

Handout 8:

---

1. (344 17SU Final)

a) You are using Twitter to study the patterns of migratory birds. A large number of volunteer citizen scientist from all over the world assist you by tweeting bird sightings using the hashtag #BirdSurvey. The response has been overwhelming and you now have a collection of over 5 million tweets in JSON format each geotagged and with specific information about the number and type of birds seen. The tweets are in a large NoSQL document store where the key is the tweetID and the value is the JSON representation of the tweet. For each of the following MapReduce programs give an explanation of what it calculates. Assume that the function `getSightings(tweet)` returns a list of (bird,count) pairs within that tweets text.

<pre>map(id, tweet):   total = 0   for bird, count in getSightings(tweet):     total += count   EmitIntermediate(tweet.country, total)</pre>	<pre>reduce(key, values):   Emit(key, sum(values))</pre>
--	--

Description of output:

<pre>map(id, tweet):   EmitIntermediate(tweet.country, tweet.username)</pre>	<pre>reduce(key, values):   unique_values = set(values)   Emit(key, unique_values.length)</pre>
--	---

Description of output:

b) Write pseudo-code for the map and reduce function below. The output of the reduce function should answer each questions (roughly equivalent SQL is also given).

Find the total number of tweets each username sent.

```
SELECT T.username, count(*)
FROM tweets T
GROUP BY T.username;
```

map(id, tweet):

reduce(key, values):

2. (344 17WI Final)

Given the following query and statistics:

```
SELECT *
```

```
FROM R, S
```

```
WHERE R.a = 10 AND R.a = S.b
```

T(R) = 100,000

V(R, a) = 300

min(R.a) = 0

max(R.a) = 1000

T(S) = 40,000

V(S, b) = 20

min(S.b) = 0

max(S.b) = 1000

Assume there are no indexes on R or S. There are 4 nodes (N1, N2, N3, N4), and each node has enough main memory to hold its partition of R and S tuples. We partition R and S using one of the following schemes:

i) block partitioned, with each node holding 25,000 tuples of R and 10,000 tuples of S.

ii) range partitioned in the following way:

N1:  $0 \leq R.a < 250$  (same for S.b)

N2:  $250 \leq R.a < 500$  (same for S.b)

N3:  $500 \leq R.a < 750$  (same for S.b)

N4:  $750 \leq R.a \leq 1000$  (same for S.b)

a) Under partition scheme i), how many tuples do you expect each partition to generate if the selection predicate was applied first?

b) Instead of selection, suppose we evaluate the join predicate first. How many tuples do you expect each node to send to other nodes if we use broadcast join on R under partition i)? "broadcast join on R" means R is broadcast to other nodes.

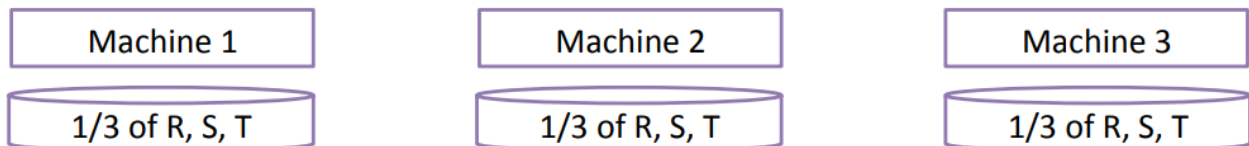
c) Under partition scheme ii), how many tuples do you expect each node to generate if the selection predicate was applied first?

d) Your roommate comes up with yet another partitioning scheme: block partition R as in i), but replicate S entirely on all nodes. She finds that the query now runs faster even when compared to hash partitioning on R.a and S.b (assume the nodes have enough memory to hold all tuples in either scheme). This is surprising because both schemes need to perform no network I/O for the join, and hash partitioning should read strictly less of S on every node, since S is partitioned instead of replicated. Provide one explanation for how this could occur.

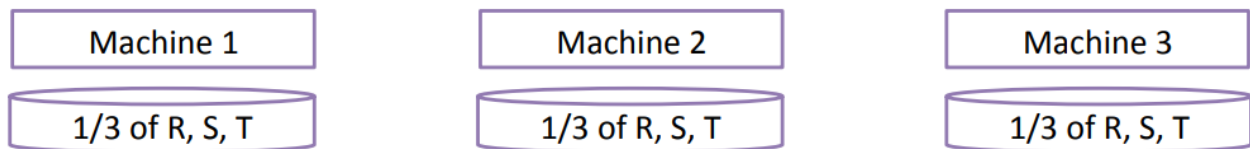
3. (344 15AU Final)

a) Consider relations  $R(a,b)$ ,  $S(c,d)$ , and  $T(e,f)$ . All three are horizontally partitioned across  $N = 3$  machines as shown in the diagram below. Each machine locally stores approximately  $1/N$  of the tuples in  $R$ ,  $S$ , and  $T$ . The tuples are randomly organized across machines (i.e.,  $R$  is block partitioned across machines). Show a relational algebra plan for the following query and how it will be executed across the  $N = 3$  machines. Use hash-join (a.k.a shuffle-join) operators. Include operators that need to re-shuffle data and add a note explaining how these operators will re-shuffle that data.

```
SELECT *  
FROM R, S, T  
WHERE R.b = S.c AND  
      S.d = T.e AND  
      (R.a - T.f) > 100
```



b) Now consider the case where the R relation is very large and both S and T are very small. Show a plan that uses broadcast joins to compute the result of the query.



c) Explain how the query would be executed in MapReduce. Make sure to specify the computations performed in the map and the reduce functions. Use hash-joins (shuffle joins). No need to give the pseudocode. Just state what each function does and what it outputs in plain text format.