
Information Systems 17/02/2016
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, Year_of_Birth)
TITLE(<u>titleID</u>, title, Cathegory, Language, Year_of_Publ, Publisher)
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
$N_{REL\_TITLE\_AUTHOR}$ = 800.000

V(<u>titleID</u>, REL_TITLE_AUTHOR) = 100.000
V(<u>authID</u>, REL_TITLE_AUTHOR) = 400.000
V(Year_of_Publ, TITLE) = 20
V(Publisher, TITLE) = 10
V(Year_of_Birth, AUTHOR) = 80

Given the query:
TitleID of books published by Springer in the last 5 years and with at least one author born in 1990.

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

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)

1) Construct a B+-tree for the following sequence of values, assuming node fanout m =4:
   (2, 6, 9, 10, 14, 18, 22, 25, 28, 29)
   The tree is initially empty and values are inserted in ascending order.

2) Show the form of the tree after the operation: Delete 10

Exercise 4 (9 points)

Let r=(A,B,C), with primary key A uniformly distributed on the interval [1; 100.000].

Assume
nr = 50.000 number of records in the relation
Lr = 20 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
Static multilevel index on a sequential file organization on A.

1. Outline the steps in answering the following queries and the cost in terms of number of block transfers from disk:

   1) select * from r where A=xxx;
   2) select * from r where 2.000 <= A < 20.000;
   3) select * from r where B=xxx