IN3020 and IN4020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 1 of 12 (9%)

Let us call a unary set relational algebra operation $Op1$ *monotonic* if, whenever every tuple in relation $R1$ is also a tuple in relation $R2$ (over the same schema), every tuple in $Op1(R1)$ is also a tuple in $Op1(R2)$. Analogously, a binary set RA operation $Op2$ is *monotonic* if, whenever every tuple in relation $R1$ is also a tuple in relation $R2$ and every tuple in relation $S1$ is also a tuple in relation $S2$, every tuple in $Op2(R1, S1)$ is also a tuple in $Op2(R2, S2)$.

Which of the following operations are not monotonic: selection, projection, union, difference, Cartesian product, renaming, natural join, theta-join, left semi-join, left anti-join, division, generalised projection, left outer join? List only those operations that are not monotonic. For each of them, explain briefly (i.e., in one sentence) why you think so (you do not need to be formal, intuition is enough).

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 2 of 12 (8%)

Consider a relation over schema $R[A, B]$. Assuming that R may be a bag (multiset), write a bag relational algebra expression that outputs the relation over the same schema that contains all the tuples that have multiplicity exactly one in the input R (and nothing more).

Note: in this and the following questions you may introduce and then use your own notation for RA operations; for example, you may write 'I use LSJ for left semi-join below' and then use LSJ in place of $\ltimes$.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 3 of 12 (9%)

Consider the relation schema $Links[Source, Target, Length]$ and the following SQL query (that uses recursion):

```
WITH RECURSIVE Rec(Source, Target, Length, Hops) AS (
    (SELECT Source, Target, Length, 0 FROM Links)
          UNION
    (SELECT R.Source, L.Target, R.Length + L.Length, R.Hops + 1
     FROM Rec R, Links L
     WHERE R.Target = L.Source AND R.Hops < 2)
) SELECT DISTINCT Source, Target, Length FROM Rec;
```

Can this query be compiled to a logical query plan in the (bag) relational algebra considered in the lectures (i.e., the one without recursion)? If your answer is no, explain why. If your answer is yes, give such a plan.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 4 of 12 (8%)

Does there exist a relation and an attribute in this relation such that it is better not to use any index of any type considered in the lectures for exact search rather than to use one? Justify your answer in no more than a few sentences.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 5 of 12 (8%)

Let the query compilation result in the query plan $\sigma_{A=a}(\sigma_{B=1}\mathsf{Table})$, and let the data consist of the following relation, where each attribute takes 2 bytes:

Table:

A	B
a	1
b	1
b	2
b	3

Optimise this query plan using the procedure considered in the lectures. In particular, give all query plans considered during the optimisation, compute their costs (no details of the computation are necessary), and determine the best plan according to the procedure.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 6 of 12 (8%)

Let us have 500 buffers in main memory and two relations, Table1 and Table2, with one common attribute, A. Let Table1 take 1000 blocks and Table2 take 500 blocks on disk. Finally, let $Vals(\mathsf{Table1.A}) = Vals(\mathsf{Table2.A}) = 1$. Which of the natural join algorithms considered in the lectures should be used to compute the natural join of Table1 and Table2, and why? Here, we assume that we should use the algorithm that is the most efficient according to our cost—that is, the one that minimises the number of disk reads and writes.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 7 of 12 (8%)

Which of the following is true?

1. Two different serial schedules of two transactions are always conflict-equivalent.
2. Two different serial schedules of two transactions are never conflict-equivalent.
3. Two different serial schedules of two transactions may or may not be conflict-equivalent.

Justify your answer in at most 3 sentences.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 8 of 12 (8%)

Consider the following two transactions:

$$T_1 : \mathtt{r}(Y), \mathtt{w}(X), \mathtt{w}(Z), \mathtt{w}(U), \mathtt{r}(Y)$$
$$T_2 : \mathtt{r}(Y), \mathtt{w}(X), \mathtt{r}(U), \mathtt{w}(U)$$

(a) Give a schedule of these two transactions that exhibits a dirty read, or explain why it is not possible.

(b) Give a schedule of these two transactions that exhibits a dirty write, or explain why it is not possible.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 9 of 12 (9%)

Is it possible for a schedule to follow the two-phase locking concurrency protocol, but not follow the timestamp-based protocol? Justify your answer—that is, either give an example of such a schedule or explain why no such schedule exists.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 10 of 12 (8%)

Assume that we use the deferred update protocol. Consider a database where, before the relevant transactions start, items X and Y have values 3 and 5 respectively. Consider also the following (partial) schedule executed on this database:

$$\mathtt{b}_2, \mathtt{b}_1, \mathtt{r}_1(X), \mathtt{eval}_1(Y{:=}X * 2), \mathtt{r}_2(Y), \mathtt{eval}_2(Y{:=}Y + 1), \mathtt{w}_2(Y), \\ \mathtt{w}_1(Y), \mathtt{a}_1, \mathtt{b}_3, \mathtt{r}_3(Y), \mathtt{eval}_3(Y{:=}Y + 1), \mathtt{w}_3(Y), \mathtt{c}_2$$

In this schedule, $\mathtt{eval}_1$, $\mathtt{eval}_2$, $\mathtt{eval}_3$ are update operations that are evaluated within transactions T_1, T_2, T_3, respectively. Assume that the last (rigid) checkpoint occurs at the beginning of the schedule, and a system crash happens at the end of the schedule. Assuming that the log buffer flushes immediately after each log update, give the log state at the moment of the crash and discuss what the system must do to recover after the crash.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 11 of 12 (8%)

In at most a few sentences, discuss the relationships between SQL and CRUD. In particular, explain whether SQL supports all CRUD operations and whether it allows data manipulation beyond CRUD.

IN3020 – Database Systems Exam, 2026

29 May, 9:00am–1:00pm (Inspera)

Question 12 of 12 (9%)

Consider the property graph created by the following Cypher clause:

```
CREATE (p1:Person {name:"Alice", born:1990})
-[m:MarriedTo {from:2015, to:2020}]-
(p2:Person {name:"Bob", born:1988})
```

Write an RDF graph that represents the same data. You may assume existence of all relevant IRIs with prefix `http://mygraph.org/`—that is, you can start your graph with `@base http://mygraph.org/` and then use shortcuts such as `<#Person>`.

Note: your knowledge of the Turtle syntax will not be evaluated, but it should be clear what triples in your RDF graph are.