CSE 444 – Homework 6
Parallelism and Distribution

Name: _____

Question	Points	Score
1	25	
2	25	
Total:	50	

1 Parallel Data Processing

1. (25 points)

   (a) (15 points) Consider two relations `R(a,b)` and `S(c,d)` that are both horizontally partitioned across $N = 3$ nodes as shown in the diagram below. Each node locally stores approximately $\frac{1}{N}$ of the tuples in `R` and $\frac{1}{N}$ of the tuples in `S`. The tuples of `R` are *randomly* organized across machines (i.e., `R` is block partitioned across machines) while the tuples of `S` are *hash-partitioned* on `S.c`.

   Show a relational algebra plan for the following query and how it will be executed across the $N = 3$ machines. Pick an *efficient* plan that leverages the parallelism as much as possible. Include operators that need to re-shuffle data and add a note explaining how these operators will re-shuffle that data. For example, if you need to re-hash the data, add a "hash" operator into your query plan.

   Draw the parallel query plan. Indicate the edges that re-shuffle data across machines by drawing them as dashed lines:

   Note: Your plan will be more efficient if you push aggregations down. Can you compute partial aggregates before shuffling data? Can you compute partial aggregates before the join?

```
SELECT a, avg(d) as avg
FROM R, S
WHERE R.b = S.c
AND S.d > 0
GROUP BY a
```

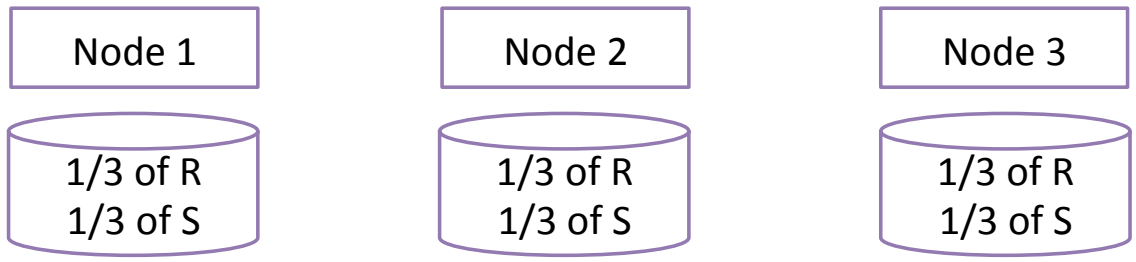

Answer:

(b) (10 points) Assume the same query as before, only now the data is distributed on 300 servers instead of just 3 servers. We expect a linear speedup, in other words we expect the runtime to be about 100 times faster. However, if the values of some attribute are skewed, then the performance of a parallel query plan can be far from a linear speedup. Indicate with of the attributes below, if skewed, can significantly prevent your query plan from achieving linear speedup.

- Skew on attribute `R.a`.
- Skew on attribute `R.b`.
- Skew on attribute `S.c`.
- Skew on attribute `S.d`.

2 Distribution and Replication

2. (25 points)

   (a) (15 points) In the two-phase commit protocol, describe what happens if a subordinate receives a PREPARE message, replies with a YES vote, crashes, and restarts.

(b) (10 points) Explain the benefits and challenges of asynchronous replication (also called lazy replication) in contrast to synchronous replication. Discuss both the configuration that uses a single master and one that uses multiple masters.