June 25th, 2021

Online Exam of Switching technologies for data centers (2020/21)

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

Problem A

Assume a Bloom filter with 10 bits used to store integer numbers. The filter uses the following three hash functions: $h_1(x) = x \mod 10$, $h_2(x) = (x + 4) \mod 10$, $h_3(x) =$ floor$(\log_{10} x)$, where $\mod$ is the modulo operator.

1. Are these 3 hash functions "good" for the considered Bloom filter? Why?
2. Add the following values, showing the state of the Bloom filter after each insertion: 10001, 2007, 35.
3. Which operation (if any) would lead to a false positive event for the above Bloom filter, after the 3 insertions?
4. Which operation (if any) would lead to a false negative event for the above Bloom filter, after the 3 insertions?
5. Explain with an example why the deletion is not allowed in the above Bloom filter, after the 3 insertions.

Problem B

Design a $20,000 \times 20,000$ rearrangeable Clos network, using a recursive construction with a basic building block being a crossbar of size 10×10.

1. Compute the total number basic building blocks, showing all the involved steps.

Problem C

Consider a data center built with the switches in the table. All the servers are equipped with ports at 10 Gbps.

Switch model	Ports
Top-of-Rack	40 @ 10 Gbps plus 4 @ 40 Gbps
Spine	8 @ 40 Gbps

1. Draw the largest leaf-and-spine data center with the above switches
2. Compute the corresponding number of servers, ToR switches and spine switches
3. Compute the corresponding oversubscription ratio
4. Assume eBGP used for routing: show a feasible AS number association for each switch
5. Assume iBGP used for routing: show a feasible AS number association for each switch

Hints for the solution

Problem A

1. $h_1(x)$ and $h_2(x)$ are not independent, so the two hash functions perform poorly for a Bloom filter
2. after inserting 10001, $BF = [0100110000]$; after inserting 2007, $BF = [0101110100]$; after inserting 35, $BF = [0101110101]$.
3. search(11) will lead to a false positive
4. delete(35) and then search(10001) would lead to a false negative, thus deletion is not allowed
5. see above

Problem B

$$C_{20000} = 4000C_{10} + 10C_{2000} = (4000 + 12000)C_{10} = 16000C_{10}$$
$$C_{2000} = 400C_{10} + 10C_{200} = (400 + 800)C_{10} = 1200C_{10}$$
$$C_{200} = 40C_{10} + 10C_{20} = (40 + 40)C_{10}$$
$$C_{20} = 4C_{10}$$

Problem C

- 8 ToR + 4 spines. Servers 320.
- oversubscription ratio $40 \times 10/(40 \times 4) = 2.5$
- for eBGP, each switch is seen as an AS. Now AS number 101,102,. . . , 108 for the ToR switches and 201, 202, 203, 204 for the spine switches.
- for iBGP, one AS is shared across all the switches. Just use AS number 101 for all the switches.