I.C.E.™

Inspiring Creativity & Endeavour

IES | GATE | PSU

DATABASE MANAGEMENT SYSTEM

For

COMPUTER SCIENCE

I.C.E.
Inspiring Creativity & Endeavour
IES | GATE | PSU

DATABASE MANAGEMENT SYSTEM

SYLLABUS

ER model. Relational model: relational algebra, tuple calculus, SQL. Integrity constraints, normal forms. File organization, indexing (e.g., B and B+ trees). Transactions and concurrency control.

ANALYSIS OF GATE PAPERS

Exam Year	1 Mark Ques.	2 Mark Ques.	Total
2003	2	3	8
2004	2	5	12
2005	3	4	11
2006	1	4	9
2007	-	6	12
2008	1	5	11
2009	-	5	10
2010	2	2	6
2011	-	2	4
2012	2	5	12
2013	-	3	6
2014 Set-1	2	3	8
2014 Set-2	2	3	8
2014 Set-3	2	3	8
2015 Set-1	2	2	6
2015 Set-2	2	2	6
2015 Set-3	2	2	6
2016 Set-1	4	1	6
2016 Set-2	2	2	6
2017 Set-1	2	3	8
2017 Set-2	2	4	10

CONTENTS

1 THE ENTITY – RELATIONSHIP MODEL

1.1 INTRODUCTION

An entity is any object, place or activity about which an enterprise keeps data. It is an object which can have instances or occurrences. An entity type is a set of objects which share common properties.
The Database Design consists of three components: Conceptual Design on the basis of user requirements, Data Modeling (Entity – Relationship Diagrams and Normalization), and Physical Design, and Implementation. The major step in conceptual design is to identify entities and relationships, which reflect the data in a natural way. The aim of this step is to specify the conceptual structure of the data. This is known as data modeling. The Entity – Relationship (E – R) model is used as an information model to develop conceptual structure.

1.2 THE E – R MODEL

The E – R data model considers the real world consisting of a set of basic objects and relationships among these objects. A number of attributes are associated with an entity and the attributes describing it. The set of all entities or relationships of the same type is called the entity set or relationship set.

1.3 DATABASE ADMINISTRATOR (DBA)

All controlling of a database system is done by database administrator.

1.3.1 FUNCTIONS OF DBA

1) Decides the storage structure and access strategy.
2) Creation of data dictionary for statistical analysis.
3) Responding to changes in requirements.
4) Performance monitoring.
5) Strategy design for backup and recovery.
6) Authorization checks and validation procedures.
7) Decides the information content.

1.4 THE ENTITY RELATIONSHIP DIAGRAM

1.4.1 RELATIONSHIPS AND RELATIONSHIP SETS

A relationship expresses an association between entities. A relationship set is a set of relationships of the same type. A relationship may also have descriptive attributes. For example, data (last data of account access) could be an attribute of the relationship set.

1.4.1.1 MAPPING CARDINALITIES

It indicates the number of entities with which another entity can be associated via a relationship. The degree of relationship is called cardinality.

a) One-to-one: An entity in A is associated with at most one entity in B, and an entity in B is associated with at most one entity in A.

b) One-to-many: An entity in A is associated with any number in B. An entity in B is associated with at least one entity in A.

c) Many-to-one (N:1): An entity in A is associated with at most one entity in B. An entity in B is associated with any number in A.

1.4.2 KEYS

Differences between entities must be expressed in terms of attributes known as keys. These facilitate us to uniquely identify each entity in a set.

1.4.3 SUPER KEY

It is a set of one or more attributes which put together enable us to identify uniquely an entity in the entity set.

1.4.4 CANDIDATE KEY

A super key may contain extraneous attributes, and we are often interested in the smallest super key. A super key for which no subset is a super key is called a candidate key.

1.4.5 PRIMARY KEY

It is a candidate key (there may be more than one) chosen by the database designer to identify entities in an entity set. The idea of strong and weak entity sets is related to the existence dependencies such as the member of a strong entity set is a dominant entity, and the member of a weak entity set is a subordinate entity. A weak entity set does not have a primary key, but we need a means of distinguishing among the entities.

1.4.6 GENERALIZATION

Generalization hides differences and emphasizes similarities. Distinction is made through attribute inheritance.
Attributes of higher – level entity are inherited by lower - level entities.

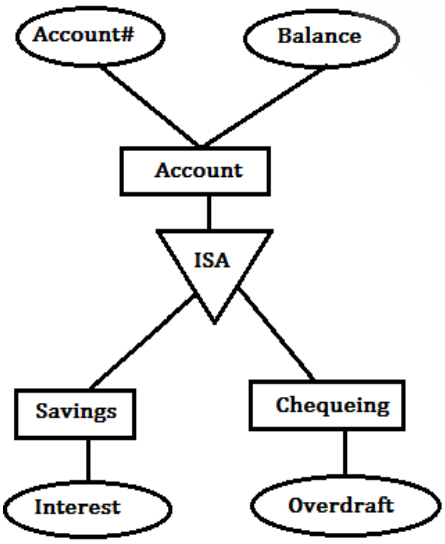

Figure : Generalization

1.4.7 AGGREGATION

The E – R model cannot express relationships among relationships. When would we need such a thing?
Consider a database with information about employees who work on a particular project and using a number of machines for doing that work.
We get the E-R diagram shown in Figure.

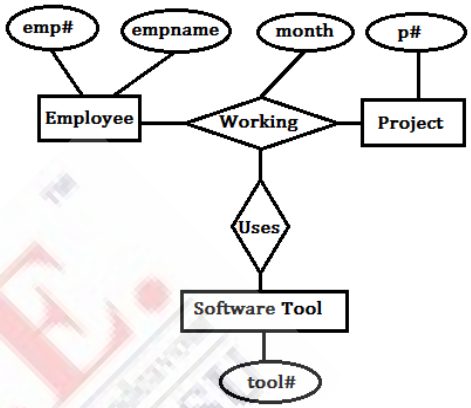

Figure : E - R diagram with redundant relationships

Relationship sets Working and Uses could be combined into a single set. However, they shouldn't be, as this would obscure the logical structure of this scheme. The solution is to use aggregation. It is an abstraction through which relationships are treated as higher – level entities. For our example, we can treat the relationship set Working and the entity sets Employee and Project as a higher –level entity set called Working. Following figure shows the E-R diagram with aggregation.

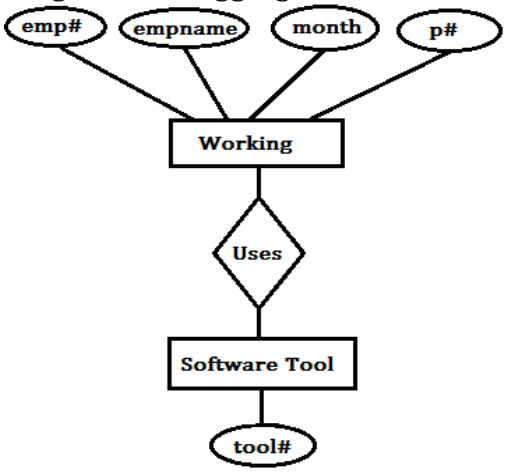

Figure : E - R diagram with aggregation

GATE QUESTIONS

Topics **Page No**

1 ER-MODEL

Q.1 It is desired to design an object-oriented employee record system for a company_ each employee has a name, unique Id and salary. Employees belong to different categories and their salary is determined by their category. The functions get Name, getld and compute salary are required. Given the class hierarchy below, possible locations for these functions are

1. getld is implemented in the super class
2. getld is implemented in the subclass
3. getld Name is an abstract function in the super class
4. getld Name is implemented in the super class
5. getld Name is implemented in the subclass
6. getld salary is an abstract function inthe super class
7. getld salary is implemented in the supercla.ss
8. getld salary is implemented in the subclass

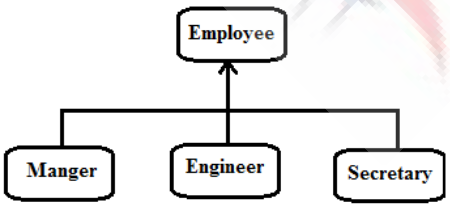

Choose the best design.

a) 1, 4, 6, 7 b) 1, 4, 7
c)1, 3, 5, 6, 8 d) 2,5,8

[GATE – 2004]

Q.2 Consider the following entity relationship diagram (ERD), where two entities E1 and E2 have a relation R of cardinality 1: m

The attributes of E1 are A11, A12 and A13 where A11 is the key attribute. The attributes of E2 are A21, A22 and A23 is the key attribute and A23 is a multi-valied attribute. Relation R does not have any attribute. A relational database containing minimum number of table with each table satisfying the requirements of the third normal form (3NF) is designed from the above ERD. The number of table in the database is

a) 2 b) 3
c) 5 d) 4

[GATE-2004]

Q.3 Consider the entities 'hotel room', and 'person' with a many to many relationship 'lodging' as shown below:

If we wish to store information about the rent payment to be made by person (s) occupying different hotel rooms, then this information should appear as an attribute of

a) Person b) Hotel Room
c) Lodging d) None of these

[GATE – 2005]

Q.4 Let E_1 and E_2 be two entities in an E/R diagram with simple single valued attributes. R_1 and R_2 are two relationships between E_l and E_2, where R_1 is one-to-many and R_2 is many-to-many. R_1, and R_2 do not have any attributes of their own. What is the minimum number of tables required to represent this situation in the relational model?

a) 2 b) 3
c) 4 d) 5

[GATE – 2005]

Q.5 The following table has two attributes A and C where A is the primary key

ANSWER KEY:

1	2	3	4	5	6	7	8	9	10
(a)	(b)	(c)	(b)	(c)	(b)	(a)	(c)	4	(c)

EXPLANATIONS

Q.1 (a)

ID is unique so; can be implemented in the super class.

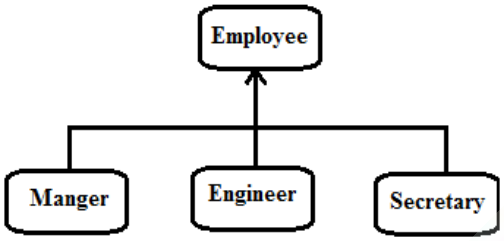

Name is a function of super class.
Also, Name can be common/ duplicate, so must be implemented in super class. Salary is a function of super class. Also, salary is dependent on the category and hence is implemented in the subclass.

Q.2 (b)

One table for E1, two table for E2(A21, A22 and A23). We need to make a separate table for multi-valued attribute to satisfy minimum 1NF condition.
Now, Relation table can be merged with (A21, A23). Tables are:
a) E1(A11, A12, A13)
b) E21(A11,A21,A22) and
c) E22(A21, A23).
Number of tables = 3

Q.3 (c)

It is many to many. Rent cannot be an attribute of room or person entities lone. If depending on number of persons sharing a room the rent for each person for the room will be different.
Otherwise rent can be attribute of room.
Hence attribute is lodging.

Q.4 (b)

Given that
R1 is one to many and R2 is many to many.
The one to many relationships are represented with entity set from one side.
This normally happens as each entity in the entity set can be associated with at most one entity of the other.
Therefore, the table is not formed for R1 the tables are hence, formed for R2, E1 and E2.
So, there are total of 3 tables.

Q.5 (c)

If (2, 4) is deleted then 2 is the primary key but in (5, 2) and (7, 2), 2 is the foreign key so these must be deleted. The primary key for (5, 2) and (7, 2) is 5 and 7 respectively but in (9, 5), 5 is the foreign key so it is also deleted.

Q.6 (b)

Minimum No of tables required = 3 = {M, P, N} Since Cardinality of R1= N:1→ It will never become a table.
And Cardinality of R2 = 1:N→ It will never become a table.

Q.7 (a)

ASSIGNMENT QUESTIONS

Q.1 Which normal form is considered adequate for relational database design?
a) 2 NF b) 3 NF
c) 4 NF d) BCNF

Q.2 The concept of locking can be used to solve the problem of
a) Lost update
b) uncommitted dependency
c) Inconsistent data
d) deadlock

Q.3 Given relations R (w, x) and S(x, y), the result of
SELECT DISTINCT w, x
FROM R, S
Is guaranteed to be the same as R, if
a) R has no duplicates and S is non-empty
b) R and S have no duplicates
c) S has no duplicates and R is non-empty
d) R and S have the same number of tuples

Q.4 A functional dependency of the form $X \rightarrow Y$ is trivial if
a) $Y \subseteq X$ b) $Y \subset X$
c) $X \subseteq Y$ d) $X \subset Y$ and $Y \subset X$

Q.5 If every non-key attribute is functionally dependent on the primary key, then the relation will be in
a) First normal form
b) Second normal form
c) Third normal form
d) Fourth normal form

Q.6 The column of a table is referred to as the
a) tuple b) attribute
c) Entity d) degree

The next four questions are based on the following details. Consider the given schemes.
Branch_scheme = (Branch_name , assets , Branch_city)
Customer_scheme = (Customer_name , street , Customer_city)
Deposit_scheme = (Branch_name ,account_number, Customer_name,Balance)
Branch_scheme = (Branch_name, loan_number , Customer_name, amount)
Client_scheme = (Customer_name , Banker , banker_name)

Q.7 Using relational algebra, the query that finds customers who have a balance of over 1000 is
a) $\pi_{\text{customer\_name}}(\sigma_{\text{balance}}>1000(\text{Deposit}))$
b) $\sigma_{\text{customer\_name}}(\pi_{\text{balance}} > 1000\ (\text{Deposit}))$
c) $\pi_{\text{customer\_name}}(\sigma_{\text{balance}}> 1000\ (\text{Borrow}))$
d) $\sigma_{\text{customer\_name}}(\pi_{\text{balance}}> 1000\ (\text{Borrow}))$

Q.8 Which of the following queries finds the clients of banker Agassi and the city they live in?
a) πclient_customer_name,customer_city (σ client.customer_name = customer.customer_name (σ banker_name = “Agassi” (Client × Customer)))
b) π customer_name.customer_city(σ banker_name = “Agassi” (Client × Customer))
c) π client.customer_name.customer_city (σ banker_name=”Agassi”(σ client.customer_name=customer.customer_name (Client × Customer)))
d) π customer_name.customer_city (σ banker_name=”Agassi” (Client × Customer))

Q.9 Which of the following tuple relational calculus finds all customers who have a loan amount of more than 1200?

EXPLANATIONS

Q.1 (b)
Adequate level of database design is 3NF.

Q.4 (a)
If Y $\subseteq$ X, then X→Y is trivial.

Q.5 (c)
If all non-key attributes depends on primary key, it is 3NF definition.

Q.7 (a)
First selection then projection must be applied.

Q.11 (b)
a, c can be get with transitive property**.**

Q.14 (c)
The join returns a record even if we have one matching.

Q.16 (c)
The same flight number flies on different dates. So the combination is key.

Q.17 (a)
ab combination is the key. Since c is depending on part of the key, it is not in 2NF.

Q.18 (a)
Maximum and Minimum would be only the students=(8,8) Since student table is parent and enroll is child.

Q.21 (c)
When parent is deleted, then child is deleted. The converse is not true.

Q.23 (d)
Subquery returns salary of the company and main query gets average salary of male employees department wise.

Q.24 (c)
Relational Algebra is procedural language, Relational Calculus is non-procedural language.

Q.25 (c)
Hierarchical model is based on graphs and Network model is based on trees.

Q.26 (d)
One student can enroll for more than one course. And one course is opted by more than one student.

Q.28 (a)
Alternate key is secondary key. It is a candidate key other than primary key.

Q.31 (a)
The problems due to redundancy are consistency and integrity.

Q.33 (a)
The primary key of one table becomes foreign key in another table.

Q.36 (a)
Check constraint is integrity constraint.

Q.39 (d)
Trigger automatically executes when condition associated with it matches when DB modification is done.

Q.42 (b)
Emp_no is primary key which can identify project_no.

Q.43 (a)
Since emp-name behaves like a key, all attributes are dependent on it.