Hwk 4 Answers

Chapter 4

1. **KR, #R2**
   The main function of the data plane is packet forwarding, which is to forward datagrams from their input links to their output links. For example, the data plane's input ports perform physical layer function of terminating an incoming physical link at a router, perform link-layer function to interoperate with the link layer at the other side of the incoming link, and perform lookup function at the input ports.

2. **KR, #R6**
   Input port, switching fabric, and output ports are implemented in hardware, because their datagram-processing functionality is far too fast for software implementation. A routing processor inside a traditional router uses software for executing routing protocols, maintaining routing tables and attached link state information, and computing the forwarding table of a router. In addition, a routing processor in a SDN router also relies on software for communication with a remote controller in order to receive forwarding table entries and install them in the router's input ports.
   Data plane is usually implemented in hardware due to the requirement of fast processing, e.g., at nanosecond time scale. Control plane is usually implemented in software and operates at the millisecond or second timescale, for example, for executing routing protocols, responding to attached links that go up or down, communicating with remote controllers, and performing management functions.

3. **KR, #R14**
   (A typo in this question: the first question mark should be replaced by a period). Only FIFO can ensure that all packets depart in the order in which they arrived.

4. **KR, #R25**
   50% overhead.

5. **KR, #P2**
   a) No, you can only transmit one packet at a time over a shared bus.
   b) No, as discussed in the text, only one memory read/write can be done at a time over the shared system bus.
   c) No, in this case the two packets would have to be sent over the same output bus at the same time, which is not possible.

6. **KR, #P7**
   Link Interface 0
   11000000
   through (32 addresses)
   11011111

   Link Interface 1
   10000000
   through(64 addresses)
   10111111

   Link Interface 2
   11100000
   through (32 addresses)
   11111111

   Link Interface 3
   00000000
   through (128 addresses) 01111111

7. **KR, #P8**
   223.1.17.0/26 223.1.17.128/25 223.1.17.192/28

8. **KR, #P18**
   It is not possible to devise such a technique. In order to establish a direct TCP connection between Arnold and Bernard, either Arnold or Bob must initiate a connection to the other. But the NATs covering Arnold and Bob drop SYN packets arriving from the WAN side. Thus neither Arnold nor Bob can initiate a TCP connection to the other if they are both behind NATs.

Chapter 5

1. **KR, #R2**
   Logically centralized control means that a logically central routing controller computes and distributes the forwarding tables to be used by each and every router, and each router does not compute its forwarding table, unlike the per-router control. In the case of logically centralized control, the data plane and control plane are implemented in separate devices; the control plane is implemented in a central server or multiple servers, and the data plane is implemented in each router.

2. KR, #R5

   The count-to-infinity problem refers to a problem of distance vector routing. The problem means that it takes a long time for a distance vector routing algorithm to converge when there is a link cost increase. For example, consider a network of three nodes x, y, and z. Suppose initially the link costs are c(x,y)=4, c(x,z)=50, and c(y,z)=1. The result of distance-vector routing algorithm says that z's path to x is z!y!x and the cost is 5(=4+1). When the cost of link (x,y) increases from 4 to 60, it will take 44 iterations of running the distance-vector routing algorithm for node z to realize that its new least-cost path to x is via its direct link to x, and hence y will also realize its least-cost path to x is via z.

   KR, #R6

   No. Each AS has administrative autonomy for routing within an AS.

3. KR, #P7

   a) Dx(w) = 2, Dx(y) = 4, Dx(u) = 7

   b) First consider what happens if c(x,y) changes. If c(x,y) becomes larger or smaller (as long as c(x,y) >=1) , the least cost path from x to u will still have cost at least 7. Thus a change in c(x,y) (if c(x,y)>=1) will not cause x to inform its neighbors of any changes.
   If c(x,y)= $\delta<1$, then the least cost path now passes through y and has cost $\delta+6$.
   Now consider if c(x,w) changes. If c(x,w) = $\varepsilon \leq 1$, then the least-cost path to u continues to pass through w and its cost changes to 5 + ε; x will inform its neighbors of this new cost. If c(x,w) = $\delta > 6$, then the least cost path now passes through y and has cost 11; again x will inform its neighbors of this new cost.

   c) Any change in link cost c(x,y) (and as long as c(x,y) >=1) will not cause x to inform its neighbors of a new minimum-cost path to u .

4. KR, #P14

   a) eBGP
   b) iBGP
   c) eBGP
   d) iBGP

5. KR, #P15

   a) I1 because this interface begins the least cost path from 1d towards the gateway router 1c.

   b) I2. Both routes have equal AS-PATH length but I2 begins the path that has the closest NEXT-HOP router.

   c) I1. I1 begins the path that has the shortest AS-PATH.