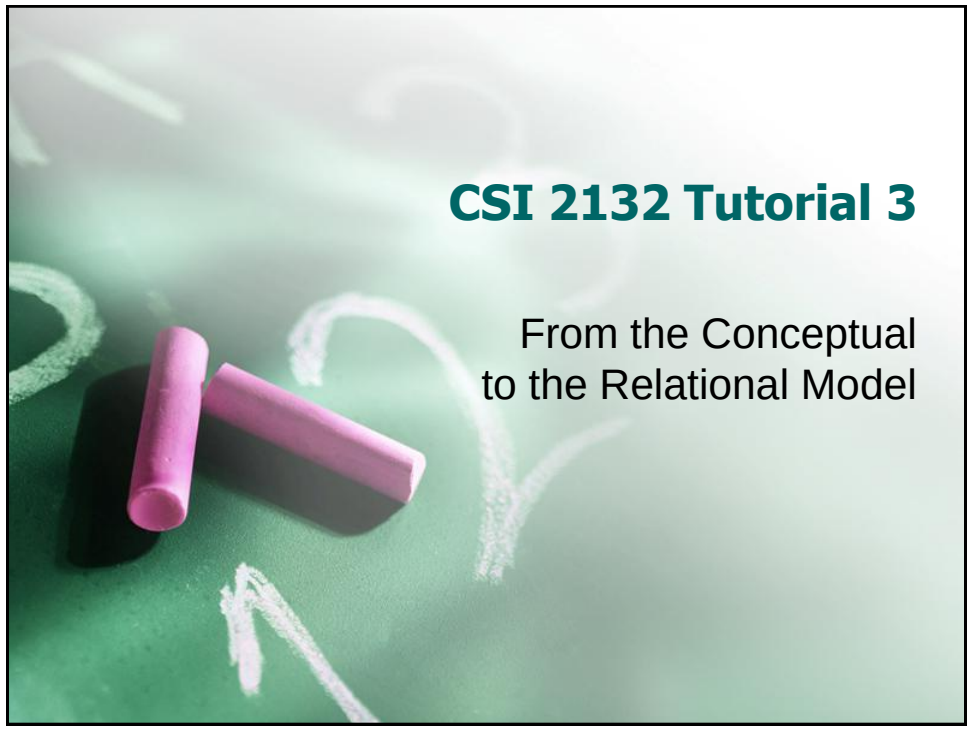

CSI 2132 Tutorial 3

From the Conceptual
to the Relational Model

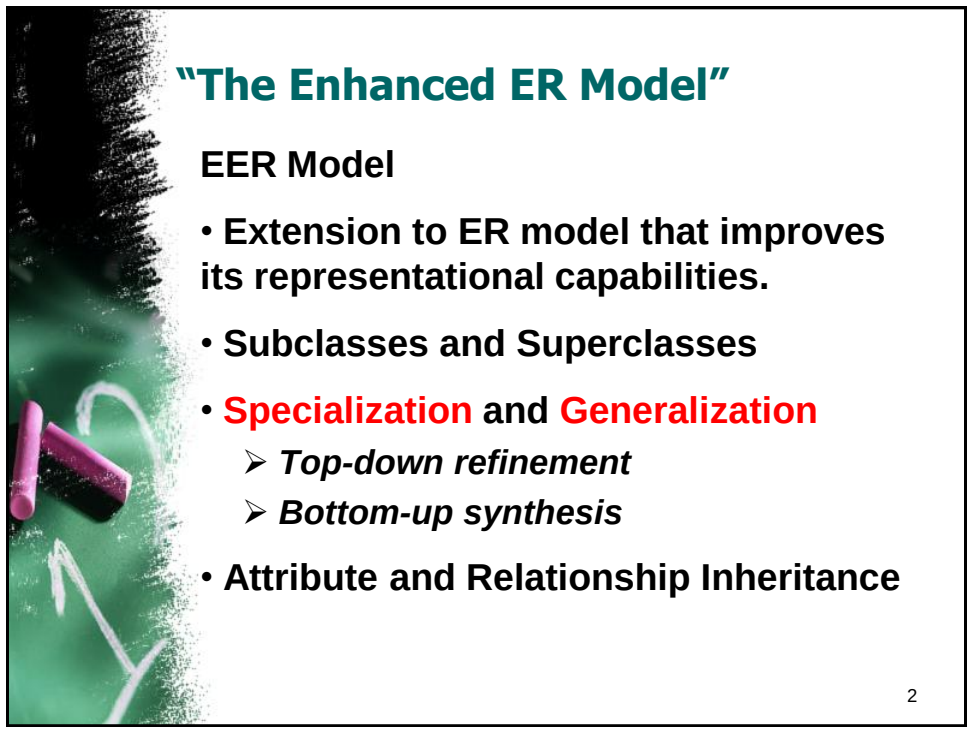

"The Enhanced ER Model"

EER Model

- **Extension to ER model that improves its representational capabilities.**
- **Subclasses and Superclasses**
- **Specialization and Generalization**
  - ***Top-down refinement***
  - ***Bottom-up synthesis***
- **Attribute and Relationship Inheritance**

2

"The Enhanced ER Model"

Exercise 1:

Consider following BANK ER schema:

3

"The Enhanced ER Model"

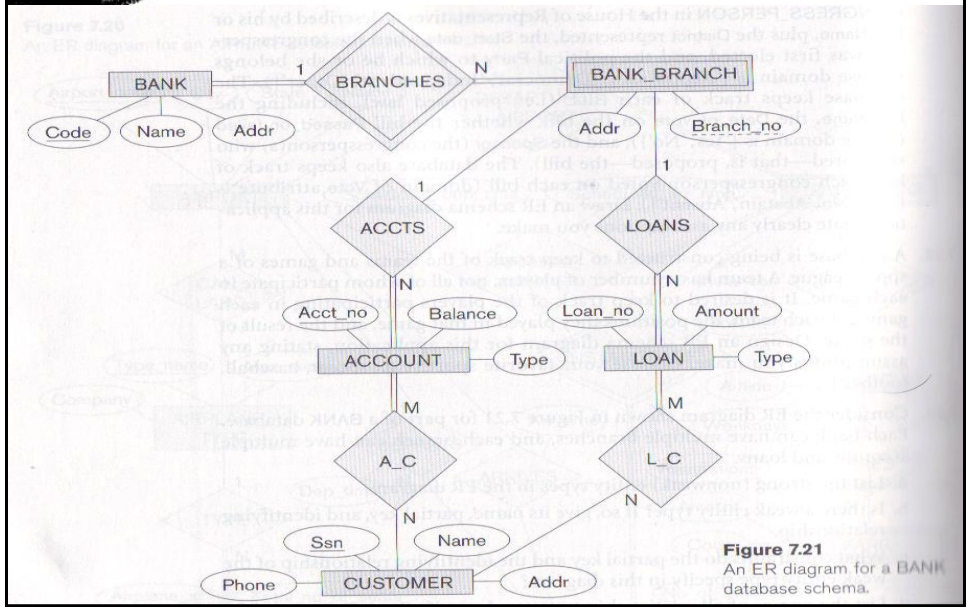

Figure 7.21
An ER diagram for a BANK database schema.

"The Enhanced ER Model"

Exercise 1:

Suppose that it is necessary to keep track of different types of accounts (SAVINGS, CHECKING, ...) and LOANS (CAR_LOANS, HOME_LOANS, ...).

5

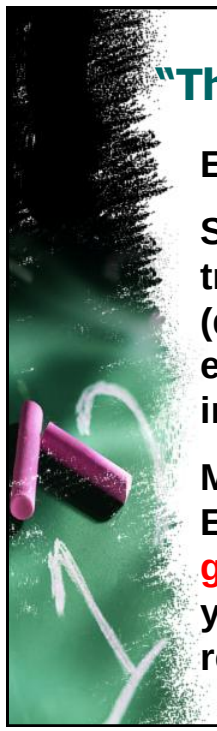

"The Enhanced ER Model"

Exercise 1:

Suppose that it is also desirable to keep track of each account's transactions (deposits, withdrawals, checks, ...) and each loan's payment; both of these include the amount, date, time, ...

Modify the BANK schema, using ER and EER concepts of specialization and generalization. State any assumptions you make about the additional requirements.

6

"The Relational Data Model and Relational Database Constraints"

Types of Integrity Constraints (Ensure Integrity of our data):

- **Entity Integrity (or key) constraint.**
  - ***No PK has a NULL value***
  - ***No instances with PK repeated***
    - **Maintains identifiable each tuple**
- **Referential integrity constraint.**
  - ***FK either NULL or matches a PK***
    - **Maintains consistency in data**
- **General Constraints.**
  - ***Functional Dependency constraints***
  - ***State constraints***
  - ***Transition Constraints***

7

"The Relational Data Model and Relational Database Constraints"

Exercise 2:

Suppose each of the following Update operations is applied directly to the following database. Discuss *all* integrity constraints violated by each operation, if any, and the different ways of enforcing these constraints.

8

"The Relational Data Model and Relational Database Constraints"

EMPLOYEE

Fname	Minit	Lname	Ssn	Bdate	Address	Sex	Salary	Super_ssn	Dno
John	B	Smith	123456789	1965-01-09	731 Fondren, Houston, TX	M	30000	333445555	5
Franklin	T	Wong	333445555	1955-12-08	638 Voss, Houston, TX	M	40000	888665555	5
Alicia	J	Zelaya	999887777	1968-01-19	3321 Castle, Spring, TX	F	25000	987654321	4
Jennifer	S	Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
Ramesh	K	Narayan	666884444	1962-09-15	975 Fire Oak, Humble, TX	M	38000	333445555	5
Joyce	A	English	453453453	1972-07-31	5631 Rice, Houston, TX	F	25000	333445555	5
Ahmad	V	Jabbar	987987987	1969-03-29	980 Dallas, Houston, TX	M	25000	987654321	4
James	E	Borg	888665555	1937-11-10	450 Stone, Houston, TX	M	55000	NULL	1

DEPARTMENT

Dname	Dnumber	Mgr_ssn	Mgr_start_date
Research	5	333445555	1988-05-22
Administration	4	987654321	1995-01-01
Headquarters	1	888665555	1981-06-19

DEPT_LOCATIONS

Dnumber	Dlocation
1	Houston
4	Stafford
5	Bellaire
5	Sugarland
5	Houston

WORKS_ON

Essn	Pno	Hours
123456789	1	32.5
123456789	2	7.5
666884444	3	40.0
453453453	1	20.0
453453453	2	20.0
333445555	2	10.0
333445555	3	10.0
333445555	10	10.0
333445555	20	10.0
999887777	30	30.0
999887777	10	10.0
987987987	10	35.0
987987987	30	5.0
987654321	30	20.0
987654321	20	15.0
888665555	20	NULL

PROJECT

Pname	Pnumber	Plocation	Dnum
ProductX	1	Bellaire	5
ProductY	2	Sugarland	5
ProductZ	3	Houston	5
Computerization	10	Stafford	4
Reorganization	20	Houston	1
Newbenefits	30	Stafford	4

DEPENDENT

Essn	Dependent_name	Sex	Bdate	Relationship
333445555	Alice	F	1986-04-05	Daughter
333445555	Theodore	M	1983-10-25	Son
333445555	Joy	F	1958-05-03	Spouse
987654321	Abner	M	1942-02-28	Spouse
123456789	Michael	M	1988-01-04	Son
123456789	Alice	F	1988-12-30	Daughter
123456789	Elizabeth	F	1967-05-05	Spouse

9

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(A) .Insert:

< 'Robert', 'F', 'Scott', '943775543', '21-JUN-42', '2365 Newcastle Rd, Bellaire, TX', M, 58000, '888665555', 1 >

into EMPLOYEE.

10

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(B) .Insert

< 'ProductA', 4, 'Bellaire', 2 >

into PROJECT

11

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(C) .Insert

< 'Production', 4,
'943775543', '01-OCT-88' >

into DEPARTMENT

12

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(D) .Insert

< '677678989', null, '40.0' >

into WORKS_ON

13

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(E) .Insert

< '453453453', 'John', M,

'12-DEC-60', 'SPOUSE' >

into DEPENDENT

14

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(F) .Delete the WORKS_ON tuples with

ESSN= '333445555'

15

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(G) .Delete the EMPLOYEE tuple with

SSN= '987654321'

16

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(H) .Delete the PROJECT tuple with

PNAME= 'ProductX'

17

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(I) .Modify the MGRSSN and MGRSTARTDATE of the DEPARTMENT tuple with DNUMBER=5 to '123456789' and '01-OCT-88', respectively

18

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(J) .Modify the SUPERSSN attribute of the EMPLOYEE tuple with SSN= '999887777' to '943775543'

19

"The Relational Data Model and Relational Database Constraints"

Exercise 2

(K) .Modify the HOURS attribute of the WORKS_ON tuple with ESSN= '999887777' and PNO= 10 to '5.0'

20

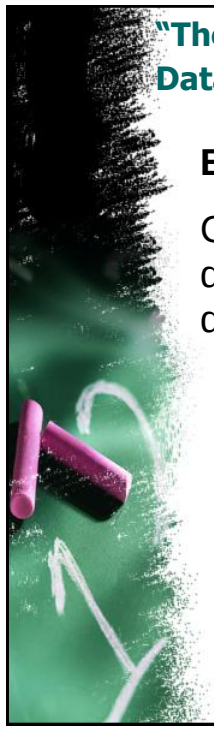

"The Relational Data Model and Relational Database Constraints"

Exercise 3

Consider the following AIRLINE relational database schema. Which describes a database for airline flight information.

21

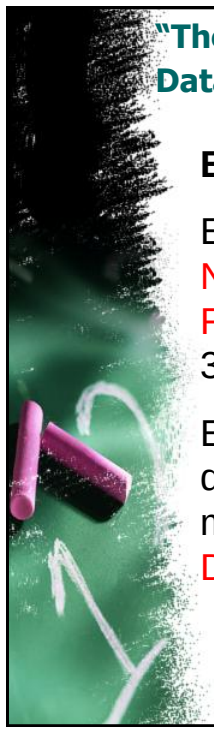

"The Relational Data Model and Relational Database Constraints"

Exercise 3

Each FLIGHT is identified by a flight NUMBER and consists of one or more FLIGHT_LEGs with LEG_NUMBERs 1, 2, 3, etc.

Each leg has scheduled arrival and departure times and airports, and has many LEG_INSTANCEs--one for each DATE on which the flight travels.

22

"The Relational Data Model and Relational Database Constraints"

Exercise 3

FARES are kept for each flight.

For each leg instance, SEAT_RESERVATIONs are kept, as is the AIRPLANE used in the leg, and the actual arrival and departure times and airports.

An AIRPLANE is identified by an AIRPLANE_ID, and is of a particular AIRPLANE_TYPE.

23

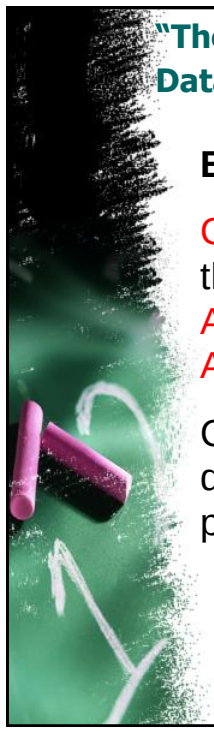

"The Relational Data Model and Relational Database Constraints"

Exercise 3

CAN_LAND relates AIRPLANE_TYPEs to the AIRPORTs in which they can land. An AIRPORT is identified by an AIRPORT_CODE.

Consider an update for the AIRLINE database to enter a reservation on a particular flight or flight leg on a given date.

24

"The Relational Data Model and Relational Database Constraints"

AIRPORT

Airport_code	Name	City	State

FLIGHT

Flight_number	Airline	Weekdays

FLIGHT_LEG

Flight_number	Leg_number	Departure_airport_code	Scheduled_departure_time
		Arrival_airport_code	Scheduled_arrival_time

LEG_INSTANCE

Flight_number	Leg_number	Date	Number_of_available_seats	Airplane_id
	Departure_airport_code	Departure_time	Arrival_airport_code	Arrival_time

FARE

Flight_number	Fare_code	Amount	Restrictions

AIRPLANE_TYPE

Airplane_type_name	Max_seats	Company

CAN_LAND

Airplane_type_name	Airport_code

AIRPLANE

Airplane_id	Total_number_of_seats	Airplane_type

SEAT_RESERVATION

Flight_number	Leg_number	Date	Seat_number	Customer_name	Customer_phone

Exercise 3

(A) Give the operations for this update.

Enter a reservation on a particular flight on a given date

25

"The Relational Data Model and Relational Database Constraints"

AIRPORT

Airport_code	Name	City	State

FLIGHT

Flight_number	Airline	Weekdays

FLIGHT_LEG

Flight_number	Leg_number	Departure_airport_code	Scheduled_departure_time
		Arrival_airport_code	Scheduled_arrival_time

LEG_INSTANCE

Flight_number	Leg_number	Date	Number_of_available_seats	Airplane_id
	Departure_airport_code	Departure_time	Arrival_airport_code	Arrival_time

FARE

Flight_number	Fare_code	Amount	Restrictions

AIRPLANE_TYPE

Airplane_type_name	Max_seats	Company

CAN_LAND

Airplane_type_name	Airport_code

AIRPLANE

Airplane_id	Total_number_of_seats	Airplane_type

SEAT_RESERVATION

Flight_number	Leg_number	Date	Seat_number	Customer_name	Customer_phone

Exercise 3

(B) What types of constraints would you expect to check?.

26

"The Relational Data Model and Relational Database Constraints"

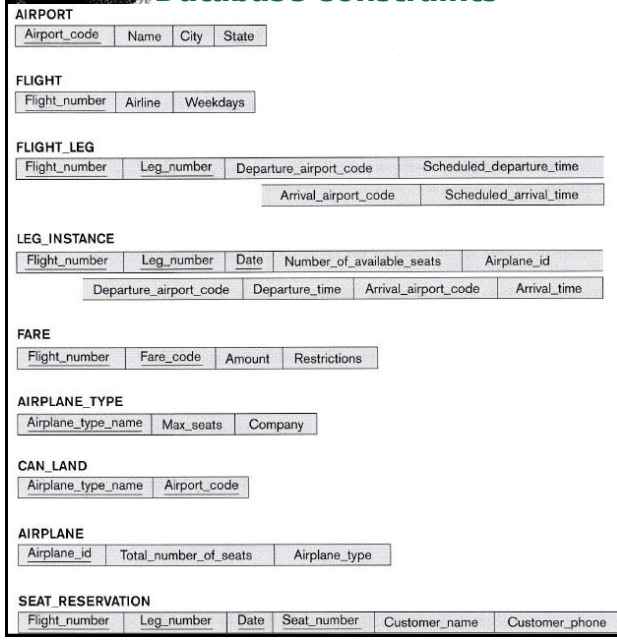

Exercise 3

(C) Which of these constraints are key, entity integrity, and referential integrity constraints and which are not?.

27

"The Relational Data Model and Relational Database Constraints"

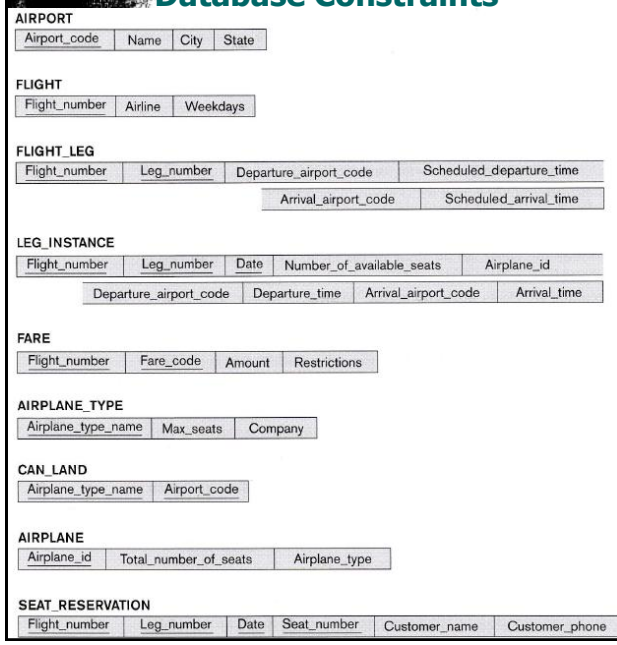

Exercise 3

(D) Specify all the referential integrity constraints.

28

"The Relational Data Model and Relational Database Constraints"

Exercise 4

- Consider the following relations for a database that keeps track of business trips of salespersons in a sales office:
  - *SALESPERSON (SSN, Name, Start_Year, Dept_No)*
  - *TRIP (SSN, From_City, To_City, Departure_Date, Return_Date, Trip_ID)*
  - *EXPENSE (Trip_ID, Account#, Amount)*
- Specify the foreign keys for this schema, stating any assumptions you make.

29

"The Relational Data Model and Relational Database Constraints"

Exercise 5

Database design often involves decisions about the storage of attributes.

For example a Social Security Number can be stored as a one attribute or split into three attributes(one for each of the three hyphen-deliniated groups of numbers in a Social Security Number—XXX-XX-XXXX)

(Area code – group number – serial number).

However, Social Security Number is usually stored in one attribute. The decision is usually based on how the database will be used. This exercise asks you to think about specific situations where dividing the SSN is useful.

30

"The Relational Data Model and Relational Database Constraints"

Exercise 6

Recent changes in privacy laws have disallowed organizations from using SSN to identify individuals unless certain restrictions are satisfied.

As a result, most US universities cannot use SSNs as primary keys (except for financial data).

In practice, StudentID, a unique ID identifier, assigned to every student, is likely to be used as the primary key rather than SSN

Since StudentID is usable across all aspects of the system.

31

"The Relational Data Model and Relational Database Constraints"

Exercise 6 (A)

Some database designers are reluctant to use generated keys (*surrogate* keys) for theStudentID because they are artificial.

Can you propose any natural choices of keys that can be used to store the student record in a UNIVERSITY database?.

32

"The Relational Data Model and Relational Database Constraints"

Exercise 6 (B)

Suppose that you were able to guarantee uniqueness of a natural key that included last name.

Are you guaranteed that the last name will not change during the lifetime of the database?

If the last name can change, what solutions can you propose for creating a primary key that still includes last name but remains unique?

33

"The Relational Data Model and Relational Database Constraints"

Exercise 6 (C)

What are the advantages and disadvantages of using generated (surrogate) keys?

34