
Information Systems ***January 14, 2015***
Master of Science in Computer Engineering

Exercise 1 (9 points)

Let's consider the following relational schema of a Library:

AUTHOR(<u>authID</u>, AuthName, Nationality, Age)
TITLE(<u>titleID</u>, title, Cathegory, Language, Year)
REL_TITLE_AUTHOR(<u>titleID, authID</u>)

Primary keys are underlined. Moreover, <u>titleID</u> in REL_TITLE_AUTHOR is foreign key of TITLE and <u>authID</u> in REL_TITLE_AUTHOR is foreign key of AUTHOR.
There is a many to many relations between TITLE and AUTHOR.

Assume that:

N_{AUTHOR} =400.000
N_{TITLE}= 100.000
$N_{REL\_TITLE\_AUTHOR}$ = 600.000

V(<u>titleID</u>, REL_TITLE_AUTHOR) = 100.000
V(<u>authID</u>, REL_TITLE_AUTHOR) = 400.000
V(Year, TITLE) = 20
V(Nationality, AUTHOR) = 5

Given the query:
Title of books published in the year "YYY" by at least one author whose nationality is "XXX".

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) r2(y) w3(x) w3(y) r1(y) w2(x) r3(z) w1(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:
2 26 13 10 18 5 9 55 45 27

2) Show the form of the B+-tree after each operation of the sequence:
Insert 20; Insert 25; Insert 19; Insert 22; Delete 10

Exercise 4 (9 points)

Let r=(A,B,C), with A a key.
Assume
nr = 100.000 number of records in the relation
Lr = 50 byte size of a record (fixed length records)
LA = 10 byte size of attribute A
Lp = 4 byte size of a pointer
Lb = 1024 byte size of a block
Sequential file organization on search-key A. We assume that blocks are not chained.

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

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 < 70.000;
assuming A uniformly distributed on the interval [1; 200.000]

3) select * from R where B=xxx;