
Information Systems ***July 3, 2015***
Master of Science in Computer Engineering

Exercise 1 (9 points)

Let's consider the following relational schema of a one day booking system:

SAILOR(Code, Name, Skill, City, DateOfBirth)
BOAT(ID, BoatName, Harbour, Category)
RESERVATION(Code, ID, Date)

Primary keys are underlined. Moreover, Code in RESERVATION is foreign key of SAILOR and ID in RESERVATION is foreign key of BOAT.

Assume that:

N_{SAILOR} =2.000
N_{BOAT}= 50
$N_{RESERVATION}$= 1.500.000

V(Code, RESERVATION) = 2000
V(ID, RESERVATION) = 50
V(Category, BOAT) = 5
V(Harbour, BOAT) = 5
V(Skill, SAILOR) = 3
V(City, SAILOR) = 10

Given the query:
Reservations made by sailors from Pisa for 'A' or 'B' category boats.

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) w3(z) w4(y) w2(x) r2(y) r3(x) w5(x) r3(y) r1(z)

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 schedule S accepted?
3) Apply the timestamp-ordering protocol to the schedule, assuming that aborted transactions are immediately restarted. Is schedule S accepted?

Exercise 3 (6 points) Let us consider the following extendable hashing organization with blocking factor **F = 2**

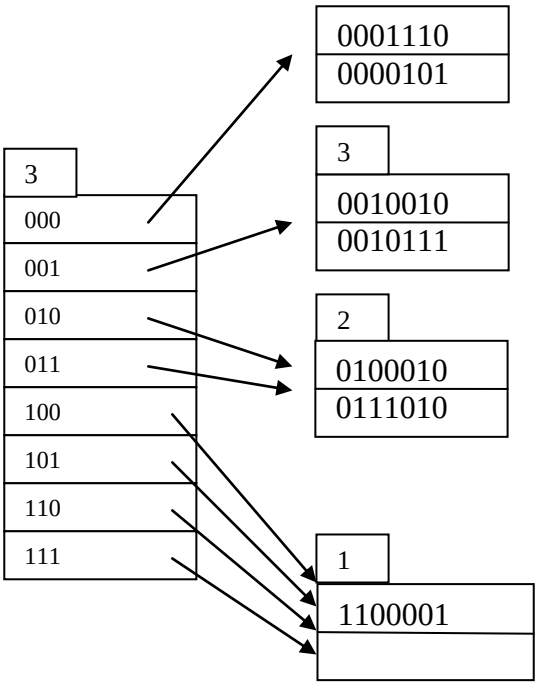

Show how the organization changes after the following operations:

- insert of **0100000**
- insert of **1110000**
- insert of **1011111.**

Exercise 4 (9 points)

Let r=(A,B,C), with A a key uniformly distributed on the interval [1; 10.000.000]
Assume

nr = 2.000.000	number of records in the relation
Lr = 100 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

Sequential file organization on search key A.

1. Show (a) the minimum and (b) the maximum number of blocks of a B+-tree 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 10.000 <= A < 40.000;