TECHNISCHE UNIVERSITEIT DELFT
Faculteit Elektrotechniek, Wiskunde en Informatica

TI2506 Information and Data Modeling (2015-2016 Q2)

Monday, 25 January 2016
13.30-15.30

INSTRUCTIONS:

- This exam consists of 4 parts with in total 20 multiple-choice questions. All questions are worth an equal number of points.
- The usage of books, notes, old exams, and other written resources is explicitly allowed during the exam. The use of electronic aids such as smart-phones, laptops, et cetera, is not allowed.
- There is only one right answer for each question. If you think there are more, pick the best one.
- You are not allowed to make corrections on the multiple-choice answer form (MAF). You are therefore advised to first mark the answers on this exam and later copy them to the MAF. If you need to make corrections anyway, ask for a new form and copy all your answers to it.
- Note that the order of the answers on your form is not always A-B-C-D.
- Enter your *studentnumber* on the form with digits as well as by filling the boxes.
- Sign the MAF. Without your signature, the form is not valid. Since you might forget this at the end, you are advised to do this at the start of the exam.

Good Luck!

Part 2 – Concurrency Controllers

Consider the three transactions shown below. Transaction T_1 is transferring 100€ from one account with balance bal_x to another account with balance bal_y . T_2 is increasing the balance of these two accounts by 10%, while T_3 is withdrawing 50€ from one account with balance bal_z .

T_1	T_2	T_3
begin_transaction		
read(bal_x)		
bal_x=bal_x - 100		
write(bal_x)		
	begin_transaction	
	read(bal_x)	
	bal_x=bal_x * 1.1	
		begin_transaction
		read(bal_z)
		bal_z=bal_z - 50
		write(bal_z)
		commit_transaction
	write(bal_x)	
	read(bal_y)	
	baly=bal_y * 1.1	
	write(bal_y)	
	commit_transaction	
read(bal_y)		
bal_y=bal_y + 100		
write(bal_y)		
commit_transaction		

7) Is the schedule serializable?

(a)	True	(b)	False

8) In what order should the operations of the schedule be executed if it is fed to a strict two-phase locking concurrency controller to ensure a serializable strict schedule? (Assume that locks are requested one by one when needed and released once after commit)

(a)	R1(X), W1(X), R2(X), R3(Z), W3(Z), Commit3, R1(Y), W1(Y), Commit1, W2(X), R2(Y), W2(Y), Commit2
(b)	R1(X), W1(X), R3(Z), W3(Z), Commit3, R1(Y), W1(Y), Commit1, R2(X), W2(X), R2(Y), W2(Y), Commit2
(c)	R1(X), W1(X), R1(Y), W1(Y), Commit1, R2(X), W2(X), R2(Y), W2(Y), Commit2, R3(Z), W3(Z), Commit3

9) Considering a validation concurrency controller with deferred update approach, which transactions will commit and which will abort? (consider the grey shaded cells as the duration of the validation period).

(a)	T1, T2 and T3 will all commit.
(b)	T3 will commit while T1 and T2 will abort.
(c)	T2 and T3 will commit while T1 will abort.

Part 4 – RDF SPARQL

16) Which query should be used to list all the classes in a dataset

(a)
```
SELECT DISTINCT ?type
WHERE {
  ?s rdf:type ?type .
}
```

(b)
```
SELECT DISTINCT ?s
WHERE {
  ?s rdf:type ?type .
}
```

(c)
```
SELECT DISTINCT ?type
WHERE {
  ?s ?type owl:Thing .
}
```

17) Which query should be used to know whether or not the Amazon river is longer than the Nile river?

(a)
```
ASK
{
  Dbpedia:Amazon_River prop:length ?amazon .
  Dbpedia:Nile prop:length ?nile .
  FILTER(?amazon < ?nile) .
}
```

(b)
```
ASK
{
  Dbpedia:Amazon_River prop:length ?amazon .
  Dbpedia:Nile prop:length ?nile .
  FILTER(?amazon > ?nile) .
}
```

(c)
```
Select ?amazon ?nile
{
  Dbpedia:Amazon_River prop:length ?amazon .
  Dbpedia:Nile prop:length ?nile .
}
```

19) Given the following RDF dataset:

```
@prefix : <http://inf.ed.ac.uk/ont#> .
@prefix dbp: <http://dbpedia.org/property/> .
@prefix dc: <http://purl.org/dc/elements/1.1/> .
@prefix rdf:  <http://www.w3.org/1999/02/22-rdf-syntax-ns#> .
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#> .
@prefix xsd: <http://www.w3.org/2001/XMLSchema#> .

:aroast dc:title "Artisan Roast" .
:aroast :rating "5"^^xsd:decimal .

:pyard dc:title "Peter's Yard" .
:pyard :rating "4"^^xsd:decimal .

:hacraft dc:title "Himalaya Art & Craft" .
:hacraft :rating "3"^^xsd:decimal .

:vittoria dc:title "Vittoria" .
:vittoria :rating "4"^^xsd:decimal .
:vittoria :hasCuisine :italian .
```

What is the answer of the following SPARQL query?

```
SELECT ?title WHERE
{
  ?restaurant dc:title ?title .
  OPTIONAL {?restaurant :hasCuisine ?Cuisine}
  FILTER {!Bound(?Cuisine)}
}
```

(a)	"Artisan Roast" "Peter's Yard" "Himalaya Art & Craft"
(b)	"Artisan Roast" "Peter's Yard" "Himalaya Art & Craft" "Vittoria"
(c)	"Vittoria"