
Information Systems ***January 14, 2014***
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, Age, City_cust)
INSURANCE_COMPANY(Id_company, Director_name, nEmployee, City)
POLICY(Policy_number, Id_cust, Id_company, 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:

$n_{CUSTOMER}$ =2100
$n_{INSURANCE\_COMPANY}$ = 20
n_{POLICY} = 100.000

V(Id_cust, POLICY) = 2100
V(Id_company, POLICY) = 20
V(Age, CUSTOMER) = 15
V(City, INSURANCE_COMPANY) = 4

Given the query:
"Policy_number of policies stipulated by customers who are 18 with companies located in Milan."

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

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 the following B+-tree with m= 4.

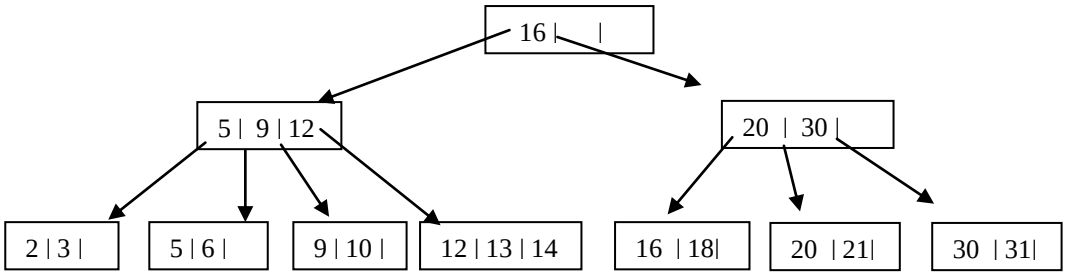

Show the form of the B+-tree after each operation of the sequence:
Delete 18; Delete 31; Insert 15.

Exercise 4 (9 points)

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

nr = 400.000	number of records in the relation
Lr = 200 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.

1. Show the number of blocks of the file.

2. Show the number of blocks of a **static multilevel index** on search-key A.

3. Outline the steps in answering the following queries, 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]