PADERBORN UNIVERSITY

Advanced Networked Systems SS24

Exercise 4: Programmable Networks – A

1 Basic Concepts

Question 1: What is the core idea of software-defined networking (SDN)?

Question 2: What are the main abstractions provided by SDN?

Question 3: How does FlowVisor leverage SDN for network slicing?

Question 4: How does CoVisor enable the co-existence of different SDN control programs in a network?

Question 5: Explain what limitations of SDN are addressed by the programmable data plane.

Question 6: Explain the basic idea of the P4 language.

Question 7: Explain the Reconfigurable Match Table (RMT) architecture.

Question 8: What is the resource efficiency issue in RMT and how does dRMT solve it?

Question 9: How does NetCache store variable-length values on the switch?

Question 10: How does NetCache obtain query statistics for updating the cached keys?

Question 11: How does NetChain ensure strong consistency?

Question 12: How does programmable switches handle large input data for in-network aggregation?

2 Bloom Filters

Assume we are given a set of 2000 unique records and we do not have enough space to store all these records. We use different hash-based approximate data structures for membership detection.

Question 1: What is the false positive rate of a hash table of 8000 slots with binary indicators?

Question 2: What is the false positive rate of a Bloom filter of 8000 slots using 10 hash functions?

Question 3: If you compare the above two results, you may notice that adding more hash functions does not decrease the false positive rate in this case. Could you explain why? How can we improve the false positive rate?

Question 4: What is the false negative rate of the above Bloom filter?

3 Count-Min Sketch

Consider a count-min sketch with three hash functions and 10 slots for each hash function. The hash functions are: $h_1(x) = 2x\%10$, $h_2(x) = 3x\%10$, and $h_3(x) = 7x\%10$. The current state of the count-min sketch is shown in the following.

Hash	**0**	**1**	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**
$h_1(x)$	12	83	49	16	0	22	17	8	91	12
$h_2(x)$	8	3	11	77	23	42	83	18	69	31
$h_3(x)$	65	48	29	14	70	15	74	6	52	78

Question 1: What is the estimated count for $x = 75$, $x = 12$, $x = 111$, respectively?

Question 2: Assume records $36, 9, 23$ have arrived. What is the state of the hash table?

Question 3: What is the relationship between the estimated count and the real count?