## Data Management - exam of 12/07/2023 (B)

## Problem 1

A schedule $S$ is called dichotomic if it contains only write actions and every element of the database is written by at most two transactions in $S$. Prove or disprove that every dichotomic schedule is view serializable if and only if it is conflict serializable.

## Solution 1

As usual, we rule out the possibility that a schedule contains two different write actions on the same element issued by the same transaction. Therefore, if an element of the database is written by two actions in a dichotomic schedule $S$, such actions belong to different transactions. Also, since it is well known that conflict serializability implies view serializability, in order to prove the statement it is sufficient to prove that for dichotomic schedules, view serializability implies conflict serializability. We proceed by showing that if a serial schedule $S^{\prime}$ is view equivalent to a dichotomic schedule $S$, then $S^{\prime}$ is also conflict equivalent to $S$. Let $S^{\prime}$ be a serial schedule view equivalent to $S$. Obviously, since $S$ is dichotomic, $S^{\prime}$ is dichotomic too, and the read-from relations for both is empty. Also, they have the same final-write set, and this allows us to show that $S^{\prime}$ is conflict equivalent to $S$. Indeed, consider any pair of conflicting actions $w_{i}(x), w_{j}(x)$ such that $w_{j}(x)$ comes after $w_{i}(x)$ in $S$. Since $S$ is dichotomic, there is no other write action on $x$ in $S$, and therefore, $w_{j}(x)$ is a final write in $S$. It follows that $w_{j}(x)$ is also a final write in $S^{\prime}$ (otherwise $S^{\prime}$ would not be view equivalnent to $S$ ), which means that $w_{i}(x), w_{j}(x)$ appear in the same order in $S$ and $S^{\prime}$. We then conclude that $S^{\prime}$ is conflict equivalent to the serial schedule $S^{\prime}$, implying that $S$ is conflict serializable.

## Problem 2

Let $S(A, B, C)$ and $T(D, E, F)$ be two relations, each stored in a heap. A tuple $t$ of $S$ is said to be an upper $h$-match with T if there are at most $h$ tuples of T whose value in the attribute D is equal to the value of $t$ in the attribute B. Let $M$ be the number of buffer frames available, let $Q$ be the query that, given a value for $h$, computes all the tuples in S that are upper $h$-matches with T , and consider the following questions.
2.1 Describe the conditions under which we can process $Q$ using the one-pass, the two-passes, and the block-nested loop methods, respectively.
2.2 For each of the methods mentioned above, illustrate the corresponding algorithm and tell which is the cost of such algorithm in terms of the number of pages of S and T .

## Solution 2

In order to process $Q$ in one pass, a possible condition is that the relation T fits in $M-2$ buffer frames. Under this condition the algorithms would be as follows: we load T in the buffer, and then use one buffer frame $F$ for reading S one page at a time, and for each tuple $t$ in $F$, we check whether there are at most $h$ tuples of T in the buffer whose value in the attribute D is equal to the value of $t$ in the attribute B. If yes, we write $t$ in the output frame, otherwise, we ignore $t$. As usual, we write the output frame in secondary storage when full. The cost is simply $B(\mathrm{~T})+B(\mathrm{~S})$. Note that another condition that would allow the one pass strategy is that the relation $S$ fit in the buffer, but in this case we would need another field to associate to each tuple of $S$ a counter and this must be taken into account in order to establish the condition under which we can apply the one pass strategy (see the paragraph below on the block-nested loop technique).
In order to process $Q$ in two passes, a possible condition is that $B(\mathrm{~S})+B(\mathrm{~T}) \leq M \times(M-1)$. Under this condition the algorithm would be as follows: we produce the sorted sublists (each of size M) of S , sorted on the sorting key B; we produce the sorted sublists (each of size $M$ ) of T , sorted on the sorting key D; we do the merging phase by simply checking for each tuple $t$ of S , if there are at most $h$ tuples of T whose value in the attribute D is equal to the value of $t$ in the attribute $C$. If yes, we write $t$ in the output frame, otherwise, we ignore $t$. As usual, we write the output frame in secondary storage when full. The cost is simply $3 \times(B(\mathrm{~T})+B(\mathrm{~S}))$.
In order to process $Q$ using the block nested loop technique, we just need 3 buffer frames ( $M \geq 3$ ) and we proceed as follows: we read S in block of $M-2$ pages, by also reserving in the buffer enough space for a new integer field $N$ storing for each tuple of $S$ the count of joining tuples of $T$, we read one by one the pages of T in one frame, and while reading such pages we register, for each tuple $t$ of S in the $M-2$ pages, in
the field $N$ associated to $t$ the number of tuples of T joining with $t$. After reading all the pages of T we use the field $N$ to figure out which are the tuples of S to be written in the output frame, and then we consider the next block of $S$. As usual, we write the output frame in secondary storage when full. If we consider the space used to store $N$ as irrelevant, the cost is simply $B(\mathrm{~S})+B(\mathrm{~T}) \times B(\mathrm{~S}) /(M-2)$. If we want to be more precise, then we can assume that the size of $N$ is the same as the size of the other fields of S , so that we consider the size of the outer relation of the block nested loop algorithm $4 \times B(\mathrm{~S}) / 3$ instead of $B(\mathrm{~S})$, and obtain the total cost of $4 / 3 \times(B(\mathrm{~S})+B(\mathrm{~T}) \times B(\mathrm{~S}) /(M-2))$.

## Problem 3

Consider the following schedule $S$ :

$$
w_{1}(Z) r_{2}(X) w_{3}(X) r_{3}(Y) w_{4}(Y) w_{4}(X) r_{2}(Y) r_{1}(Y) w_{2}(Z)
$$

and answer (with suitable motivations) the following questions: (i) Is $S$ conflict serializable? (ii) Is $S$ a 2PL schedule (with shared and exclusive locks)? (iii) Can we insert the commit commands in such a way that the resulting schedule is accepted by the timestamp-based scheduler (assuming that the timestamp associated to a transaction coincides with the physical time when the transaction starts)? (iv) In all the three cases above where the answer is no, provide the answer to the following further question: is there any action $\alpha$ in $S$ such that, if we delete $\alpha$ from $S$, then the answer to the question would switch to yes?

## Solution 3

(i) Since the precedence graph associated to $S$ is cyclic, $S$ is not conflict serializable.
(ii) Since transaction $T_{2}$ should enter the shrinking phase when $w_{3}(x)$ is executed, and we cannot anticipate all the future locks required by $T_{2}$ (in particular, we cannot anticipate the shared lock on $Y$ required for $r_{2}(Y)$ ), we have that $S$ is not a 2PL schedule with both shared and exclusive locks.
(iii) Since $r_{2}(Y)$ comes after $w_{4}(Y)$ and the timestamp of $T_{4}$ is higher than the one of $T_{2}$, we have that $r_{2}(Y)$ is a "read too late action". Analogously, since $r_{1}(Y)$ comes after $w_{4}(Y)$ and the timestamp of $T_{4}$ is higher than the one of $T_{1}$, we have that $r_{1}(Y)$ is a "read too late action". Therefore, we cannot insert the commit commands in such a way that the resulting schedule is accepted by the timestamp-based scheduler.
(iv) We analyze the three cases separately.

- Looking at the precedence graph, we can realize that all cycles involve node $T_{4}$ and deleting action $w_{4}(Y)$ eliminates both the edge from $T_{4}$ to $T_{2}$ and the edge from $T_{4}$ to $T_{1}$, thus eliminating all cycles. So, if we delete $w_{4}(Y)$ from $S$, then the answer to question $(i)$ would switch to yes.
- It can be verified that if we delete $w_{4}(Y)$ from $S$, then the answer to question $(i i)$ would switch to yes.
- Finally, if we delete action $w_{4}(Y)$, insert the commit commands in $S$ as follows:

$$
w_{1}(Z) r_{2}(X) w_{3}(X) r_{3}(Y) c_{3} w_{4}(X) c_{4} r_{2}(Y) r_{1}(Y) c_{1} w_{2}(Z) c_{2}
$$

and use the Thomas rule when executing $w_{2}(Z)$, the resulting schedule is accepted by the timestampbased scheduler.

## Problem 4

The friendship between a city $C_{1}$ and a city $C_{2}$ (different from $C_{1}$ ) is strong if they have signed an agreement to be twin cities, and is solid if there is a third city $C_{3}$ such that both $C_{1}$ and $C_{2}$ are twin cities with $C_{3}$. Let $\mathrm{S}(\mathrm{A}, \mathrm{B}, \mathrm{C}$ ) be a relation (with key $\langle\mathrm{A}, \mathrm{B}\rangle$ ), storing all the tuples $\langle x, y, z\rangle$ such that city $x$ and $y$ are twin cities and $z$ is the year of the agreement between them. Obviously, S is symmetric. We know that ( $i$ ) no city has more than 400 twins, ( $i$ i ) associated to S there is a clustering sorted index on attribute A whose cost for retrieving the first tuple with a certain value for A is 5 page accesses, (iii) every page has space for 900 tuples of R, and (iv) the buffer has 100 frames available. Considering the query $Q$ that, given a specific city $x$, computes the set (therefore, without duplicates) of cities with whom $x$ has a strong or solid friendship, illustrate the most efficient algorithm you can think of for answering $Q$, and tell which is the cost of the algorithm in terms of number of page accesses.

## Solution 4

We can use the index for retrieving the tuples of S whose first component is $x$, telling us which are the twin cities of $x$ (i.e., which are the cities with which $x$ ha a string friendship) and store them in one frame $F$ of the buffer, relying on the fact that their maximum number is 400 ( 900 tuples fit in one frame). We can then use the index again for each of the twins of $x$ so as to retrieve all the twins of the twins of $x$ (i.e., the cities with which $x$ has a solid friendship), whose maximum number is $400 \times 400=160.000$. Since each page contains at most 900 tuples, we need to retrieve $160.000 / 900=178$ pages. Since 900 tuples of three values fit in one page, we have that 2.700 values fit in one page and therefore we can store the codes of all the twins of $x$ (i.e., the codes of the cities with which $x$ has a strong friendship) and the codes of the twins of the twins of $x$ (i.e., the codes of the cities with which $x$ has a solid friendship) using $160.400 / 2.700=60$ frames in the buffer, eliminate all duplicates by working in the buffer, and then write all the pages of the result. As for the cost, we count 5 page accesses for retrieving the first twin city of $x$, plus one page for retrieving, in the worst case, all the twin cities of $x$. Similarly, we need $(5+1) \times 400$ page accesses for retrieving all the cities with which $x$ has a solid friendship. Therefore, the cost is simply 2.406 page accesses.

Problem 5 (only for students enrolled in an A.Y. before 2021/22 who do not do the project)
Parallel algorithms for the implementation of relational operators are based on the idea of horizontal data partitioning, by which we can split a relation in various chunks, each one stored at a different node.
5.1 Describe the techniques that can be used for partitioning.
5.2 Among the described techniques, tell which is the most common one and explain why.
5.3 Illustrate the idea for designing a parallel algorithm that, given a table S , based on the partitioning technique mentioned for item 2) above, and depending on the number of buffer frames available at the various nodes, computes the relation obtained from S by the "group by" operator. Also, discuss the cost of the algorithm.
5.4 Illustrate the idea for designing a parallel algorithm that, based on the partitioning technique mentioned for item 2) above, and depending on the number of buffer frames available at the various nodes, computes the set intersection of two relations given in input, also discussing the cost of the algorithm.

## Solution 5

Please, refer to the slides of the course.

