
Information Systems *February 20, 2014*

Master of Science in Computer Engineering

Exercise 1 (9 points)

Let's consider the following relational schema for an insurance company, customers and policies:

CUSTOMER(Cust_Id, Name, Age, City)
BRANCH(Branch_Id,_Branch_name, nEmployee, Branch_city)
POLICY(Policy_number, Cust_Id, Branch_Id, expiry_date)
Primary keys are underlined in the relations. Moreover, Cust_Id in POLICY is foreign key of CUSTOMER and Branch_Id in POLICY is foreign key of BRANCH.
A customer can have more than one policy in the same or in different branches.

Assume that:

$n_{CUSTOMER}$ =50.000	V(Cust_Id, POLICY) = 50.000
n_{BRANCH}= 10	V(Branch_Id, POLICY) = 10
n_{POLICY} = 100.000	V(City, CUSTOMER) = 15
	V(Branch_city, BRANCH) = 5
	V(expiry_date, POLICY) = 20

Given the query:
"Branch name of branches located in Florence holding policies with expiry date = 2018 stipulated by customers leaving in Pisa."

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: r1(x) r1(v) r3(z) w2(z) r3(x) w3(y) w1(y) w2(v)

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= 4.

1) Show the B+-tree after the insertion of the following values of the search key:
36 12 18 25 3 70 19 5 40 30

2) Show the form of the B+-tree after the following operation:
Delete 70.

Exercise 4 (9 points)

Let r=(A,B,C), with A a key.
Assume

nr = 1.000.000	number of records in the relation
Lr = 50 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

Heap file organization.

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; 10.000.000]

3) select * from R where B=xxx;