


CS5300

Homework #3

- 1) Problems 18.13, 19.15, and 19.18 PP. 689 and 740.
- 2) Consider the relations $r_1(A, B, C)$, $r_2(C,D,E)$, and $r_3(E,F)$, with primary keys A, C, and E, respectively. Assume that r_1 has 1,000 tuples, r_2 has 1,500 tuples, and r_3 has 750 tuples. Estimate the size (cardinality) of $r_1 \bowtie r_2 \bowtie r_3$, and give an efficient strategy for computing the join.
- 3) Let relations $r_1(A, B, C)$ and $r_2(C,D,E)$ have the following properties: r_1 has 20,000 tuples, r_2 has 45,000 tuples, 25 tuples of r_1 fit on one block, and 30 tuples of r_2 fit on one block:

- a. Devise an efficient nested loop algorithm to perform $r_1 \bowtie r_2$, and
- b. Estimate the number of block accesses.

- 4) Consider the relations $r_1(A, B, C)$, $r_2(C,D,E)$, and $r_3(E,F)$. Assume that there are no primary keys, except the entire schema. Let $V(C, r_1)$ be 900, $V(C, r_2)$ be 1,100, and $V(E, r_3)$ be 100. Assume that r_1 has 1,000 tuples, r_2 has 1,500 tuples, and r_3 has 750 tuples. Estimate the size (cardinality)



,and give an efficient strategy for computing the join. Note $V(A, r)$ is the number of distinct values of A in r.