
Information Systems ***June 12, 2014***
Master of Science in Computer Engineering

Exercise 1 (9 points)

Let's consider the following relational schema:

EMPLOYEE(EName, Phone, Address, Date_of_birth)
DEPARTMENT(DName, _Manager, Area, Budget)
CONTRACTS(Contract_Id, Empl, Depart, Salary, Year)

Primary keys are underlined in the relations. Moreover, Empl in CONTRACTS is foreign key of EMPLOYEE and Depart in CONTRACTS is foreign key of DEPARTMENT.
An employee can have one or more contracts with different departments in a year.

Assume that:

$n_{EMPLOYEE}$ =10.000
$n_{DEPARTMENT}$= 50
$n_{CONTRACTS}$ = 1.000.000

V(Depart, CONTRACTS) = 50
V(Empl, CONTRACTS) = 10.000
V(Area, DEPARTMENT) = 10
V(Year, CONTRACTS) = 20

Given the query:
Name and Address of employees holding a contract in the year 2012 with a department whose Area is "Computer Science"

1) express the query as a relational-algebra expression;
2) show the basic steps of the query optimization process in terms of relational-algebra expression transformations
3) give an efficient strategy for computing the query.

Exercise 2 (6 points)
Consider the following schedule of concurrent transactions:

S: `w3(z) w2(y) r2(z) w3(y) w3(x) r2(x) w1(y) w2(x)`

1) Show if it is conflict serializable (CSR) or view serializable (VSR). Explain why. If serializable, show equivalent serial schedules.
2) Apply the rigorous two-phase locking protocol to the schedule. Is the schedule accepted?
3) Apply the timestamp-ordering protocol to the schedule, assuming that aborted transactions are immediately restarted. Is the schedule accepted?

Exercise 3 (6 points)
Consider an empty B+-tree with m= 5.

1) Show the B+-tree after the insertion of the following values of the search key:
16 21 8 5 13 30 19 50 40 12

2) Show the form of the B+-tree after each operation of the sequence:
Delete 5; Delete 13

Exercise 4 (9 points)

Let r=(A,B,C), with A a key.
Assume
nr = 500.000 number of records in the relation
Lr = 100 byte size of a record (fixed length records)
LA = 6 byte size of attribute A
Lp = 4 byte size of a pointer
Lb = 1000 byte size of a block
Sequential file organization on search-key A.

1. Show the number of blocks of a **B+-tree index** on search-key A, assuming that each block contains 70% of the records that can fit in a block.

2. Outline the steps in answering the following queries, showing **the best strategy** and **the cost** in terms of number of block transfers from disk:

1) select * from r where A=xxx;

2) select * from r where 50.000 <= A < 250.000;
assuming A uniformly distributed on the interval [1; 1.000.000]

3) select * from R where B=xxx;