RDBMS and SQL, Aug–Nov 2024

Assignment 4, 19 November 2024, due 30 November 2024

1. Consider relations $r_1(A, B, C)$ and $r_2(C, D, E)$. We wish to compute $r_1 \bowtie r_2$, the natural join of r_1 and r_2. Explain the difference in efficiency in computing this join under the following two assumptions:

   (a) A is a primary key for r_1, there is no primary key for r_2.

   (b) C is a primary key for r_2, there is no primary key for r_1.

2. Let relations $r_1(A, B, C)$ and $r_2(C, D, E)$ have the following properties: r_1 has 6×10^6 tuples, r_2 has 3×10^6 tuples, 15 tuples of r_1 fit in one block, 30 tuples of r_2 fit in one block and the memory can hold approximately 100 blocks.

   Estimate the number of block accesses required using the following strategies for $r_1 \bowtie r_2$. Assume that neither r_1 nor r_2 is already sorted with respect to any of its attributes. Explain your answer, including your choice of outer and inner relations.

   (a) Nested-loop join

   (b) Block nested-loop join

3. Consider the following schedules of concurrent read and write operations.

Schedule A

T_a	T_b	T_c	T_d	T_e
	$w(x_2)$			
$w(x_1)$				
			$w(x_3)$	
	$r(x_1)$			
				$r(x_1)$
			$r(x_2)$	
				$w(x_4)$
		$r(x_4)$		
$r(x_3)$				
		$r(x_3)$		
			$r(x_4)$	
	$r(x_4)$			

Schedule B

T_a	T_b	T_c	T_d	T_e
	$w(x_2)$			
$w(x_1)$				
			$r(x_1)$	
	$r(x_1)$			
				$r(x_1)$
			$r(x_2)$	
				$w(x_4)$
		$r(x_4)$		
$r(x_3)$				
		$r(x_3)$		
			$r(x_4)$	
	$r(x_4)$			

   (a) One of these schedules is conflict-serializable and the other is not. Explain which is which.

   (b) For the schedule that is conflict-serializable, list out all possible serial schedules that are consistent with this schedule.

4. Describe a concurrent schedule for the following two transactions T_1 and T_2 that is serializable, but not *conflict serializble*. Explain your answer.

```
T1: read(A);
    A := A − 50;
    write(A);
    read(B);
    B := B + 50;
    write(B).
```

```
T2: read(A);
    temp := A * 0.1;
    A := A − temp;
    write(A);
    read(B);
    B := B + temp;
    write(B).
```