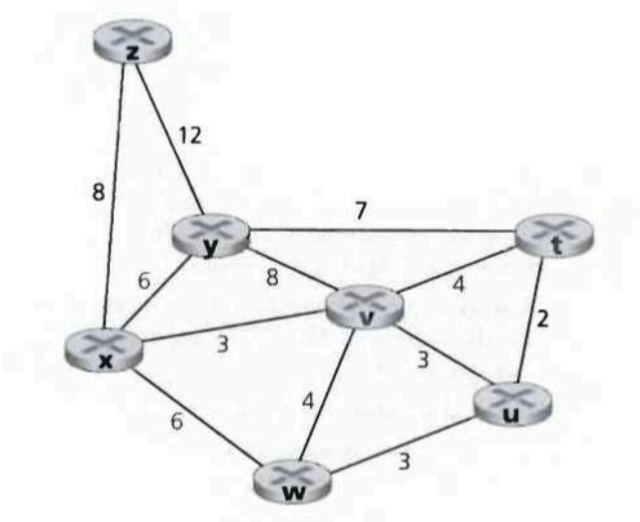

Figure 4: Network topology.

Exercise 6: distance-vector routing

Consider the network in Figure 5. Execute the distance-vector (Bellman-Ford) algorithm we saw in class and show the information that router *z* knows after each iteration.

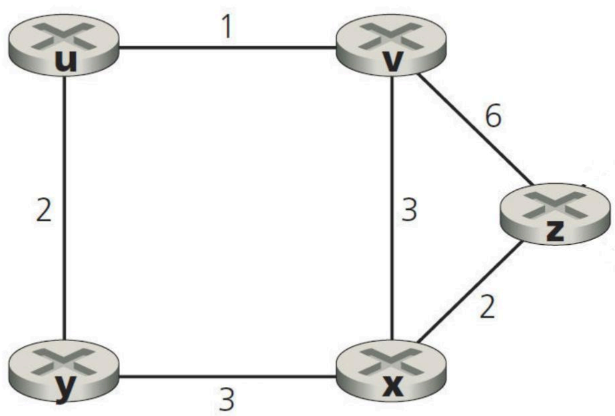

Figure 5: Network topology.

Exercise 7: convergence

What is the maximum number of iterations required for the distance-vector (Bellman Ford) algorithm that we saw in class to converge (i.e., to finish, assuming no change occurs in the network graph and link costs)? Justify your answer.