# Data Management - exam of 12/07/2021 <br> Solutions 

## Problem 1

A schedule $S$ is called "one-writer" if at most one transaction appearing in $S$ has write actions. Prove or disprove the following two statements:

1. A "one-writer" schedule $S$ is conflict-serializable if and only if it is view-serializable.
2. A "one-writer" schedule $S$ is conflict-serializable if and only if it is a 2 PL schedule.

The solution must be provided under the usual assumption that no transaction reads or writes the same element more than once.

## Solution to problem 1

1. It is well-known that if $S$ is conflict-serializable, then it is also view-serializable. So, it remains to show that if $S$ is a view-serializable one-writer schedule, then it is also conflict-serializable. Assume that $S$ is view-serializable, implying that there exists a serial schedule $S^{\prime \prime}$ that is viewequivalent to $S$. Our aim is to prove that $S$ and $S^{\prime}$ are conflict-equivalent. Suppose, by contradiction, that $S$ and $S^{\prime}$ are not conflict-equivalent. This means that there is a pair of conflicting actions, say $a_{i}(X)$ and $a_{j}(X)$, appearing in different order in the two schedules. Since one of these actions is a write and one is a read, we conclude that $S^{\prime}$ has a different "read-from" relation wrt $S$, which implies that $S$ and $S^{\prime}$ are not view-equivalent.
2. It is well-known that if $S$ is a 2PL schedule, then it is conflict-serializable. So, it remains to check whether if $S$ is a conflict-serializable one-writer schedule, then it is also a 2 PL schedule. This is disproved by the following counterexample:

$$
S=w_{1}(x) r_{2}(x) r_{3}(y) w_{1}(y)
$$

Indeed, it is easy to see that $S$ is conflict serializable, but it is not a 2 PL schedule.

## Problem 2

Consider the following schedule $S$ :

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

where the action $B$ means "begin transaction", and every write action performed by transaction $T_{i}$ writes the value $i$ on the element corresponding to the argument. Suppose that $S$ is executed by PostgreSQL with all the transactions defined with the isolation level "read committed", and for each action different from $B$ tell what is the effect of the action and what is the behavior of the system when it executes the action.

## Solution to problem 2

We should remember the following important points regarding SQL and PostgreSQL:

- The lock-based concurrency control implementation of SQL keeps write locks until the end of the transaction, but read locks are released as soon as the SELECT operation is performed.
- The concurrency control strategy of PostgreSQL is a sort of multiversion control, where reads are never blocked (they read the value written by the last committed transaction), and write actions are performed on a local store (snapshot), and their effects are transferred to the database at commit.
- PostgreSQL adopts a recognition approach to deadlock.

The effect of each action and the behavior of the system is specified as follows:

- $w_{0}(A)$ writes 0 on $A$ in the local store associated to $T_{0}$
- $c_{0}$ commits $T_{0}$, and the value 0 for $A$ is now in the database
- $w_{1}(A)$ writes 1 on $A$ in the local store associated to $T_{1}$
- $r_{2}(A)$ reads the value 0 , the value written by $T_{0}$, the last committed transaction that wrote on A
- $w_{2}(B)$ writes 2 on $B$ in the local store associated to $T_{2}$
- $c_{1}$ commits $T_{1}$, and the value 1 for $A$ is now in the database
- $r_{3}(A)$ reads the value 1 , the value written by $T_{1}$, the last committed transaction that wrote on A
- $w_{3}(C)$ writes 3 on $C$ in the local store associated to $T_{3}$
- when $T_{3}$ tries to execute the action $w_{3}(B)$, such action is not executed because $T_{2}$ holds the exclusive lock on $B$. So, $T_{3}$ is suspended waiting for the commit of $T_{2}$
- since $C$ was written by $T_{3}$, when $T_{2}$ tries to execute the action $w_{2}(C)$, the system recognizes a deadlock, because $T_{3}$ is waiting for the end of $T_{2}$, and $T_{2}$ should now wait for the end of $T_{3}$. Therefore, action $w_{2}(C)$ is not executed, $T_{2}$ is aborted in order to resolve the deadlock, and $T_{3}$ becomes active again
- $c_{3}$ commits $T_{3}$, and the value 3 for $C$ and the value 3 for $B$ are now in the database
- $c_{2}$ is ignored, because $T_{2}$ was aborted


## Problem 3

Let $R(A, B)$ be a relation with 100.000 pages, and $S(C, D)$ a relation with 500.000 pages. We know that (i) R has 2.000 values in A , uniformly distributed in the various tuples, (ii) there is a clustered, sparse $\mathrm{B}^{+}$-tree index on R with search key A , (iii) 60 tuples of R fit in one page, (iv) we have 260 frames available in the buffer, and $(v)$ every value and pointer occupy the same amount of space. Consider the query:

```
select A,D
from R, S
where B = C and A >= 10 and A <= 15
```

Show the logical query plan associated to the query, as well as the logical query plan and the physical query plan you would choose for executing the query efficiently. Also, tell which is the cost (in terms of number of page accesses) of executing the query according to the chosen physical query plan.

## Solution to problem 3

The logical query plan associated to the query is shown below:


## Logical query plan

The chosen logical query plan is shown below:


The chosen logical query plan

The selection operation can be performed by using the index. Indeed, this is a range selection, and can be well supported by a clustered index. So, let us evaluate the cost of computing all the pages of $R$ satisfying the selction condition using the index, assuming that the index uses alternative 1 . Since in every page we have space for 60 tuples of $R(A, B)$, it follows that we have space for 60 data entries, and taking into account the $67 \%$ occupancy rule, we have 40 data entries per page. Since the index is sparse, we have one data entry for each page of the relation, and so the number of leaves is $100.000 / 40$ $=2.500$. Since 60 tuples of $R$ fit in one page, we have that the number of tuples of $R(A, B)$ is 100.000 $\times 60=6.000 .000$, and the number of tuples with the same value for A is $6.000 .000 / 2.000=3.000$.

So, the number of pages with the same value of A is $3 \cdot 000 / 60=50$. Since we have 6 values for A to consider, we have to access $50 \times 6=300$ pages of $R(A, B)$. To access them, we use the index searching for the value 10 of A in $\mathrm{R}(\mathrm{A}, \mathrm{B})$ and then we move to the data file in order to access the 300 pages of $R(A, B)$. The fan-out of the tree is $(60+30) / 2=45$, and therefore the cost of accessing the first data record is $\log _{45} 2.500+1$.

The set of tuples in the pages of the data file accessed by the index form one operand of the join with S (C,D). Since $300>260$, the join cannot be executed in one pass. One possibility is using a two pass algorithm, but the cost would be $(300+500.000) \times 3$. Instead, it is no difficult to see that it is more efficient to choose a block-nested loop algorithm for the join, requiring to scan the relation S(C,D) only twice.

The resulting physical query plan is shown below:
1-pass projection


## Physical query plan

The cost of the execution of the whole physical query plan is $\log _{45} 2.500+1+300+2 \times 500.000=$ 1.000 .304 page accesses.

## Problem 4

Let Building(bcode,floors, area, value,cat,city) be stored in a heap file with 840.000 tuples, and City (ccode, nation, nab) be stored in a heap file with 9.000 .000 tuples. We assume that every value has the same size, that every page has room for 600 values, that $V$ (Building,floors) $=100$, $V$ (Building, city) $=300$, and that we have 85 free buffer frames available. Consider the query:
select distinct floors, city, nab
from Building, City
where city $=$ ccode and nab > 10000
Show the logical query plan associated to the query, as well as the logical query plan and the physical query plan you would choose for executing the query efficiently. Also, tell which is the cost (in terms of number of page accesses) of executing the query according to the chosen physical query plan.

## Solution to problem 4

The logical query plan associated to the query is shown below.


After pushing selection, duplicate elimination and projections we have the following logical query plan:


Note that duplicate elimination has not been pushed through the branch of City, because the right-hand side operand of the join contains a key of the relation, and therefore does not have duplicates. Note also that after duplicate elimination we cannot have more that $V$ (Building,floors) $\times$ $V$ (Building, city) $=30.000$ tuples. The number of tuples of the various nodes of the plan is shown here:


The number of pages is shown here:


Note that $2.800+3.334<85 \times 84$, and therefore we can use a two pass algorithm for the join. Also, we can combine the join with the duplicate elimination for Building if we use an algorithm based on sorting. Indeed, we can produce the sorted sublists for Building based on a sorting on〈city,floors〉, and we can produce the sorted sublists for City based on a sorting on ccode, so that
we can carry out the merging phase based on city and ccode, and simply eliminate duplicates during such merging phase of the two pass join algorithm.

Here is the physical query plan:


The cost is: 8.400 (reading of Building) +45.000 (reading of City) +2.800 (writing of the $2.800 /$ $85=33$ sublists for Building during the first phase of the two pass join algorithm) +3.334 (writing of the $3.334 / 85=40$ sublists for City during the first phase of the two pass join algorithm) +2.800 (reading of the 33 sublists during the merging phase of the two pass join algorithm) +3.334 (reading of the 40 sublists during the merging phase of the two pass join algorithm) $=65.668$ page accesses.

## Problem 5

If $R_{1}(\underline{A}, B)$ and $R_{2}(\underline{A}, B)$ are two relations each with key $A$, then the disjoint left-union of $R_{1}$ and $R_{2}$, indicated as $R_{1} \oslash R_{2}$, is the relation with attributes $A, B$ and with the set of tuples specified as follows: (i) for each tuple $t \in \mathrm{R}_{1}$ such that $t$.B is not null, we have $t \in \mathrm{R}_{1} \oslash \mathrm{R}_{2}$; (ii) for each tuple $t \in \mathrm{R}_{1}$ such that $t . \mathrm{B}$ is null and such that there exists $t^{\prime} \in \mathrm{R}_{2}$ with $t . \mathrm{A}=t^{\prime}$.A, we have $t^{\prime} \in \mathrm{R}_{1} \oslash \mathrm{R}_{2}$. Suppose that we have a multiprocessor system with $N$ nodes $n_{1}, \ldots, n_{N}$, each of them with $M>N$ free frames available, and that the two relations $\mathrm{R}_{1}$ and $\mathrm{R}_{2}$ are stored in node $n_{1}$.

1. Illustrate a parallel algorithm for computing $R_{1} \oslash R_{2}$.
2. Assuming $B\left(\mathrm{R}_{1}\right)=10.000, B\left(\mathrm{R}_{2}\right)=15.000, N=10$ and $M=40$, describe the cost of the algorithm both in terms of the elapsed time, and in terms of number of page accesses.

## Solution to problem 5

1. We use a hash function on the values of $A$ for distributing the tuples of both $R_{1}$ and $R_{2}$ among the various nodes, in such a way that each node $n_{i}$ will handle one bucket of $\mathrm{R}_{1}$ and one bucket of $R_{2}$ resulted from the same value of the hash function, and therefore containing all the tuples with values of A that have given the result $i$ for the hash function. This implies that if a tuple $t$ of $\mathrm{R}_{1}$ has the same A -value as the tuple $t^{\prime}$ of $\mathrm{R}_{2}, t$ and $t^{\prime}$ will be stored in the same node.

- The distribution of the tuples is done at $n_{1}$, following the usual mechanism: we reserve one frame for each node (this is possible because $M>N$ ), and one frame for the input. We load each page of $\mathrm{R}_{1}$ in the input frame, one by one. For each tuple $t$ in the input frame such that $t$.B is not null, we write $t$ in the output of $n_{1}$, because it has to be part of the result. For each tuple $t$ in the input frame such that $t$.B is null, we apply the hash function so as to obtain the bucket (or, the processor) where the tuple should go, and we store the tuple in the frame corresponding to the "correct" node. Obviously, when a frame is full, we ship the tuples to the corresponding node. The same algorithm is carried out for $R_{2}$. If the hash function is good, each node receives at most $B\left(\mathrm{R}_{1}\right) / N$ pages with tuples of $\mathrm{R}_{1}$ and at most $B\left(\mathrm{R}_{2}\right) / N$ pages with tuples of $\mathrm{R}_{2}$, where $B(R)$, as usual, denotes the number of pages of $R$.
- The computation at each node $n_{i}$ is done as follows: the pages of $\mathrm{R}_{1}$ received are stored in a file $F_{1}$ and the pages of $\mathrm{R}_{2}$ received are stored in a different file $F_{2}$. Then, using a multi-pass strategy, we iteratively produce a suitable number of sorted sublists for $F_{1}$ and a suitable number of sorted sublists for $F_{2}$. Obviously, the sorting is done based on the attribute A. In particular, the goal is to reach a situation where the number of sorted sublists for $F_{1}$ is less than $\left(M \times B\left(\mathrm{R}_{1}\right) /\left(B\left(\mathrm{R}_{1}\right)+B\left(\mathrm{R}_{2}\right)\right)\right.$, and for $F_{2}$ is less than $\left(M \times B\left(\mathrm{R}_{2}\right) /\left(B\left(\mathrm{R}_{1}\right)+B\left(\mathrm{R}_{2}\right)\right)\right.$. At this point, we carry out a final "merge-like" phase using $M$ frames, and we compute the final result for the buckets in $n_{i}$. For doing that, we simply apply the definition of $\mathrm{R}_{1} \oslash \mathrm{R}_{2}$.

2. The worst-case cost of the algorithm in terms of total number page accesses is determined as follows.

- We have the cost of reading the pages of both relations in the tuple distribution phase: $10.000+15.000=25.000$;
- For each of the 10 nodes $n_{1}, \ldots, n_{10}$, we have the following:
- the cost $10.000 / 10=1.000$ of writing the bucket for $\mathrm{R}_{1}$, and the cost $15.000 / 10=1.500$ of writing the bucket for $R_{2}$. Since $2.500>40,2.500>40 \times 39$, and $2.500<40 \times 40$ $\times 39$, we need 3 passes to solve the problem, and therefore the additional cost is $5 \times$ 2.500 (as usual, we ignore the cost of writing the result). The total cost at each node is 15.000 .

Therefore the cost in terms of number of page accesses is: $25.000+15.000 \times 10=175.000$ page accesses, and the cost in terms of elapsed time is $25.000+15.000=40.000$ page accesses.

Notice that the above strategy could be made more efficient, by avoiding storing the pages shipped to nodes $n_{2}, \ldots, n_{10}$. However, we do not discuss such an improvement here.

