## Data Management - exam of 08/06/2022 (Compito B)

## Problem 1

Let $\tau$ indicate the ternary operator such that $\tau(R, S, T)=R-\left(S \cup_{b} T\right)$, where $R, S$ and $T$ are three relations with the same schema and without duplicates, $\cup_{b}$ indicates bag union and - indicates bag difference.
1.1 Design and describe in detail a two pass algorithm that, given $R, S, T$, each one stored as a heap, computes $\tau(R, S, T)$.
1.2 Tell under which condition the algorithm can be used and illustrate the cost of the algorithm in terms of number of page accesses.

## Solution 1

Assuming that the buffer has $M$ frames available, a two pass algorithm based on sorting can be defined as follows (similarly, we could design a two pass algorithm based on hashing, but we will not illustrate such an algorithm).

- First pass: produce $M-1$ sorted sublists for $R, S$ and $T$ (each one with $M$ pages).
- Second pass: reserve one buffer frame for the output and one buffer frame for each of the sublists produced in the first pass, and perfom the "merge" phase of algorithm, based on the fact that all the sublists are sorted and we know exactly which is the relation associated to each buffer frame. As usual, whenever an input frame is exhausted, we read in that frame the next page of the corresponding sublist, if it exists, and whenever the output frame is full, we write its content in the result. At each step of the merge phase we do the following $(\operatorname{tup}(X)$ denotes the least tuple of relation $X$ that we have in the buffer or NULL if no tuple of $X$ exists in the buffer; we assume that advancing $X$ when $\operatorname{tup}(X)$ is NULL has no effect, and we assume that any condition on NULL is always false):
- if $\operatorname{tup}(R)$ is NULL, then stop
- if $\operatorname{tup}(S)$ and $\operatorname{tup}(T)$ are NULL, then copy $R$ in the output and stop
- If $\operatorname{tup}(R)<\operatorname{tup}(S)$ and $\operatorname{tup}(R)<\operatorname{tup}(T)$, then put $\operatorname{tup}(R)$ in the output, and go the next iteration
- If $\operatorname{tup}(R)=\operatorname{tup}(S)$ or $\operatorname{tup}(R)=\operatorname{tup}(T)$, then advance $R$ and go the next iteration
- If $\operatorname{tup}(R)>\operatorname{tup}(S)$, then advance $S$ and go to the next iteration
- If $\operatorname{tup}(R)>\operatorname{tup}(T)$, then advance $T$ and go to the next iteration

The above algorithm can be used under the following condition:

$$
\lceil(B(R) / M\rceil+\lceil B(S) / M\rceil+\lceil B(T) / M)\rceil \leq M-1
$$

where $B(X)$ denotes the number of pages of relation $X$. The cost of the algorthm is, as usual for two pass algorithms (we obviousy ignore the cost of writing the result):

$$
3 \times(B(R)+B(S)+B(T))
$$

## Problem 2

Given a schedule $S$, a serial schedule $S_{1}$ on the same transactions as $S$ is said to be "begin-order preserving with respect to $S^{\prime \prime}$ if it satisfies the following property: for every pair of transactions $T_{i}, T_{j}$ in $S$, if the first action of $T_{i}$ precedes the first action of $T_{j}$ in $S$, then $T_{i}$ precedes $T_{j}$ in $S_{1}$. A schedule $S$ is called begin-order preserving view serializable if there exists a serial schedule $S_{1}$ on the same transactions that is both view equivalent to $S$ and begin-order preserving with respect to $S$.
2.1 Prove or disprove the following claim: every conflict serializable schedule is "begin-order preserving view serializable".
2.2 Is the problem of checking whether a schedule is "begin-order preserving view serializable" decidable? If the answer is negative, then motivate the answer; if the answer is positive, then exhibit an algorithm for the problem, provide evidence of the correctness of the algorithm and illustrate its computational conplexity.

## Solution 2

1.1 The intuition is that begin-order preservation is independent from conflict serializability and therefore it should be easy to disprove the claim. Indeed, it is sufficient to consider the schedule $S: w_{2}(A) r_{1}(B) w_{2}(B)$ that is clearly conflict serializable. In particular, the serial schedule $T_{1}, T_{2}$ is the only serial schedule that is conflict equivalent to $S$, and is clearly not begin-order preserving with respect to $S$.
1.2 The problem of checking whether a schedule $S$ is "begin-order preserving view serializable" is decidable. An algorithm solving the problem is based on the following property: by definition, the only serial schedule that is begin-order preserving with respect to $S$ is the serial schedule $S^{\prime}$ that is coherent with the order imposed by the first actions of the transactions, i.e., the serial schedule where $T_{i}$ comes immediately after $T_{j}$ if the first action of $T_{i}$ comes after the first action of $T_{j}$ and no transaction starts in between. Therefore, an appropriate algorithm simply checks that such serial schedule $S^{\prime}$ is view equivalent to $S$. Clearly, constructing the serial schedule $S^{\prime}$ can be done in linear time with respect to the size of $S$, and checking if $S$ and $S^{\prime}$ are view equivalent can be done in quadratic time with respect to the size of $S$ and $S^{\prime}$.

## Problem 3

Consider the following schedule $S$ (where we have relaxed the condition that no transaction contains more than one occurrence of the same action):

$$
B\left(T_{0}\right) r_{0}(E) c_{0} B\left(T_{1}\right) r_{1}(A) r_{1}(E) w_{1}(E) B\left(T_{2}\right) r_{2}(E) w_{2}(E) B\left(T_{3}\right) r_{3}(A) w_{3}(A) r_{1}(A) c_{3} r_{1}(A) c_{1} c_{2}
$$

where the action $B$ means "begin transaction", the initial values of $A$ and $E$ are 10 and 30 , respectively, and every write action increases the value of the element on which it operates by 10 . Suppose that $S$ is executed by PostgreSQL, and describe what happens when the scheduler analyzes each action (illustrating also which are the values read and written by all the "read" and "write" actions) in both the following two cases: (1) all the transactions are defined with the isolation level "read committed"; (2) all the transactions are defined with the isolation level "repeatable read".

## Solution 3

We first deal with the isolation level "read committed". We remind the reader that such isolation level does not prevent the unrepeatable read anomaly nor the lost update anomaly.

- $r_{0}(E): T_{0}$ reads 30 .
- $c_{0}: T_{0}$ commits.
- $r_{1}(A): T_{1}$ reads 10 .
- $r_{1}(E): T_{1}$ reads 30 .
- $w_{1}(E): T_{1}$ writes 40 on $E$ in the local store.
- $r_{2}(E): T_{2}$ reads 30 .
- $w_{2}(E)$ : not executed, because $T_{1}$ holds the write lock on $E$; so $T_{2}$ must wait for the end of $T_{1}$ and therefore is suspended.
- $r_{3}(A): T_{3}$ reads 10 .
- $w_{3}(A): T_{3}$ writes 20 on $A$ in the local store.
- $r_{1}(A): T_{1}$ reads 10.
- $c_{3}: T_{3}$ commits and the value 20 for $A$ is written in the database.
- $r_{1}(A): T_{1}$ reads 20.
- $c_{1}: T_{1}$ commits and the value 40 for $E$ is written in the database.
- $w_{2}(E): T_{2}$ resumes and writes 40 on $E$ in the local store.
- $c_{2}: T_{2}$ commits and the value 40 for $E$ is written in the database.

We now deal with the isolation level "repeatable read" (in bold the difference with respect to the previous case). We remind the reader that such isolation level prevents both the unrepeatable read anomaly and the lost update anomaly.

- $r_{0}(E): T_{0}$ reads 30 .
- $c_{0}: T_{0}$ commits.
- $r_{1}(A): T_{1}$ reads 10 .
- $r_{1}(E): T_{1}$ reads 30.
- $w_{1}(E): T_{1}$ writes 40 on $E$ in the local store.
- $r_{2}(E): T_{2}$ reads 30 .
- $w_{2}(E)$ : not executed, because $T_{1}$ holds the write lock on $E$; so $T_{2}$ must wait for the end of $T_{1}$ and therefore is suspended.
- $r_{3}(A): T_{3}$ reads 10.
- $w_{3}(A): T_{3}$ writes 20 on $A$ in the local store.
- $r_{1}(A): T_{1}$ reads 10.
- $c_{3}: T_{3}$ commits and the value 20 for $A$ is written in the database.
- $r_{1}(A): T_{1}$ reads 10.
- $c_{1}: T_{1}$ commits and the value 40 for $E$ is written in the database.
- $w_{2}(E): T_{2}$ aborted with message: "ERROR: could not serialize access due to concurrent update".
- $c_{2}$ : ignored, because $T_{2}$ has been aborted.


## Problem 4

Consider the relations Bus(code,stopcode,passengers) with 15.000 .000 tuples, and Stop(stopcode, district,city) occupying 290 pages, each page with 100 tuples (keys are underlined). We assume that 150 values are in the attribute city, that each value (regardless of the type) requires the same number of bytes and that we have 300 frames available in the buffer. Consider the query select city, sum(passengers)
from Bus b, Stop s where b.stopcode = s.stopcode
group by city
and describe the algorithm you would use to execute the query, illustrating the number of page accesses required by the execution of the algorithm.

## Solution 4

We immediately notice that each page can contain 300 values (100 tuples of three values each). Also, we immediately notice that the whole relation Stop fits in the buffer, and that the 150 tuples of the form (city, sum(passengers)) in the result of the query fits in one frame $F$ of the buffer (because 300 values fit in one page, and therefore also in one frame of the buffer). It follows that we can use a one pass algorithm as follows:

- load the relation Stop in 290 frames of the buffer;
- load the relation Bus one page $P$ at a time in one frame of the buffer and for each tuple $t$ of Bus in $P$ single out the tuple $t^{\prime}$ of Stop (at most one) joining with $t$. Consider the value $c$ that the tuple $t^{\prime}$ has in the attribute city, and update the value in $F$ of sum(passengers) corresponding to $c$, making use of the value that $t$ has in the attribute passengers.

The number of page accesses required by the algorithm is simply $B$ (Bus) $+B$ (Stop). Since each page has space for 100 tuples of Stop, and since Bus and Stop have the same number of attributes, and since we assume that all values requires the same number of bytes, we conclude that each page has space for 100 tuples of Bus as well, and therefore $B($ Bus $)=15.000 .000 / 100=150.000$. Therefore, the number of page accesses required by the algorithm is $150.000+290=150.290$.

## Problem 5

Consider a graph database with nodes of type Shop with properties name (identifying the shop) and city, nodes of type Product with property category, and edges of type Sold connecting each shop with the products sold in that shop. In such database, for example, one may represent the node s1 of type Shop with name: "Bakery100" and city:"Roma" connected to node p1 of type Product with category:"book" by means of an edge of type Sold.
5.1 Illustrate how you would represent the above database as a schema in the relational model.
5.2 Assuming that the three most relevant queries are (1) given the name of a shop, compute the products sold by the shop, (2) produce the sorted list of 〈name of shop, name of product〉 for shops and the products they have sold and (3) compute all the shops of a given city, describe the file organizations you would choose for the relations in the relational database mentioned above, and then describe the algorithms for executing the queries on the basis of such representation.

## Solution 3

The conceptual schema representing the graph database is as follows:


Consequently，the correct relational schema to represent such database is as follows：

```
Shop(name, city)
Product(code, category)
Sold(sname, pcode)
    foreign key Sold(sname) \subseteq Shop(name)
    foreign key Sold(pcode) \subseteq Product(pcode)
```

We immediately notice that in order to support the second query it is essential to have the sorted list of 〈name of shop，code of product〉 in secondary storage．Notice that we get such a list simply by building a $\mathrm{B}^{+}$－tree index on Sold with search key 〈sname，pcode〉 using alternative 1：the leaves of the tree will be exactly the sorted list of interest．The nice feature of such an index is that it also supports the first query， because a $\mathrm{B}^{+}$－tree index well supports all queries searching for a value corresponding to a prefix of the search key．As for the third query，we can simply build a hash－based index on Shop with search key city．Based on the above decisions，the algorithms for the various queries are as follows：
－Query 1：given the name $n$ of a shop，use the $\mathrm{B}^{+}$－tree index on Sold to search for the value $n$ in the attribute sname，and then return all the values in the attribute pcode associated to $n$ ．
－Query 2：simply return all the records in the leaves of the $\mathrm{B}^{+}$－tree index．
－Query 3：given a city $c$ ，use the hash index to retrieve all the records in Shop having $c$ in the attribute city，and return all the values in the attribute sname of such records．

