


1. The operations a router perform are typically divided into a control plane and a data plane.
  - (a) What is the difference between forwarding and routing? On which plane does each belong?
  - (b) To which plane does the cross-bar (more generally called the interconnect) belong?
  - (c) What is the purpose of a routing table? What plane does the routing table belong to?
2. Routers typically employ many buffers/queues in their design.
  - (a) Describe a situation when buffers at the input port of a router might fill up to its capacity.
  - (b) Describe a situation when buffers at the output port of a router might fill up to its capacity.
  - (c) Give one reasonable course of action for a router to do when a buffer is completely full.
3. Layer 2 switches and Layer 3 routers have similar functionality in that they take a chunk (e.g., frame and packet, respectively) of data coming in on one port and send it out on a different one. Where there is a switch in a network, it is usually possible to replace it with a router and vice-versa. Let's examine some of the reasons influencing a network engineer's choice to use a switch (and make the LAN bigger) versus a router (to connect two separate LANs).
  - (a) What are some pros and cons of switches versus routers when considering management and set-up?
  - (b) What are some pros and cons of switches versus routers when considering network size?
  - (c) What are some pros and cons of switches versus routers when considering how efficiently the packet traverses the network?
4. Consider the following network:



The following table above shows a routing table on node A before any distance vector updates have been received (assume A knows the cost of reaching its immediate neighbours). As you can see, each row of the table stores the distance to a particular node through one of A's neighbours.

to/via	B	C	F
B	1	$\infty$	$\infty$
C	$\infty$	3	$\infty$
D	$\infty$	$\infty$	$\infty$
E	$\infty$	$\infty$	$\infty$
F	$\infty$	$\infty$	2

- (a) What does the table look like when A receives the initial distance vector update from node C?
  - (b) What does the table look like when A receives the initial distance vector update from node B?
  - (c) What does the table look like in the steady state, after several updates?
5. For this question you will design a custom link-state routing system for a network. You can assume that all links are bi-directional and every node has a unique network identifier of four bytes. You can also assume that the link layer provides a mechanism for a node to exchange arbitrary messages with its immediate neighbours.
- (a) Design the messages that nodes will exchange to obtain information about the state of the link between itself and its immediate neighbours.
  - (b) Describe the steps (using the messages you designed above) that nodes should follow in order to collect link state information from their immediate neighbours.

- (c) Design an appropriate link-state message for communicating link-state information across the entire network.
- (d) Describe the steps nodes should employ to share out the link-state messages in designed in the previous part.
- (e) A network using your system has  $n$  nodes,  $l$  links, the maximum degree of any node is  $k$  and there is a path between any two nodes of not more than  $d$  hops. Give a bound on the total amount of information which must be transmitted (in bytes) to ensure that every node acquires complete topology information, in terms of  $n$ ,  $l$ ,  $k$  and  $d$ .