January 28th, 2020

Exam of Switching technologies for data centers (2019/20)

Rules for the exam. It is **forbidden** to use notes, books or calculators. Use only draft paper provided by the professor. When needed, use approximations. The answers must be provided in correct English. Any notation must be defined.
Time available: 70 minutes.

Problem A

Consider the design of a data center based on a leaf-spine topology, with an oversubscription 3:1, connecting 24,576 servers with each server interface running at 10 Gbit/s. Each switch is equipped with 64 ports running at 10 Gbit/s.

1. Design the interconnection topology, showing all the steps in the design construction.
2. Compute the total number of switches and cables required to build the data center.

Problem B

Consider the new network paradigm denoted as “Software Defined Networking” (SDN).

1. What are the flow tables?
2. Show three examples of flow rules corresponding to (i) an Ethernet switch, (ii) an IP router and (iii) a firewall.
3. Explain the differences between Single Match Tables, Multiple Match Tables and Reconfigurable Match Tables.
4. Show an example in which the three kind of tables are different for the same set of flow rules.
5. What is the difference between Openflow and P4, referring to the flow tables?

Problem C

Consider a $N \times M$ input queued switch with Virtual Output Queueing and supporting Strict Priority Queueing based on C classes. For each (input-i,output-j) pair there exist C queues: $[VOQ_{ij}^{c}]_{c=1}^{C}$, with decreasing priority level, where $c = 1$ refers to the highest priority traffic and $c = C$ refers to the lowest priority traffic (i.e., best effort traffic). Therefore, $C \times N \times M$ queues are present in total. We assume that an input traffic classifier sends the incoming packets to the correct queue.

1. Write in pseudocode a greedy algorithm to schedule the transmissions across the switching fabric, to be maximal and to maximize the number of higher priority packets that are selected at each timeslot. Define all the data structures adopted in your code.

2. Is it possible that low priority traffic will be starved indefinitely by high priority traffic? Motivate your answer with an example.

Hints for the solution

Problem A

The largest pod that can designed connects 1536 servers to 512 spine switches and it is built with $32 + 16 = 48$ switches. To connects all the servers, $24576/1536 = 16$ pods are necessary. Now the total number of leave switches connecting the pods is 128, with 4 parallel links connected to each pod.
Thus, the total number of switches is $16 \times 48 + 128 = 896$. The total number of cables is

$$24576 + 16 \times 32 \times 16 + 512 \times 16 = 24576 \times (1 + 2/3) = 40960$$

.

Problem B

See the class notes.

Problem C

See the class notes.