COMPSCI 314 S1T Assignment 3
2008

Department of Computer Science
The University of Auckland

*This assignment contributes **5%** of your overall course mark. Submit your assignment in **PDF** format to the **Assignment Drop Box**. Include all **workings** and **explanations**. Marks will be deducted for ambiguous solutions**.** Zero marks are awarded if the answers contain no explanation. Also, refer to the Departmental Policy on Cheating on Assignments.*

Assignment Drop Box *(*https://adb.ec.auckland.ac.nz/adb/*).*
Departmental Policy on Cheating on Assignments *(*http://www.cs.auckland.ac.nz/CheatingPolicy.php*)*

[Total: 40 marks]

Sample Solution(v.1.1 2pm-04-May-2008)

Q1. Bridging and routing [20 marks]

a) Consider the following network of LAN bridges (in practice, these would be multiport switches, with only the links *between* switches shown).

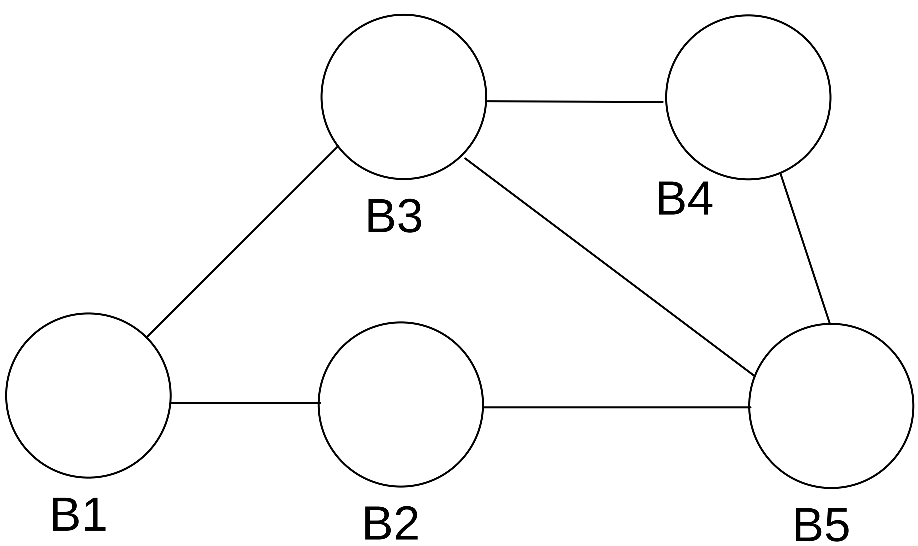

Assume B1 is elected as the root bridge. Assuming the link metrics are all the same, list the links in a valid spanning tree (e.g. B1-B2, etc.). Which link(s) will not be used? **[4 marks]**

B1-B2
B1-B3
B3-B4
B3-B5
(B2-B5 and B4-B5 not used)

or

B1-B2
B2-B5
B1-B3
B3-B4
(B3-B5 and B4-B5 not used)

b) Consider the following network of routers, where all the paths have equal weight (1).

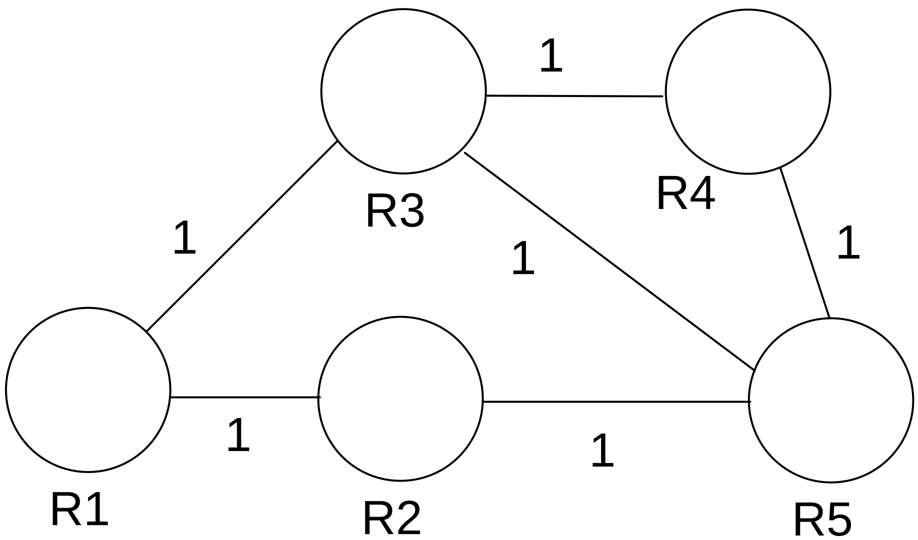

List the shortest distances from **R1** to each other R, as the Dijkstra (shortest path first) algorithm would find them. Then repeat for the distances from **R5** to each other R. **[4 marks]**

For R1:
R1-R2=1
R1-R3=1
R1-R4=2
R1-R5=2

For R5:
R5-R1=2
R5-R2=1
R5-R3=1
R5-R4=1

(ok to write B1-B3-B4, etc)

c) What will change if the link between R3 and R4 stops working? **[4 marks]**

For R1:
R1-R4 becomes 3 instead of 2. The shortest path is now via R5.
Two choices: R1-R2-R5-R4=3, or R1-R3-R5-R4=3

For R5:
No changes of the distance or path.

d) Repeat part b) but with different weights on the links, as shown below. **[4 marks]**

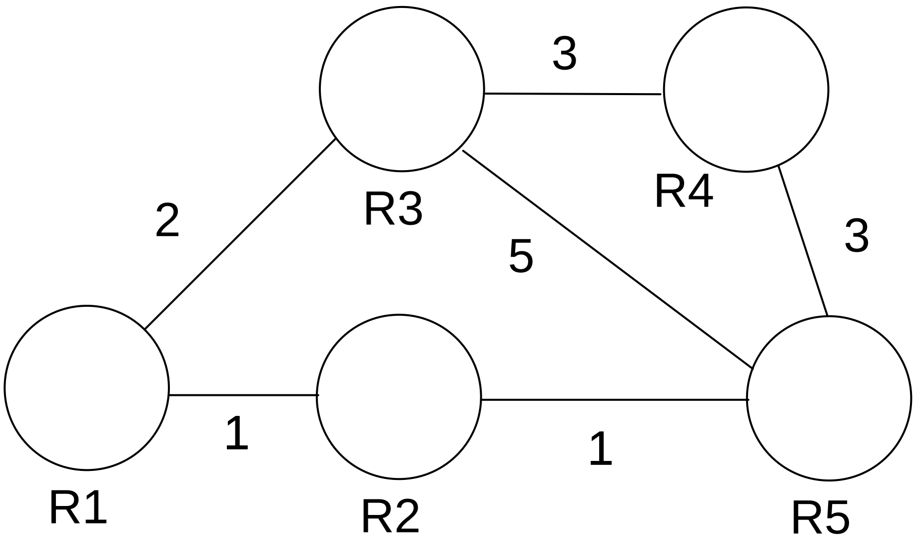

For R1:
R1-R2=1
R1-R3=2
R1-R4=5
R1-R5=2

For R5:
R5-R1=2
R5-R2=1
R5-R3=4
R5-R4=3

Note: we can get R1-R4=5 by going through R1-R3-R4=5, or R1-R2-R5-R4 = 5

e) What will change this time, if the link between R3 and R4 stops working? **[4 marks]**

For R1:
Distance (cost) does not change. The path may or may not change for R1-R4.

If the shortest path were from R1-R3-R4=5, then the route changes to R1-R2-R5-R4=5
But if your shortest path were already from R1-R2-R5-R4=5, then no route changes.

For R5:
No changes of the distance or routes.

Q2. IP fragmentation [20 marks]

a) Explain briefly what is meant by 'fragmentation and reassembly' of IP packets. **[2 marks]**
Fragmentation: breaking a large (> link MTU) into smaller (<= MTU) parts
Reassembly: collecting incoming fragments at a receiving (destination) host and joining them together again so as to reconstruct the original datagram.

b) List and explain the fields in the IPv4 header that are used to implement fragmentation. [**3 marks]**
Identification – identifies the original datagram
Flags – MF = more fragments follow, DF = Don't fragment
Fragment Offset (in 8-byte units) = byte position for first byte of fragment in original datagram.

c) How does a host know that an incoming Ipv4 packet contains a fragment? **[2 marks]**
MF = 1, or MF = 0 and Frag Offset !=0

d) Why are IPv4 fragment offsets specified in 8-byte units? **[3 marks]**
So that we can have 3 flag bits.
That leaves 13 bits for offset, enough to reach ~65536 (max IP datagram size).

e) Assume that you are attempting to send a large Ipv4 packet (6000 bytes in size including 20 bytes IP header) to an Ethernet switch that has 2000 MTU. How many IP fragments would you need to send? Assume Ethernet header/trailer size is 18 bytes. Also, how would that affect the link's effective data transport rate? **[6 marks]**
We are first limited by the 2000-18 = 1982 bytes for first IP packet. But, we need to take off 20 bytes for IP header, leaving 1962 for the IP payload. Still, 1962 bytes is not divisible by 8; the next lowest multiple of 8 is thus 1960 bytes. Therefore, we have IP payload size of 1960 + 1960 + 1960 + 100 bytes, sending four IP fragments.
Illustration below (warning: not to scale)

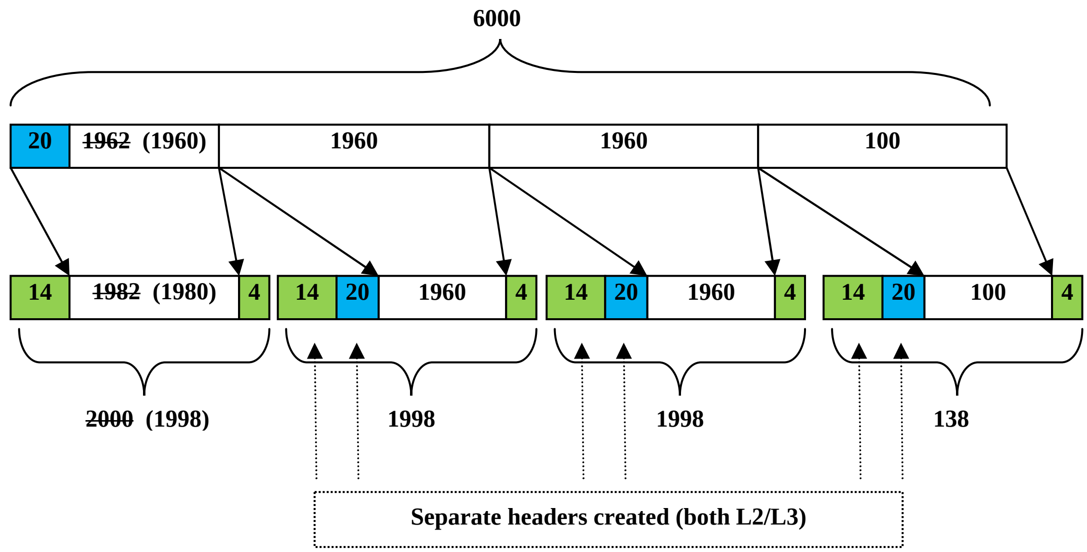

Overhead will increase: Instead of one packet, we are now sending four packets, each with 20 bytes + 18 bytes (of headers). In other words, 3 x 38 = 114 bytes more data transmitted; lowering the link effective data transport rate.

f) In IPv6, what devices perform fragmentation? How does a fragment appear in an IPv6 packet? **[2 marks]**
Sending host
In a fragmentation header (with Identification. MF and Frag Offset as in Ipv4).

g) Give one advantage and one disadvantage of the IPv6 approach to fragmentation. **[2 marks]**
Advantages:
No multiple fragmentation (1500 to 1000 to 576, etc) or Frag fields in sub-header, not present in most packets. Less computational resources required for the routers since they do not perform the fragmentation.
Disadvantages:
Hosts must discover minimum link MTU (of the entire end-to-end path) (hard).
Or use 1280 (minimum MTU for IPv6) which can be inappropriate.