

Data Science and Database Technology

Iniziato mercoledì, 3 febbraio 2021, 14:29

Stato Completato

Terminato mercoledì, 3 febbraio 2021, 14:31

Tempo impiegato 1 min. 33 secondi

Valutazione 14,00 su un massimo di 31,00 (45%)

Domanda 1

Risposta corretta

Punteggio ottenuto

1,00 su 1,00

Theory (Recovery) (1 point, -15% penalty for a wrong answer)

A log file contains the following list of operations

B(T1) B(T2) U(T1) U(T1) I(T2) B(T3) A(T2) CK(T1,T3) I(T3) B(T4) D(T4) C(T1) B(T5)
U(T5) A(T4) **failure**

Notation:

T_i: transaction i

B(T_i): Transaction begin

U(T_i): Update - I(T_i): Insert - D(T_i): Delete (the other parameters have been omitted for the sake of brevity)

CK(T_i, ..., T_j, ...): Checkpoint with a list of active transactions

C(T_i): Commit transaction

A(T_i): Abort transaction

What are the final Undo/Redo lists built during the warm restart phase?

- (a) Undo {T1, T5}, Redo {T3, T4}
- (b) Undo {T3, T5}, Redo {T1, T4}
- (c) Undo {T3, T4, T5}, Redo {T1} ✓
- (d) Undo {T1, T4}, Redo {T3, T5}
- (e) Undo {T3, T5}, Redo {T1, T2, T4}
- (f) Undo {T3,T4, T5}, Redo {T1, T2}

Risposta corretta.

La risposta corretta è: Undo {T3, T4, T5}, Redo {T1}

Domanda 2

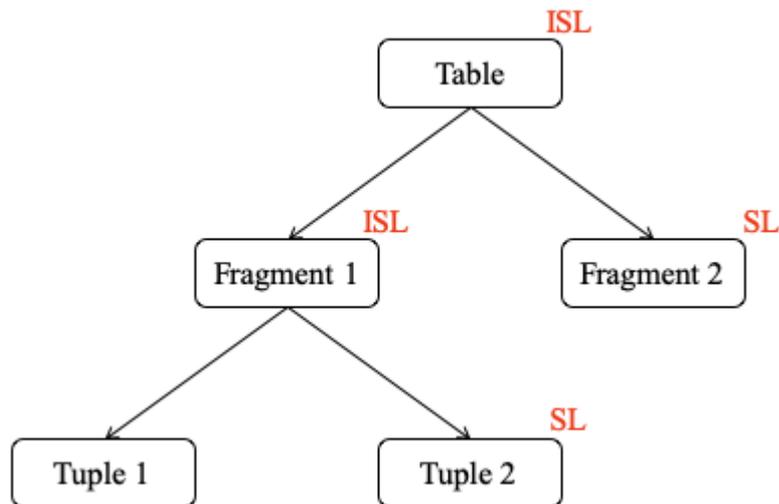
Risposta corretta

Punteggio ottenuto

1,00 su 1,00

Theory (Hierarchical locking) (1 point, -15% penalty for a wrong answer)

The following image shows (in red) the locks acquired by a transaction TA for a table, at various granularity levels.



A transaction TB needs to modify Tuple 1 (among other operations). Which of the following statements is correct?

- (a) Transaction TB can obtain an SIXL on Table, IXL on Fragment 1, XL on Tuple 1 ✓
- (b) Transaction TB can obtain an ISL on Table, IXL on Fragment 1, XL on Tuple 1
- (c) Transaction TB can obtain an IXL on Table, IXL on Fragment 1, IXL on Tuple 1
- (d) Transaction TB cannot obtain any lock that allows modifying Tuple 1
- (e) Transaction TB can obtain an XL on Table, XL on Fragment 1, XL on Tuple 1
- (f) Transaction TB can obtain an IXL on Table, XL on Fragment 1, XL on Tuple 1

Risposta corretta.

La risposta corretta è: Transaction TB can obtain an SIXL on Table, IXL on Fragment 1, XL on Tuple 1

Domanda 3

Risposta corretta

Punteggio ottenuto
1,00 su 1,00**Theory (MongoDB) (1 point, -15% penalty for a wrong answer)**

The following is a document from a MongoDB collection called "stayings". This collection contains data on reservations for an hotel.

```
{
  checkin: 2021-01-10,
  cost_per_night: 120.0,
  total_nights: 4,
  payment_method: "cash",
  room_type: "deluxe"
}
```

Based on the structure that can be inferred from this document, which of the following MongoDB queries extracts, for each room type, the earnings of the hotel?

-
- (a) `db.stayings.mapReduce(
 function() { emit(null, this.cost_per_night *
this.total_nights)}, // map
 function(key, values) { return Array.sum(values) }, //
reduce
 ...
)`
 - (b) `db.stayings.mapReduce(
 function() { emit(this.room_type, this.cost_per_night
)}, // map
 function(key, values) { return Array.sum(values) }, //
reduce
 ...
)`
 - (c) `db.stayings.mapReduce(
 function() { emit(null, this.cost_per_night)}, // map
 function(key, values) { return Array.sum(values) }, //
reduce
 ...
)`
 - (d) `db.stayings.mapReduce(
 function() { emit(this.room_type, this.cost_per_night *
this.total_nights)}, // map
 function(key, values) { return Array.sum(values) }, //
reduce
 ...
) ✓`

- (e) `db.stayings.mapReduce(
 function() { emit(null, this.cost_per_night)}, // map
 function(key, values) { return Array.avg(values) }, //
 reduce
 ...
)`
- (f) `db.stayings.mapReduce(
 function() { emit(null, this.cost_per_night *
 this.total_nights)}, // map
 function(key, values) { return Array.avg(values) }, //
 reduce
 ...
)`
- (g) `db.stayings.mapReduce(
 function() { emit(this.room_type, this.cost_per_night
)}, // map
 function(key, values) { return Array.avg(values) }, //
 reduce
 ...
)`
- (h) `db.stayings.mapReduce(
 function() { emit(this.room_type, this.cost_per_night *
 this.total_nights)}, // map
 function(key, values) { return Array.avg(values) }, //
 reduce
 ...
)`

Risposta corretta.

La risposta corretta è: `db.stayings.mapReduce(
 function() { emit(this.room_type, this.cost_per_night *
 this.total_nights)}, // map
 function(key, values) { return Array.sum(values) }, //
 reduce
 ...
)`

Domanda 4

Risposta corretta

Punteggio ottenuto

1,00 su 1,00

Theory (Classification) (1 point, -15% penalty for a wrong answer)

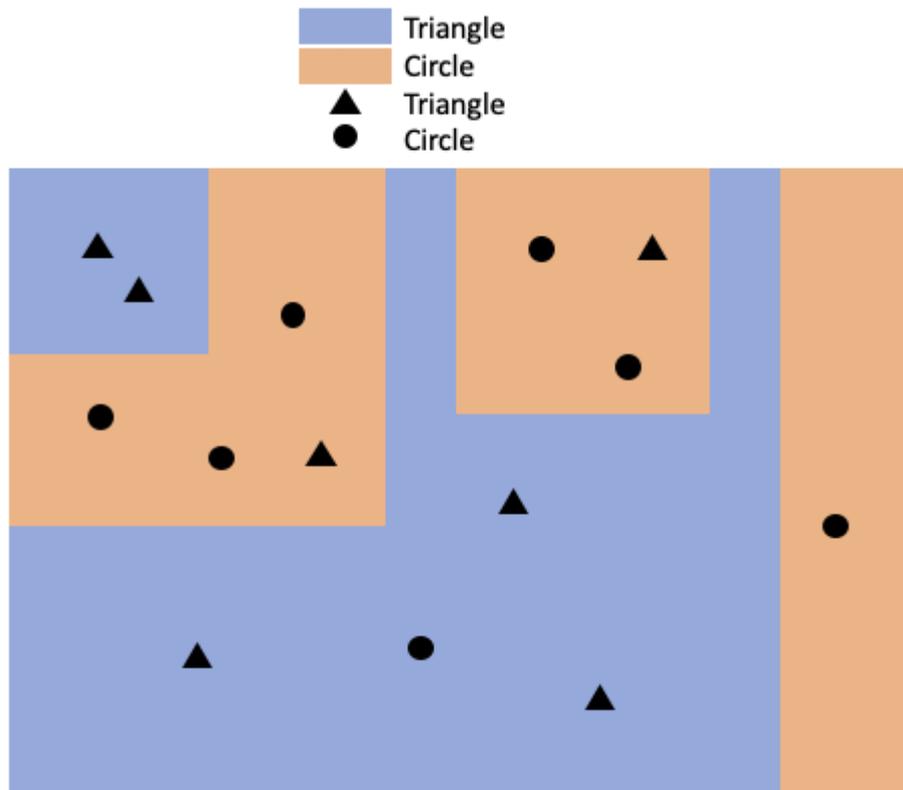
Precision(C): fraction of elements correctly classified in class C, among all elements assigned to class C

Recall(C): fraction of elements correctly classified in class C, among all elements that belong to class C

Accuracy: Fraction of elements correctly classified, among all elements that have been classified

Given a decision tree, we want to classify a set of points described by two numerical variables that can be displayed in a 2-dimensional space.

The following figure shows the 5 regions delimited by the decision boundaries of the decision tree. The color of each region identifies the class label assigned by the decision tree to each region ("blue" for class "triangle", "orange" for class "circle").



The test set includes 14 points whose class label is known. These points are shown in the 2-dimensional space above with their correct class label (either a triangle, or a circle).

Which of the following statements is correct?

- (a) The accuracy is 11/13
- (b) The precision for class "triangle" is 5/7
- (c) The recall for class "circle" is 6/7 ✓
- (d) The recall for class "triangle" is 5/6
- (e) The precision for class "circle" is 6/7

Risposta corretta.

Domanda 5

Risposta corretta

Punteggio ottenuto

2,00 su 2,00

Theory (Frequent itemsets) (1 point, -15% penalty for a wrong answer)

The Apriori principle states that, if an itemset is frequent, then all of its subsets must also be frequent.

You are given the following transactional dataset.

| |
|---------|
| A C E |
| A B C D |
| C D |
| A C D |
| A B |
| A B E |
| B C E |
| A D E |
| A B E |
| A C D E |

You need to apply the Apriori algorithm to extract the frequent itemsets with minsup = 2 (an itemset is frequent if it is contained in at least 2 transactions).

During the generation of the candidate itemsets of length 3, after the "prune step" with the application of the "Apriori Principle", what is the number of candidates extracted for counting their support in the dataset?

-
- (a) 7 ✓
 - (b) 4
 - (c) 9
 - (d) 6
 - (e) 5
 - (f) 8

Risposta corretta.

La risposta corretta è: 7

Domanda 6

Risposta corretta

Punteggio ottenuto

1,50 su 1,50

Cardinalities (1.5 points, -15% penalty for each wrong answer)

The following tables are provided:

VEHICLE(CodV, Model, Brand, Colour, RegistrationDate)
USER(CodU, FirstName, Surname, BirthDate, Address, City, Region)
REPAIR-HISTORY(CodV, StartDate, EndDate, ProblemType)
VEHICLE-RENTAL(CodV, StartDate, CodU, EndDate, PaymentMode)

Assume the following cardinalities:

- card(VEHICLE) = 10^4 tuples,
MIN(RegistrationDate) = 1/1/2001,
MAX(RegistrationDate) = 31/12/2020
distinct values of Brand = 5
- card(USER) = 10^5 tuples,
MIN(BirthDate) = 1/1/1930, MAX(BirthDate) = 31/12/2000
distinct values of City = 10
- card(REPAIR-HISTORY) = 10^6 tuples,
distinct values of ProblemType = 10
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020
- card(VEHICLE-RENTAL) = $2 \cdot 10^6$ tuples,
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020

Furthermore, assume the following reduction factor for the having clauses:

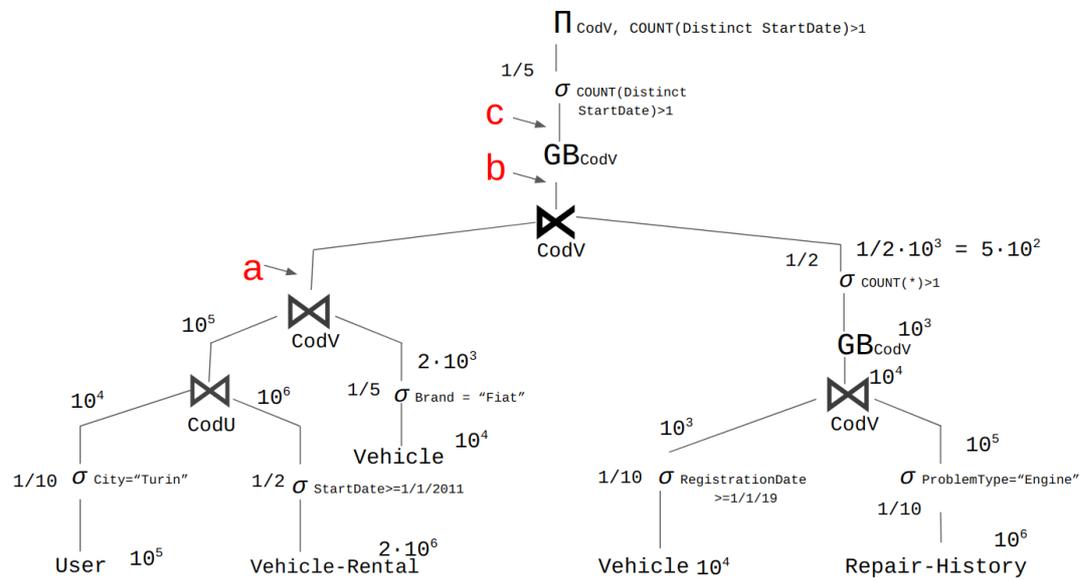
Having COUNT(*) > 1 = 1/2

Having COUNT(Distinct StartDate) > 1 = 1/5

Consider the following query:

```
select VE.CodV, count(Distinct VR.StartDate)
from VEHICLE-RENTAL VR, VEHICLE VE, USER U
where VR.CodV=VE.CodV and U.CodU=VR.CodU and VR.Start
Date>=1/1/2011
and VE.Brand='Fiat'
and U.City='Turin'
and VE.CodV IN ( SELECT CodV
                  FROM VEHICLE V2, REPAIR-HISTORY RH
                  WHERE V2.CodV=RH.CodV
                    and ProblemType='Engine'
                    and RegistrationDate>=1/1/2019
                  GROUP BY V2.CodV
                  HAVING COUNT(*)>1 )
group by VE.CodV
having COUNT(Distinct VR.StartDate)>1
```

The figure below represents the query tree for the query above.



Specify the cardinality of each node indicated by the red letters (a,b,c) in the figure. There is one right answer for each node a,b,c.

Scegli una o più alternative:

- (a) c: $\approx 5 \cdot 10^3$
- (b) a: $2 \cdot 10^3$
- (c) b: 10^3 ✓
- (d) c: $\approx 10^3$
- (e) c: $\approx 2 \cdot 10^3$
- (f) a: 10^4
- (g) c: $\approx 10^2$ ✓
- (h) b: $2 \cdot 10^3$
- (i) a: $5 \cdot 10^4$
- (j) b: $5 \cdot 10^3$
- (k) b: $2 \cdot 10^2$
- (l) a: $2 \cdot 10^4$ ✓

Risposta corretta.

La risposta corretta è: a: $2 \cdot 10^4$, b: 10^3 , c: $\approx 10^2$

Domanda 7

Risposta corretta

Punteggio ottenuto

1,50 su 1,50

Indexes (1.5 points, -15% penalty for each wrong answer)

The following tables are provided:

VEHICLE(CodV, Model, Brand, Colour, RegistrationDate)
USER(CodU, FirstName, Surname, BirthDate, Address, City, Region)
REPAIR-HISTORY(CodV, StartDate, EndDate, ProblemType)
VEHICLE-RENTAL(CodV, StartDate, CodU, EndDate, PaymentMode)

Assume the following cardinalities:

- card(VEHICLE) = 10^4 tuples,
MIN(RegistrationDate) = 1/1/2001,
MAX(RegistrationDate) = 31/12/2020
distinct values of Brand = 5
- card(USER) = 10^5 tuples,
MIN(BirthDate) = 1/1/1930, MAX(BirthDate) = 31/12/2000
distinct values of City = 10
- card(REPAIR-HISTORY) = 10^6 tuples,
distinct values of ProblemType = 10
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020
- card(VEHICLE-RENTAL) = $2 \cdot 10^6$ tuples,
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020

Furthermore, assume the following reduction factor for the having clauses:

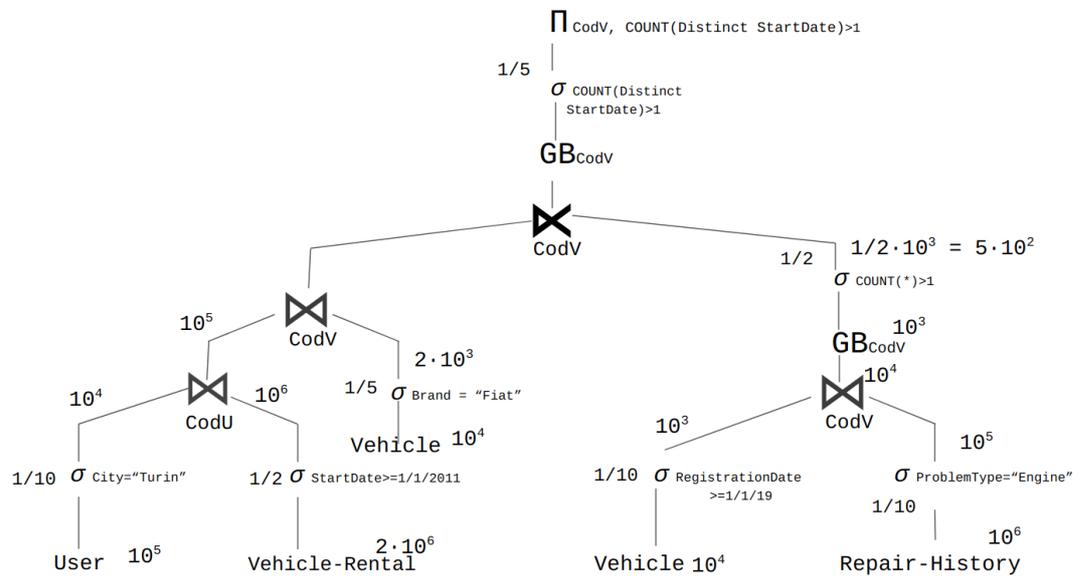
Having COUNT(*) > 1 = 1/2

Having COUNT(Distinct StartDate) > 1 = 1/5

Consider the following query:

```
select VE.CodV, count(Distinct VR.StartDate)
from VEHICLE-RENTAL VR, VEHICLE VE, USER U
where VR.CodV=VE.CodV and U.CodU=VR.CodU and VR.Start
Date>=1/1/2011
and VE.Brand='Fiat'
and U.City='Turin'
and VE.CodV IN ( SELECT CodV
                  FROM VEHICLE V2, REPAIR-HISTORY RH
                  WHERE V2.CodV=RH.CodV
                    and ProblemType='Engine'
                    and RegistrationDate>=1/1/2019
                  GROUP BY V2.CodV
                  HAVING COUNT(*)>1 )
group by VE.CodV
having COUNT(Distinct VR.StartDate)>1
```

The figure below represents the query tree for the query above.



Select one or more secondary physical structures to increase query performance (if possible) among the options below. You can select multiple correct answers.

Scegli una o più alternative:

- (a) CREATE INDEX IndexF ON Vehicle(Brand) - HASH
- (b) CREATE INDEX IndexB ON User(City) - B+-Tree
- (c) CREATE INDEX IndexE ON Vehicle(Brand) - B+-Tree
- (d) CREATE INDEX IndexA ON User(City) - HASH ✓
- (e) CREATE INDEX IndexJ ON Repair-History(ProblemType) - B+-Tree
- (f) CREATE INDEX IndexI ON Repair-History(ProblemType) - HASH ✓
- (g) CREATE INDEX IndexD ON Vehicle-Rental(StartDate) - B+-Tree
- (h) None - secondary physical structures would not increase query performance.
- (i) CREATE INDEX IndexG ON Vehicle(RegistrationDate) - B+-Tree ✓
- (j) CREATE INDEX IndexC ON Vehicle-Rental(StartDate) - HASH
- (k) CREATE INDEX IndexH ON Vehicle(RegistrationDate) - HASH

Risposta corretta.

La risposta corretta è: CREATE INDEX IndexA ON User(City) - HASH, CREATE INDEX IndexG ON Vehicle(RegistrationDate) - B+-Tree, CREATE INDEX IndexI ON Repair-History(ProblemType) - HASH

Domanda 8

Risposta corretta

Punteggio ottenuto
2,00 su 2,00**Group by anticipation (2 points, -15% penalty for a wrong answer)**

The following tables are provided:

VEHICLE(CodV, Model, Brand, Colour, RegistrationDate)
USER(CodU, FirstName, Surname, BirthDate, Address, City, Region)
REPAIR-HISTORY(CodV, StartDate, EndDate, ProblemType)
VEHICLE-RENTAL(CodV, StartDate, CodU, EndDate, PaymentMode)

Assume the following cardinalities:

- card(VEHICLE) = 10^4 tuples,
MIN(RegistrationDate) = 1/1/2001,
MAX(RegistrationDate) = 31/12/2020
distinct values of Brand = 5
- card(USER) = 10^5 tuples,
MIN(BirthDate) = 1/1/1930, MAX(BirthDate) = 31/12/2000
distinct values of City = 10
- card(REPAIR-HISTORY) = 10^6 tuples,
distinct values of ProblemType = 10
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020
- card(VEHICLE-RENTAL) = $2 \cdot 10^6$ tuples,
MIN(StartDate) = 1/1/2001, MAX(StartDate) = 31/12/2020

Furthermore, assume the following reduction factor for the having clauses:

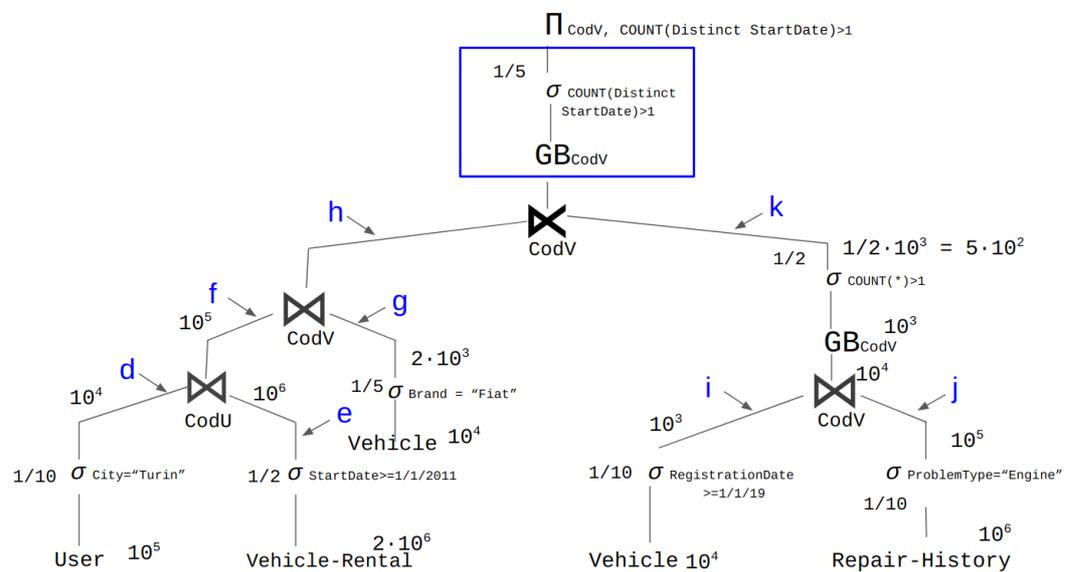
Having COUNT(*) > 1 = 1/2

Having COUNT(Distinct StartDate) > 1 = 1/5

Consider the following query:

```
select VE.CodV, count(Distinct VR.StartDate)
from VEHICLE-RENTAL VR, VEHICLE VE, USER U
where VR.CodV=VE.CodV and U.CodU=VR.CodU and VR.Start
Date>=1/1/2011
and VE.Brand='Fiat'
and U.City='Turin'
and VE.CodV IN ( SELECT CodV
                  FROM VEHICLE V2, REPAIR-HISTORY RH
                  WHERE V2.CodV=RH.CodV
                    and ProblemType='Engine'
                    and RegistrationDate>=1/1/2019
                  GROUP BY V2.CodV
                  HAVING COUNT(*)>1 )
group by VE.CodV
having COUNT(Distinct VR.StartDate)>1
```

The figure below represents the query tree for the query above.



Analyze the group by anticipation of **group by CodV Having COUNT(DISTINCT StartDate) > 1** represented in the box. Select the solution that allows the greatest reduction in cardinality (if any).

- (a) It is possible to anticipate it in branch **f** ✓
- (b) It is possible to anticipate it in branch **j**
- (c) It is possible to anticipate it in branch **g**
- (d) It is not possible to anticipate the Group BY **group by CodV Having COUNT(DISTINCT StartDate) > 1**
- (e) It is possible to anticipate it in branch **k**
- (f) It is possible to anticipate it in branch **i**
- (g) It is possible to anticipate it in branch **e**
- (h) It is possible to anticipate it in branch **h**
- (i) It is possible to anticipate it in branch **d**

Risposta corretta.

La risposta corretta è: It is possible to anticipate it in branch **f**

Domanda 9

Risposta corretta

Punteggio ottenuto

1,00 su 1,00

Conceptual schema (1 point, -15% penalty for each wrong answer)

A company that manufactures cameras is interested in analyzing the sales history. The analysis should be conducted separately for the different camera models. Each model is characterized by a category (e.g., "compact", "reflex", ...) and by a list of technical specs. The technical specs vary depending on the model and their number is not known during the design phase of the data warehouse. Cameras are sold through two sale channels: e-commerce and traditional stores. For each purchase, some customer information is recorded.

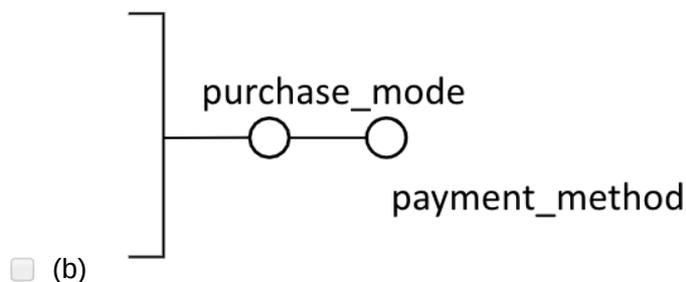
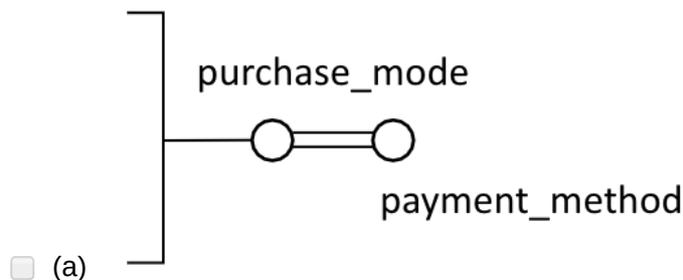
The company wants to analyze the average revenue (in euros) per purchased item based on:

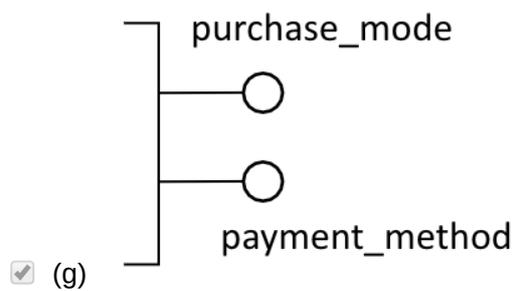
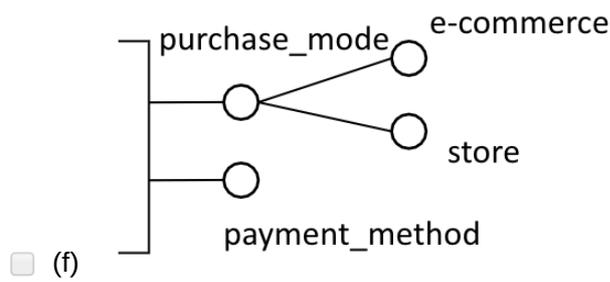
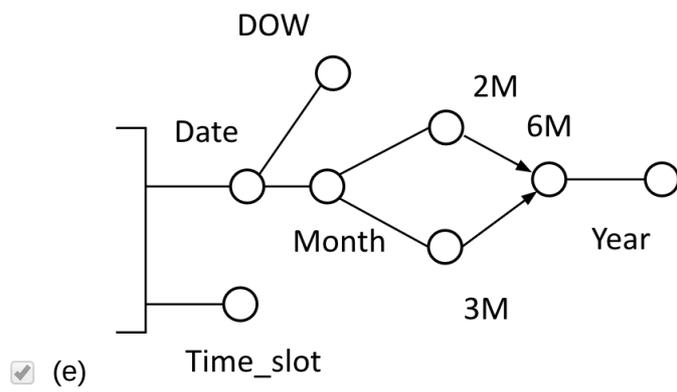
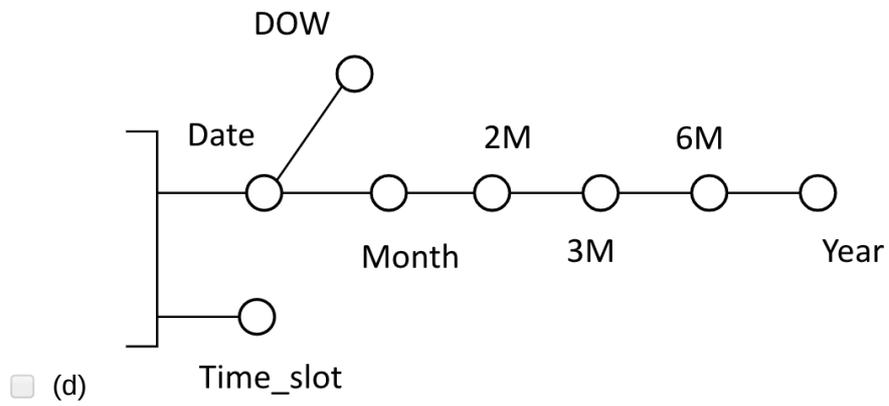
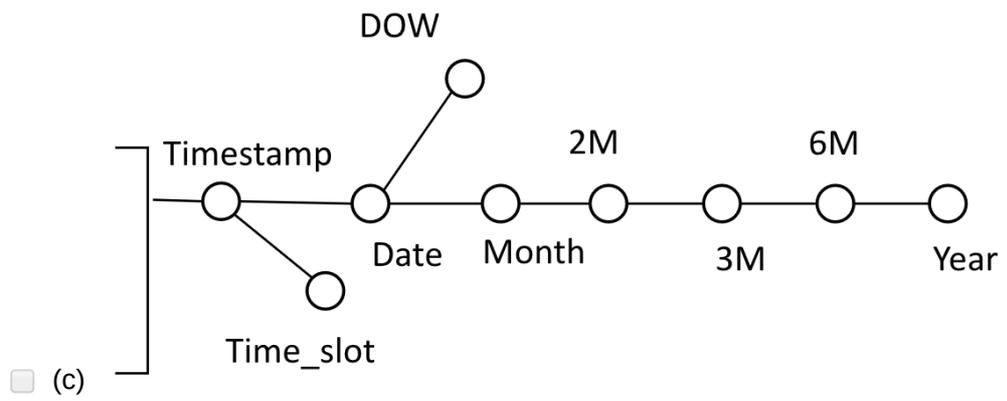
- purchase date, day of the week, month, 2-month period, 3-month period, semester and year
- time slot (in 4-hour blocks, from 00:01 to 04:00, then 04:01-08:00, 8:01-12:00, 12:01-16:00, 16:01-20:00, 20:01-00:00)
- camera model
- model technical specs (e.g., lens, resolution, zoom, etc.)
- sale channel (e-commerce, traditional store)
- payment method (e.g., credit card, debit card, check, etc.)
- city, province, region and country of the customer
- age range of the customer (18-30, 31-50, 51-65, >65 years old)

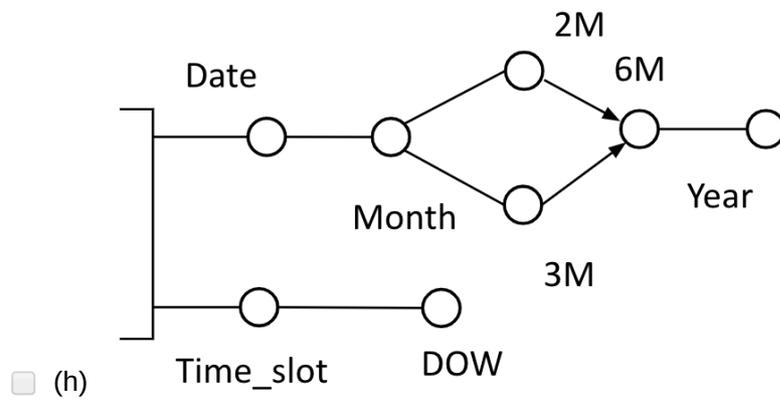
Select, among the following proposed dimensions, those that meet the requirements described in the problem specification.

Note: you should select all and only the right answers; multiple answers are correct. There is only one right answer for each dimension group (2 groups in total).

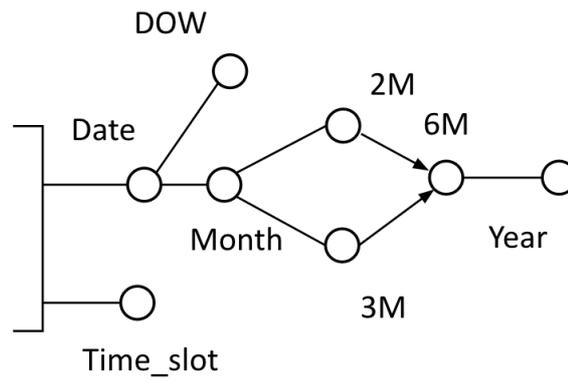
Scegli una o più alternative:



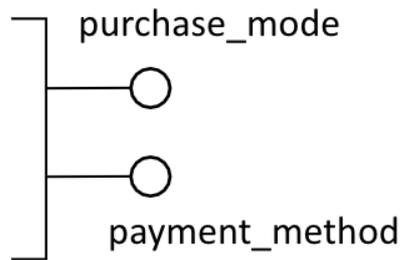




Risposta corretta.



La risposta corretta è:



Domanda 10

Risposta corretta

Punteggio ottenuto
1,00 su 1,00**Conceptual schema (1 point, -15% penalty for each wrong answer)**

A company that manufactures cameras is interested in analyzing the sales history. The analysis should be conducted separately for the different camera models. Each model is characterized by a category (e.g., "compact", "reflex", ...) and by a list of technical specs. The technical specs vary depending on the model and their number is not known during the design phase of the data warehouse. Cameras are sold through two sale channels: e-commerce and traditional stores. For each purchase, some customer information is recorded.

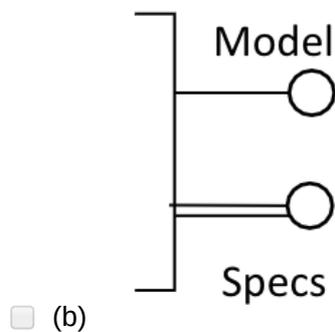
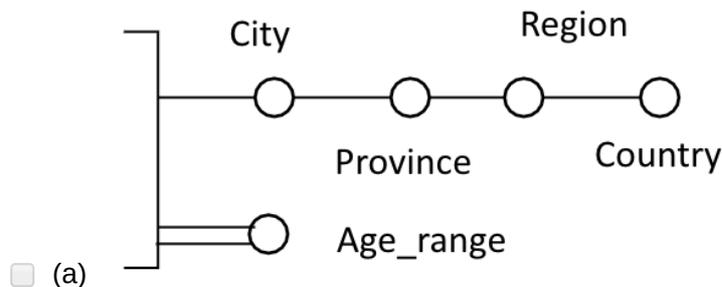
The company wants to analyze the average revenue (in euros) per purchased item based on:

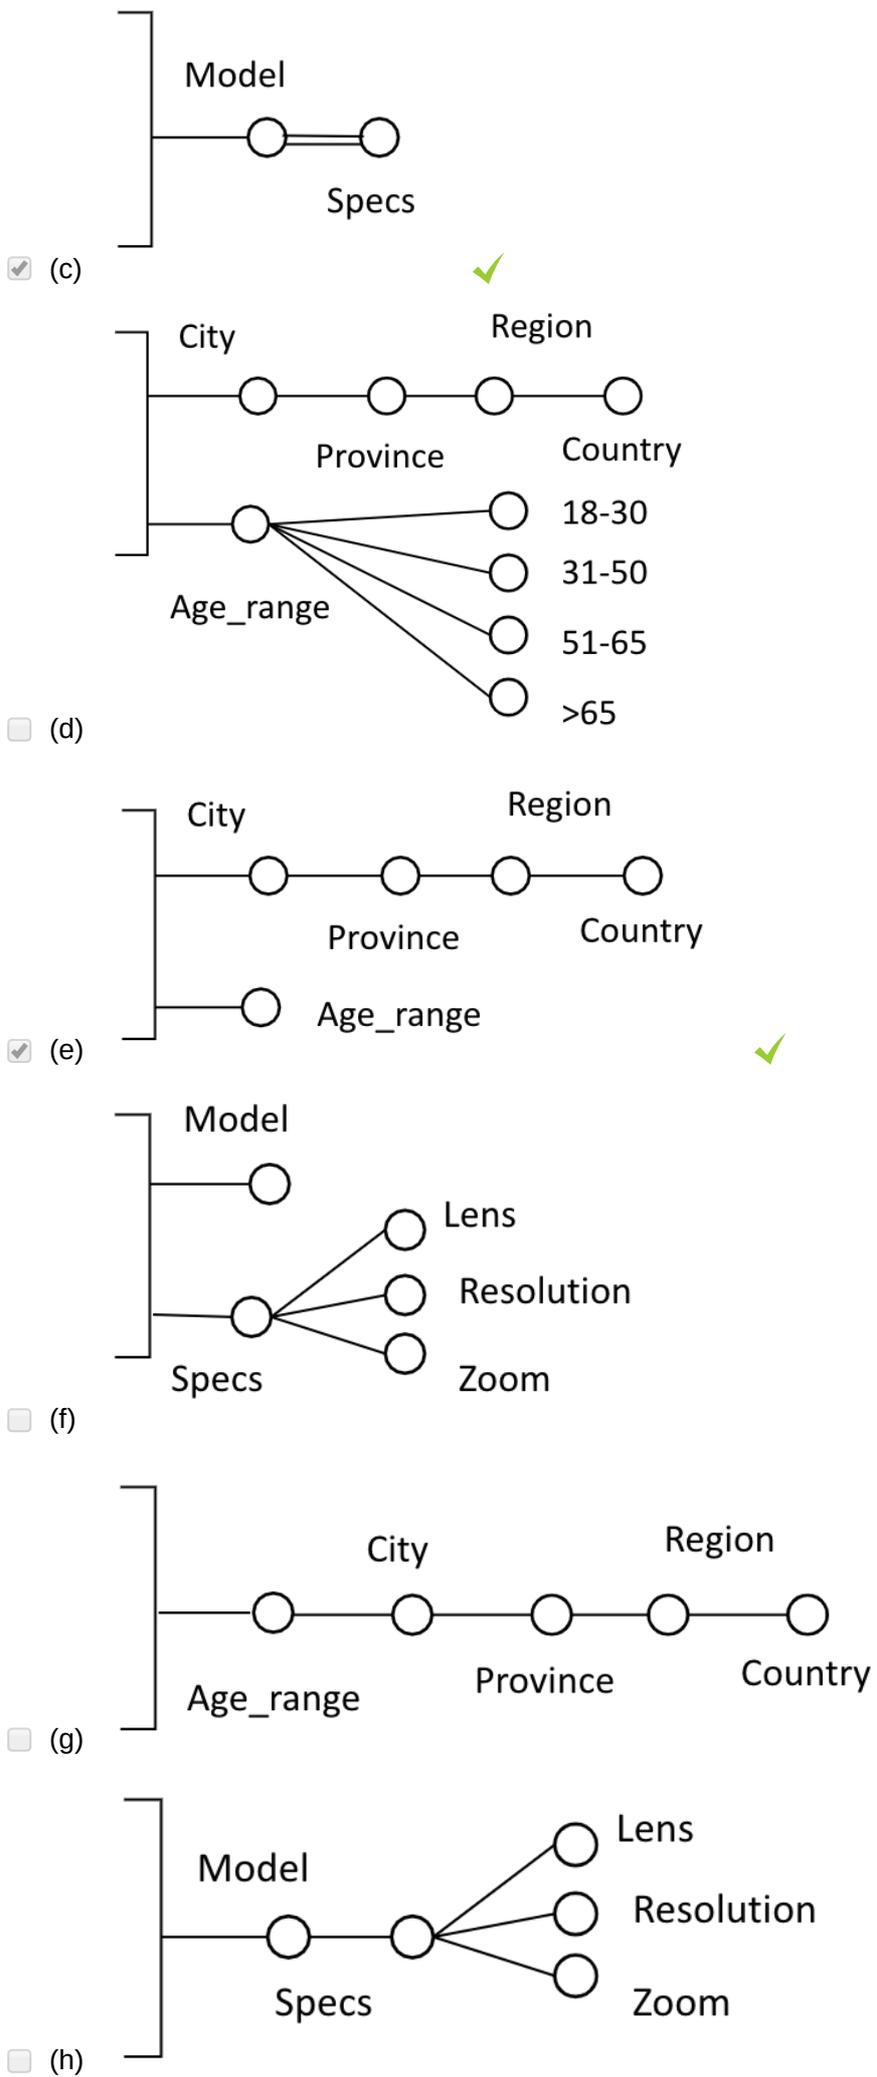
- purchase date, day of the week, month, 2-month period, 3-month period, semester and year
- time slot (in 4-hour blocks, from 00:01 to 04:00, then 04:01-08:00, 8:01-12:00, 12:01-16:00, 16:01-20:00, 20:01-00:00)
- camera model
- model technical specs (e.g., lens, resolution, zoom, etc.)
- sale channel (e-commerce, traditional store)
- payment method (e.g., credit card, debit card, check, etc.)
- city, province, region and country of the customer
- age range of the customer (18-30, 31-50, 51-65, >65 years old)

Select, among the following proposed dimensions, those that meet the requirements described in the problem specification.

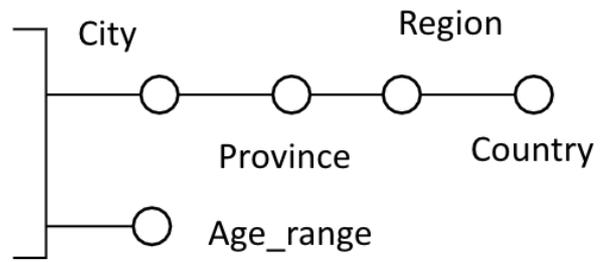
Note: you should select all and only the right answers; multiple answers are correct. There is only one right answer for each dimension group (2 groups in total).

Scegli una o più alternative:

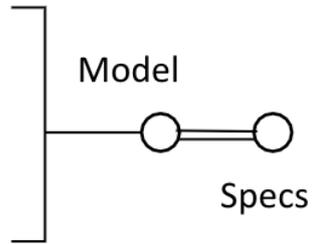




Risposta corretta.



La risposta corretta è:



Domanda 11

Risposta corretta

Punteggio ottenuto
1,00 su 1,00**Measurements (1 point, -15% penalty for each wrong answer)**

A company that manufactures cameras is interested in analyzing the sales history. The analysis should be conducted separately for the different camera models. Each model is characterized by a category (e.g., "compact", "reflex", ...) and by a list of technical specs. The technical specs vary depending on the model and their number is not known during the design phase of the data warehouse. Cameras are sold through two sale channels: e-commerce and traditional stores. For each purchase, some customer information is recorded.

The company wants to analyze the average revenue (in euros) per purchased item based on:

- purchase date, day of the week, month, 2-month period, 3-month period, semester and year
- time slot (in 4-hour blocks, from 00:01 to 04:00, then 04:01-08:00, 8:01-12:00, 12:01-16:00, 16:01-20:00, 20:01-00:00)
- camera model
- model technical specs (e.g., lens, resolution, zoom, etc.)
- sale channel (e-commerce, traditional store)
- payment method (e.g., credit card, debit card, check, etc.)
- city, province, region and country of the customer
- age range of the customer (18-30, 31-50, 51-65, >65 years old)

Select from the list below all and only the attributes required to correctly model the requests in the specifications for the fact table (multiple answers can be correct, since multiple measurements can be required).

Scegli una o più alternative:

- (a) Average number of sales
- (b) Total number of purchased items ✓
- (c) Average number of purchased items
- (d) Total number of customers
- (e) Average revenue per purchase
- (f) Total number of sales transactions
- (g) Total revenue ✓

Risposta corretta.

La risposta corretta è: Total number of purchased items, Total revenue

Domanda 12

Risposta non data

Punteggio max.:
4,00**Extended SQL (4 points)**

The following relational schema is used to analyze the information collected from a hotel review site. For each hotel, the available amenities are known (e.g., wifi, swimming pool, etc.).

```
HOTEL (CodH, Name, CodL)
AMENITIES (CodH, Amenity)
USER (CodU, Nickname, Genre, CodL)
LOCATION (CodL, City, Province, Region)
TIME (CodT, Date, Month, Month_of_Year, Year)
STATS (CodH, CodU, CodT, TotalRatings, NumberOfRatings)
```

Considering the hotels that have a swimming pool, for each month and province, compute:

- The average ratings
 - The ratio between the number of ratings and the total monthly ratings in the region
 - The number of ratings received per day on average (daily average)
-

```
SELECT Month, Province,
       SUM(TotalRatings)/SUM(NumberOfRatings),
       SUM(NumberOfRatings)/SUM(SUM(NumberOfRatings))
OVER (PARTITION BY Month, Region),
       SUM(NumberOfRatings)/COUNT(DISTINCT Date)
FROM STATS S, HOTEL H, AMENITIES A, LOCATION L, TIME T
WHERE S.CodH=H.CodH AND H.CodL=L.CodL AND H.CodH=A.CodH AND
S.CodT=T.CodT AND Amenity = 'Swimming pool'
GROUP BY Month, Province, Region
```

Domanda 13

Risposta non data

Punteggio max.:
4,00**Extended SQL (4 points)**

The following relational schema is used to analyze the information collected from a hotel review site. For each hotel, the available amenities are known (e.g., wifi, swimming pool, etc.).

```
HOTEL (CodH, Name, CodL)
AMENITIES (CodH, Amenity)
USER (CodU, Nickname, Genre, CodL)
LOCATION (CodL, City, Province, Region)
TIME (CodT, Date, Month, Month_of_Year, Year)
STATS (CodH, CodU, CodT, TotalRatings, NumberOfRatings)
```

Considering users resident in the Piedmont region, for each month, compute:

- The total number of ratings
 - The monthly cumulative of the number of ratings since the beginning of each year
 - the rank based on the average rating (the highest value is the first)
-

```
SELECT Month
       SUM(NumberOfRatings),
       SUM(SUM(NumberOfRatings)) OVER
(PARTITION BY Year
 ORDER BY Month
 ROWS UNBOUNDED PRECEDING),
       RANK() OVER (ORDER BY SUM(TotalRatings)/SUM(NumberOfRatings) DE
SC)
FROM STATS S, USER U, LOCATION L, TIME T
WHERE S.CodU=U.CodU AND U.CodL=L.CodL AND S.CodT=T.CodT AND Region = 'Piedmont'
GROUP BY Month, Year
```

Domanda 14

Risposta non data

Punteggio max.:

6,00

Trigger 1 (6 points, no penalty for a wrong answer)

The following relations are given (primary keys are underlined).

VACCINES(VaccineCode, Description, NumberOfAvailableDoses)

NURSE (NurseCode, Name, Surname, BirthDate)

VACCINATION_SHIFT (NurseCode, Date, TimeSlot, NumberOfAssignedPatients)

BOOKINGS (PatientSSN, Date, TimeSlot, VaccineCode, NurseCode)

VACCINE_REQUEST(PatientSSN, Date, TimeSlot, VaccineCode)

We want to automatically manage some activities of a vaccination center. Write the trigger to handle the following activity.

Vaccine booking by a patient.

A patient requests to book a specific vaccine on a given date and time slot (insertion of a new record in the VACCINE_REQUEST table). The trigger must perform the following activities.

(1) It must check if the requested vaccine is administered at the centre and there are still doses available. The VACCINE table stores the vaccines administered at the center and the number of doses still available for each vaccine (when no dose is available, the NumberOfAvailableDoses attribute is equal to 0).

(2) If the vaccine is not administered at the vaccination center, or no dose is available, the trigger ends with an error.

(3) If there are available doses, the nurse who will carry out the vaccination must be selected. The VACCINATION_SHIFT table stores, for each nurse, the number of assigned patients on each date and time slot. The nurse with the lowest number of assigned patients on the requested date and time slot must be selected. Assume that there is at least one nurse on duty on the requested date and time slot and that at most one nurse meets all conditions above.

(4) The booking for the requested vaccine must be inserted into the BOOKINGS table. Furthermore, the number of patients assigned to the selected nurse on the requested date and time slot must be updated. The number of available doses for the requested vaccine in the VACCINES table must also be updated. If there are no available doses left, the value of the NumberOfAvailableDoses attribute must be set to 0.

Draft solution

```
CREATE TRIGGER Vaccination
AFTER insert on VACCINE_REQUEST
FOR EACH ROW
DECLARE
X number;
MyNurseCode number;

BEGIN
```

---- check if the requested vaccine is administered at the centre and there are still doses available

```
SELECT count(*) into X
FROM VACCINES
WHERE VaccineCode = :NEW.VaccineCode and NumberOfAvailableDoses > 0;
```

If (X=0) then

```
----- The vaccine is not available
      Raise_application_error (...);
end if;
```

--- select the nurse who will carry out the vaccination. The nurse with the lowest number of assigned patients on the requested date and time slot must be selected

```
SELECT NurseCode into MyNurseCode
FROM VACCINATION_SHIFT
WHERE Date = :NEW.Date and TimeSlot =:NEW.TimeSlot
and NumberOfAssignedPatients =
      (SELECT MIN(NumberOfAssignedPatients)
      FROM VACCINATION_SHIFT
      WHERE Date = :NEW.Date and TimeSlot =:NEW.TimeSlot);
```

-----insert the reservation

```
INSERT INTO BOOKINGS (PatientSSN, Date, TimeSlot, VaccineCode, NourseCode)
VALUES(:NEW.PatientSSN,:NEW.Date,:NEW.TimeSlot,:NEW.VaccineCode,
myNurseCode);
```

--- update the number of patients assigned to the selected nurse on the requested date and time slot

```
UPDATE VACCINATION_SHIFT
SET NumberOfAssignedPatients = NumberOfAssignedPatients +1
WHERE NurseCode = myNurseCode and Date =:NEW.Date
and TimeSlot =:NEW.TimeSlot;
```

----- update the number of available doses

```
UPDATE VACCINES
SET NumberOfAvailableDoses = NumberOfAvailableDoses - 1
WHERE VaccineCode = :NEW.VaccineCode;
```

END;

Domanda 15

Risposta non data

Punteggio max.:
3,00**Trigger 2 (3 points, no penalty for a wrong answer)**

The following relations are given (primary keys are underlined).

VACCINES(VaccineCode, Description, NumberOfAvailableDoses)

NURSE (NurseCode, Name, Surname, BirthDate)

VACCINATION_SHIFT (NurseCode, Date, TimeSlot, NumberOfAssignedPatients)

BOOKINGS (PatientSSN, Date, TimeSlot, VaccineCode, NurseCode)

VACCINE_REQUEST(PatientSSN, Date, TimeSlot, VaccineCode)

We want to automatically manage some activities of a vaccination center. Write the trigger to handle the following activity.

Integrity constraint on the total number of patients assigned to each nurse on a daily basis.

The total number of patients assigned *per day* to each nurse must not exceed the maximum threshold MAXPATIENTS. Any modification of the VACCINATION_SHIFT table that causes the constraint violation must not be executed. Carefully evaluate all the triggering events on the VACCINATION_SHIFT table.

Draft solution

```
CREATE TRIGGER check_max_reservation
AFTER insert or update of Date, NurseCode, NumberOfAssignedPatients on
VACCINATION_SHIFT
DECLARE
X number;
BEGIN

SELECT count(*) into X
FROM NURSE
```

```
WHERE NurseCode in
  (SELECT NurseCode
   FROM VACCINATION_SHIFT
   GROUP BY NurseCode, Date
   HAVING SUM(NumberOfAssignedPatients) > MAXPATIENTS)
```

```
IF (X <>0) THEN
  Raise_application_error(...)
END IF;
```

```
END;
```

Domanda 16

Risposta non data

Non valutata

This question is not a part of the exam

You can use the text area below to write any note or draft (e.g. intermediate steps of an exercise).

Any text written below will not be considered toward the correction of the exam.