## Data Management－exam of 08／06／2023（B）

## Problem 1

Let Trainer（id，nation）be a relation with 500 pages stored in a file sorted on 〈id，nation〉， Team（name，city）a relation with 100 pages sorted on 〈name，city〉 and Coached（id，nation，name，city） a relation with 35.000 pages sorted on 〈id，nation，name，city〉．With the goal of knowing，for each trainer， the teams that（s）he has not coached，we want to compute the set difference between the cartesian product of Trainer and Team and the relation Coached．Tell which is the best algorithm that a query engine with 107 free buffer frames should use for this task，indicating also the cost of executing such algorithm．

## Problem 2

Answer the following two questions，providing a suitable motivation for each answer．
3．1 Give an example of three transactions which have the following properties：$(i)$ there exists a schedule $S$ on the three transactions such that when $S$ is given in input to a timestamp－based scheduler，a deadlock occurs；（ii）for each pair of the three transactions and for any schedule $S$ on such pair，no deadlock occurs when $S$ is given in input to a timestamp－based scheduler．
3．2 Try to generalize the above example to the case of $n$ transactions．More precisely，try to come up with an example of $n$ transactions having the following properties：（i）There exists a schedule $S$ on the $n$ transactions such that when $S$ is given in input to a timestamp－based scheduler，a deadlock occurs． （ii）For each $n-1$ transactions of the $n$ chosen transactions and for any schedule $S$ on such $n-1$ transactions，no deadlock occurs when $S$ is given in input to a timestamp－based scheduler．

## Problem 3

Let $R(A, B, C, D)$（with 52.000 pages）and $S(E, F, G, H)$（with 20.000 pages）be two relations stored in a heap at processor $P_{0}$ ．We know that 300 free buffer frames are available at $P_{0}$ and that 100 tuples of each relation fit in one page．Also，we know that attribute C contains 500 values uniformly distributed in the tuples of R，and the same holds for attribute H in S ．Finally，we know that there are 10 processors $P_{1}, \ldots, P_{10}$ that we can use for processing queries，each with 80 free frames available．Consider the query that computes the equi－join between $R$ and $S$ on the condition $C=H$ ，and compare the following two cases：
1．1 The query is executed at processor $P_{0}$ ．
1．2 The query is executed in parallel after sending the tuples of the two relations to processors $P_{1}, \ldots, P_{10}$ ． For each of the above cases，illustrate the algorithm you would use and discuss its cost in terms of number of page accesses（case 1．1）and elapsed time（case 1．2）．

## Problem 4

Consider the following schedule $S$ ：

$$
B_{1} w_{1}(X) B_{2} w_{2}(Y) B_{3} w_{3}(Z) c_{3} B_{4} r_{4}(Z) r_{4}(Y) w_{4}(W) B_{5} r_{5}(W) c_{4} w_{5}(W) c_{5} w_{1}(Y) c_{1} w_{2}(X) c_{2}
$$

where the action $B_{i}$ means＂begin transaction $T_{i}$＂，the initial value of every item $X, Y, Z, W$ is 200 and every write action increases the value of the element on which it operates by 200 ．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＂．
Problem 5 （only for students enrolled in an A．Y．before 2021／22 who do not do the project）
Consider a database $B$ about musicians and concerts，where：$(i)$ for each musician we are interested in the id，the age，the sex，the city of birth，the city（s）he is living in，the music academy where（s）he graduated and the concerts to which（s）he participated；（ii）for each concert we are interested in the average price of the ticket，in the month and the year when it was held，with the decade of the year，and the city where the concert was held；（iii）for each city we are interested in the province，the region and the number of inhabitants；（iv）for each music academy we are interested in the name and the founder．We want to build a data warehouse on all the above data for various analyses on participation of musicians to concerts in the last 50 years．You are asked to show（1）the ER schema of the database，（2）the DFM schema of the data warehouse，（3）the corresponding star schema（optionally，with the queries used to populate the star schema tables），and（4）based on such tables，the SQL query that computes the number of musicians who participated in concerts，for each region of the city of birth of the musicians，and for each music academy where the musicians graduated．

