CS W186 - Spring 2020
Guerrilla Section 4
Parallel Query Processing, Transactions and Concurrency, Database Design

Sunday, April 5, 2020

1   Types of Parallelism

For each of the following scenarios, state whether it is an example of:

- Inter-query parallelism
- Intra-query, inter-operator parallelism
- Intra-query, intra-operator parallelism
- No parallelism

1. A query with a selection, followed by a projection, followed by a join, runs on a single machine with one thread.

2. Same as before, but there is a second machine and a second query, running independently of the first machine and the first query.

3. A query with a selection, followed by a projection, runs on a single machine with multiple threads; one thread is given to the selection and one thread is given to the projection.

4. We have a single machine, and it runs recursive hash partitioning (for external hashing) with one thread.

5. We have a multi-machine database, and we are running a join over it. For the join, we are running parallel sort-merge join.

2   Partitioning for Parallelism

1. Suppose we have a table of size 50,000 KB, and our database has 10 machines. Each machine has 100 pages of buffer, and a page is 4 KB.

   We would like to perform parallel sorting on this table, so first, we perfectly range partition the data. Then on each machine, we run standard external sorting.

   How many passes does this external sort on each machine take?

2. Suppose we were doing parallel hash join. The first step is to partition the data across the machines, and we usually use hash partitioning to do this.

   Would range partitioning also work? What about round-robin partitioning?

3. Suppose we have a table of 1200 rows, perfectly range-partitioned across 3 machines in order.

   We just bought a 4th machine for our database, and we want to run parallel sorting using all 4 machines.

   The first step in parallel sorting is to repartition the data across all 4 machines, using range partitioning. (The new machine will get the last range.)

   For each of the first 3 machines, how many rows will it send across the network during the repartitioning? (You can assume the new ranges are also perfectly uniform.)

3 Transactions and Concurrency

In this question, we will explore the key topics of transactions and concurrency: serializability, types of locks, two-phase locking, and deadlocks.

We will do this by actually running multiple transactions at the same time on a database, and seeing what happens. You may find it helpful (but not necessary) to draw some graphs:

- For the lock type questions, you may wish to draw a graph representing the whole database and which resources are being locked.

- For serializability questions, you may wish to draw a graph with a node for each transaction, and arrows if there are *conflicts* between transactions.

- For deadlock questions, you may wish to draw a graph with a node for each transaction, and arrows if a transaction is *waiting* for a lock held by another transaction.

We will use a database with tables A, B, C, ... and table A holds rows A1, A2, A3, ... and so on.

Consider the following sequence of operations:

$$\text{Txn 1:} \quad \text{IX-Lock(Database)} \tag{1}$$
$$\text{Txn 1:} \quad \text{IX-Lock(Table A)} \tag{2}$$
$$\text{Txn 1:} \quad \text{X-Lock(Row A1)} \tag{3}$$
$$\text{Txn 1:} \quad \text{Write(Row A1)} \tag{4}$$
$$\text{Txn 1:} \quad \text{Unlock(Row A1)} \tag{5}$$
$$\text{Txn 1:} \quad \text{S-Lock(Row A2)} \tag{6}$$
$$\text{Txn 2:} \quad \text{IX-Lock(Database)} \tag{7}$$
$$\text{Txn 2:} \quad \text{IX-Lock(Table A)} \tag{8}$$
$$\text{Txn 2:} \quad \text{X-Lock(Row A1)} \tag{9}$$
$$\text{Txn 2:} \quad \text{Write(Row A1)} \tag{10}$$
$$\text{Txn 2:} \quad \text{S-Lock(Row A2)} \tag{11}$$
$$\text{Txn 2:} \quad \text{Read(Row A2)} \tag{12}$$

1. Is Transaction 1 doing Two-Phase Locking so far?

2. Is Transaction 1 doing Strict Two-Phase Locking?

3. Is this schedule conflict-serializable so far? If not, what is the cycle?

4. Is this schedule serial so far?

Continuing with all operations so far:

$$\text{Txn 2:} \quad \texttt{SIX-Lock(Table B)} \tag{13}$$
$$\text{Txn 2:} \quad \texttt{X-Lock(Row B1)} \tag{14}$$
$$\text{Txn 2:} \quad \texttt{Write(Row B1)} \tag{15}$$
$$\text{Txn 2:} \quad \texttt{Unlock(Row B1)} \tag{16}$$
$$\text{Txn 2:} \quad \texttt{Unlock(Table B)} \tag{17}$$
$$\text{Txn 1:} \quad \texttt{SIX-Lock(Table B)} \tag{18}$$
$$\text{Txn 1:} \quad \texttt{X-Lock(Row B1)} \tag{19}$$
$$\text{Txn 1:} \quad \texttt{Read(Row B1)} \tag{20}$$

5. Is Transaction 2 doing Two-Phase Locking so far?

6. Is Transaction 2 doing Strict Two-Phase Locking?

7. Is this schedule conflict-serializable so far? If not, what is the cycle?

8. Suppose we start a new transaction, Transaction 3. What kind of locks can Transaction 3 acquire on the whole database?

9. Given the above answers, what kind of locks can Transaction 3 acquire on Table A?

10. Given the above answers, what kind of locks can Transaction 3 acquire on Row A3?

11. Given the above answers, what kind of locks can Transaction 3 acquire on Table B?

12. Given the above answers, what kind of locks can Transaction 3 acquire on Row B2?

13. What kind of locks can Transaction 1 acquire on Row B2?

$$\text{Txn 2: } \texttt{IX-Lock(Table B)} \tag{21}$$
$$\text{Txn 3: } \texttt{IX-Lock(Database)} \tag{22}$$
$$\text{Txn 3: } \texttt{X-Lock(Table C)} \tag{23}$$
$$\text{Txn 3: } \texttt{IX-Lock(Table A)} \tag{24}$$
$$\text{Txn 3: } \texttt{X-Lock(Row A1)} \tag{25}$$
$$\text{Txn 1: } \texttt{IX-Lock(Table C)} \tag{26}$$

14. We have now entered a deadlock. What is the waits-for cycle between the transactions?

15. We can end this deadlock by aborting the youngest transaction. Which transaction do we abort?

    Alternatively, we could have avoided this deadlock in the first place by using *wound-wait* or *wait-die*. Recall from lecture that these methods cause transactions to sometimes abort, according to a priority order, when they try to acquire locks.

    Let's say that the priority of the transaction is its number (Txn 1 is highest priority).

16. If we were using *wound-wait*, what is the **first operation** in this sequence that would cause a transaction to get aborted, and which transaction gets aborted?

17. If we were using *wait-die*, what is the **first operation** in this sequence that would cause a transaction to get aborted, and which transaction gets aborted?

4 Database Design

As chancellor of UC Berkeley, you are tasked with designing a new final exam scheduling system to make it easier on students. Using the following assumptions, fill in the ER diagram we give you.

- A student may take any number of exams, and every exam is taken by at least one student.

- An exam is uniquely identified by the combination of a course and a semester.

- Every exam has at least one supervisor. A supervisor oversees exactly one exam.

- There is at least one question on every exam, and a question appears on at most one exam.

- A question on an exam may be answered by any number of students, and a student may answer any number of questions on an exam.
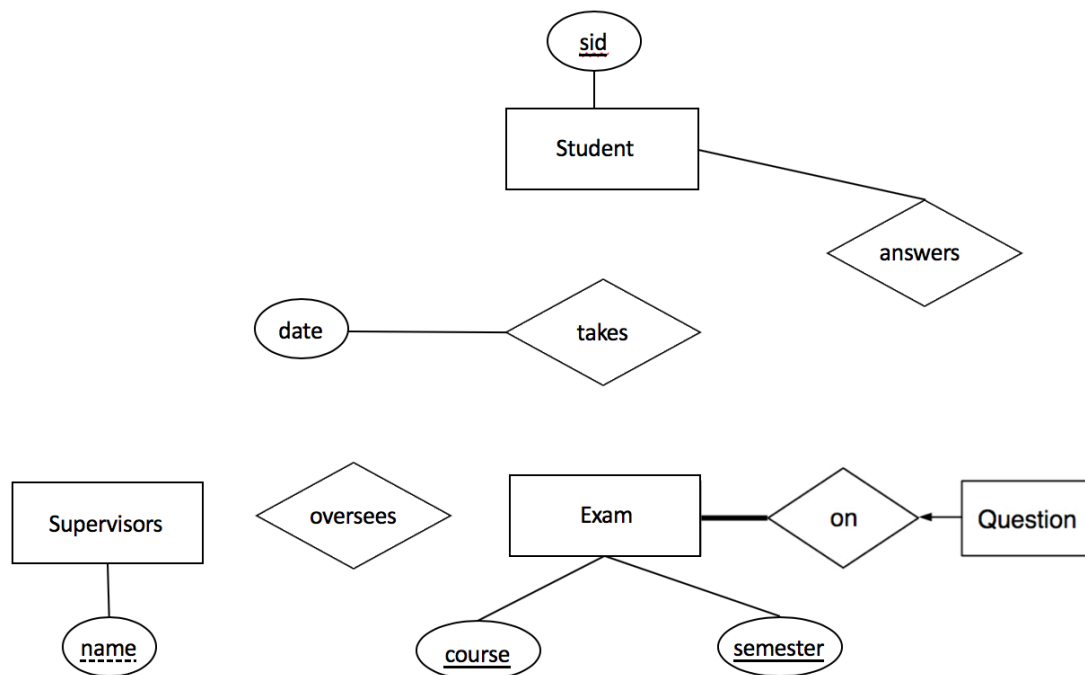
1. What type of edge should be drawn between the **Supervisors** entity and the **oversees** relationship set?

2. What type of edge should be drawn between the **Exam** entity and the **oversees** relationship set?

3. What type of edge should be drawn between the **Student** entity and the **takes** relationship set?

4. What type of edge should be drawn between the **Exam** entity and the **takes** relationship set?

5. What type of edge should be drawn between the **Questions** entity and the **answers** relationship set?

6. Consider the attribute set R = ABCDEF and the functional dependency set
   F = {BE → C, B → F, D → F, AEF → B, A → E}. Which of the following are candidate keys of R?
   Mark all that apply

   (A) ACD

   (B) AD

   (C) FC

   (D) BF

7. Given Attribute Set R = ABCDEFGH and functional dependencies set
   F = {CE → GH, F → G, B → CEF, H → G}. What relations are included in the final decomposition
   when decomposing R into BCNF in the order of functional dependencies set F?

8. True or False: The decomposition of attribute set R = ABCDEF, given the functional dependency set
   F = {B → D, E → F, D → E, D → B, F → BD}, into ABDE, BCDF is lossless.