Section 8 Worksheet: Parallel DBMS

Problem 1

Create a Parallel Relational Algebra plan using the following database setup and query. Tuples of R are block-partitioned evenly across all 3 machines.

```
SELECT a, count(b) as countb
FROM R
WHERE a > 0
GROUP BY a
```

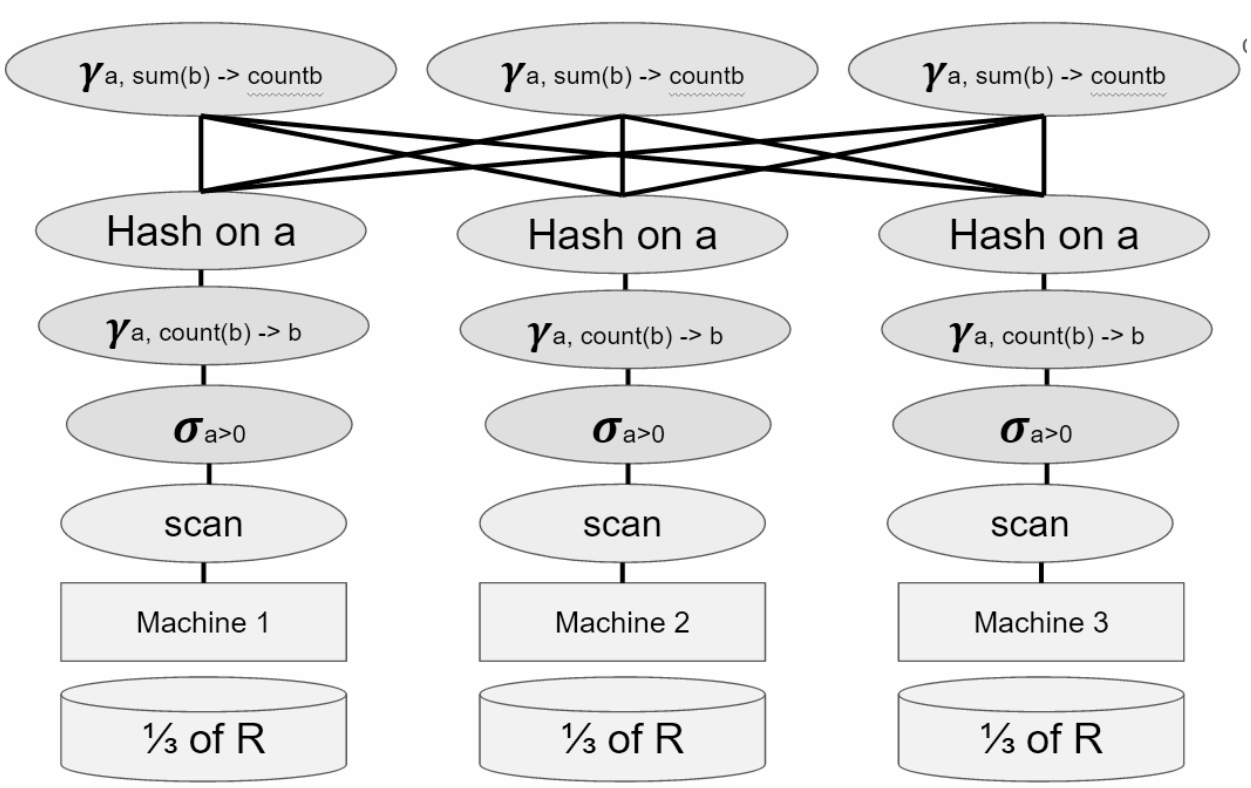

Problem 2

Create a Parallel Relational Algebra plan using the following database setup and query. Tuples of R is block-partitioned evenly across all 3 machines and S is hash partitioned on S.b.

SELECT R.a, avg(S.c) as myavg
FROM R, S
WHERE R.b = S.b
AND R.a <= 10 and S.c > 20
GROUP BY R.a;

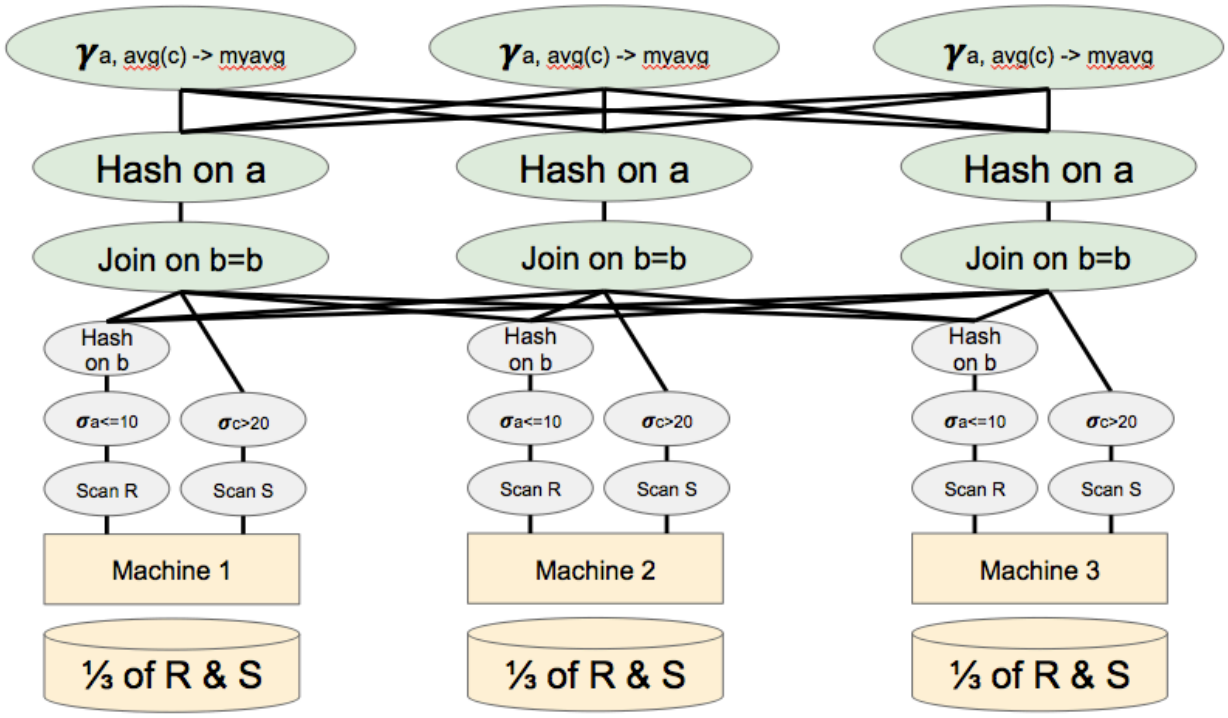

Problem 3

Create a Parallel Relational Algebra plan using the following database setup and query. Tuples of A and B are hash partitioned on y.

```
SELECT A.x
FROM A,
WHERE NOT EXISTS (
        SELECT *
        FROM B
        WHERE A.y = B.y)
HAVING sum(A.y) > 100
GROUP BY A.x;
```