

There are 10 questions totalling 100 points.

1. (5 points) A unix router with two ethernet interfaces, *fa0* and *fa1*, and a loopback interface, *lo0*, is issued the following commands at boot time. Show the combined forwarding/arp table. Make up MAC addresses if you need to.


```
ifconfig fa0 inet 207.140.138.1 netmask 255.255.0.0
ifconfig fa1 inet 207.140.172.1 netmask 255.255.255.0
ifconfig lo0 inet 127.0.0.1
route add default 207.140.1.1
ping -c 1 207.140.138.1
ping -c 1 207.140.171.254
ping -c 1 207.140.172.255
ping -c 1 10.0.0.1
```

Destination	Gateway	Flags	Netif
default	207.140.1.1	UG	fa0
127.0.0.1	127.0.0.1	UH	lo0
207.140.0.0/16	link#1	U	fa0
207.140.138.1	00:02:04:06:08:0a	UHL	fa0
207.140.171.254	04:06:08:0a:0c:0e	UHL	fa0
207.140.172.255	ff:ff:ff:ff:ff:ff	UBH	fa1
207.140.172.0/24	link#2	U	fa1
207.140.255.255	ff:ff:ff:ff:ff:ff	UB	fa0

2. (5 points) What is the fundamental difference between a *multi-homed host* and a *router*?
A router forwards packets; a multi-homed host does not.
3. (a) (5 points) What is the *forwarding table* or *forwarding information base* (FIB)?
The FIB is a data structure consulted every time a node sends out an IP packet; for each destination prefix (including the default prefix) that the node knows how to reach, it shows the next-hop node and the output interface.
(b) (5 points) What is the *routing table* or *routing information base* (RIB)?
The RIB is a data structure maintained by the routing code that contains information about every reachable prefix; the information kept is typically the next-hop router, the cost to the destination, etc. The RIB is updated whenever new routes are computed (for example, when the link-state database gets modified and the SPF algorithm is run again), and it is used to compute the FIB.
(c) (5 points) How does the FIB implementation in a host usually differ (if it does) from the FIB implementation in a router?
It doesn't have to, but general-purpose hosts have software implementations of their FIB, whereas routers, because they need to forward packets very fast, keep their FIB in specialized hardware (ternary CAMs).
(d) (5 points) Does a single-homed host need a FIB?
Of course.

4. (5 points) Name three data structures that are used in software implementations of FIBs.
Hash table, trie, p-trie (Patricia tree).
5. (5 points) A network with no redundant links (*i.e.*, with only one path between each pair of routers) is running a rudimentary D-V routing protocol (no SH or PR). Can it ever exhibit counting-to-infinity behavior? Give a brief justification for your answer.
Yes. Any network running a Distance-Vector protocol without Split Horizon may exhibit counting-to-infinity.
6. (5 points) A network with no redundant links (*i.e.*, with only one path between each pair of routers) is running a D-V routing protocol with split horizon and poison reverse. Can it ever exhibit counting-to-infinity behavior? Give a brief justification for your answer.
No, since routing messages are only propagated away from the advertising router.

Questions 7 to 10 refer to the following network:



7. (5 points) The network is running a Distance-Vector protocol without split horizon or poison reverse. Show the routing table of **D** after routing has converged.

Simply by inspection:

Destination	Next hop	Distance
d	D	0
c	C	1
e	E	1
b	C	2
a	E	2

8. (5 points) The network is running a Link-State protocol. Assume that all links have equal cost. Show the routing table of **D** after routing has converged.

The table is identical to the previous one, obviously!

Destination	Next hop	Cost
d	D	0
c	C	1
e	E	1
b	C	2
a	E	2

9. (5 points) The network is running a Link-State protocol again; links DE and DC have cost 1; links AE and BC have cost 2; link AB has cost 3. Show the routing table of **D** after routing has converged.

Again, by inspection:

Destination	Next hop	Cost
d	D	0
c	C	1
e	E	1
b	C	3
a	E	3

10. The network is back to running a Distance-Vector protocol, only this time with split horizon. Link AB becomes unidirectional (traffic from A can flow to B over the link, but not the other way around).

First, let us compute all the routing tables.

Each router starts with a route to itself...

t = 0														
A			B			C			D			E		
a	A	0	b	B	0	c	C	0	d	D	0	e	E	0

... which it advertises to its neighbors. Note that A did not hear a route from B over the AB link, since that link is only passing packets from A to B and not the other way around.

t = 1														
A			B			C			D			E		
a	A	0	b	B	0	c	C	0	d	D	0	e	E	0
e	E	1	a	A	1	b	B	1	c	C	1	d	D	1
			c	C	1	d	D	1	e	E	1	a	A	1

Each router passes along routes it heard from its counterclockwise neighbor to its clockwise neighbor, and vice versa. Again, A is not hearing any routes from B.

t = 2														
A			B			C			D			E		
a	A	0	b	B	0	c	C	0	d	D	0	e	E	0
e	E	1	a	A	1	b	B	1	c	C	1	d	D	1
d	E	2	c	C	1	d	D	1	e	E	1	a	A	1
			e	A	2	a	B	2	b	C	2	c	D	2
			d	C	2	e	D	2	a	E	2			

If link AB were working properly, this is as far as we would get, and indeed, the routing tables of B, C, and D are complete. But E does not have a route to b yet, nor does A have routes

to b or c . This lets the route to c that E sends to A to be installed with a distance of 3 (if AB were working, A would have a distance-2 route to c). Similarly, the route to b that D sends to E also gets installed with a metric of 3.

t = 3														
A			B			C			D			E		
a	A	0	b	B	0	c	C	0	d	D	0	e	E	0
e	E	1	a	A	1	b	B	1	c	C	1	d	D	1
d	E	2	c	C	1	d	D	1	e	E	1	a	A	1
c	E	3	e	A	2	a	B	2	b	C	2	c	D	2
			d	C	2	e	D	2	a	E	2	b	D	3

Finally, E passes along to A its route to b , which A installs.

t = 4														
A			B			C			D			E		
a	A	0	b	B	0	c	C	0	d	D	0	e	E	0
e	E	1	a	A	1	b	B	1	c	C	1	d	D	1
d	E	2	c	C	1	d	D	1	e	E	1	a	A	1
c	E	3	e	A	2	a	B	2	b	C	2	c	D	2
b	E	4	d	C	2	e	D	2	a	E	2	b	D	3

Observe that all routers think that they can reach all other routers.

- (5 points) Show the routing table of **A** after routing has converged.
See the $t = 4$ table above.
- (5 points) Show the routing table of **B** after routing has converged.
See the $t = 4$ table above.
- (5 points) Show the routing table of **D** after routing has converged.
See the $t = 4$ table above.
- (25 points) Fill out the reachability matrix for the network.

The only paths that will have a problem are paths that have to go over the B -to- A link.

- A can reach all destinations, since all its routes go through E ; note, however, that even though packets to b and c could have gone through B , they take a suboptimal path.
- B thinks it can get to a and e through A , but of course it cannot.
- C thinks it can get to a through B , but again it cannot.
- D can reach all nodes.
- E can also reach all nodes.

The reachability table is thus:

src \ dst	A	B	C	D	E
A	✓	✓	✓	✓	✓
B		✓	✓	✓	
C		✓	✓	✓	✓
D	✓	✓	✓	✓	✓
E	✓	✓	✓	✓	✓

Observe that if the network were running an OSPF-like link-state protocol, the fact that link *AB* is not bidirectional would have been detected (*A* would not be able to hear *B*'s HELLO messages) and no router would have attempted to use that link in its SPF calculations.