

Database Management Systems

Written Exam

25.06.2010

First name		Last name	
Student number		Signature	

Instructions for Students

- Write your name, student number, and signature on the exam sheet and on every solution sheet you hand in.
- This is a closed book exam: the only resources allowed are blank paper, pens, and your head. Use a pen, not a pencil.
- Write neatly and clearly. The clarity of your explanations affects your grade.
- You have 2 hours for the exam.

Good luck!

Reserved for the Teacher

Exercise	Max. points	Points
1	20	
2	8	
3	20	
4	24	
5	8	
6	8	
7	12	
Total	100	

Exercise 1 (20 pt) Answer the following questions:

- a. What is a pinned block?
- b. For which operation is a clustering file organisation advantageous?
- c. In a multilevel index, what is indexed by the outer index?
- d. Describe a uniform hash function for a search-key value with range $[1, 1000]$ and 20 buckets.
- e. What are two index structures that can efficiently handle multiple-key queries?
- f. What is always applicable: materialized evaluation or pipelined evaluation?
- g. How many different join orders exist for $r_1 \bowtie r_2 \bowtie r_3$?
- h. What are the 3 steps of query processing?
- i. What is a cascading rollback?
- j. What is the main idea of multiversion concurrency protocols to increase concurrency?

Exercise 2 (8 pt) Consider the following file organization using fixed-length records and a free list.

header					
record 0	BMW	1990	red	10	
record 1					
record 2	BMW	1991	red	2	
record 3	Fiat	1990	white	3	
record 4					
record 5	Fiat	1991	blue	3	
record 6					
record 7	Ford	1990	blue	1	

Show the structure of the file after each of the following operations (in that order):

- a. Insert(BMW,1991,blue,6)
- b. Delete(record 2)
- c. Insert(Ford,1990,white,7)

Exercise 3 (20 pt) Consider the following relation r :

	<i>Name</i>	<i>Course</i>	<i>Grade</i>
r_0	Tom	ITP	30
r_1	Tom	DMS	21
r_2	Aron	CSA	18
r_3	Ann	OS	18
r_4	Ann	DMS	25
r_5	Nick	ITP	27
r_6	Nick	DSA	23
r_7	Nick	IDB	26
r_8	Sue	ITP	28
r_9	Sue	CSA	19

Show the following index structures and file organisations:

- A primary dense B⁺-tree index on *Course*. Assume $n = 3$ for the B⁺-tree. The tuples are inserted in the order r_0, \dots, r_9 . Show the tree after inserting the first five tuples and after inserting all tuples.
- An extensible hash table on *Grade* with the hash function $h(n) = n \bmod 8$. Each bucket holds at most 2 tuples. The tuples are inserted in the order r_0, \dots, r_9 . Show the structure after inserting the first five tuples and after inserting all tuples.
- Assume that you have a B⁺-tree index on *Course* and a (separate) B⁺-tree index on *Grade*. Then consider the following query:

```
SELECT * FROM r WHERE Course = 'ITP' AND Grade = 30
```

Describe 3 different evaluation strategies for this query that take advantage of the indexes, using one of the indexes or both.

Exercise 4 (24 pt) Let $r(A, B)$ and $s(A, C)$ be two relations with the following characteristics: $|r| = 45.000$, $|s| = 20.000$, A is primary key in both relations and equally distributed between 1 and 1.000.000, and s has a primary B⁺-tree index on attribute A with 100 search-key/pointer pairs per node. A single block can contain 25 tuples of r , 30 tuples of s , or 1 node of the index.

- Determine the number of blocks needed for r , s , and the index, respectively.
- Determine the access strategy and determine the number of block IOs for the following selection queries:
 - $\sigma_{A=100.000}(s)$
 - $\sigma_{A<100.000}(s)$
 - $\sigma_{A>100.000}(s)$
- Determine the number of block IOs for the following evaluation plans for $r \bowtie s$ when 3 buffer blocks in memory are available:
 - Plan p1: Block nested loop join
 - Indexed nested loop join using the index in a)
 - Plan p3: Hash join
- For the hash join in plan p3 above a partition of s need to fit entirely in main memory. Assume a main memory buffer size of 12 blocks. How should the buffer blocks be used and what would be a useful hash function such that the number of s -partitions is minimal, i.e., the partitions are maximal. (Assume that the A -values are perfectly distributed)

Exercise 5 (8 pt) Proof that the following expressions do not hold:

- a. $\sigma_\theta(E_1 \cup E_2) = \sigma_\theta(E_1) \cup E_2$
- b. $\pi_A(E_1 - E_2) = \pi_A(E_1) - \pi_A(E_2)$

Exercise 6 (8 pt) Given is the following schedule over transactions T_1, T_2, T_3 :

T_1	T_2	T_3
	read(Z) read(Y) write(Y)	
read(X) write(X)		read(Y) read(Z)
		write(Y) write(Z)
read(Y) write(Y)	read(X)	
	write(X)	

Answer the following questions and explain your answers:

- a. Draw the conflict graph of this schedule and show whether the schedule is conflict serializable or not.
- b. Is the schedule view serializable to $\langle T_1, T_2, T_3 \rangle$?

Exercise 7 (12 pt) Consider the following two transactions:

T_1 : read(A);
 read(B);
if A=0 **then** B := B + 1;
 write(B).
 T_2 : read(B);
 read(A);
if B=0 **then** A := A + 1;
 write(A).

- a. Add lock and unlock instructions to T_1 and T_2 so that they observe the two-phase locking protocol.
- b. Show a concurrent schedule of T_1 and T_2 that results in a deadlock? Show also the evolution of the wait-for graph.
- c. For the schedule in b.) what happens under the wait-die deadlock prevention protocol?

Solution 1

- a. Memory block that is not allowed to be written back to disk as long as it is pinned.
- b. Join
- c. The inner (primary) index
- d. $h = n \text{ mod } 20$
- e. Bitmap index, grid file index
- f. Materialized evaluation
- g. 12
- h. Parsing/translation, optimization, evaluation
- i. A single transaction failure leads to a series of transaction rollbacks
- j. Keep old versions of data items such that reads are always successful

Solution 2

- a. Insert(BMW,1991,blue,6)

header					
record 0	BMW	1990	red	10	
record 1	BMW	1991	blue	6	
record 2	BMW	1991	red	2	
record 3	Fiat	1990	white	3	
record 4					
record 5	Fiat	1991	blue	3	
record 6					
record 7	Ford	1990	blue	1	

- b. Delete(record 2)

header					
record 0	BMW	1990	red	10	
record 1	BMW	1991	blue	6	
record 2					
record 3	Fiat	1990	white	3	
record 4					
record 5	Fiat	1991	blue	3	
record 6					
record 7	Ford	1990	blue	1	


- c. Insert(Ford,1990,white,7)

header					
record 0	BMW	1990	red	10	
record 1	BMW	1991	blue	6	
record 2	Ford	1990	white	7	
record 3	Fiat	1990	white	3	
record 4					
record 5	Fiat	1991	blue	3	
record 6					
record 7	Ford	1990	blue	1	


Solution 3

a. B⁺-tree index

– after inserting r_0, \dots, r_4 :



– after inserting all tuples:




b. Extensible hash file organization:

• The hash function gives:

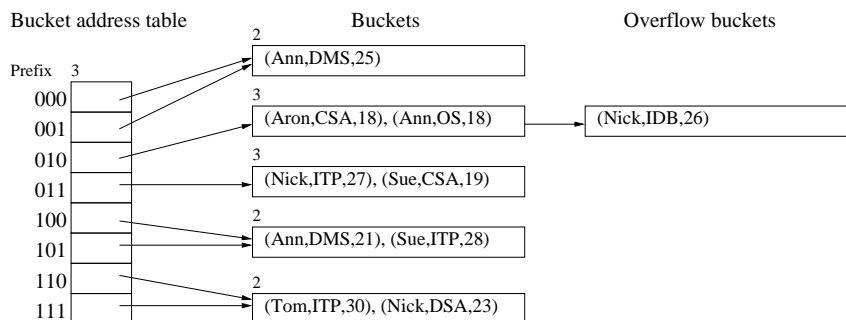
$$h(18) = 010, h(19) = 011, h(21) = 101, h(23) = 111, h(25) = 001, \\ h(26) = 010, h(27) = 011, h(28) = 100, h(30) = 110$$

• Overflow buckets are used, if a bucket is already full.

– after inserting r_0, \dots, r_4 :



– after inserting all tuples:



c. The 3 evaluation strategies are:

1. Use index on *Course* to find all tuples with *Course* = 'ITA'; then test for *Grade* = 30.
2. Use index on *Grade* to find all tuples with *Grade* = 30; then test for *Course* = 'ITA'.
3. Use index on *Course* to find pointers to all records with a *Course* = 'ITA'. Similarly, use index on *Grade* to find pointers to all records with a *Grade* = 30. Take the intersection of the two pointer sets.

Solution 4


- a. Data blocks for r : $b_r = \lceil 45.000/25 \rceil = 1.800$ blocks
Data blocks for s : $b_s = \lceil 20.000/30 \rceil = 667$ blocks
Index on s :
– level 3: $\lceil 20.000/100 \rceil = 200$ nodes
– level 2: $\lceil 200/100 \rceil = 2$ nodes
– level 1: $\lceil 2/100 \rceil = 1$ node
Total for index: 203 blocks
- b. $\sigma_{A=100.000}(s)$:
– Traverse the B^+ -tree to locate the matching tuple
– 3 index block IOs + 1 data block IO = 4 block IOs
 $\sigma_{A<100.000}(s)$:
– Scan the data file from the beginning; the index is not needed.
– Avg. distance between A -values: $1.000.000/20.000 = 50$
– Tuples that match the selection predicate: $100.000/50 = 2.000$
– Thus, $\lceil 2.000/30 \rceil = 67$ data block IOs
 $\sigma_{A>100.000}(s)$:
– Traverse the B^+ -tree to locate the first matching tuple: 3 index blocks
– Scan the data file sequentially from that tuple
– Avg. distance between A -values: $1.000.000/20.000 = 50$
– Tuples that match the selection predicate: $900.000/50 = 18.000$
– Thus, $\lceil 18.000/30 \rceil = 600$ data block IOs
– Total block IOs: $3 + 600 = 603$
- c. Plan p1: Block nested loop join (with r as outer relation):
– $C = b_r * b_s + b_r = 1.800 * 667 + 1.800 = 1.202.400$
Plan p2: Indexed nested loop join:
– Use the index to access matching tuples in s
– Cost c to access a matching tuple: $c = 3 + 1 = 4$ block IOs
– Cost for p2: $C = n_r * c + b_r = 45.000 * 4 + 1.800 = 181.800$
Plan p3: Hash join (partially filled blocks are ignored):
– $C = 3 * (b_r + b_s) = 3 * (1.800 + 667) = 7.401$
- d. – Use 1 block for the result, 1 block for r -partitions, 10 blocks for s -partitions
– An s -partition can hold at most $30 * 10 = 300$ tuples
– Avg. distance between A -values in s : $\lceil 1.000.000/20.000 \rceil = 50$
– The range of A -values that fit in a partition is $300 * 50 = 15.000$
– A hash function that assigns 1.500 tuples to a partition: $h = A \text{ div } 15.000$

Solution 5

- a. Proof by counter-example: Assume (for E_1 and E_2) two relations with schema (A, B) and instances $E_1 = \{(a, 1)\}$ and $E_2 = \{(b, 1)\}$, and let θ be the condition $A = 'a'$.
 Then on the right-hand side we get $E_1 \cup E_2 = \{(a, 1), (b, 1)\}$ and $\sigma_{A='a'}(E_1 \cup E_2) = \{(a, 1)\}$.
 On the left-hand side we get $\sigma_{A='a'}(E_1) \cup E_2 = \{(a, 1)\} \cup \{(b, 1)\} = \{(a, 1), (b, 1)\}$, which is different from the result of the left-hand side.
- b. Proof by counter-example: Assume (for E_1 and E_2) two relations with schema (A, B) and instances $E_1 = \{(1, 2), (1, 5)\}$ and $E_2 = \{(1, 2)\}$. The left-hand side is empty, whereas the right-hand side gives (1) .

Solution 6

- a. Conflict graph:



The schedule is not conflict serializable, since the conflict graph contains cycles.

- b. No.
 Example of violating a condition for view serializability: In the concurrent schedule, T_2 reads the initial value of Y ; in $\langle T_1, T_2, T_3 \rangle$, transaction T_2 reads the value of Y which is produced by T_1 (but should read the initial value).

Solution 7

- a. Lock and unlock instructions:

T_1 : lock-S(A); read(A); lock-X(B); read(B); if A=0 then B := B + 1; write(B). unlock(A); unlock(B);	T_2 : lock-S(B); read(B); lock-X(A); read(A); if B=0 then A := A + 1; write(A). unlock(B); unlock(A);
--	--

- b. The following schedule results in a deadlock at step 6:

	T_1	T_2	Wait-for graph
1	lock-S(A);		
2		lock-S(B);	
3		read(B);	
4	read(A)		
5	lock-X(B)		$T_1 \longrightarrow T_2$ (T_1 waits for T_2)
6		lock-X(B);	$T_1 \rightleftarrows T_2$ (T_1 waits for T_2 and vice versa)

- c. Wait-die deadlock prevention protocol: We assume that T_1 is the older transaction and T_2 is the younger transaction. Then at step 6, T_2 (the younger transaction) will not wait for T_1 (the older transaction) to release the lock. Instead, T_2 is rolled back, and the lock on B is released. T_1 can now continue.