
Information Systems July 1, 2016
Master of Science in Computer Engineering

Exercise 1 (9 points)
Let's consider the following relational schema for a group of insurance companies located in different cities:

CUSTOMER(Id_cust, Name, Gender, YearOfBirth)
INSURANCE_COMPANY(Id_company, _Director_name, Budget, City)
POLICY(Policy_number, Id_cust, Id_company, Policy_type, expiry_date)

Primary keys are underlined in the relations. Moreover, Id_cust in POLICY is foreign key of CUSTOMER and Id_company in POLICY is foreign key of INSURANCE_COMPANY.
A customer can have more than one policy in the same company or in different companies.

Assume that:

$n_{CUSTOMER}$ =5000	V(Id_cust, POLICY) = 5000
$n_{INSURANCE\_COMPANY}$ = 28	V(Id_company, POLICY) = 28
n_{POLICY} = 800.000	V(City, INSURANCE_COMPANY) = 7

Given the query:
Name of female customers with policies stipulated with companies located in Pisa or Livorno.

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(y) r2(y) w2(x) r3(z) r1(y) w3(x) r2(z) w1(z) w1(x)

1) Show if S 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.
3) Apply the timestamp-ordering protocol to the schedule, assuming that aborted transactions are immediately restarted.

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:
72 5 8 29 19 24 22 9 13 11 6

2) Show the form of the B+-tree after each operation of the sequence:
Insert 2; Insert 3 ; Delete 19.

Exercise 4 (9 points)

Let r=(A,B,C), with primary key A uniformly distributed on the interval [1; 1.000.000].

Assume
nr = 100.000 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 = 1024 byte size of a block
Heap file organization on A.

1. Show (a) the minimum and (b) the maximum height of a B+tree tree index on search-key A.

2. Outline the steps in answering the following queries, the **best strategy** and the **cost** in terms of number of block transfers from disk in case (a):
1) select * from r where A=xxx;
2) select * from r where 20.000 <= A < 120.000;
3) select * from r where B=xxx