Birla Institute of Technology & Science, Pilani - K.K. Birla Goa Campus
Department of CSIS

Due Date: 25.10.2012

1. Write the reviews for the technical papers

   (a) **[Jain]** Albert Greenberg, Srikanth Kandula, David A. Maltz, James R. Hamilton, Changhoon Kim, Parveen Patel, Navendu Jain, Parantap Lahiri, and Sudipta Sengupta, **VL2: A Scalable and Flexible Data Center Network**, SIGCOMM'09.

   (b) **[Lua]** Eng Keong Lua, Jon Crowcroft, Marcelo Pias, Ravi Sharma, and Steven Lim, **A Survey and Comparison of Peer-to-Peer Overlay Network Schemes**, IEEE Communications Survey and Tutorial, MARCH 2004.

   The review must include your analysis of the new concepts introduced in the paper, basic assumptions of the paper, limitations of the proposals and possible advantages of the proposals. The review size for the first paper must be in the range of 1500 to 2000 words. A more detailed review is expected of you for the second paper.

2. Consider distributing a file of F = 15 Gbits to N peers. The server has an upload rate of u_s = 30 Mbps, and each peer has a download rate of d_i = 2 Mbps and an upload rate of u. For N = 10, 100, and 1,000 and u = 300 Kbps, 700 Kbps, and 2Mbps, prepare a chart giving the minimum distribution time for each of the combinations of N and u for both client-server distribution and P2P distribution.

3. (a) In the circular DHT given, suppose that peer 3 learns that peer 5 has left. How does peer 3 update its successor state information? Which peer is now its first successor? Its second successor?

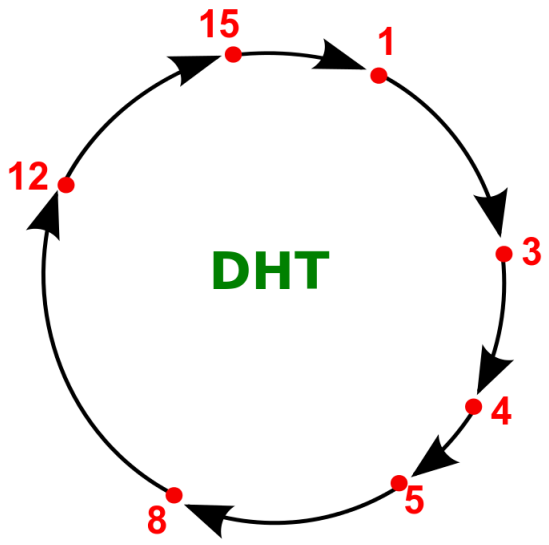

   (b) In the circular DHT given, suppose that a new peer 6 wants to join the DHT and peer 6 initially only knows peer 15's IP address. What steps are taken?

4. Consider a circular DHT with node and key identifiers in the range [0, 63]. Suppose there are eight peers with identifiers 0, 8, 16, 24, 32,40,48, and 56.

   (a) Suppose each peer can have one shortcut peer. For each of the eight peers, determine its shortcut peer so that the number of messages sent for any query (beginning at any peer) is minimized.

   (b) Repeat (a) but now allow each peer to have two shortcut peers.

5. For the given data center architecture, answer the following.

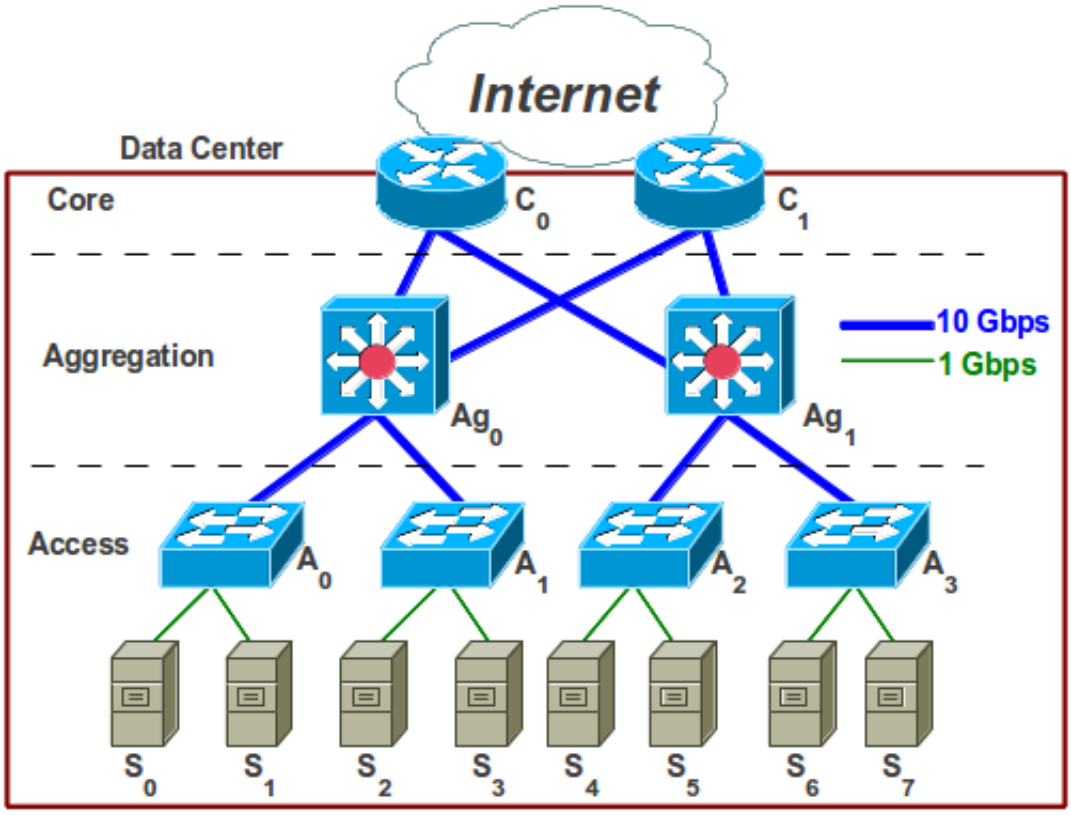

   (a) Is there any oversubscription in the architecture? According to the definition of oversubscription, what are the levels of oversubscription in this architecture?

   (b) The modulo-N hashing is to be used for ECMP. Assuming that the traffic between all the server pairs is equal, suggest the forwarding tables for the core routers, aggregate routers and the access routers.

   (c) What will happen to the traffic balancing and forwarding tables if S_1 and S_6 exchange places?

6. For the 4-ary (k=4) fat tree given in [Al-Fares], answer the following questions.

   (a) Show the IP addresses and forwarding tables of the switches in Pod_1 and Pod_3.

   (b) Show the forwarding table of the core switch with IP address 10.4.2.2.

(c) Suggest a modification to the address allocation and forwarding tables for allowing the servers to run the virtual machines.
(apart from Portland idea described in [Mysore])

(d) For 8-ary (k=8) fat tree, generate the PMACs for all the servers attached to Pod_1. Assume that each server runs only one virtual machine.

(e) How does the packet forwarding work in the Portland architecture? Is it same / different from the 2-level routing introduced in [Al-Fares]? Justify your answer.

7. (a) The SEATTLE architecture uses consistent hashing to hash the MAC address to a router. Use SHA-1 for consistent hashing and consider the last 7-bits for building the 7-bit DHT. For example, the MAC address 24:b6:fd:15:2c:7e (hashed along with the colons) produces a SHA-1 hash of *8adfaac1ddadd0fefd5444f86b2a477cd0d02e5a* in hexadecimal representation; Extracting the last 7-bits of the hash result gives us *5a* in hexadecimal or *90* in decimal.
There are two routers in the SEATTLE architecture with the following MAC addresses: R1(00:5f:02:b8:eb:14) and R2(00:5e:04:e3:a9:51). There are seven hosts in the network and their MAC addresses are given below.
H1(00:3d:02:93:c2:72), H2(00:8d:00:1a:0b:17), H3(00:4e:0c:be:e6:ed), H4(00:81:08:34:a7:07), H5(00:12:06:b8:a4:74), H6(00:71:07:fd:10:3d), and H7(00:6d:06:fc:df:99).
Point out the locations of storage (R1/R2) for the <key,value> pairs of each of these hosts based on the hash scheme described above. A new router R3(00:5e:02:db:73:11) is added to the network; Show the new locations for all the <key,value> pairs of all the hosts. Clearly indicate any changes in the location.

(b) Comment on the suitability of the the Cisco data center architecture and fat tree architecture for implementing the SEATTLE architectural idea?

8. A proposed network for Arvind Eye Hospital is given here.

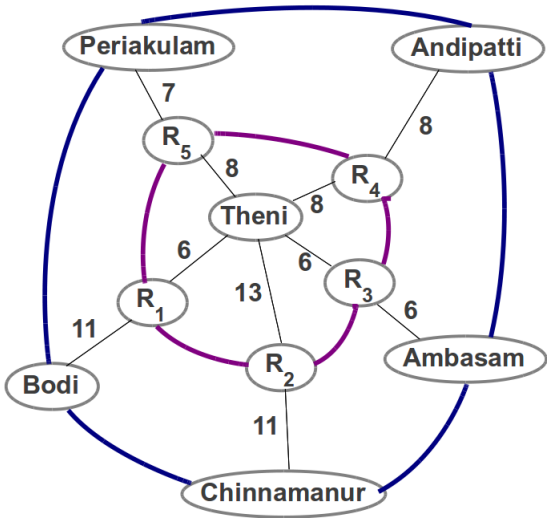

The black lines connect the remote locations (ex: Andipatti) to the hospital in Theni using relays; the black lines indicate wireless connections and the cost of using the wireless link

is shown on the network. The cost of using the blue line during the day is 6 + number of hops; the cost of using the blue line during the night is 10. The cost of using the violet line 4 + number of hops.
Using the given information, answer the following questions.

(a) On the network, use the DTLSR[1] to indicate the best paths to go from all the end nodes to Theni during the day. The resulting network is called a sink tree for Theni. How does the sink tree change during the night?

(b) The node in Chinnamanur prefers to use tree-based flooding with n=2. What is the average cost of going to Theni using the tree-based flooding? Clearly indicate all the replicas on the tree.

(c) The remote locations want to change their routing protocol to location-based forwarding. The distance metric to be used is the Eulerian distance as indicated by the formula given below.

$$\text{d} = \sqrt{(\text{lattitude difference})^2 + (\text{longitude difference})^2}$$

Assume that the relays are half-way between the remote locations and the hospital in Theni, and hence have half the cost for the distance metric for each of their links. For example, if the distance metric between Periakulam and Theni is 10, then cost of Periakulam-R5 link is 5 and the cost of R5-Theni link is 5. Now compute and illustrate the sink tree for Theni.

Other Useful URLs

(a) Two simplified explanations of consistent hashing are available at the following URLs.
   i. *http://www.tomkleinpeter.com/2008/03/17/programmers-toolbox-part-3-consistent-hashing/*
   ii. *http://www.lexemetech.com/2007/11/consistent-hashing.html*

(b) Explanation of SHA-1 is available at the following URL.
http://en.wikipedia.org/wiki/SHA-1

(c) You can use the **sha1sum** command to compute the SHA1 hash of a MAC address; Optionally, you can use the following URL for computing the SHA1 hash.
http://www.sha1hash.com/

[1] DTLSR: Delay-Tolerant Link State Routing