

Homework 5 – Solution

18.7 The text of exercise is somewhat ambiguous or not intuitive. First, there is no description on how to find manager in accounting department. Second, age is something that is rarely greater than 100 and salary – as of now – should be more than 100. The exercise would be more appropriate if it asks for the manager whose salary is greater than 1000 times his age. Despite of this ambiguity or unintuitiveness, let us stick to what is given and let us find all manager whose salary is greater than their age. The following plan is an appropriate one:

1. At site C performs a SELECT statement that returns the SSnum and Age of all employees for which Title='Manager'. Let the result be T_1 .
2. Send T_1 to B. Since there are 50 managers, the total amount of information that must be sent is $50 * (9 + 2) = 550$ bytes.
3. At site B, perform the join of T_1 with EMP21, select those names whose salary is greater than their age, and send the result to site A.

Because we do not have information to estimate the result, we cannot estimate the communication costs for the third step. However, this cannot be more than $50 * 15 = 750$ bytes (at most, there are 50 managers whose salary is greater than their age).

Notice that even though we cannot estimate the cost of the third step, the above plan is still the best possible plan because the other plans will cost more. Indeed, the two alternatives are:

1. send both to A; this plan requires us to send 100,000 tuples from B to A and 100,000 tuples from C to A; the cost is much higher than 550 bytes plus a maximal of 750 bytes
2. select SSnum, Name, and Salary at B, send to C, join then send to A is another plan); here we need to send 100,000 tuples from B to C; again, the cost is much higher than 550 bytes plus 750 bytes.

18.14. The best alternative is the one given in the text.

1. At each of the warehouse sites, select all tuples for which the title is manager.
2. Send these tables to the headquarter. Then send the same to site A.

The number of bytes to be sent is exactly the same as in the text - 1300 bytes from all the warehouses to the headquarter plus 675 bytes from the headquarter to A - which is 1975 bytes.