T–106.5221 Transaction Management in Databases

Exam, May 18th, 2010

Write the following clearly on top of every paper you submit: "T–106.5221, May 18th, 2010", your full name, student ID and study programme, and the **total number of papers** you submit.

1. (6p) Explain the following concepts briefly:
   (a) savepoint
   (b) WAL (write-ahead logging) protocol
   (c) "steal" buffering policy
   (d) unrepeatable read
   (e) conditional lock request
   (f) update-mode lock (U-lock)

2. (6p) We apply the key-range locking protocol. What locks are acquired by the transactions in the following histories and when are these locks released? Are the histories possible under the key-range locking protocol? What isolation anomalies (dirty writes, dirty reads, unrepeatable reads) do the histories contain? In all cases, the database is initially empty.
   (a) $B_1 I_1[5] B_2 I_2[6] C_2 I_1[3] C_1$
   (b) $B_1 I_1[5] C_1 B_2 R_2[5, >2] B_3 I_3[3] C_3 I_2[8] C_2$
   (c) $B_1 I_1[5] C_1 B_2 D_2[5] B_3 I_3[2] C_3 I_2[5] C_2$

3. Relation $R(A, B, C, D)$ has a sparse index (a B$^+$-tree) on the pair of attributes AB (a unique key of R) and a dense static hash-table index on attribute C. (a–c) (3p) How can these index structures be utilized in the implementation of the following operations?
   (a) select D from R where $A = 12$ and $B < 8$;
   (b) select $*$ from R where $A \geq 59$ and $C \leq 400$;
   (c) select $*$ from R where $C = 35$;
   (d) (3p) Let us execute the following operation:
   insert into R values (15,20,35,40);
   Assume that the B$^+$-tree mentioned above has height 4, and all of its nodes are as full as possible before the insertion – thus, some nodes need to be split. What structure modifications are done to the tree, and in what order? Which pages of the tree are latched, and when is each latch released? The course discussed several possible answers to these questions; you can choose any one of the methods described on the course for your answer.

4.(a–c) (3p) The transaction generated from the following program fragment is run concurrently with other transactions.
set transaction isolation level L;
begin transaction;
select avg(V) into :a from r;
select sum(V) into :s from r;
select count($*$) into :c from r;
insert into q values (:a, :s / :c);
commit transaction.
How can the result produced by the transaction vary when (a) L = "read committed", (b) L = "repeatable read", (c) L = "serializable"? We assume that all other transactions that possibly run concurrently with T have isolation level "read uncommitted", at least.
   (d) (3p) Describe briefly the Commit-LSN mechanism. What are its advantages? On which isolation levels can it be used?

(problem 5 is on the next page)

5. (6p) The contents of the log on disk at the time of a system crash are the following:

101: ⟨*begin-checkpoint*⟩
102: ⟨*transaction-table*, {}⟩
103: ⟨*page-table*, {}⟩
104: ⟨*end-checkpoint*⟩
105: $\langle T_1, B\rangle$
106: $\langle T_1, I, p, 6, 9, 105\rangle$
107: $\langle T_2, B\rangle$
108: $\langle T_2, I, p, 19, 4, 107\rangle$
109: $\langle T_2, D, p, 12, 2, 108\rangle$
110: $\langle T_3, B\rangle$
111: $\langle T_3, D, p, 23, 5, 110\rangle$
112: $\langle T_3, C\rangle$
113: $\langle T_2, A\rangle$
114: $\langle T_2, D^{-1}, p, 12, 2, 108\rangle$
115: $\langle T_1, I, p, 25, 1, 106\rangle$

What actions are involved in restart recovery when using the ARIES algorithm? What log records are generated and when is the log forced to disk? Assume that Page-LSN = 108 in the disk version of p. We assume that all inverse operations in the undo pass can be performed physically.