

Database Systems

- Database Management System (DBMS) contains information about a particular enterprise (interrelated data), that is both convenient and efficient to use.

Ex: Universities (new students, instructors, ~~gpa~~ courses, assign GPA's to students, etc.)

- Disadvantages of file systems

- Multiple file formats, duplication of different files.
- Hard to add new constraints or change existing ones.
- Difficult to access data w/o writing new programs to carry out that task.
- Uncontrolled concurrent accesses (by multiple users) leading to inconsistencies.
- Security problems.

- Levels of Abstraction

- Physical level (describes how a record is stored)
 - Logical level (describes data stored in database/record and relationships among data)
 - View level (application programs hide details of datatypes, can also hide information for security purposes)
-
- Logical Schema (overall logical structure of database)
(analogous to type of information of variable)
 - Physical Schema (overall physical structure of database)
 - Instance (actual content of database at a particular point)

in time) (analogy is value of variable)

* Data Models

- A collection of tools for describing data, data relationships, data semantics, data constraints.

- Relational Model

- All the data is stored in various tables

Ex:-

<u>ID</u>	<u>name</u>	<u>dept_name</u>	<u>salary</u>
22222	Einstein	Physics	95000
12345	Wu	Finance	75000
32154	Gold	Music	120000
67891	Singh	History	87000
⋮	⋮	⋮	⋮

(INSTRUCTOR
TABLE)

<u>dept_name</u>	<u>building</u>	<u>budget</u>
Comp Sci.	Taylor	100000
Biology	Watson	90000
Music	Taylor	220000
Finance	Packard	57000
History	Packard	70000

(DEPT. TABLE)

- Data Definition Language (DDL) is a specific notation used for defining the database schema.

Ex:-

```
create table instructor (  
    ID char(5),  
    name varchar(20),  
    dept_name varchar(20),  
    salary numeric(8,2));
```

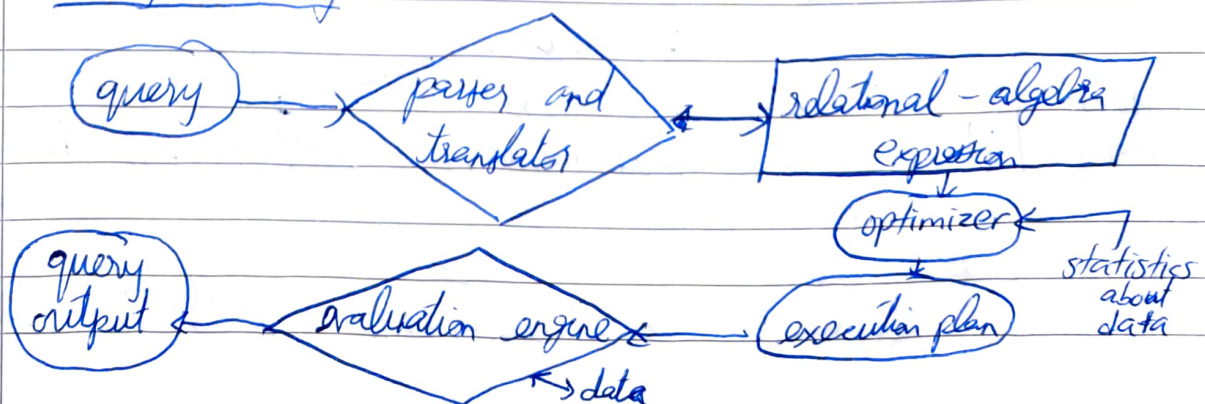
- DDL compiler ~~uses~~ generates a set of table templates stored in data dictionary which contains metadata about the database schema, primary key (ID uniquely identifies instructions) and authorization.
- Data Manipulation Language (DML) is the language for accessing and manipulating data organized by appropriate data model. (aka query language) (Ex: SQL)
- Extensible Markup Language (XML), originally intended as document markup language and not a database language, provides a way to specify new tags and create nested tag structures.

★ Database Engine

— Storage Management

- Program module that provides the interface between low-level data stored in database and application programs and queries submitted to the system.
- It is responsible for interaction with OS file managers and efficient storing, retrieving and updating of data.

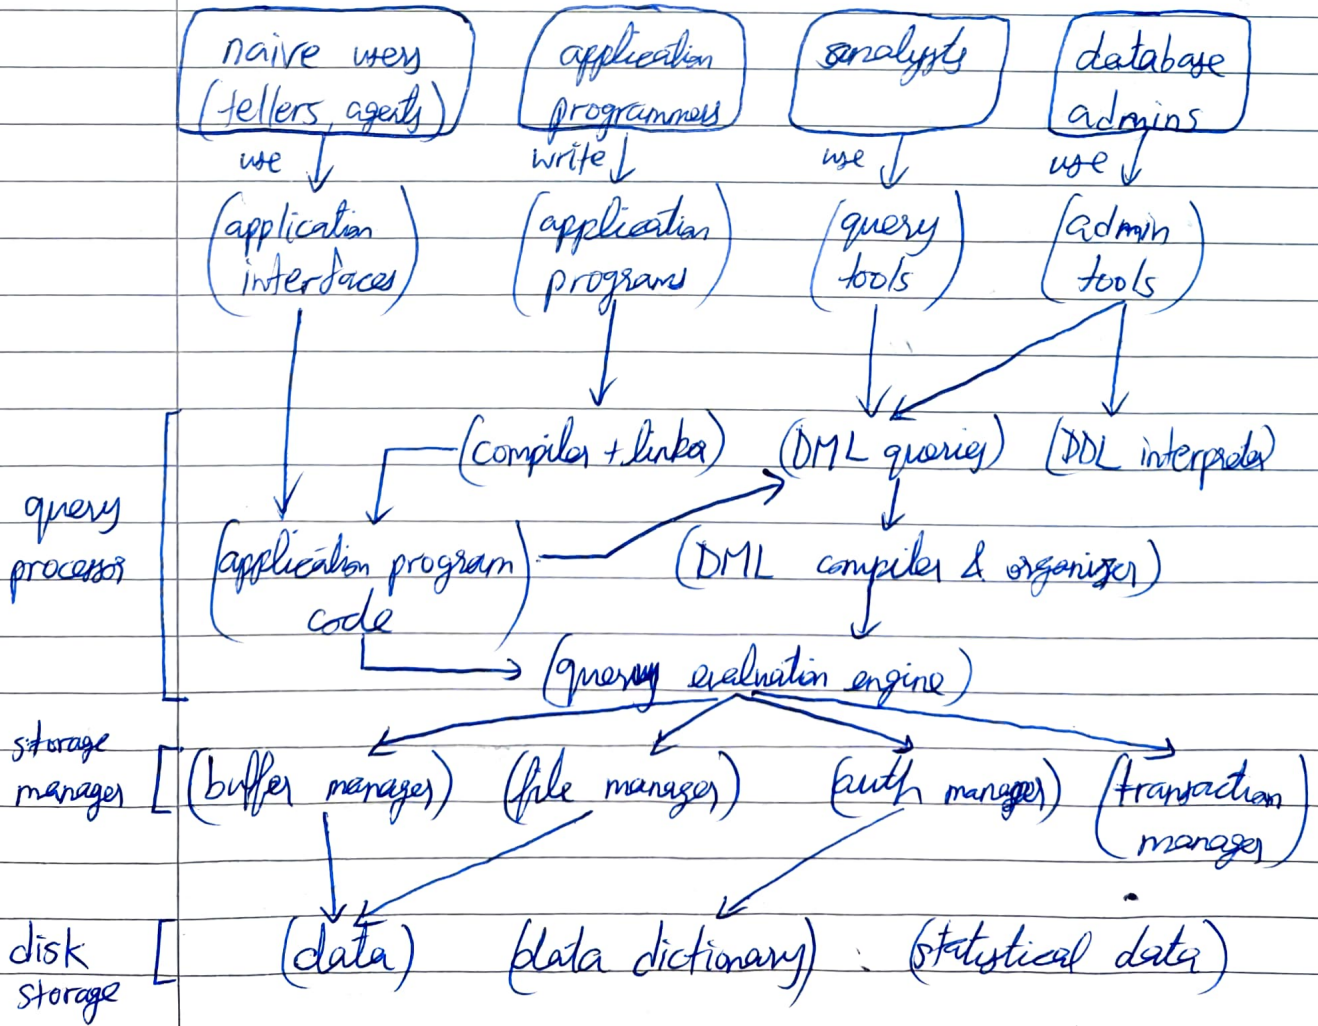
— Query Processing



Transaction Management

- Transaction is a collection of operations that perform a single logical function in a database application.
- Transaction-management component ensures that database remains correct/consistent state during system/power failures.
- Concurrency-control manager controls interaction among concurrent transactions ensuring consistency.

Note: Architecture of database systems is greatly influenced by underlying computer system on which database is running; centralized, client-server, parallel or distributed.



- Cartesian Product (X): Output pairs of rows from two input relations having same value on all attributes having same name.

Ex:- $instructors \times department$

- Union (U): Returns union of tuples from two input relations.

Ex:- $\Pi_{name}(instructors) \cup \Pi_{name}(student)$
(tuples in instructors + tuples in student)

- Set Difference (-): Returns set difference of tuples from two input relations.

Ex:- $\Pi_{name}(instructors) - \Pi_{name}(student)$
(tuples in instructors - tuples in student)

- Natural Join (\bowtie): Returns pairs of rows from two input relations where attributes are of same name and same value.

Ex:

	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>	<u>B</u>	<u>D</u>	<u>E</u>
	2	1	2	a	1	a	2
	3	2	3	a	3	a	3
(r)	3	4	3	b	1	a	3
	2	1	3	a	2	b	3
	3	2	3	b	3	b	3

$r \bowtie s$ \Rightarrow

	<u>A</u>	<u>B</u>	<u>C</u>	<u>D</u>	<u>E</u>
	2	1	2	a	2
	2	1	3	a	2
	2	1	3	a	3
	3	2	3	b	3

$\Pi_{A, r.B, C, r.D, E} (\sigma_{r.B=s.B \wedge r.D=s.D} (r \times s))$

- Intersection (\cap): Returns common tuples from two input relations.

Ex: $r \cap s = r - (r - s)$

- Renaming a Table: For a relation r , $\rho_s(r)$ would rename the table to s .

- x -

INTRODUCTION TO SQL

- Domain Types

- char(n) (Fixed length character string, with user-specified length n)
- varchar(n) (Variable length character string, with user-specified max length n)
- int (integer, machine dependent)
- numeric(p,d) (fixed point number, with precision p digits and d decimal digits)
- float(n) (precision of atleast n digits)

- Create Table Construct

```

create table instructs (
  ID char(5),
  name varchar(20) not null,
  dept-name varchar(20),
  salary numeric(8,2));
primary key (ID),
foreign key (dept-name) references department);

```

★ Updates to tables

- Insert: insert into instructor values ('10211', 'Smith', 'Biology', 66000)
- Delete: removes all tuples.
→ delete from student
- Drop: drop table r
- Alter: alter table r add gender varchar(1)
(for the new attribute, all tuples are assigned with null values.)

alter table r drop gender

★ Basic Query Structure (select... from... where...)

— select

- list the attributes desired in the result of query (more like projection operation (Π))

select name from instructor

- To remove duplicates
select distinct dept_name from instructor
Otherwise,
select all dept_name from instructor.

- To select all attributes / entire table

select * from instructor.

- select '437' → table with 1 column and 1 row with value 437.
- select 'A' from instructors → table with 1 column and n rows, each having value 'A'.
- You can rename columns using 'as' clause:

```
select ID, name, salary/12 as monthly_salary;
```

- where

- Specifies conditions that result must satisfy.

```
select name from instructors where
dept_name = 'Comp Sci' and salary > 80000
(or, not)
```

- from

- Lists the relations involved in the query.

→ select * from instructors, teachers (Cartesian product) (instructors X teachers)

Exs Find names of all instructors in Art department who have taught some course and course-id.

```
select name, course_id
from instructors, teachers
where instructors_id = teachers_id and instructors.dept_name = 'Art';
```

- String Operations

- To check for matching substrings, operator 'like' is used in the where clause.

- % character matches any substring
_ character matches any character.

Ex: select name
from instructors Sunday
where name like '%dar%';
(name that includes substring 'dar')

- Escape character (\)

- 'Intro%' → string beginning with 'Intro'
- '_____' → matches string of exactly 3 characters
- '_____%' → matches string of at least 3 characters

- Concatenation (||)

- Ordering tuples

select distinct name
from instructors
order by name desc/asc (descending/ascending)
(alphabetically)

- where (Contd.)

select name
from instructors
where salary between 90000 and 100000

select name, course_id
from instructors, teaches
where (instructors.ID, dept_name) = (teaches.ID, 'Biology');
(tuple comparison)

Note: Set operations such as union, intersect or except can be used in between ~~where~~ selected clauses.

- To return duplicates, use 'all'
- Suppose a tuple occurs m times in r and n times in s then
 - $(m+n)$ ^{times} in r union all s
 - $\min(m, n)$ ^{times} in r intersect all s
 - $\max(0, m-n)$ times in r except all s

• null signifies an unknown value that doesn't exist.

Ex: select name from instructors where salary is null.

- Three-valued logic:
 - OR : (null or true) = true
 - (null or false) = null
 - (null or null) = null
 - AND : (null and true) = null
 - (null and false) = false
 - (null and null) = null
 - NOT : (not null) = null

— Aggregate Functions

- ① avg : average value
- ② min : minimum value
- ③ max : maximum value
- ④ sum : sum of values
- ⑤ count : no. of values.

Ex: No. of tuples in 'course' relation.

select count (*) from course

⑥ group by : `select dept_name, avg(salary) as avg_salary
from instructors
group by dept_name;`

⑦ having : It is like where clause for 'group by' clause
`→ select dept_name, avg(salary)
from instructors
group by dept_name
having avg(salary) > 42000;`

Note: All aggregate operations except count(*) ignore tuples with null values on attributes.

★ Nested Subqueries :

• Subquery is a select-from-where expression nested within another query.

- Subqueries in 'where' clause

```
select distinct course_id  
from section  
where semester = 'Fall' and year = 2009 and  
course_id in (select course_id  
from section  
where semester = 'Spring' and year = 2010);  
(Courses offered in Fall 2009 and Spring 2010)
```

Ex: Find total no. of students who have taken course taught by instructor with ID 10101 (set membership)

```

select count (distinct ID)
from takes
where (course-id, sec-id, semester, year) in
      (select course-id, sec-id, semester, year
       from teaches
       where teaches.ID = 10101);

```

Ex: (For comparison) find names of instructors with salary greater than atleast one / all from Biology dept.

(Classic)

```

select distinct T.name
from instructors as T, instructors as S
where T.salary > S.salary and S.dept.name = 'Biology';

```

```

select name
from instructors
where salary > some/all (select name salary
                          (ANY) from instructors
                          where dept.name = 'Biology');

```

Ex: (Empty Relation)

```

select course-id
from section as S
where semester = 'Fall' and year = 2009 and
exists (select *
        from section as T
        where semester = 'Spring' and year = 2010 and
        S.course-id = T.course-id);

```

Note: 'unique' clause tests whether subquery has any duplicate tuples in it's result.

Subqueries in 'from' clause

Ex Find avg instructors salaries to those depts where average salary is greater than \$2000.

```
select dept_name, avg_salary
from (select dept_name, avg(salary) as avg_salary
      from instructors
      group by dept_name)
where avg_salary > 2000;
```

• 'with' clause provides a way to define a temporary relation. (table)

Ex

```
with max_budget
with dept_total (dept_name, value) as
  (select dept_name, sum(salary)
   from instructors
   group by dept_name),
dept_total_avg (value) as
  (select avg(value)
   from dept_total)
select dept_name
from dept_total, dept_total_avg
where dept_total.value > dept_total_avg.value;
```

Subqueries in 'select' clause

```
select dept_name,
  (select count(*)
   from instructors
   where department = dept_name)
as num_instructors
```

from department;

★ Modification of DB

- Deletion

delete from 'instructs' where dept-name = 'Finance';

- Insertion

insert into course (course-id, title, dept-name, credits)
values ('CS-437', 'DBMS', 'CompSci', 4);

- Updates

(i) update instructs
set salary = salary * 1.03
where salary > 100000;

update instructs
set salary = salary * 1.05
where salary <= 100000;

(ii) update instructs
set salary = case
when salary < 100000 then salary * 1.05
else salary * 1.03
end

* Intermediate SQL

* JOIN operations

• It is a Cartesian product which requires that tuples in both relations match (under some condition). Used as subquery in from clause.

- (i) Left Outer Join (all tuples in left relation + attributes) (if ~~does~~ not value on right, null)
- (ii) Right Outer Join (all tuples in right relation + attributes) (if no value on left, null)
- (iii) Full Outer Join (Left + Right Outer Join)
- (iv) Inner Join (strict on condition) (more like intersection)

* Views

• A virtual relation that hides certain data from view of certain users.

```

Ex: create view faculty as
      select ID, name, dept_name (w/o salary)
      from instructs
select name
from faculty
where dept_name = 'Biology'

```

- Expanded View

```

create view physics_fall_2009_naton as
(select course_id, room_number
from (select course.course_id, building, room_number

```

//_

```
from course, section
where course.course_id = section.course_id
and course.dept_name = 'Physics'
and section.semester = 'Fall'
and section.year = '2009'
where building = 'Watson';
```

- View relation V_1 is said to be directly dependent on view relation V_2 if V_2 is used in expression defining V_1 .
- View relation V_2 is said to be recursive if it depends on itself.
- Views can be updated using insert clause. However SQL prefers updates on simple views where (a) FROM clause has only one base table (b) ~~SELECT clause uses~~ Uses only attribute names from table in SELECT clause (c) missing attributes set to NULL (d) no group by/having.

Ex create view instructs_budget as
select instructs.ID, instructs.Name, instructs.dept_name,
department.Budget
from instructs JOIN department
ON instructs.dept_name = department.dept_name.

```
update instructs_budget set Budget = 700000 where
dept_name = 'History'; (Error)
```

(Since view is derived from multiple tables, database does not know whether to update one row/multiple rows)
(What if multiple instructs belong to History dept)

Excr

create view history_instructors as
select *
from instructors
where dept_name = 'History';

insert ~~into~~ into history_instructors values ('25566', 'Brown',
'Biology', 100000) (Error)
(It will give error, since it violates the condition that
only allows rows ~~that~~ where dept_name = History)
(Even though base table does not have any restriction)

Materialized View

- Database object that stores result of a query in a table as a physical table for faster access (saved)
- Since data can become outdated when original table changes, this view can be manually, periodically or automatically refreshed.

Excr

```
CREATE MATERIALIZED VIEW sales_summary AS  
SELECT region, SUM(sales) AS total_sales  
FROM sales GROUP BY region;
```

```
REFRESH MATERIALIZED VIEW sales_summary;
```

★ Transactions

- Unit of work in a database that ensures that a series of operations are fully completed or rolled back if something goes wrong (atomic transaction)

- In most databases, transactions start automatically.
- COMMIT (saves changes permanently)
- ROLLBACK (cancels transaction, undoing all changes)

★ Integrity Constraints

- Guard against accidental damage to database, ensuring that changes do not result in loss of data consistency.

Exi- Customers must have non-null phone number.

- (a) not null
- (b) unique (states that attributes form candidate key)
(permitted to be null unlike primary key)
- (c) primary key
- (d) check (P) where P is predicate
(check (semester in ('Fall', 'Winter', 'Spring', 'Summer')))
- (e) foreign key (referential integrity)

Excr (Violation)

```
CREATE TABLE person (
  ID char (10),
  name char (40),
  mother char (10),
  father char (10),
  primary key ID,
  foreign key father references person,
  foreign key mother references person);
```

To insert a tuple, either insert father and mother before inserting person, or set them to null initially.

_ / _ / _

★ Built-in Data types in SQL

- (a) date '2005-03-17'
- (b) time '9:00:30'
- (c) timestamp '2005-03-17 9:00:30.75'
- (d) interval '2 hours' (usually +/- with time/date)

★ User-defined Datatypes

```
CREATE TYPE Dollars AS NUMERIC(12,2) FINAL;  
(Use)  
CREATE TABLE DEPARTMENT (  
    . . . . .  
    budget Dollars);
```

- Domains (restricts allowed values) (even not null)

```
CREATE DOMAIN degree-level Varchar(10)  
CONSTRAINT degree-level-test  
check (value in ('Bachelors', 'Masters', 'Doctorate'));  
(Use)  
CREATE TABLE students (  
    . . . . .  
    degree degree-level);
```


Note \wedge (and), \vee (or), \neg (not) in propositional calculus.

Entity - Relationship (ER) Model

- Models an enterprise as a collection of entities & relationships
- Entity is a thing/object in the enterprise that is distinguishable from other objects. Described by set of attributes
- Relationship is an association among several entities. Represented by ER diagram.
- ER data model was developed to facilitate database design by allowing specification of enterprise schema, representing overall logical structure of database.
- Entity set is set of all entities of some type sharing same properties.

Ex course = (course_id, title, credits)

- Relationship set is a mathematical relation among $n \geq 2$ entities, each taken from entity sets.
- Binary relationship involves two entity sets (degree two)

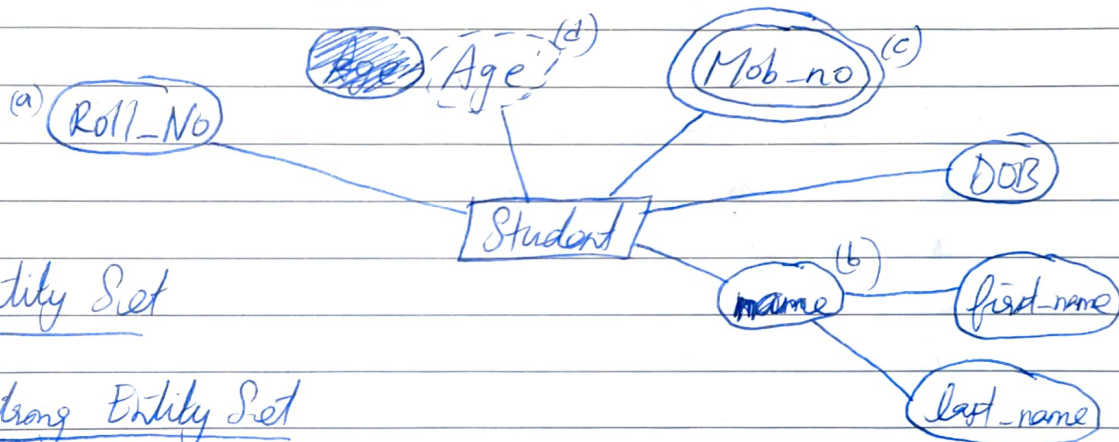


- For binary relationships set, mapping cardinality must be of:
 - one - to - one
 - one - to - many
 - many - to - one
 - many - to - many

★ Complex Attributes

- (a) Simple attribute (can't be classified further)
- (b) Composite attribute : name (first_name, middle_name)
- (c) Single-valued and multi-valued (multiple phone numbers)
- (d) Derived attribute (age from date-of-birth)


• Domain is set of permitted values for each attribute.



★ Entity Set

- Strong Entity Set

• Entity set that contains sufficient attributes to uniquely identify all its entities (primary key exists)




• Here two strong entities are Student and Course are related to each other. Here, primary keys are underlined.

• Double line between entities signifies total participation (every student must enrol in atleast one course)

• Single line means partial participation (there might be some courses for which no enrolments are made)

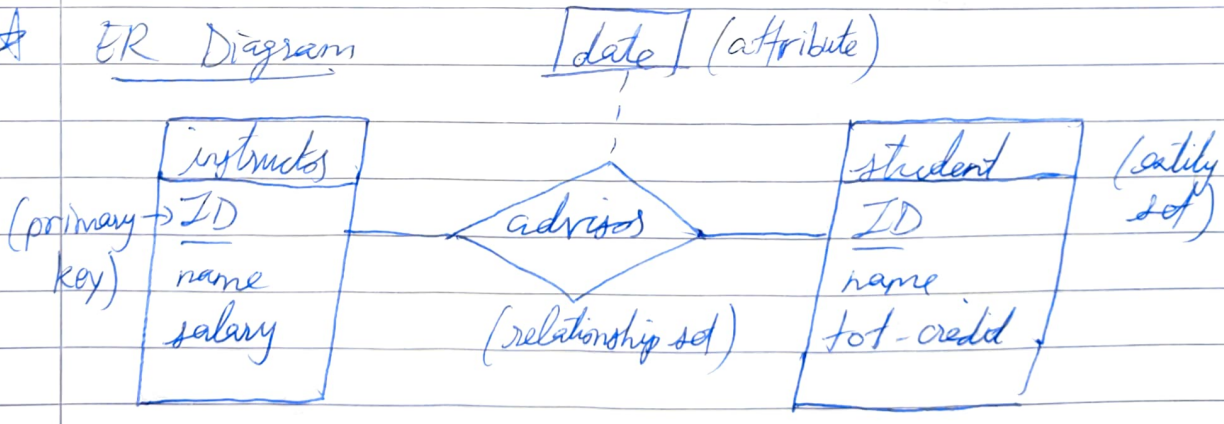
- Weak Entity Set

- Entity set does not contain sufficient attributes to uniquely identify entities. No primary key. However it contains a partial key called as discriminators that can identify a group of entities from entity set.

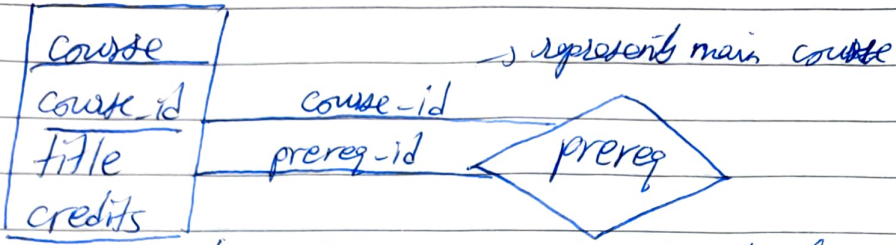


- Here ~~apartment~~ Apartment is weak entity set. Door-no is discriminator which cannot uniquely identify an apartment alone since there can be many other buildings with same door number.
- A weak entity set is one whose existence is dependent on another entity called identifying entity, and discriminators.

* ER Diagram



- Roles (function / part an entity plays in a relationship when same entity set appears more than once)



(both from same entity set but play diff roles)

- Cardinality Constraints (→ one, — many)

(i) One-to-one relationship



A student is associated with at most one instructor via advisor and vice versa

(ii) Many-to-one relationship



Instructor is associated with at most one student
Student is associated with several instructors

(iii) One-to-many relationship



Instructor is associated with several students
Student is associated with at most one student.

(iv) Many-to-many relationship.



Instructors is associated with several students and vice versa.

- Total participation implies that every entity in entity set participates in atleast one relationship.



Here min value of 1 → total participation
 max value of 1 → entity participates in atleast one relationship.
 max value of * → no limit.

Note Primary key for weak entity set = (primary key of related strong entity set + discriminators)

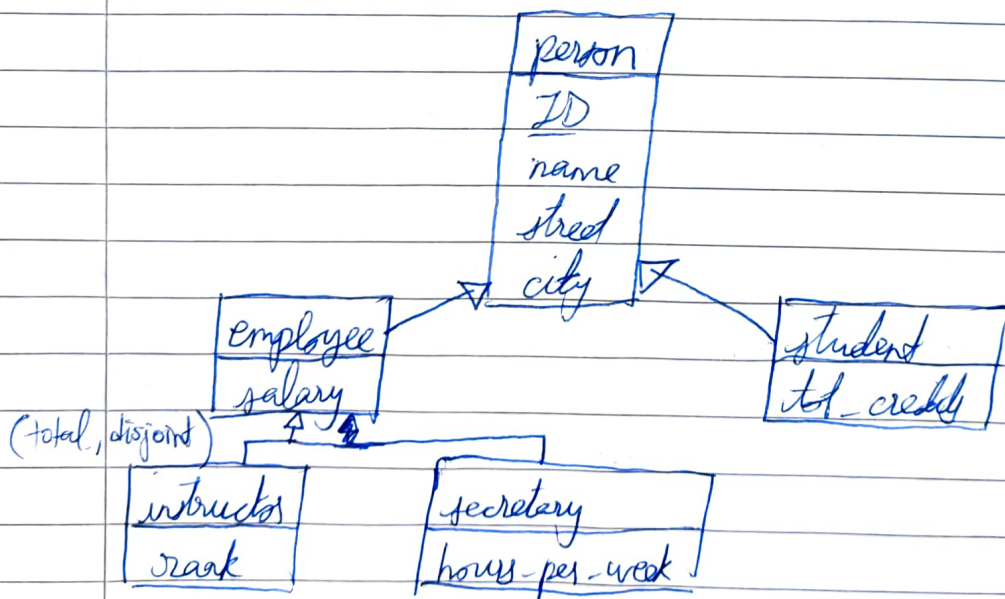
<ul style="list-style-type: none"> Example of entity with Complex Attributes: 	<p><u>instructor</u></p> <p><u>ID</u></p> <p>name</p> <p>first-name</p> <p>middle-initial</p> <p>last-name</p>
<ul style="list-style-type: none"> To represent strong entity set, student (<u>ID</u>, name, tot-cred) 	
<ul style="list-style-type: none"> To represent weak entity set, section (<u>course-id</u>, sec-id, sem, year) <p>↳ from identifying strong entity set</p>	<p>address</p> <p>street</p> <p>street-number</p> <p>street-name</p> <p>apt-number</p>
<ul style="list-style-type: none"> To represent relationship set (advisor) <p>advisor = (student-id, instructor-id)</p>	
<ul style="list-style-type: none"> For entity sets with composite attributes, either prefix is omitted (name-first-name) or all component attributes are listed out. (w/o multivalued attributes) 	<p>city</p> <p>state</p> <p>zip</p> <p>{ phone number }</p> <p>date-of-birth</p> <p>age ()</p>
<p>instructors (ID, first-name, middle-name, last-name, street-number, street-name, apt-number, city, state, zip-code, date-of-birth)</p>	

- For multivalued attributes, we use a separate schema having attributes corresponding to primary key of entity set and that of multivalued attribute.

instructor-phone = (ID, phone-number)
 = (22222, 987654321)
 = (22222, 123456789)

★ Specialization

- It is a Top-Down design approach in ER modelling. It involves dividing an entity set into subsets (low level entity sets) having distinct attributes or participate in unique relationships.
- Attribute inheritance means that low level entities inherit all attributes and relationships of high-level entity.



- Overlapping: A person can be both employee and student at the same time.
- Disjoint: An employee can either be instructor or secretary not both.

- To represent specialization via schemas,
- (i) For schema for each low entity set, include primary key of higher level entity set and local attributes.

person (ID, name, city, street) (Cons: Getting info about an employee requires accessing both higher & lower level schema)
 student (ID, st-cred)
 employee (ID, salary)

- (ii) Form a schema for each entity set with all local and inherited attributes

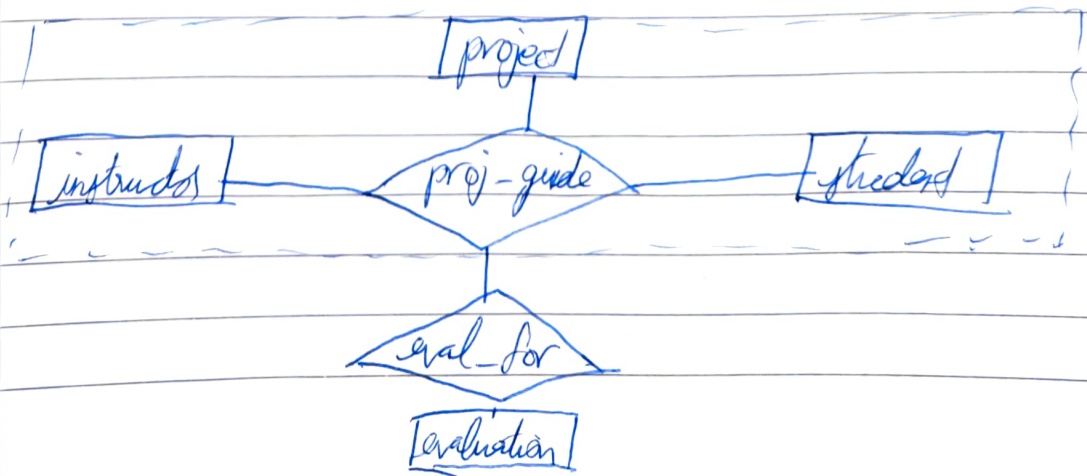
person (ID, name, street, city) (For people being both student and employee, storing name, street and city would be redundant)
 student (ID, name, street, city, st-cred)
 employee (ID, name, street, city, salary)

★ Generalization

- Opposite of Specialization, it is a Bottom-Up design process combining entity sets sharing same features into higher level entity set.

★ Aggregation

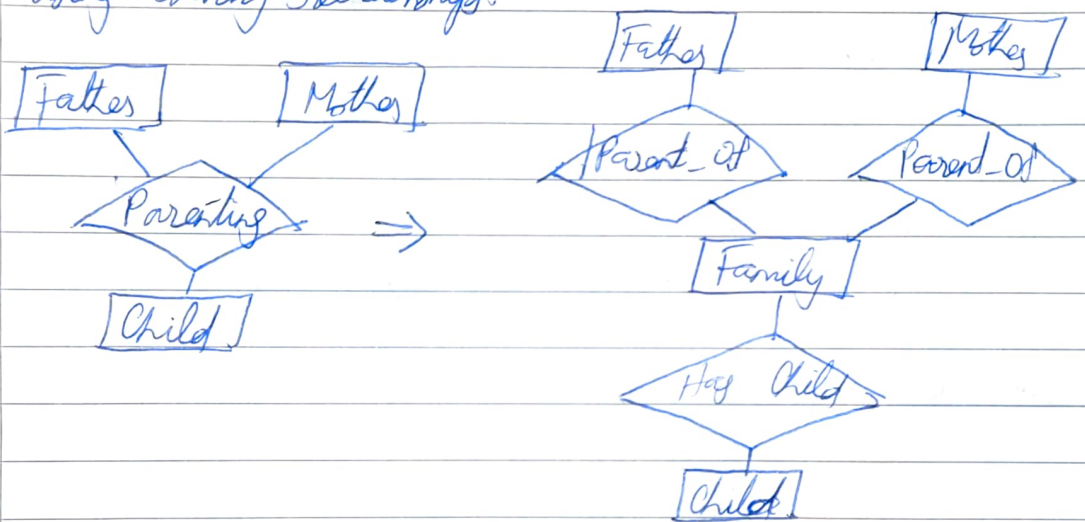
- Given the ternary relationship proj-guide, suppose we want to record evaluations of student by guide on a project



- Instead of connecting project, instructor and student to eval_for, causing redundancy, we treat the ternary relationship as an abstract entity, thus allowing relationships between relationships.

eval_for (s-ID, project-ID, i-ID, evaluation-ID)


Note Some non-binary relationships may be better represented using binary relationships.




RELATIONAL DATABASE DESIGN

- Functional dependency (FD) is a rule of the form $X \rightarrow Y$ where values of attribute set X can uniquely determine values of attribute set Y
- To avoid redundancy, ~~data~~ relations are decomposed into smaller relations / tables.

(i) Lossy Decomposition: Occurs when relation is divided into smaller sub-relations, and original relation cannot be reconstructed without introducing extra / incorrect tuples.



(ii) Lossless-Join Decomposition: Original relation is reconstructed



- First Normal Form

• Domain is atomic if elements are indivisible.

→ R is in 1NF if domains of all attributes of R are atomic.

Note: Set of names are composite attributes (non-atomic)
These complicate storage and lead to repeated/redundant storage of data.

Ex: Consider $r(A, B)$

	A	B
Here, $A \rightarrow B$ does not hold (A is not unique)	1	5
but $B \rightarrow A$ does hold.	3	7

• K is ~~super~~ superkey for relational schema R if $K \rightarrow R$.

• K is candidate key iff $K \rightarrow R$ and for no $\alpha \subset K$
 $\alpha \rightarrow R$.

• A functional dependency is trivial if $\beta \subseteq \alpha$
($\alpha \rightarrow \beta$)

- BCNF (Boyce-Codd Normal Form)

• $\alpha \rightarrow \beta$ is trivial ($\beta \subseteq \alpha$) (or)

• α is superkey for R

• If violated, we can decompose R into

$R_1 (\alpha \cup \beta)$

$R_2 (R - (\beta - \alpha))$

Ex: instructor_dept (ID, name, salary, dept_name, building, budget)

$F: \text{dept_name} \rightarrow \text{building, budget}$

Not in BCNF since dept_name is not superkey in R.

Decompose R into:

- $R_1 (R - (B - A)) = R_1 (\text{dept_name, building, budget})$
- $R_2 (R - (A - B)) = R_2 (\text{ID, name, salary, dept_name})$

Note: A decomposition is called dependency preserving, if we can check all functional dependencies by looking at individual tables, w/o joining them. (if all FD's are available in respective sub-relations)

Third Normal Form

- $A \rightarrow B$ is trivial ($B \subseteq A$) (OK)
- A is superkey for R (OK)
- Each attribute A in $(B - A)$ contained in candidate key for R.

Closure of Set of FD's (F^+) (using Armstrong's axioms)

Set of all FD's logically implied by F is closure of F .

- If $B \subseteq A$, then $A \rightarrow B$ (reflexivity)
- If $A \rightarrow B$, then $\gamma A \rightarrow \gamma B$ (augmentation)
- If $A \rightarrow B$ and $B \rightarrow \gamma$, then $A \rightarrow \gamma$ (transitivity)
- $A \rightarrow B, A \rightarrow \gamma$, then $A \rightarrow B\gamma$ (union)
- $A \rightarrow B\gamma \Rightarrow A \rightarrow B$ and $A \rightarrow \gamma$ (decomposition)
- $A \rightarrow B, \gamma B \rightarrow \delta$, then $A \rightarrow \delta$ (pseudo-transitivity)

Ex: $R = (A, B, C, G, H, I)$
 $F = \{A \rightarrow B, A \rightarrow C, CG \rightarrow H, CG \rightarrow I, B \rightarrow H\}$

Some members of F^+ :

$A \rightarrow H$ ($A \rightarrow B, B \rightarrow H$)

$AG \rightarrow I$ ($AG \rightarrow CG$ (augment), $CG \rightarrow I$)

$CG \rightarrow HI$ ($CG \rightarrow H, CG \rightarrow I$)

$(AG)^+ = \{AG\}$
 $= \{ABCG\}$ ($A \rightarrow C, A \rightarrow B$)
 $= \{ABCGHI\}$ ($CG \rightarrow H, CG \rightarrow I$)

AG is superkey

For AG to be candidate key, AG should be minimal

$(A)^+ = \{A\} = \{ABC\} = \{A, B, C, H\}$ \times

$(G)^+ = \{G\} = \{G\}$ \times

Since AG is minimal, it is candidate key.

Note From attribute closure, we can

- (i) Test for superkey (Compute α^+ , check if α^+ contains all attributes of R .)
- (ii) Test FD (To check if $\alpha \rightarrow \beta$ holds, check if $\beta \subseteq \alpha^+$. Compute α^+ and check if it contains β .)
- (iii) Compute F^+ ($\forall X \in R$, find X^+ and $\forall S \subseteq X^+$ obtain FD: $X \rightarrow S$.)

Canonical Cover

- Minimal set of FD's equivalent to F with no redundant dependencies nor redundant parts.

To check if any attribute in $\alpha \rightarrow \beta$ in F is extraneous

(i) If $A \in \alpha$ is extraneous in α .

Compute $(\alpha - A)^+$ using FD's and check if it contains β , then A is extraneous in α

(ii) If $A \in \beta$ is extraneous in β

Compute $(\alpha)^+$ using only FD's in $F' = (F - \{\alpha \rightarrow \beta\}) \cup \{\alpha \rightarrow (\beta - A)\}$

Check if $(\alpha)^+$ contains A , then A is extraneous in β .

Ex $F = \{A \rightarrow C, AB \rightarrow C\}$

B is extraneous since $AB - B = (A)^+$ does contain $C(B)$

$F = \{A \rightarrow C, AB \rightarrow CD\}$

C is extraneous in $AB \rightarrow CD$, since in $F' = (A \rightarrow C, AB \rightarrow D)$

$(AB)^+ = \{ABDC\}$, since C is there, it is extraneous.

Ex $R = (A, B, C)$

$F = \{A \rightarrow BC, B \rightarrow C, A \rightarrow B, AB \rightarrow C\}$

$= \{A \rightarrow BC, B \rightarrow C, AB \rightarrow C\}$ (Combine $A \rightarrow B$)

$= \{A \rightarrow BC, B \rightarrow C\}$

(A is extraneous in $AB \rightarrow C$)

$= \{A \rightarrow B, B \rightarrow C\}$

(C is extraneous, can be obtained by transitivity)

(Canonical Covers)

Ex (Lossless-Join Decomposition)

Decomposition of R into R_1 and R_2 is lossless if $R_1 \bowtie R_2 \rightarrow R$

$R_1 \cap R_2 \rightarrow R_1$

$R_1 \cap R_2 \rightarrow R_2$

$R(A, B, C)$

$F = \{A \rightarrow B, B \rightarrow C\}$

(i) $R_1 = (A, B), R_2 = (B, C)$ (no dependency lost)
 $R_1 \cap R_2 = \{B\}$ and $B \rightarrow BC$
 Since B is superkey for R_2 , lossless join ✓
 (Dependency preserved) ✓

(ii) $R_1 (A, B), R_2 (A, C)$
 $R_1 \cap R_2 = \{A\}$ and $A \rightarrow AB$
 Since A is superkey of R_1 , lossless join ✓
 But since $B \rightarrow C$ is not directly available, dependency
 is not preserved directly. ✗

Ex: (BCNF Decomposition)
 class (course-id, title, dept-name, credits, sec-id,
 semester, year, building, room-number,
 capacity, time-slot-id)
 FD: course-id \rightarrow title, dept-name, credits
 building, room-number \rightarrow capacity
 course-id, sec-id, semester, year \rightarrow building,
 room-number, time-slot-id
 Candidate key {course-id, sec-id, semester, year}

(i) course-id \rightarrow title, dept-name, credits
 But course-id is not a superkey in R.

Decompose class into:

course (course-id, title, dept-name, credits) \rightarrow in BCNF
 class1 (course-id, sec-id, semester, year, building,
 room-number, capacity, time-slot-id)

(ii) building, room-number \rightarrow capacity
 \rightarrow not superkey of class1.

Decompose into:

classroom (building, room-number, capacity)
 section (course-id, sec-id, semester, year, building,
 room-number, time-slot-id) \rightarrow in BCNF

//_

- Multivalued Dependency (MVD)

- One attribute A determining multiple independent values of another attribute B ($A \twoheadrightarrow B$)

Ex: An instructor can have multiple students.
If $\alpha \twoheadrightarrow \beta$ then $\alpha \rightarrow \beta$.

- Fourth Normal Form (4NF)

- $\alpha \twoheadrightarrow \beta$ is trivial ($\beta \subseteq \alpha, \alpha \cup \beta = R$) (or)
- α is superkey for schema R .

Ex: $R = (A, B, C, G, H, I)$

$$F = \{ A \twoheadrightarrow B, B \twoheadrightarrow HI, CG \twoheadrightarrow HI \}$$

R is not in 4NF since $A \twoheadrightarrow B$ and A is not superkey of R .

$$R_1 = (A, B) \quad \checkmark$$

$$R_2 = (A, C, G, H, I) \quad \times$$

$$R_3 = (C, G, H) \quad \checkmark$$

$$R_4 = (A, C, G, I) \quad \times$$

~~R_5~~ Since $A \twoheadrightarrow B, B \twoheadrightarrow HI, A \twoheadrightarrow HI, A \twoheadrightarrow I$

$$R_5 = (A, I) \quad \checkmark$$

$$R_6 = (A, C, G) \quad \checkmark \quad (\text{no non-trivial MVD left})$$