February 10th, 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 equal to 2:1, with each server interface running at 10 Gbit/s. Each switch is equipped with 60 ports running at 10 Gbit/s.

1. Design the largest possible two level (leaf-and-spine) data center.
2. Design the largest possible two level (leaf-and-spine) POD and connect 6 of them to build the largest 3 levels data center.

For each of the two data centers, show the interconnection topology and compute (i) the total number of supported servers, (ii) the total number of switches, (iii) the total number of required cables (including the ones required to connect the servers).

Problem B

Consider a detection system for Distributed Deny of Service (DDoS) attacks that tracks the number of IP packets transferred for each traffic flow. A flow x is identified by the source-destination IP addresses and the port numbers at transport layer (only if present). We assume that the data structure to store the number of packets for each flow is a hash table with 65,536 buckets, specifically designed to exploit the "power of 2 random choices" result for random load balancing.

1. What is the claim of the "power of 2 random choices" result? Why is it relevant for hash tables?
2. Describe in details two hash functions $h_1(x)$ and $h_2(x)$ specifically designed for the flow identifier considered in the problem and for the given size of the hash table.
3. Describe in pseudocode the operation to update the hash table when a new packet of flow x is processed.

Problem C

Design a Clos network, strictly non blocking, of size 27×27, with three inputs for each module of the first stage. Use only modules 3×3 through recursive factorization.

1. Draw the network at each factorization level, with all modules and links.
2. Draw the final network, with all modules and links.
3. Compute formally the complexity of the final network, in function of the complexity $C(3)$ of module 3×3. Show all the involved mathematical steps.
4. Which routing algorithm can be used to find a path between an input and an output port?

Hints for the solution

Problem A

Data center	Number of servers	Number of switches	Number of cables
2 layers	$40 \times 60 = 2400$	$60 + 20 = 80$	$2400 + 20 \times 60 = 3600$
3 layers	$1200 \times 6 = 7200$	$(30 + 20) \times 6 + 60 = 360$	$7200 + (30 \times 20 \times 6) + 3600 = 14400$

Problem B

The problem is almost identical to Ex. 140.
In this case, each hash function must map a sequence x of 96 bits (i.e., $32 + 32$ bits for the IP addresses and $16 + 16$ bits for the ports), to a sequence $h(x)$ of 16 bits (i.e., $\log_2 65536$, to index each bucket of the hash table). One possible way is to sum (using EXOR) 6 sequences of 16 bits, each of them obtained by partitioning x into blocks of 16 bits each. Notably, in the case the port is not present in the packet, all the corresponding bits in x will be set equal to zero.

Problem C

See exercise 17. The path can be computed using either BFS or Dijkstra algorithm.