February 8th, 2019

Exam of Switching technologies for data centers (2018/19)

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.

1. What are the main advantages and disadvantages of managing the addressing at level 2 and at level 3?
2. Explain the difference between End-of-Row (EOR) and Top-of-Rack (TOR) architectures, highlighting the corresponding advantages and disadvantages.
3. What is the purpose of the technology denoted as “multi-chassing link aggregation”?
4. What is the difference between overlay and underlay in a data center?
5. What is the motivation to use an overlay in a data center?
6. What is the role of BGP in a data center?

Problem B

Consider a $N \times M$ input queued switch with Virtual Output Queueing, running in a network where each packet is tagged with the entrance time in the network (i.e. the packet is tagged at the first switch traversed in the network with a timestamp corresponding to the current time). Each VOQ is not FIFO, but stores the packets in increasing order of timestamp, such that the head-of-line packet at each VOQ is the oldest in the network. Let VOQ_{ij} be the queue corresponding to input i and output j and let $H[i][j]$ be its occupancy. Let T_{ij} be the timestamp of the head-of-line packet corresponding to VOQ_{ij}, represented by an integer number corresponding to the absolute time when the packet entered the network.

1. Write in pseudocode a greedy algorithm to schedule the transmissions across the switching fabric, to maximize the number of oldest packets that are selected at each timeslot.
2. Discuss the scheduler performance in terms of throughput.
3. Is it possible that a packet will be starved indefinitely in the switch? Motive your answer.

Problem C

Design a rearrangeable switch of size 2000×4000 using only basic modules of size 10×10, with the aim of minimizing the number of modules.

1. Describe the architecture.
2. Compute the total number of required basic modules.
3. Describe the configuration algorithm.
4. What would have been the approximated number of modules if the switch was designed to be strictly-non-blocking?

Hints for the solution

Problem C

In short:

$$C_{4000\times2000} = 600C_{10} + 10C_{400\times200} = 2300C_{10}$$

$$C_{400\times200} = 60C_{10} + 10C_{40\times20} = 170C_{10}$$

$$C_{40\times20} = 11C_{10}$$

The corresponding strictly non-blocking network is expected to have roughly twice the number of basic modules.