X SEMINAR ON SOFTWARE FOR NUCLEAR, SUBNUCLEAR AND APPLIED PHYSICS 2 – 8 JUNE 2013, PORTO CONTE, ALGHERO

### Databases in experimental Physics: theory and practice L. Tomassetti University of Ferrara and INFN

## Notice!

# I'm teaching the course
"Databases and Laboratory"
to Computer Science Bachelor's students

and "Advanced Databases" to Computer Science Master's students

## Summary

#### \* Introduction

\* Why Databases are important for Physicists

\* A 1-semester database course in half-an-hour

Some applications/use cases/examples:

# Good design (?)

\* Problems arising from database design issues

- \* Database System are widely used in several fields
- \* We are interested in how they can make our work easier:
  - \* Administrative tasks
  - \* Organizational tasks
  - \* Data store and management
  - Metadata store and management
  - **\* Backend** for services in use by experiments

Please, let me ask you a couple of questions:

- How many of you are working in HEP experiments?
- \* How many of you have ever used the Grid Infrastructure (directly)?
- # ... ever used the Grid Infrastructure through a 'custom' service?
- \* ... ever used (a query, for example) a database management system (directly)?
- … ever used a Wiki, Tracking software, …?
- No, I'll not ask you:
  - # ... ever used Amazon/eBay/AppStores/...?
  - # ... ever used Facebook/Twitter/...?

# Just a few examples to start with:

- \* Publications databases (WoS, Scopus, INFN, ...)
- \* Organizational databases (experiments, ...)
- \* Bookkeeping databases
- \* Conditions databases
- # Grid access services (Dirac, ...)

\* As 'simple' users we may consider services as black-boxes and ignore the internals

\* As soon as we need something more (in Physics applications of course) it can be useful to know some details...

\* As soon our supervisors ask for writing analysis/ core software/utility code we have to know (almost) all details...

\* Database are relatively (maybe too) simple, we just need to understand a few concepts and apply them correctly

\* We do not need to be database experts (unless we have to develop core software tightly connected to dbms)

**\*** I'll limit the discussion to Relational Databases

#### **DATABASE IN THEORY**

## Content

- \* Introduction & Conceptual Modeling
- \* Database System: Concepts and Architecture
- \* Entity-Relationship Model
- Relational Model
- \* E-R to Relational Model Mapping
- \* Relational Algebra
- **\*** SQL

You can decide: (a) my speed (b) which parts we should discuss

l would suggest at least: - Relational Model + SQL - E-R + Algebra

# Introduction & Conceptual Modeling

## **Basic Definitions**

- **\* Database:** A collection of related data.
- **\* Data:** Known facts that can be recorded and have an implicit meaning.
- \* Mini-world: Some part of the real world about which data is stored in a database. For example, student grades and transcripts at a university.
- \* Database Management System (DBMS): A software package/ system to facilitate the creation and maintenance of a computerized database.
- \* Database System: The DBMS software together with the data itself. Sometimes, the applications are also included.

## Typical DBMS Functionality

- \* Define a database: in terms of data types, structures and constraints
- \* Construct or Load the Database on a secondary storage medium
- \* Manipulating the database: querying, generating reports, insertions, deletions and modifications to its content
- \* Concurrent Processing and Sharing by a set of users and programs – yet, keeping all data valid and consistent

Typical DBMS Functionality

**\*** Other features:

\* Protection or Security measures to prevent unauthorized access

\* "Active" processing to take internal actions on data

**\*** Presentation and Visualization of data

# Main Characteristics of the Database Approach

\* Self-describing nature of a database system: A DBMS catalog stores the description of the database (the description is called meta-data). This allows the DBMS software to work with different databases.

\* Insulation between programs and data: Called program-data independence. Allows changing data storage structures and operations without having to change the DBMS access programs.

# Main Characteristics of the Database Approach

- \* Data Abstraction: A data model is used to hide storage details and present the users with a conceptual view of the database.
- \* Support of multiple views of the data: Each user may see a different view of the database, which describes only the data of interest to that user.

# Main Characteristics of the Database Approach

\* Sharing of data and multiuser transaction processing : allowing a set of concurrent users to retrieve and to update the database. Concurrency control within the DBMS guarantees that each transaction is correctly executed or completely aborted. OLTP (Online Transaction Processing) is a major part of database applications. Advantages of Using the Database Approach

- \* Controlling redundancy in data storage and in development and maintenence efforts.
- \* Sharing of data among multiple users.
- \* Restricting unauthorized access to data.
- \* Providing persistent storage for program Objects
- \* Providing Storage Structures for efficient Query Processing

Advantages of Using the Database Approach

\* Providing backup and recovery services.

\* Providing multiple interfaces to different classes of users.

\* Representing complex relationships among data.

\* Enforcing integrity constraints on the database.

\* Drawing Inferences and Actions using rules

Additional Implications of Using the Database Approach

\* Potential for enforcing standards: this is very crucial for the success of database applications in large organizations Standards refer to data item names, display formats, screens, report structures, meta-data (description of data) etc.

\* Reduced application development time: incremental time to add each new application is reduced.

### Additional Implications of Using the Database Approach

- \* Flexibility to change data structures: database structure may evolve as new requirements are defined.
- \* Availability of up-to-date information very important for on-line transaction systems such as airline, hotel, car reservations.
- \* Economies of scale: by consolidating data and applications across departments wasteful overlap of resources and personnel can be avoided.

## When not to use a DBMS

#### **\*** When no DBMS may suffice:

If the database system is not able to handle the complexity of data because of modeling limitations

\* If the database users need special operations not supported by the DBMS.

## Database System: Concepts and Architecture

## Data Models

\* Data Model: A set of concepts to describe the structure of a database, and certain constraints that the database should obey.

\* Data Model Operations: Operations for specifying database retrievals and updates by referring to the concepts of the data model. Operations on the data model may include basic operations and user-defined operations.

## Categories of data models

- \* Conceptual (*high-level*, semantic) data models: Provide concepts that are close to the way many users perceive data. (Also called entity-based or object-based data models.)
- \* Physical (low-level, internal) data models: Provide concepts that describe details of how data is stored in the computer.
- \* Implementation (representational) data models: Provide concepts that fall between the above two, balancing user views with some computer storage details.

## History of Data Models

- Relational Model: proposed in 1970 by E.F. Codd (IBM), first commercial system in 1981-82. Now in several commercial products (DB2, ORACLE, SQL Server, SYBASE, INFORMIX) and several opensource products (MySQL, PostgreSQL, ...).
- \* Network Model: the first one to be implemented by Honeywell in 1964-65 (IDS System). Adopted heavily due to the support by CODASYL (CODASYL - DBTG report of 1971). Later implemented in a large variety of systems - IDMS (Cullinet - now CA), DMS 1100 (Unisys), IMAGE (H.P.), VAX -DBMS (Digital Equipment Corp.).
- Hierarchical Data Model: implemented in a joint effort by IBM and North American Rockwell around 1965. Resulted in the IMS family of systems. The most popular model. Other system based on this model: System 2k (SAS inc.)

## History of Data Models

- \* Object-oriented Data Model(s): several models have been proposed for implementing in a database system. One set comprises models of persistent O-O Programming Languages such as C++ (e.g., in OBJECTSTORE or VERSANT), and Smalltalk (e.g., in GEMSTONE). Additionally, systems like O2, ORION (at MCC - then ITASCA), IRIS (at H.P.- used in Open OODB).
- \* Object-Relational Models: Most Recent Trend. Started with Informix Universal Server. Exemplified in the latest versions of Oracle-10i, DB2, and SQL Server etc. systems.

### Schemas versus Instances

- \* Database Schema: The description of a database. Includes descriptions of the database structure and the constraints that should hold on the database.
- **Schema Diagram**: A diagrammatic display of (some aspects of) a database schema.
- **\* Schema Construct**: A component of the schema or an object within the schema, e.g., STUDENT, COURSE.
- \* Database Instance: The actual data stored in a database at a particular moment in time. Also called database state (or occurrence).

## Database Schema Vs. Database State

- \* Database State: Refers to the content of a database at a moment in time.
- \* Initial Database State: Refers to the database when it is loaded
- \* Valid State: A state that satisfies the structure and constraints of the database.
- \* Distinction:
  - \* The database schema changes very infrequently. The database state changes every time the database is updated.
  - \* Schema is also called intension, whereas state is called extension.

Three-Schema Architecture

\* Proposed to support DBMS characteristics of:

\* Program-data independence.

\* Support of multiple views of the data.

## Three-Schema Architecture

- \* Defines DBMS schemas at three levels:
  - Internal schema at the internal level to describe physical storage structures and access paths. Typically uses a physical data model.
  - \* Conceptual schema at the conceptual level to describe the structure and constraints for the whole database for a community of users. Uses a conceptual or an implementation data model.
  - \* External schemas at the external level to describe the various user views. Usually uses the same data model as the conceptual level.

Three-Schema Architecture

\* Mappings among schema levels are needed to transform requests and data. Programs refer to an external schema, and are mapped by the DBMS to the internal schema for execution.

## Data Independence

\* Logical Data Independence: The capacity to change the conceptual schema without having to change the external schemas and their application programs.

\* Physical Data Independence: The capacity to change the internal schema without having to change the conceptual schema.

### Data Independence

\* When a schema at a lower level is changed, only the mappings between this schema and higher-level schemas need to be changed in a DBMS that fully supports data independence. The higher-level schemas themselves are unchanged. Hence, the application programs need not be changed since they refer to the external schemas.

## DBMS Languages

\* Data Definition Language (DDL): Used by the DBA and database designers to specify the conceptual schema of a database. In many DBMSs, the DDL is also used to define internal and external schemas (views). In some DBMSs, separate storage definition language (SDL) and view definition language (VDL) are used to define internal and external schemas.

## DBMS Languages

- \* Data Manipulation Language (DML): Used to specify database retrievals and updates.
  - \* DML commands (data sublanguage) can be embedded in a general-purpose programming language (host language), such as COBOL, C or an Assembly Language.
  - \* Alternatively, stand-alone DML commands can be applied directly (query language).
#### DBMS Languages

\* High Level or Non-procedural Languages: e.g., SQL, are set-oriented and specify what data to retrieve than how to retrieve. Also called declarative languages.

\* Low Level or Procedural Languages: record-at-a-time; they specify how to retrieve data and include constructs such as looping.

#### **Entity – Relationship Model**

#### **ER Model**

- \* Example Database Application (COMPANY)
- # ER Model Concepts
  - \* Entities and Attributes
  - \* Entity Types, Value Sets, and Key Attributes
  - Relationships and Relationship Types
  - \* Weak Entity Types
  - \* Roles and Attributes in Relationship Types
- # ER Diagrams Notation
- \* ER Diagram for COMPANY Schema

#### Example COMPANY Database

Requirements of the Company (oversimplified for illustrative purposes)

- \* The company is organized into DEPARTMENTs. Each department has a name, number and an employee who manages the department. We keep track of the start date of the department manager.
- \* Each department controls a number of PROJECTs. Each project has a name, number and is located at a single location.
- \* We store each EMPLOYEE's social security number, address, salary, sex, and birthdate. Each employee works for one department but may work on several projects. We keep track of the number of hours per week that an employee currently works on each project. We also keep track of the direct supervisor of each employee.
- \* Each employee may have a number of DEPENDENTs. For each dependent, we keep track of their name, sex, birthdate, and relationship to employee.

#### ER Model Concepts

#### \* Entities and Attributes

- \* Entities are specific objects or things in the mini-world that are represented in the database. For example the EMPLOYEE John Smith, the Research DEPARTMENT, the ProductX PROJECT
- \* Attributes are properties used to describe an entity. For example an EMPLOYEE entity may have a Name, SSN, Address, Sex, BirthDate
- \* A specific entity will have a value for each of its attributes. For example a specific employee entity may have Name='John Smith', SSN='123456789', Address ='731, Fondren, Houston, TX', Sex='M', BirthDate='09-JAN-55'
- \* Each attribute has a value set (or data type) associated with it e.g. integer, string, subrange, enumerated type, …

## Types of Attributes (1)

#### Simple

- \* Each entity has a single atomic value for the attribute. For example, SSN or Sex.
- Composite
  - \* The attribute may be composed of several components. For example, Address (Apt#, House#, Street, City, State, ZipCode, Country) or Name (FirstName, MiddleName, LastName). Composition may form a hierarchy where some components are themselves composite.

#### Multi-valued

\* An entity may have multiple values for that attribute. For example, Color of a CAR or PreviousDegrees of a STUDENT. Denoted as {Color} or {PreviousDegrees}.

### Types of Attributes (2)

\* In general, composite and multi-valued attributes may be nested arbitrarily to any number of levels although this is rare. For example, PreviousDegrees of a STUDENT is a composite multi-valued attribute denoted by {PreviousDegrees (College, Year, Degree, Field)}.

#### Entity Types and Key Attributes

- \* Entities with the same basic attributes are grouped or typed into an entity type. For example, the EMPLOYEE entity type or the PROJECT entity type.
- \* An attribute of an entity type for which each entity must have a unique value is called a key attribute of the entity type. For example, SSN of EMPLOYEE.
- \* A key attribute may be composite. For example, VehicleTagNumber is a key of the CAR entity type with components (Number, State).
- \* An entity type may have more than one key. For example, the CAR entity type may have two keys:
  - \* VehicleIdentificationNumber (popularly called VIN) and
  - \* VehicleTagNumber (Number, State), also known as license\_plate number.

#### SUMMARY OF ER-DIAGRAM NOTATION FOR ER SCHEMAS



Meaning ENTITY TYPE WEAK ENTITY TYPE **RELATIONSHIP TYPE IDENTIFYING RELATIONSHIP TYPE ATTRIBUTE KEY ATTRIBUTE** MULTIVALUED ATTRIBUTE COMPOSITE ATTRIBUTE DERIVED ATTRIBUTE TOTAL PARTICIPATION OF E2 IN R CARDINALITY RATIO 1:N FOR E1:E2 IN R

STRUCTURAL CONSTRAINT (min, max) ON PARTICIPATION OF E IN R

#### ER DIAGRAM – Entity Types are: EMPLOYEE, DEPARTMENT, PROJECT, DEPENDENT



### Relationships and Relationship Types (1)

- \* A relationship relates two or more distinct entities with a specific meaning. For example, EMPLOYEE John Smith works on the ProductX PROJECT or EMPLOYEE Franklin Wong manages the Research DEPARTMENT.
- \* Relationships of the same type are grouped or typed into a relationship type. For example, the WORKS\_ON relationship type in which EMPLOYEEs and PROJECTs participate, or the MANAGES relationship type in which EMPLOYEEs and DEPARTMENTs participate.
- \* The degree of a relationship type is the number of participating entity types. Both MANAGES and WORKS\_ON are binary relationships.

## Example relationship instances of the WORKS\_FOR relationship between EMPLOYEE and DEPARTMENT



# Example relationship instances of the WORKS\_ON relationship between EMPLOYEE and PROJECT



Relationships and Relationship Types (2)

\* More than one relationship type can exist with the same participating entity types. For example, MANAGES and WORKS\_FOR are distinct relationships between EMPLOYEE and DEPARTMENT, but with different meanings and different relationship instances.

#### ER DIAGRAM – Relationship Types are: WORKS\_FOR, MANAGES, WORKS\_ON, CONTROLS, SUPERVISION, DEPENDENTS\_OF



## Weak Entity Types

- \* An entity that does not have a key attribute
- \* A weak entity must participate in an identifying relationship type with an owner or identifying entity type
- \* Entities are identified by the combination of:
  - \* A partial key of the weak entity type
  - \* The particular entity they are related to in the identifying entity type
- \* Example:
  - Suppose that a DEPENDENT entity is identified by the dependent's first name and birhtdate, and the specific EMPLOYEE that the dependent is related to. DEPENDENT is a weak entity type with EMPLOYEE as its identifying entity type via the identifying relationship type DEPENDENT\_OF

#### Weak Entity Type is: DEPENDENT Identifying Relationship is: DEPENDENTS\_OF



#### Constraints on Relationships

- Constraints on Relationship Types
  - \* (Also known as ratio constraints)
  - \* Maximum Cardinality
    - \* One-to-one (1:1)
    - One-to-many (1:N) or Many-to-one (N:1)
    - Many-to-many
  - Minimum Cardinality (also called participation constraint or existence dependency constraints)
    - \* zero (optional participation, not existence-dependent)
    - \* one or more (mandatory, existence-dependent)

### Many-to-one (N:1) RELATIONSHIP



### Many-to-many (M:N) RELATIONSHIP



### Relationships and Relationship Types (3)

- \* We can also have a recursive relationship type.
- \* Both participations are same entity type in different roles.
- \* For example, SUPERVISION relationships between EMPLOYEE (in role of supervisor or boss) and (another) EMPLOYEE (in role of subordinate or worker).
- In following figure, first role participation labeled with 1 and second role participation labeled with 2.
- In ER diagram, need to display role names to distinguish participations.

#### A RECURSIVE RELATIONSHIP SUPERVISION



#### Recursive Relationship Type is: SUPERVISION (participation role names are shown)



# Attributes of Relationship types

\* A relationship type can have attributes; for example, HoursPerWeek of WORKS\_ON; its value for each relationship instance describes the number of hours per week that an EMPLOYEE works on a PROJECT.

#### Attribute of a Relationship Type is: Hours of WORKS\_ON



Structural Constraints – one way to express semantics of relationships

\* Structural constraints on relationships:

\* Cardinality ratio (of a binary relationship): 1:1, 1:N, N:1, or M:N

\* SHOWN BY PLACING APPROPRIATE NUMBER ON THE LINK.

\* Participation constraint (on each participating entity type): total (called existence dependency) or partial.

**\*** SHOWN BY DOUBLE LINING THE LINK

\* NOTE: These are easy to specify for Binary Relationship Types.

## Alternative (min, max) notation for relationship structural constraints:

- \* Specified on each participation of an entity type E in a relationship type R
- \* Specifies that each entity e in E participates in at least min and at most max relationship instances in R
- \* Default(no constraint): min=0, max=n
- \* Must have min $\leq$ max, min $\geq$ 0, max  $\geq$ 1
- \* Derived from the knowledge of mini-world constraints
- \* Examples:
  - \* A department has exactly one manager and an employee can manage at most one department.
    - \* Specify (0,1) for participation of EMPLOYEE in MANAGES
    - \* Specify (1,1) for participation of DEPARTMENT in MANAGES
  - An employee can work for exactly one department but a department can have any number of employees.
    - \* Specify (1,1) for participation of EMPLOYEE in WORKS\_FOR
    - \* Specify (0,n) for participation of DEPARTMENT in WORKS\_FOR

### (min, max) notation



## COMPANY ER Schema Diagram using (min, max) notation



### Relationships of Higher Degree

Relationship types of degree 2 are called binary

\* Relationship types of degree 3 are called ternary and of degree n are called n-ary

\* In general, an n-ary relationship is not equivalent to n binary relationships

#### PROBLEM with ER notation

\* THE ENTITY RELATIONSHIP MODEL IN ITS ORIGINAL FORM DID NOT SUPPORT THE SPECIALIZATION/GENERALIZATION ABSTRACTIONS

#### Extended Entity-Relationship (EER) Model

\* Incorporates Set-subset relationships

**\*** Incorporates Specialization/Generalization Hierarchies

#### **Relational Model**

#### **Relational Model**

#### \* Relational Model Concepts

\* Relational Model Constraints and Relational Database Schemas

\* Update Operations and Dealing with Constraint Violations

#### Concepts

\* The relational Model of Data is based on the concept of a Relation.

\* A Relation is a mathematical concept based on the ideas of sets.

\* The strength of the relational approach to data management comes from the formal foundation provided by the theory of relations.

\* We review the essentials of the relational approach in the next slides.

#### Concepts

 The model was first proposed by Dr. E.F. Codd of IBM in 1970 in the following paper:
"A Relational Model for Large Shared Data Banks," Communications of the ACM, June 1970.

\* The above paper caused a major revolution in the field of Database management and earned Ted Codd the coveted ACM Turing Award.
- RELATION: A table of values
  - \* A relation may be thought of as a set of rows.
  - \* A relation may alternately be though of as a set of columns.
  - \* Each row represents a fact that corresponds to a real-world entity or relationship.
  - \* Each row has a value of an item or set of items that uniquely identifies that row in the table.
  - Sometimes row-ids or sequential numbers are assigned to identify the rows in the table.
  - \* Each column typically is called by its column name or column header or attribute name.

\* A Relation may be defined in multiple ways.

\* The Schema of a Relation: R (A1, A2, ..., An)

- \* Relation schema R is defined over attributes A1, A2, ..., An
- \* For Example: CUSTOMER (Cust-id, Cust-name, Address, Phone#)

\* Here, CUSTOMER is a relation defined over the four attributes Cust-id, Cust-name, Address, Phone#, each of which has a domain or a set of valid values. For example, the domain of Cust-id is 6 digit numbers.

\* A tuple is an ordered set of values

\* Each value is derived from an appropriate domain.

\* Each row in the CUSTOMER table may be referred to as a tuple in the table and would consist of four values.

\* <632895, "John Smith", "101 Main St. Atlanta, GA 30332", "(404) 894-2000"> is a tuple belonging to the CUSTOMER relation.

- \* A relation may be regarded as a set of tuples (rows).
- \* Columns in a table are also called attributes of the relation.

\* A domain has a logical definition: e.g., "USA\_phone\_numbers" are the set of 10 digit phone numbers valid in the U.S.

\* A domain may have a data-type or a format defined for it. The USA\_phone\_numbers may have a format: (ddd)-ddd-dddd where each d is a decimal digit. E.g., Dates have various formats such as monthname, date, year or yyyy-mm-dd, or dd mm,yyyy etc.

\* An attribute designates the role played by the domain. E.g., the domain Date may be used to define attributes "Invoicedate" and "Payment-date".

- \* The relation is formed over the cartesian product of the sets; each set has values from a domain; that domain is used in a specific role which is conveyed by the attribute name.
- \* For example, attribute Cust-name is defined over the domain of strings of 25 characters. The role these strings play in the CUSTOMER relation is that of the name of customers.
- \* Formally,
  - \* Given R(A1, A2, ....., An)
    - \*  $r(R) \subset dom(A1) X dom(A2) X \dots X dom(An)$
- \* R: schema of the relation
- \* r of R: a specific "value" or population of R.

**\*** Let S1 = {0,1}

**\*** Let S2 = {a,b,c}

### **\*** Let $R \subset S1 \times S2$

\* Then for example: r(R) = {<0,a> , <0,b> , <1,c> }

is one possible "state" or "population" or "extension" r of the relation R, defined over domains S1 and S2. It has three tuples.

### DEFINITION SUMMARY

Informal Terms	Formal Terms
Table	Relation
Column	Attribute/Domain
Row	Tuple
Values in a column	Domain
Table Definition	Schema of a Relation
Populated Table	Extension

## Example

	Relation na	ame		Attri	butes			*
	STUDENT	Name	SSN	HomePhone	Address	OfficePhone	Age	GPA
	×	Benjamin Bayer	305-61-2435	373-1616	2918 Bluebonnet Lane	null	19	3.21
	/*	Katherine Ashly	381-62-1245	375-4409	125 Kirby Road	null	18	2.89
	1	Dick Davidson	422-11-2320	null	3452 Elgin Road	749-1253	25	3.53
Tuples		Charles Cooper	489-22-1100	376-9821	265 Lark Lane	749-6492	28	3.93
		Barbara Benson	533-69-1238	839-8461	7384 Fontana Lane	null	19	3.25

### CHARACTERISTICS OF RELATIONS

- \* Ordering of tuples in a relation r(R): The tuples are not considered to be ordered, even though they appear to be in the tabular form.
- \* Ordering of attributes in a relation schema R (and of values within each tuple): We will consider the attributes in R(A1, A2, ..., An) and the values in t=<v1, v2, ..., vn> to be ordered.
- \* (However, a more general alternative definition of relation does not require this ordering). But... Physical level may rely on order!
- \* Values in a tuple: All values are considered atomic (indivisible). A special null value is used to represent values that are unknown or inapplicable to certain tuples.

## CHARACTERISTICS OF RELATIONS

\* Notation:

\* We refer to component values of a tuple t by t[Ai] = vi (the value of attribute Ai for tuple t).

Similarly, t[Au, Av, ..., Aw] refers to the subtuple of t containing the values of attributes Au, Av, ..., Aw, respectively.

### Relational Integrity Constraints

\* Constraints are conditions that must hold on all valid relation instances. There are three main types of constraints:

\* Key constraints

\* Entity integrity constraints

\* Referential integrity constraints

### Key Constraints

- Superkey of R: A set of attributes SK of R such that no two tuples in any valid relation instance r(R) will have the same value for SK. That is, for any distinct tuples t1 and t2 in r(R), t1[SK] ≠ t2[SK].
- **Key** of R: A "minimal" superkey; that is, a superkey K such that removal of any attribute from K results in a set of attributes that is not a superkey.
  - **\*** Example: The CAR relation schema:
    - \* CAR(State, Reg#, SerialNo, Make, Model, Year)
    - \* has two keys, Key1 = {State, Reg#}, Key2 = {SerialNo}, which are also superkeys. {SerialNo, Make} is a superkey but not a key.
- If a relation has several candidate keys, one is chosen arbitrarily to be the primary key. The primary key attributes are underlined.

### Key Constraints

Figure 7.4 The CAR relation with two candidate keys: LicenseNumber and EngineSerialNumber.

CAR	LicenseNumber EngineSerialNumber		Make	Model	Year
	Texas ABC-739	A69352	Ford	Mustang	96
	Florida TVP-347	B43696	Oldsmobile	Cutlass	99
1	New York MPO-22	X83554	Oldsmobile	Delta	95
	California 432-TFY	C43742	Mercedes	190-D	93
	California RSK-629	Y82935	Toyota	Camry	98
	Texas RSK-629	U028365	Jaguar	XJS	98

© Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

## Entity Integrity

\* Relational Database Schema: A set S of relation schemas that belong to the same database. S is the name of the database.

**\*** S = {R1, R2, ..., Rn}

- \* Entity Integrity: The primary key attributes PK of each relation schema R in S cannot have null values in any tuple of r(R). This is because primary key values are used to identify the individual tuples.
- \*  $t[PK] \neq null for any tuple t in r(R)$
- \* Note: Other attributes of R may be similarly constrained to disallow null values, even though they are not members of the primary key.

### Referential Integrity

- \* A constraint involving two relations (the previous constraints involve a single relation).
- \* Used to specify a relationship among tuples in two relations: the referencing relation and the referenced relation.
- \* Tuples in the referencing relation R1 have attributes FK (called foreign key attributes) that reference the primary key attributes PK of the referenced relation R2. A tuple t1 in R1 is said to reference a tuple t2 in R2 if t1[FK] = t2[PK].
- \* A referential integrity constraint can be displayed in a relational database schema as a directed arc from R1.FK to R2.

### Referential Integrity Constraint

- Statement of the constraint
- \* The value in the foreign key column (or columns) FK of the the referencing relation R1 can be either:
  - 1.a value of an existing primary key value of the corresponding primary key PK in the referenced relation R2, or
  - 2.a null.
- In case (2), the FK in R1 should not be a part of its own primary key.

### Other Types of Constraints

\* Semantic Integrity Constraints:

- \* based on application semantics and cannot be expressed by the model per se
- \* E.g., "the max. no. of hours per employee for all projects he or she works on is 56 hrs per week"
- \* A constraint specification language may have to be used to express these

\* SQL-99 allows TRIGGERS and ASSERTIONS to allow for some of these

### **Figure 7.7** Referential integrity constraints displayed on the COMPANY relational database schema diagram.



C Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

## Update Operations on Relations

- **\*** INSERT a tuple.
- **\*** DELETE a tuple.
- **\*** MODIFY a tuple.
- Integrity constraints should not be violated by the update operations.
- \* Several update operations may have to be grouped together.
- \* Updates may propagate to cause other updates automatically. This may be necessary to maintain integrity constraints.

## Update Operations on Relations

\* In case of integrity violation, several actions can be taken:

- \* Cancel the operation that causes the violation (REJECT option)
- \* Perform the operation but inform the user of the violation
- \* Trigger additional updates so the violation is corrected (CASCADE option, SET NULL option)

\* Execute a user-specified error-correction routine

# Conceptual – Relational model Mapping

### ER – Relational Mapping

- \* Step 1: Mapping of Regular Entity Types
- \* Step 2: Mapping of Weak Entity Types
- \* Step 3: Mapping of Binary 1:1 Relation Types
- \* Step 4: Mapping of Binary 1:N Relationship Types.
- \* Step 5: Mapping of Binary M:N Relationship Types.
- Step 6: Mapping of Multivalued attributes.
- \* Step 7: Mapping of N-ary Relationship Types.
- Mapping EER Model Constructs to Relations
  - \* Step 8: Options for Mapping Specialization or Generalization.
  - Step 9: Mapping of Union Types (Categories).

- \* Step 1: Mapping of Regular Entity Types.
  - \* For each regular (strong) entity type E in the ER schema, create a relation R that includes all the simple attributes of E.
  - Choose one of the key attributes of E as the primary key for R. If the chosen key of E is composite, the set of simple attributes that form it will together form the primary key of R.
    - \* Example: We create the relations EMPLOYEE, DEPARTMENT, and PROJECT in the relational schema corresponding to the regular entities in the ER diagram. SSN, DNUMBER, and PNUMBER are the primary keys for the relations EMPLOYEE, DEPARTMENT, and PROJECT as shown.





- Step 2: Mapping of Weak Entity Types
  - \* For each weak entity type W in the ER schema with owner entity type E, create a relation R and include all simple attributes (or simple components of composite attributes) of W as attributes of R.
  - In addition, include as foreign key attributes of R the primary key attribute(s) of the relation(s) that correspond to the owner entity type(s).
  - \* The primary key of R is the combination of the primary key(s) of the owner(s) and the partial key of the weak entity type W, if any.
    - Example: Create the relation DEPENDENT in this step to correspond to the weak entity type DEPENDENT. Include the primary key SSN of the EMPLOYEE relation as a foreign key attribute of DEPENDENT (renamed to ESSN).
  - \* The primary key of the DEPENDENT relation is the combination {ESSN, DEPENDENT\_NAME} because DEPENDENT\_NAME is the partial key of DEPENDENT.

Step 3: Mapping of Binary 1:1 Relation Types

- \* For each binary 1:1 relationship type R in the ER schema, identify the relations S and T that correspond to the entity types participating in R. There are three possible approaches:
- 1. Foreign Key approach: Choose one of the relations-S, say-and include a foreign key in S the primary key of T. It is better to choose an entity type with total participation in R in the role of S.
  - Example: 1:1 relation MANAGES is mapped by choosing the participating entity type DEPARTMENT to serve in the role of S, because its participation in the MANAGES relationship type is total.
- 2. Merged relation option: An alternate mapping of a 1:1 relationship type is possible by merging the two entity types and the relationship into a single relation. This may be appropriate when both participations are total.
- Cross-reference or relationship relation option: The third alternative is to set up a third relation R for the purpose of cross-referencing the primary keys of the two relations S and T representing the entity types.

- \* Step 4: Mapping of Binary 1:N Relationship Types.
  - \* For each regular binary 1:N relationship type R, identify the relation S that represent the participating entity type at the N-side of the relationship type.
  - Include as foreign key in S the primary key of the relation T that represents the other entity type participating in R.
  - \* Include any simple attributes of the 1:N relation type as attributes of S.
    - \* Example: 1:N relationship types WORKS\_FOR, CONTROLS, and SUPERVISION in the figure. For WORKS\_FOR we include the primary key DNUMBER of the DEPARTMENT relation as foreign key in the EMPLOYEE relation and call it DNO.

- Step 5: Mapping of Binary M:N Relationship Types.
  - \* For each regular binary M:N relationship type R, create a new relation S to represent R.
  - Include as foreign key attributes in S the primary keys of the relations that represent the participating entity types; their combination will form the primary key of S.
  - \* Also include any simple attributes of the M:N relationship type (or simple components of composite attributes) as attributes of S.
    - Example: The M:N relationship type WORKS\_ON from the ER diagram is mapped by creating a relation WORKS\_ON in the relational database schema. The primary keys of the PROJECT and EMPLOYEE relations are included as foreign keys in WORKS\_ON and renamed PNO and ESSN, respectively.
    - \* Attribute HOURS in WORKS\_ON represents the HOURS attribute of the relation type. The primary key of the WORKS\_ON relation is the combination of the foreign key attributes {ESSN, PNO}.

- \* Step 6: Mapping of Multivalued attributes.
  - \* For each multivalued attribute A, create a new relation R. This relation R will include an attribute corresponding to A, plus the primary key attribute K-as a foreign key in R-of the relation that represents the entity type of relationship type that has A as an attribute.
  - \* The primary key of R is the combination of A and K. If the multivalued attribute is composite, we include its simple components.
    - \* Example: The relation DEPT\_LOCATIONS is created. The attribute DLOCATION represents the multivalued attribute LOCATIONS of DEPARTMENT, while DNUMBER-as foreign key-represents the primary key of the DEPARTMENT relation. The primary key of R is the combination of {DNUMBER, DLOCATION}.

- \* Step 7: Mapping of N-ary Relationship Types.
  - \* For each n-ary relationship type R, where n>2, create a new relationship S to represent R.
  - Include as foreign key attributes in S the primary keys of the relations that represent the participating entity types.
  - \* Also include any simple attributes of the n-ary relationship type (or simple components of composite attributes) as attributes of S.
    - \* Example: The relationship type SUPPY in the ER below. This can be mapped to the relation SUPPLY shown in the relational schema, whose primary key is the combination of the three foreign keys {SNAME, PARTNO, PROJNAME}



#### SUPPLIER

SNAME	• • •

#### PROJECT

PROJNAME	•••
----------	-----

#### PART

PARTNO	•••
--------	-----

### SUPPLY

SNAME	PROJNAME	PARTNO	QUANTITY	

## Summary of Mapping constructs and constraints

ER Model	Relational Model
Entity type	"Entity" Relation
1:1 or 1:N Relationship type	Foreign Key (or "Relationship" relation)
M:N Relationship type	"Relationship" relation and two Foreign Keys
N-ary Relationship type	"Relationship" relation and N Foreign Keys
Simple Attribute	Attribute
Composite Attribute	Set of simple component attributes
Multivalued attribute	Relation and Foreign Key
Value set	Domain
Key Attribute	Primary Key (or Key)

### **Relational Algebra**

### Relational Algebra

\* Example Database Application (COMPANY)

\* Relational Algebra

- **\*** Unary Relational Operations
- \* Relational Algebra Operations From Set Theory
- \* Binary Relational Operations
- **\*** Additional Relational Operations
- \* Examples of Queries in Relational Algebra
## Database State for COMPANY

Figure 7.5 Schema diagram for the COMPANY relational database schema; the primary keys are underlined.

MINIT	LNAME	SSN	BUATE	ADDRESS	SEX	SALARY	SUPERSSN	
	DE	PARTMEN	п		_			
	DNAM	IE DN	UMBER	MGRSSN	MGRS	TARTDATE		
			DEPT	LOCATIONS				
			DNUMBER	DLOCA	TION			
		PROJEC	т					
		PNAME	PNUMBE	R PLOCAT		NUM		
			w	ORKS_ON				
			ESSN	PNO HOU	JRS			
Г	DEPEND	EDENDE		CEV 0	DATE	DELATIONOU		
I	ESSIN L	PENDEI	NT_INAME	SEA 0	DATE	RELATIONSHI	P	

## Relational Algebra

- \* The basic set of operations for the relational model is known as the relational algebra. These operations enable a user to specify basic retrieval requests.
- \* The result of a retrieval is a new relation, which may have been formed from one or more relations. The algebra operations thus produce new relations, which can be further manipulated using operations of the same algebra.
- \* A sequence of relational algebra operations forms a relational algebra expression, whose result will also be a relation that represents the result of a database query (or retrieval request).

## Unary Relational Operations

- SELECT Operation
- SELECT operation is used to select a subset of the tuples from a relation that satisfy a selection condition. It is a filter that keeps only those tuples that satisfy a qualifying condition – those satisfying the condition are selected while others are discarded.
  - \* Example: To select the EMPLOYEE tuples whose department number is four or those whose salary is greater than \$30,000 the following notation is used:
    - \* σ<sub>DNO = 4</sub> (EMPLOYEE)
- \* In general, the select operation is denoted by  $\sigma$  <selection condition>(R) where the
- \* symbol σ (sigma) is used to denote the select operator, and the selection condition is a Boolean expression specified on the attributes of relation R

## Unary Relational Operations

- SELECT Operation Properties
- \* The SELECT operation σ <selection condition>(R) produces a relation S that has the same schema as R
  - \* The SELECT operation  $\sigma$  is commutative; i.e.,
    - \*  $\sigma_{\text{condition1>}}(\sigma_{\text{condition2>}}(R)) = \sigma_{\text{condition2>}}(\sigma_{\text{condition1>}}(R))$
  - \* A cascaded SELECT operation may be applied in any order; i.e.,
    - \*  $\sigma_{\text{condition1>}}(\sigma_{\text{condition2>}}(R)) = \sigma_{\text{condition2>}}(\sigma_{\text{condition3>}}(R)))$
  - \* A cascaded SELECT operation may be replaced by a single selection with a conjunction of all the conditions; i.e.,
    - \*  $\sigma_{\text{condition1>}}(\sigma_{\text{condition2>}}(R)) = \sigma_{\text{condition1>}} \text{ AND } \text{-condition2>} \text{ AND } \text{-condition3>}(R)$

Figure 7.8 Results of SELECT and PROJECT operations. (a)  $\sigma_{(DNO=4 \text{ AND SALARY>25000}) \text{ OR } (DNO=5 \text{ AND SALARY>30000})}$  (EMPLOYEE). (b)  $\pi_{\text{LNAME, FNAME, SALARY}}$  (EMPLOYEE). (c)  $\pi_{\text{SEX, SALARY}}$  (EMPLOYEE)

a)	FNAME	MINIT	LNAME	<u>SSN</u>	BDATE	ADDRESS	SEX	SALARY	SUPERSSN	DNO
	Franklin	Т	Wong	3334455555	1955-12-08	638 Voss, Houston, TX	М	40000	888665555	5
	Jennifer		Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
	Ramesh		Narayan	666884444	1962-09-15	975 FireOak,Humble,TX	М	38000	333445555	5

(c)

(b)

LNAME	FNAME	SALARY
Smith	John	30000
Wong	Franklin	40000
Zelaya	Alicia	25000
Wallace	Jennifer	43000
Narayan	Ramesh	38000
English	Joyce	25000
Jabbar	Ahmad	25000
Borg	James	55000

SEX	SALARY
М	30000
М	40000
F	25000
F	43000
М	38000
М	25000
М	55000

C Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

## Unary Relational Operations (cont.)

- **PROJECT** Operation
- \* This operation selects certain columns from the table and discards the other columns. The PROJECT creates a vertical partitioning – one with the needed columns (attributes) containing results of the operation and other containing the discarded Columns.
  - \* Example: To list each employee's first and last name and salary, the following is used:
    - \* πLNAME, FNAME, SALARY (EMPLOYEE)
- \* The general form of the project operation is π<sub><attribute list>(</sub>R) where π (pi) is the symbol used to represent the project operation and <attribute list> is the desired list of attributes from the attributes of relation R.
- \* The project operation removes any duplicate tuples, so the result of the project operation is a set of tuples and hence a valid relation.

Unary Relational Operations (cont.)

**\* PROJECT Operation Properties** 

\* The number of tuples in the result of projection π <list> (R)is always less or equal to the number of tuples in R.

If the list of attributes includes a key of R, then the number of tuples is equal to the number of tuples in R.

\*  $\pi_{<list1>}(\pi_{<list2>}(R)) = \pi_{<list1>}(R)$ as long as <list2> contains the attributes in <list2> Figure 7.8 Results of SELECT and PROJECT operations. (a)  $\sigma_{(DNO=4 \text{ AND SALARY>25000}) \text{ OR } (DNO=5 \text{ AND SALARY>30000})}$  (EMPLOYEE). (b)  $\pi_{\text{LNAME, FNAME, SALARY}}$  (EMPLOYEE). (c)  $\pi_{\text{SEX, SALARY}}$  (EMPLOYEE)

a)	FNAME	MINIT	LNAME	<u>SSN</u>	BDATE	ADDRESS	SEX	SALARY	SUPERSSN	DNO
	Franklin	Т	Wong	3334455555	1955-12-08	638 Voss, Houston, TX	М	40000	888665555	5
	Jennifer		Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
	Ramesh		Narayan	666884444	1962-09-15	975 FireOak,Humble,TX	М	38000	333445555	5

(c)

(b)

LNAME	FNAME	SALARY
Smith	John	30000
Wong	Franklin	40000
Zelaya	Alicia	25000
Wallace	Jennifer	43000
Narayan	Ramesh	38000
English	Joyce	25000
Jabbar	Ahmad	25000
Borg	James	55000

SEX	SALARY
М	30000
М	40000
F	25000
F	43000
М	38000
М	25000
М	55000

C Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

## Unary Relational Operations (cont.)

- Rename Operation
- We may want to apply several relational algebra operations one after the other. Either we can write the operations as a single relational algebra expression by nesting the operations, or we can apply one operation at a time and create intermediate result relations. In the latter case, we must give names to the relations that hold the intermediate results.
  - \* Example: To retrieve the first name, last name, and salary of all employees who work in department number 5, we must apply a select and a project operation. We can write a single relational algebra expression as follows:
    - \* πfname, lname, salary(σdno=5 (EMPLOYEE))
  - \* OR We can explicitly show the sequence of operations, giving a name to each intermediate relation:
    - \* DEP5\_EMPS  $\leftarrow \sigma_{DNO=5}$  (EMPLOYEE)
    - **\*** RESULT ←  $\pi$  FNAME, LNAME, SALARY (DEP5\_EMPS)

## Unary Relational Operations (cont.)

- Rename Operation (cont.)
- \* The rename operator is ρ
- \* The general Rename operation can be expressed by any of the following forms:
  - \* ρ S (B1, B2, ..., Bn ) (R) is a renamed relation S based on R with column names B1, B1, ..., Bn.
  - \* ρ S (R) is a renamed relation S based on R (which does not specify column names).
  - \* ρ (B1, B2, ..., Bn ) (R) is a renamed relation with column names B1, B1, ..., Bn which does not specify a new relation name.

## Relational Algebra Operations From Set Theory

#### \* UNION Operation

- ★ The result of this operation, denoted by R ∪ S, is a relation that includes all tuples that are either in R or in S or in both R and S. Duplicate tuples are eliminated.
  - Example: To retrieve the social security numbers of all employees who either work in department 5 or directly supervise an employee who works in department 5, we can use the union operation as follows:
    - \* DEP5\_EMPS  $\leftarrow \sigma_{DNO=5}$  (EMPLOYEE)
    - \* RESULT1  $\leftarrow \pi_{SSN}$  (DEP5\_EMPS)
    - **\*** RESULT2(SSN) ←  $\pi$  superssn(DEP5\_EMPS)
    - **\*** RESULT ← RESULT1 ∪ RESULT2
- \* The union operation produces the tuples that are in either RESULT1 or RESULT2 or both. The two operands must be "type compatible".

## Relational Algebra Operations From Set Theory

### \* Type Compatibility

- \* The operand relations R1(A1, A2, ..., An) and R2(B1, B2, ..., Bn) must have the same number of attributes, and the domains of corresponding attributes must be compatible; that is, dom(Ai)=dom(Bi) for i=1, 2, ..., n.
- ★ The resulting relation for R1∪R2,R1 ∩ R2, or R1-R2 has the same attribute names as the first operand relation R1 (by convention).

Figure 7.11 Illustrating the set operations union, intersection, and difference. (a) Two union compatible relations. (b) STUDENT  $\cup$  INSTRUCTOR. (c) STUDENT  $\cup$  INSTRUCTOR. (d) STUDENT - INSTRUCTOR. (e) INSTRUCTOR - STUDENT.

STUDENT	FN	LN
	Susan	Yao
	Ramesh	Shah
	Johnny	Kohler
	Barbara	Jones
	Amy	Ford
	Jimmy	Wang
	Ernest	Gibert

Jones

Ford

Wang

Gibert

Smith

Browne

Johnson

INSTRUCTOR	FNAME	LNAME
	John	Smith
	Ricardo	Browne
	Susan	Yao
	Francis	Johnson
	Ramesh	Shah



Barbara

Amy

Jimmy

Ernest

John

Ricardo

Francis

(C)	FN	LN
	Susan	Yao
	Ramesh	Shah

1.10	_
6711	
1041	

(a)

FN	LN
Johnny	Kohier
Barbara	Jones
Amy	Ford
Jimmy	Wang
Ernest	Gibert

FNAME	LNAME
John	Smith
Ricardo	Browne
Francis	Johnson

C Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

(e)

### Relational Algebra Operations From Set Theory (cont.)

\* CARTESIAN (or cross product) Operation

- This operation is used to combine tuples from two relations in a combinatorial fashion. In general, the result of R(A1, A2, ..., An) x S(B1, B2, ..., Bm) is a relation Q with degree n + m attributes Q(A1, A2, ..., An, B1, B2, ..., Bm), in that order. The resulting relation Q has one tuple for each combination of tuples—one from R and one from S.
- \* Hence, if R has nR tuples (denoted as |R| = nR), and S has nS tuples, then
  - R x S | will have nR \* nS tuples.
- \* The two operands do NOT have to be "type compatible"
- \* Example:
  - **\*** FEMALE\_EMPS ←  $\sigma_{SEX='F'}$ (EMPLOYEE)
  - \* EMPNAMES  $\leftarrow \pi$  FNAME, LNAME, SSN (FEMALE\_EMPS)
  - ★ EMP\_DEPENDENTS ← EMPNAMES × DEPENDENT

#### Figure 7.12 An illustration of the CARTESIAN PRODUCT operation.

FEMALE_ EMPS	FNAME	MINT	LNAME	SSN	BDATE	ADDRESS	SEX	SALARY	SUPERSSN	DNO
	Akis	1	Zeleyn	999887777	1968-07-19	3321 Castle, Spring, TX	F	25000	987654321	4
	Jerviler	S	Wallace	987654321	1941-06-20	291 Berry Bellaire, TX	F	43000	888965555	4
	Joyce	A	English	453453453	1972-07-31	5631 Rice,Houston,TX	F	25000	333445555	5

EMPNAMES	FNAME	LNAME	SSN
	Alicia	Zoloya	9998877777
	Jerstifer	Wallape	987654321
	Jayse	English	453453453

EMP_DEPENDENTS	FNAME	LNAME	SSN	ESSN	DEPENDENT_NAME	SEX	BDATE	
	Alicia	Zekaya	999887777	333445555	Alce	F	1986-04-05	
	Alicin	Zelaya	999867777	333445555	Theodore	M	1983-10-25	
	Alicia	Zelaya	999887777	333445555	Joy	F	1958-05-05	
	Alicia	Zelaya	999887777	987654321	Abeer	M	1942-02-28	
	Alicin	Zekiya	999887777	123456780	Michael	M	1988-01-04	
	Alicia	Zelaya	999867777	123456780	Alce	P.	1988-12-30	
	Alicin	Zelaya	990887777	123456789	Ekabeth	F	1967-05-05	
	Jenniler	Walace	987654321	333445555	Alce	F	1986-04-05	
	Jenniler	Wallace	987654321	333445555	Theodore	M	1983-10-25	
	Jerniler	Wallace	987654321	333445555	Joy	F	1958-05-03	
	Jermiler	Walkce	987654321	967654321	Abner	M	1942-02-28	
	Jenniler	Wallace	987654321	123456780	Michael	M	1968-01-04	
	Jerniler	Walace	987054321	123456789	Alce	1	1988-12-30	
	Jenniler	Wallace	987654321	123456789	Elizabeth	F	1967-05-05	
	Joyce	English	453453453	333445555	Alce	p.	1986-04-05	
	Joyce	Erglish	453453453	333445555	Theodore	M	1983-10-25	
	Jayoe	Erglish	453453453	333445555	Joy	F	1958-05-03	
	Jayoe	Brglish	453453453	967654321	Abner	M	1942-02-28	
	Joyce	Briglish	453453453	123456780	Michael	M	1988-01-04	
	Joyce	Erglish	453453453	123456789	Alce	1	1968-12-30	
	Joyce	Erglish	453453453	123456780	Elizabeth	F	1967-05-05	

ACTUAL DEPENDENTS	FNAME	LNAME	SSN	ESSN	DEPENDENT_NAME	SEX	BDATE
	Jernifer	Walace	987654321	967654321	Abner	M	1942-02-28

RESULT	FNAME	LNAME	DEPENDENT_NAME
	Jenniler	Walace	Abner

Γ

C Addison Wesley Longman, Inc. 2000, Elmasri/Navathe, Fundamentals of Database Systems, Third Edition

## Binary Relational Operations

- **\*** JOIN Operation
  - ★ The sequence of cartesian product followed by select is used quite commonly to identify and select related tuples from two relations, a special operation, called JOIN. It is denoted by a
  - \* This operation is very important for any relational database with more than a single relation, because it allows us to process relationships among relations.
  - \* The general form of a join operation on two relations R(A1, A2, ..., An) and S(B1, B2, ..., Bm) is:
    - \* R <-join condition> S
    - \* where R and S can be any relations that result from general relational algebra expressions.

## Binary Relational Operations (cont.)

- \* Example: Suppose that we want to retrieve the name of the manager of each department. To get the manager's name, we need to combine each DEPARTMENT tuple with the EMPLOYEE tuple whose SSN value matches the MGRSSN value in the department tuple. We do this by using the join operation.
  - \* DEPT\_MGR ← DEPARTMENT MGRSSN=SSN EMPLOYEE

DEPT_MGR	DNAME	DNUMBER	MGRSSN	 FNAME	MINIT	LNAME	SSN	
	Research	5	333445555	 Franklin	Т	Wong	333445555	
	Administration	4	987654321	 Jennifer	S	Wallace	987654321	
	Headquarters	1	888665555	 James	E	Borg	888665555	

FIGURE 6.6 Result of the JOIN operation DEPT\_MGR ← DEPARTMENT MGRSSN=SSN EMPLOYEE.

## Binary Relational Operations (cont.)

- **EQUIJOIN** Operation
- \* The most common use of join involves join conditions with equality comparisons only. Such a join, where the only comparison operator used is =, is called an EQUIJOIN. In the result of an EQUIJOIN we always have one or more pairs of attributes (whose names need not be identical) that have identical values in every tuple.
- \* The JOIN seen in the previous example was EQUIJOIN.

#### \* NATURAL JOIN Operation

- Because one of each pair of attributes with identical values is superfluous, a new operation called natural join—denoted by \*—was created to get rid of the second (superfluous) attribute in an EQUIJOIN condition.
- \* The standard definition of natural join requires that the two join attributes, or each pair of corresponding join attributes, have the same name in both relations. If this is not the case, a renaming operation is applied first.

## Binary Relational Operations (cont.)

\* Example: To apply a natural join on the DNUMBER attributes of DEPARTMENT and DEPT\_LOCATIONS, it is sufficient to write:

### ★ DEPT\_LOCS ← DEPARTMENT \* DEPT\_LOCATIONS

PROJ_DEPT	PNAME	PNUMBER	PLOCATION	DNUM	DNAM	1E	MGRSSN	MGRSTARTDATE
	ProductX	1	Bellaire	5	Research		333445555	1988-05-22
	ProductY	2	Sugarland	5	Research		333445555	1988-05-22
	ProductZ	3	Houston	5	Research	tion of	333445555	1988-05-22
	Computerization	10	Stafford	4	Administra	ation	987654321	1995-01-01
	Reorganization	20	Houston	1	Headquar	rters	888665555	1981-06-19
	Marchanofile	20	Ctational		A	ation	007051001	1005 01 01
	Newbenefits		Station		Administra	abon	987654321	1995-01-01
	Newbenefits	30	Station		Administra	abon	987654321	1995-01-01
DEPT_LOCS	DNAME	DNUMBER	MGRSSN	4 MGRSTA	RTDATE	LOC	987654321	1995-01-01
DEPT_LOCS	DNAME Headquarters	DNUMBER 1	MGRSSN 888665555	MGRSTAI 1981-0	RTDATE	LOC	CATION uston	1995-01-01
DEPT_LOCS	DNAME Headquarters Administration	DNUMBER 1 4	MGRSSN 888665555 987654321	4 MGRSTAI 1981-0 1995-0	RTDATE 06-19 01-01	LOC Hou Sta	CATION uston	1995-01-01
DEPT_LOCS	DNAME Headquarters Administration Research	DNUMBER 1 4 5	MGRSSN 888665555 987654321 333445555	4 MGRSTAI 1981-1 1995-1 1988-0	RTDATE 06-19 01-01 05-22	LOC Hou Sta Bell	CATION uston fford laire	1995-01-01
DEPT_LOCS	DNAME Headquarters Administration Research Research	30 DNUMBER 1 4 5 5 5	MGRSSN 888665555 987654321 333445555 333445555	4 MGRSTAI 1981-4 1995-1 1988-4 1988-4	Administra RTDATE 06-19 01-01 05-22 05-22	LOC Hou Sta Bell	CATION uston flord laire garland	1995-01-01

### Recap of Relational Algebra Operations

#### TABLE 6.1 OPERATIONS OF RELATIONAL ALGEBRA

Operation	Purpose	Notation
SELECT	Selects all tuples that satisfy the selection condition from a relation <i>R</i> .	$\sigma_{\text{$
PROJECT	Produces a new relation with only some of the attributes of <i>R</i> , and removes duplicate tuples.	$\pi_{< \text{ATTRIBUTE LIST>}}(R)$
THETA JOIN	Produces all combinations of tuples from $R_1$ and $R_2$ that satisfy the join condition.	$R_1^{M} \leq \text{ODIN CONDITIONS} R_2$
EQUIJOIN	Produces all the combinations of tuples from $R_1$ and $R_2$ that satisfy a join condition with only equality compar- isons.	$R_1^{\bowtie} <_{\text{JOIN CONDITION>}} R_2$ , OR $R_1^{\bowtie} (_{\text{JOIN ATTRIBUTES 1>}})$ , $(_{\text{JOIN ATTRIBUTES 2>}}) R_2$
NATURAL JOIN	Same as EQUIJOIN except that the join attributes of $R_2$ are not included in the resulting relation; if the join attributes have the same names, they do not have to be specified at all.	$R_1^*_{}R_2$ , OR $R_1^*(_{})$ , $(_{})R_2$ OR $R_1^*R_2$
UNION	Produces a relation that includes all the tuples in $R_1$ or $R_2$ or both $R_1$ and $R_2$ ; $R_1$ and $R_2$ must be union compatible.	$R_1 \cup R_2$
INTERSECTION	Produces a relation that includes all the tuples in both $R_1$ and $R_2$ ; $R_1$ and $R_2$ must be union compatible.	$R_1 \cap R_2$
DIFFERENCE	Produces a relation that includes all the tuples in $R_1$ that are not in $R_2$ ; $R_1$ and $R_2$ must be union compatible.	$R_1 - R_2$
CARTESIAN	Produces a relation that has the attributes of $R_1$ and $R_2$	$R_1 \times R_2$
PRODUCT	and includes as tuples all possible combinations of tuples from $R_1$ and $R_2$ .	
DIVISION	Produces a relation $R(X)$ that includes all tuples $t[X]$ in $R_1(Z)$ that appear in $R_1$ in combination with every tuple from $R_2(Y)$ , where $Z = X \cup Y$ .	$R_1(Z) \div R_2(Y)$

# Additional Relational Operations

- \* Aggregate Functions and Grouping
  - \* A type of request that cannot be expressed in the basic relational algebra is to specify mathematical aggregate functions on collections of values from the database.
  - \* Examples of such functions include retrieving the average or total salary of all employees or the total number of employee tuples. These functions are used in simple statistical queries that summarize information from the database tuples.
  - \* Common functions applied to collections of numeric values include SUM, AVERAGE, MAXIMUM, and MINIMUM. The COUNT function is used for counting tuples or values.

(a)	R	DNO	NO_OF_EMPLOYEES	AVERAGE_SAL
		5	4	33250
		4	3	31000
		1	1	55000

(b)

DNO	COUNT_SSN	AVERAGE_SALARY
5	4	33250
4	3	31000
1	1	55000

11		L	
"	٠	1	

COUNT_SSN	AVERAGE_SALARY	
8	35125	

- \* Use of the Functional operator F
  - \* MAX Salary (Employee) retrieves the maximum salary value from the Employee relation
  - \* MIN Salary (Employee) retrieves the minimum Salary value from the Employee relation
  - SUM Salary (Employee) retrieves the sum of the Salary from the Employee relation
  - \* DNO FCOUNT SSN, AVERAGE Salary (Employee) groups employees by DNO (department number) and computes the count of employees and average salary per department.[Note: count just counts the number of rows, without removing duplicates]

- The OUTER JOIN Operation
  - In NATURAL JOIN tuples without a matching (or related) tuple are eliminated from the join result. Tuples with null in the join attributes are also eliminated. This amounts to loss of information.
  - \* A set of operations, called outer joins, can be used when we want to keep all the tuples in R, or all those in S, or all those in both relations in the result of the join, regardless of whether or not they have matching tuples in the other relation.
  - \* The left outer join operation keeps every tuple in the first or left relation R in R S; if no matching tuple is found in S, then the attributes of S in the join result are filled or "padded" with null values.
  - \* A similar operation, right outer join, keeps every tuple in the second or right relation S in the result of R S.
  - \* A third operation, full outer join, denoted by \_\_\_\_\_keeps all tuples in both the left and the right relations when no matching tuples are found, padding them with null values as needed.

RESULT	FNAME	MINIT	LNAME	DNAME
	John	В	Smith	null
	Franklin	Т	Wong	Research
	Alicia	J	Zelaya	null
	Jennifer	S	Wallace	Administration
	Ramesh	К	Narayan	null
	Joyce	Α	English	null
	Ahmad	V	Jabbar	null
	James	E	Borg	Headquarters



## Data Definition, Constraints, and Schema Changes

\* Used to CREATE, DROP, and ALTER the descriptions of the tables (relations) of a database

## CREATE TABLE

\* Specifies a new base relation by giving it a name, and specifying each of its attributes and their data types (INTEGER, FLOAT, DECIMAL(i,j), CHAR(n), VARCHAR(n))

\* A constraint NOT NULL may be specified on an attribute

CREATE TABLE DEPARTMENT ( DNAME VARCHAR(10) NOT NULL, DNUMBER INTEGERNOT NULL, MGRSSN CHAR(9), MGRSTARTDATE CHAR(9) );

## CREATE TABLE

- In SQL2, can use the CREATE TABLE command for specifying the primary key attributes, secondary keys, and referential integrity constraints (foreign keys).
- \* Key attributes can be specified via the PRIMARY KEY and UNIQUE phrases

```
* CREATE TABLE DEPT
```

DNAME VARCHAR(10) NOT NULL, DNUMBER INTEGER NOT NULL, MGRSSN CHAR(9), MGRSTARTDATE CHAR(9), PRIMARY KEY (DNUMBER), UNIQUE (DNAME), FOREIGN KEY (MGRSSN) REFERENCES EMP );

## DROP TABLE

\* Used to remove a relation (base table) and its definition

\* The relation can no longer be used in queries, updates, or any other commands since its description no longer exists

\* Example:

### DROP TABLE DEPENDENT;

## ALTER TABLE

- \* Used to add an attribute to one of the base relations
- \* The new attribute will have NULLs in all the tuples of the relation right after the command is executed; hence, the NOT NULL constraint is not allowed for such an attribute
- \* Example:

```
ALTER TABLE EMPLOYEE ADD JOB VARCHAR(12);
```

\* The database users must still enter a value for the new attribute JOB for each EMPLOYEE tuple. This can be done using the UPDATE command.

## REFERENTIAL INTEGRITY OPTIONS

\* We can specify RESTRICT, CASCADE, SET NULL or SET DEFAULT on referential integrity constraints (foreign keys)

CREATE TABLE DEPT ( DNAME VARCHAR(10) NOT NULL, DNUMBER INTEGER NOT NULL, MGRSSN CHAR(9), MGRSTARTDATE CHAR(9), PRIMARY KEY (DNUMBER), UNIQUE (DNAME), FOREIGN KEY (MGRSSN) REFERENCES EMP ON DELETE SET DEFAULT ON UPDATE CASCADE );

## REFERENTIAL INTEGRITY OPTIONS (continued)

### **\*** CREATE TABLE EMP ( ENAME VARCHAR(30) NOT NULL, ESSN CHAR(9), BDATE DATE, DNO INTEGER DEFAULT 1, SUPERSSN CHAR(9), PRIMARY KEY (ESSN), FOREIGN KEY (DNO) REFERENCES DEPT ON DELETE SET DEFAULT ON UPDATE CASCADE, FOREIGN KEY (SUPERSSN) REFERENCES EMP ON DELETE SET NULL ON UPDATE CASCADE );

## Additional Data Types in SQL2 and SQL-99

- \* Has DATE, TIME, and TIMESTAMP data types
- **\*** DATE:
  - Made up of year-month-day in the format yyyy-mm-dd
- **\*** TIME:
  - \* Made up of hour:minute:second in the format hh:mm:ss
- # TIME(i):
  - Made up of hour:minute:second plus i additional digits specifying fractions of a second
  - \* format is hh:mm:ss:ii...i
- **\*** TIMESTAMP:
  - \* Has both DATE and TIME components

Additional Data Types in SQL2 and SQL-99 (cont.)

### **\*** INTERVAL:

Specifies a relative value rather than an absolute value

\* Can be DAY/TIME intervals or YEAR/MONTH intervals

\* Can be positive or negative when added to or subtracted from an absolute value, the result is an absolute value

## Retrieval Queries in SQL

- \* SQL has one basic statement for retrieving information from a database; the SELECT statement
- \* This is not the same as the SELECT operation of the relational algebra
- Important distinction between SQL and the formal relational model; SQL allows a table (relation) to have two or more tuples that are identical in all their attribute values
- # Hence, an SQL relation (table) is a multi-set (sometimes called a bag) of tuples; it is not a set of tuples
- \* SQL relations can be constrained to be sets by specifying PRIMARY KEY or UNIQUE attributes, or by using the DISTINCT option in a query
# Retrieval Queries in SQL (cont.)

- Basic form of the SQL SELECT statement is called a mapping or a SELECT-FROM-WHERE block
  - SELECT <attribute list>
  - FROM
  - WHERE <condition>
- « attribute list> is a list of attribute names whose values are to be retrieved by the query
  »
- \* is a list of the relation names required to process the query
- \* <condition> is a conditional (Boolean) expression that identifies the tuples to be retrieved by the query

FNAME	MINIT	LNAME	<u>SSN</u>	BDATE	ADDRES	S SEX	SALARY	SUPERSSN	DNC
		D	EPARTMEN	л					
		DNA		UMBER	MGRSSN	MGRS	STARTDATE		
			Г	DEPT_		;			
				DNUMBER	DLOC	ATION			
		r	PROJEC	ст 					
			PNAME	PNUMBE	R PLOCA		DNUM		
				wo	ORKS_ON				
				ESSN	PNO HO	URS			
DEPENDENT									

EMPLOYEE	FNAME	MINIT	LNAME	SSN	BDATE	ADDRESS	SEX	SALARY	SUPERSSN	DNO
	John	В	Smith	123456789	1965-01-09	731 Fondren, Houston, TX	M	30000	333445555	5
	Franklin	Т	Wong	333445555	1955-12-08	638 Voss, Houston, TX	M	40000	888665555	5
	Alicia	J	Zelaya	999887777	1968-07-19	3321 Castle, Spring, TX	F	25000	987654321	4
	Jennifer	S	Wallace	987654321	1941-06-20	291 Berry, Bellaire, TX	F	43000	888665555	4
	Ramesh	к	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
										-

#### DEPT\_LOCATIONS DNUMBER DLOCATION

1	Houston
4	Stafford
5	Bellaire
5	Sugarland
5	Houston

DEPARTMENT	DNAME	DNUMBER	MGRSSN	MGRSTARTDATE
	Research	5	333445555	1988-05-22
	Administration	4	987654321	1995-01-01
	Headquarters	1	888665555	1981-06-19

WORKS_ON	ESSN	PNO	HOURS
6	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
	999687777	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

### Simple SQL Queries

- Basic SQL queries correspond to using the SELECT, PROJECT, and JOIN operations of the relational algebra
- \* All subsequent examples use the COMPANY database
- \* Example of a simple query on one relation
- \* Query 0: Retrieve the birthdate and address of the employee whose name is 'John B. Smith'.
  - © Q0: SELECT BDATE, ADDRESS
    FROM EMPLOYEE
    WHERE FNAME='John' AND MINIT='B'
    AND LNAME='Smith'
    - Similar to a SELECT-PROJECT pair of relational algebra operations; the SELECT-clause specifies the projection attributes and the WHERE-clause specifies the selection condition
    - \* However, the result of the query may contain duplicate tuples

# Simple SQL Queries (cont.)

- \* Query 1: Retrieve the name and address of all employees who work for the 'Research' department.
  - \* Q1: SELECT FNAME, LNAME, ADDRESS FROM EMPLOYEE, DEPARTMENT WHERE DNAME='Research' AND DNUMBER=DNO
  - \* Similar to a SELECT-PROJECT-JOIN sequence of relational algebra operations
  - (DNAME='Research') is a selection condition (corresponds to a SELECT operation in relational algebra)
  - (DNUMBER=DNO) is a join condition (corresponds to a JOIN operation in relational algebra)

# Simple SQL Queries (cont.)

- Query 2: For every project located in 'Stafford', list the project number, the controlling department number, and the department manager's last name, address, and birthdate.
  - \* Q2: SELECT PNUMBER, DNUM, LNAME, BDATE, ADDRESS FROM PROJECT, DEPARTMENT, EMPLOYEE WHERE DNUM=DNUMBER AND MGRSSN=SSN AND PLOCATION='Stafford'
  - \* In Q2, there are two join conditions
  - \* The join condition DNUM=DNUMBER relates a project to its controlling department
  - \* The join condition MGRSSN=SSN relates the controlling department to the employee who manages that department

### Aliases, \* and DISTINCT, Empty WHERE-clause

\* In SQL, we can use the same name for two (or more) attributes as long as the attributes are in different relations

A query that refers to two or more attributes with the same name must qualify the attribute name with the relation name by prefixing the relation name to the attribute name

**\*** Example:

#### **\*** EMPLOYEE.LNAME, DEPARTMENT.DNAME

#### ALIASES

- Some queries need to refer to the same relation twice
- \* In this case, aliases are given to the relation name
- \* Query 8: For each employee, retrieve the employee's name, and the name of his or her immediate supervisor.
  - Q8: SELECT E.FNAME, E.LNAME, S.FNAME, S.LNAME FROM EMPLOYEE E S WHERE E.SUPERSSN=S.SSN
  - In Q8, the alternate relation names E and S are called aliases or tuple variables for the EMPLOYEE relation
  - \* We can think of E and S as two different copies of EMPLOYEE; E represents employees in role of supervisees and S represents employees in role of supervisors

ALIASES (cont.)

\* Aliasing can also be used in any SQL query for convenience Can also use the AS keyword to specify aliases

Q8: SELECT E.FNAME, E.LNAME, S.FNAME, S.LNAME FROM EMPLOYEE AS E, EMPLOYEE AS S WHERE E.SUPERSSN=S.SSN

### UNSPECIFIED WHERE-clause

- \* A missing WHERE-clause indicates no condition; hence, all tuples of the relations in the FROM-clause are selected
- \* This is equivalent to the condition WHERE TRUE
- \* Query 9: Retrieve the SSN values for all employees.
  - \* Q9: SELECT SSN FROM EMPLOYEE
- If more than one relation is specified in the FROM-clause and there is no join condition, then the CARTESIAN PRODUCT of tuples is selected

### UNSPECIFIED WHERE-clause (cont.)

\* Example:

Q10: SELECT SSN, DNAME FROM EMPLOYEE, DEPARTMENT

It is extremely important not to overlook specifying any selection and join conditions in the WHEREclause; otherwise, incorrect and very large relations may result

#### USE OF \*

\* To retrieve all the attribute values of the selected tuples, a \* is used, which stands for all the attributes Examples:

\* Q1C: SELECT \* FROMEMPLOYEE WHERE DNO=5

Q1D: SELECT \* FROMEMPLOYEE, DEPARTMENT WHERE DNAME='Research' AND DNO=DNUMBER

### USE OF DISTINCT

- \* SQL does not treat a relation as a set; duplicate tuples can appear
- \* To eliminate duplicate tuples in a query result, the keyword DISTINCT is used
- \* For example, the result of Q11 may have duplicate SALARY values whereas Q11A does not have any duplicate values

<b>*</b> Q11:	SELECT	SALARY
	FROM	EMPLOYEE
Q11A:	SELECT	DISTINCT SALARY
	FROM	EMPLOYEE

#### SET OPERATIONS

\* SQL has directly incorporated some set operations

- \* There is a union operation (UNION), and in some versions of SQL there are set difference (MINUS) and intersection (INTERSECT) operations
- \* The resulting relations of these set operations are sets of tuples; duplicate tuples are eliminated from the result
- \* The set operations apply only to union compatible relations; the two relations must have the same attributes and the attributes must appear in the same order

### SET OPERATIONS (cont.)

- \* Query 4: Make a list of all project numbers for projects that involve an employee whose last name is 'Smith' as a worker or as a manager of the department that controls the project.
  - Q4: (SELECT PNAME FROM PROJECT, DEPARTMENT, EMPLOYEE WHERE DNUM=DNUMBER AND MGRSSN=SSN AND LNAME='Smith') UNION (SELECT PNAME FROM PROJECT, WORKS\_ON, EMPLOYEE WHERE PNUMBER=PNO AND ESSN=SSN AND LNAME='Smith')

### NESTING OF QUERIES

- \* A complete SELECT query, called a nested query, can be specified within the WHERE-clause of another query, called the outer query
- \* Many of the previous queries can be specified in an alternative form using nesting
- \* Query 1: Retrieve the name and address of all employees who work for the 'Research' department.

Q1:	SELECT	FNAME, LNAME, ADDRESS
	FROM	EMPLOYEE
	WHERE	DNO IN (SELECT DNUMBER
	FROM	DEPARTMENT
	WHERE	DNAME='Research')

### NESTING OF QUERIES (cont.)

- \* The nested query selects the number of the 'Research' department
- \* The outer query select an EMPLOYEE tuple if its DNO value is in the result of either nested query
- \* The comparison operator IN compares a value v with a set (or multiset) of values V, and evaluates to TRUE if v is one of the elements in V
- \* In general, we can have several levels of nested queries
- \* A reference to an unqualified attribute refers to the relation declared in the innermost nested query
- \* In this example, the nested query is not correlated with the outer query

### CORRELATED NESTED QUERIES

- If a condition in the WHERE-clause of a nested query references an attribute of a relation declared in the outer query, the two queries are said to be correlated
- \* The result of a correlated nested query is different for each tuple (or combination of tuples) of the relation(s) in the outer query
- \* Query 12: Retrieve the name of each employee who has a dependent with the same first name as the employee.

Q12: SELECT	E.FNAME, E	.LNAME	
FROM	EMPLOYEE	ASE	
WHERE	E.SSN IN	(SELECT	ESSN
		FROM	DEPENDENT
		WHERE	ESSN=E.SSN AND
			E.FNAME=DEPENDENT_NAME)

### CORRELATED NESTED QUERIES (cont.)

- \* In Q12, the nested query has a different result for each tuple in the outer query
- \* A query written with nested SELECT... FROM... WHERE... blocks and using the = or IN comparison operators can always be expressed as a single block query. For example, Q12 may be written as in Q12A

Q12A: SELECT E.FNAME, E.LNAME FROM EMPLOYEE E, DEPENDENT D WHERE E.SSN=D.ESSN AND E.FNAME=D.DEPENDENT\_NAME

### THE EXISTS FUNCTION

- \* EXISTS is used to check whether the result of a correlated nested query is empty (contains no tuples) or not
- \* We can formulate Query 12 in an alternative form that uses EXISTS as Q12B below
- \* Query 12: Retrieve the name of each employee who has a dependent with the same first name as the employee.

Q12B: SELECT FNAME, LNAME FROM EMPLOYEE WHERE EXISTS (SELECT \* FROM DEPENDENT WHERE SSN=ESSN AND FNAME=DEPENDENT\_NAME)

# THE EXISTS FUNCTION (cont.)

\* Query 6: Retrieve the names of employees who have no dependents.

Q6: SELECT FNAME, LNAME FROM EMPLOYEE WHERE NOT EXISTS (SELECT \* FROM DEPENDENT WHERE SSN=ESSN)

In Q6, the correlated nested query retrieves all DEPENDENT tuples related to an EMPLOYEE tuple. If none exist, the EMPLOYEE tuple is selected

\* EXISTS is necessary for the expressive power of SQL

#### EXPLICIT SETS

It is also possible to use an explicit (enumerated) set of values in the WHERE-clause rather than a nested query

\* Query 13: Retrieve the social security numbers of all employees who work on project number 1, 2, or 3.

\* Q13: SELECT DISTINCT ESSN FROM WORKS\_ON WHERE PNO IN (1, 2, 3)

### NULLS IN SQL QUERIES

- \* SQL allows queries that check if a value is NULL (missing or undefined or not applicable)
- \* SQL uses IS or IS NOT to compare NULLs because it considers each NULL value distinct from other NULL values, so equality comparison is not appropriate.

 Query 14: Retrieve the names of all employees who do not have supervisors.
 Q14: SELECT FNAME, LNAME FROM EMPLOYEE
 WHERESUPERSSN IS NULL

Note: If a join condition is specified, tuples with NULL values for the join attributes are not included in the result

### Joined Relations Feature in SQL2

\* Can specify a "joined relation" in the FROM-clause

\* Looks like any other relation but is the result of a join

\* Allows the user to specify different types of joins (regular "theta" JOIN, NATURAL JOIN, LEFT OUTER JOIN, RIGHT OUTER JOIN, CROSS JOIN, etc)

### Joined Relations Feature in SQL2 (cont.)

\* Examples:

Q8: SELECT E.FNAME, E.LNAME, S.FNAME, S.LNAME FROM EMPLOYEE E S WHERE E.SUPERSSN=S.SSN

can be written as:

Q8: SELECT E.FNAME, E.LNAME, S.FNAME, S.LNAME FROM (EMPLOYEE E LEFT OUTER JOIN EMPLOYEES ON E.SUPERSSN=S.SSN)

### Joined Relations Feature in SQL2 (cont.)

\* Q1: SELECT FNAME, LNAME, ADDRESS
 FROM EMPLOYEE, DEPARTMENT
 WHERE DNAME='Research' AND DNUMBER=DNO

could be written as:

Q1: SELECT FNAME, LNAME, ADDRESS FROM (EMPLOYEE JOIN DEPARTMENT ON DNUMBER=DNO) WHERE DNAME='Research'

or as:

Q1: SELECT FNAME, LNAME, ADDRESS FROM (EMPLOYEE NATURAL JOIN DEPARTMENT AS DEPT(DNAME, DNO, MSSN, MSDATE) WHERE DNAME='Research'

### Joined Relations Feature in SQL2 (cont.)

\* Another Example;

- \* Q2 could be written as follows; this illustrates multiple joins in the joined tables
  - Q2: SELECT PNUMBER, DNUM, LNAME, BDATE, ADDRESS FROM (PROJECT JOIN DEPARTMENT ON DNUM=DNUMBER) JOIN EMPLOYEE ON MGRSSN=SSN WHERE PLOCATION='Stafford'

### AGGREGATE FUNCTIONS

\* Include COUNT, SUM, MAX, MIN, and AVG

\* Query 15: Find the maximum salary, the minimum salary, and the average salary among all employees.

Q15: SELECT MAX(SALARY), MIN(SALARY), AVG(SALARY) FROM EMPLOYEE

\* Some SQL implementations may not allow more than one function in the SELECT-clause

### AGGREGATE FUNCTIONS (cont.)

\* Query 16: Find the maximum salary, the minimum salary, and the average salary among employees who work for the 'Research' department.

Q16: SELECT MAX(SALARY), MIN(SALARY), AVG(SALARY) FROM EMPLOYEE, DEPARTMENT WHERE DNO=DNUMBER AND DNAME='Research'

### AGGREGATE FUNCTIONS (cont.)

\* Queries 17 and 18: Retrieve the total number of employees in the company (Q17), and the number of employees in the 'Research' department (Q18).

- Q17: SELECT COUNT (\*) FROM EMPLOYEE
- Q18: SELECT COUNT (\*) FROM EMPLOYEE, DEPARTMENT WHERE DNO=DNUMBER AND DNAME='Research'

#### GROUPING

In many cases, we want to apply the aggregate functions to subgroups of tuples in a relation

- \* Each subgroup of tuples consists of the set of tuples that have the same value for the grouping attribute(s)
- \* The function is applied to each subgroup independently
- \* SQL has a GROUP BY-clause for specifying the grouping attributes, which must also appear in the SELECT-clause

### GROUPING (cont.)

Query 20: For each department, retrieve the department number, the number of employees in the department, and their average salary.

Q20: SELECT DNO, COUNT (\*), AVG (SALARY) FROM EMPLOYEE GROUP BY DNO

- In Q20, the EMPLOYEE tuples are divided into groups--each group having the same value for the grouping attribute DNO
- \* The COUNT and AVG functions are applied to each such group of tuples separately
- \* The SELECT-clause includes only the grouping attribute and the functions to be applied on each group of tuples
- \* A join condition can be used in conjunction with grouping

### GROUPING (cont.)

\* Query 21: For each project, retrieve the project number, project name, and the number of employees who work on that project.

Q21: SELECT PNUMBER, PNAME, COUNT (\*) FROM PROJECT, WORKS\_ON WHERE PNUMBER=PNO GROUP BY PNUMBER, PNAME

\* In this case, the grouping and functions are applied after the joining of the two relations

#### THE HAVING-CLAUSE

Sometimes we want to retrieve the values of these functions for only those groups that satisfy certain conditions

\* The HAVING-clause is used for specifying a selection condition on groups (rather than on individual tuples)

### THE HAVING-CLAUSE (cont.)

\* Query 22: For each project on which more than two employees work, retrieve the project number, project name, and the number of employees who work on that project.

Q22: SELECT PNUMBER, PNAME, COUNT(\*) FROM PROJECT, WORKS\_ON WHERE PNUMBER=PNO GROUP BY PNUMBER, PNAME HAVING COUNT (\*) > 2

### SUBSTRING COMPARISON

\* The LIKE comparison operator is used to compare partial strings

\* Two reserved characters are used: '%' (or '\*' in some implementations) replaces an arbitrary number of characters, and '\_' replaces a single arbitrary character
### SUBSTRING COMPARISON (cont.)

\* Query 25: Retrieve all employees whose address is in Houston, Texas. Here, the value of the ADDRESS attribute must contain the substring 'Houston,TX'.

Q25: SELECT FNAME, LNAME FROM EMPLOYEE WHERE ADDRESS LIKE '%Houston,TX%'

### SUBSTRING COMPARISON (cont.)

\* Query 26: Retrieve all employees who were born during the 1950s. Here, '5' must be the 8th character of the string (according to our format for date), so the BDATE value is '\_\_\_\_5\_', with each underscore as a place holder for a single arbitrary character.

Q26: SELECT FNAME, LNAME FROM EMPLOYEE WHERE BDATE LIKE '\_\_\_\_\_

5'

#### ARITHMETIC OPERATIONS

- \* The standard arithmetic operators '+', '-'. '\*', and '/' (for addition, subtraction, multiplication, and division, respectively) can be applied to numeric values in an SQL query result
- \* Query 27: Show the effect of giving all employees who work on the 'ProductX' project a 10% raise.

Q27: SELECT FNAME, LNAME, 1.1\*SALARY FROM EMPLOYEE, WORKS\_ON, PROJECT WHERE SSN=ESSN AND PNO=PNUMBER AND PNAME='ProductX'

#### ORDER BY

- \* The ORDER BY clause is used to sort the tuples in a query result based on the values of some attribute(s)
- \* Query 28: Retrieve a list of employees and the projects each works in, ordered by the employee's department, and within each department ordered alphabetically by employee last name.
  - Q28: SELECT DNAME, LNAME, FNAME, PNAME FROM DEPARTMENT, EMPLOYEE, WORKS\_ON, PROJECT WHERE DNUMBER=DNO AND SSN=ESSN AND PNO=PNUMBER ORDER BY DNAME, LNAME

# ORDER BY (cont.)

\* The default order is in ascending order of values

\* We can specify the keyword DESC if we want a descending order; the keyword ASC can be used to explicitly specify ascending order, even though it is the default

#### Summary of SQL Queries

\* A query in SQL can consist of up to six clauses, but only the first two, SELECT and FROM, are mandatory. The clauses are specified in the following order:

SELECT<attribute list>FROM[WHERE<condition>][GROUP BY<grouping attribute(s)>][HAVING<group condition>][ORDER BY<attribute list>]

### Summary of SQL Queries (cont.)

- \* The SELECT-clause lists the attributes or functions to be retrieved
- \* The FROM-clause specifies all relations (or aliases) needed in the query but not those needed in nested queries
- \* The WHERE-clause specifies the conditions for selection and join of tuples from the relations specified in the FROM-clause
- # GROUP BY specifies grouping attributes
- # HAVING specifies a condition for selection of groups
- \* ORDER BY specifies an order for displaying the result of a query
- \* A query is evaluated by first applying the WHERE-clause, then GROUP BY and HAVING, and finally the SELECT-clause

# Specifying Updates in SQL

\* There are three SQL commands to modify the database; INSERT, DELETE, and UPDATE

#### INSERT

\* In its simplest form, it is used to add one or more tuples to a relation

\* Attribute values should be listed in the same order as the attributes were specified in the CREATE TABLE command

#### \* Example:

- U1: INSERT INTO EMPLOYEE VALUES ('Richard','K','Marini', '653298653', '30-DEC-52', '98 Oak Forest,Katy,TX', 'M', 37000,'987654321', 4)
- \* An alternate form of INSERT specifies explicitly the attribute names that correspond to the values in the new tuple
- \* Attributes with NULL values can be left out
- \* Example: Insert a tuple for a new EMPLOYEE for whom we only know the FNAME, LNAME, and SSN attributes.
  - U1A: INSERT INTO EMPLOYEE (FNAME, LNAME, SSN) VALUES ('Richard', 'Marini', '653298653')

Important Note: Only the constraints specified in the DDL commands are automatically enforced by the DBMS when updates are applied to the database

\* Another variation of INSERT allows insertion of multiple tuples resulting from a query into a relation

Example: Suppose we want to create a temporary table that has the name, number of employees, and total salaries for each department. A table DEPTS\_INFO is created by U3A, and is loaded with the summary information retrieved from the database by the query in U3B.

U3A: CREATE TABLE DEPTS\_INFO ( DEPT\_NAME VARCHAR(10), NO\_OF\_EMPS INTEGER, TOTAL\_SAL INTEGER);

U3B: INSERT INTO DEPTS\_INFO (DEPT\_NAME, NO\_OF\_EMPS, TOTAL\_SAL) SELECT DNAME, COUNT (\*), SUM (SALARY) FROM DEPARTMENT, EMPLOYEE WHERE DNUMBER=DNO GROUP BY DNAME ;

\* Note: The DEPTS\_INFO table may not be up-to-date if we change the tuples in either the DEPARTMENT or the EMPLOYEE relations after issuing U3B. We have to create a view (see later) to keep such a table up to date.

#### DELETE

Removes tuples from a relation

Includes a WHERE-clause to select the tuples to be deleted

\* Tuples are deleted from only one table at a time (unless CASCADE is specified on a referential integrity constraint)

\* A missing WHERE-clause specifies that all tuples in the relation are to be deleted; the table then becomes an empty table

\* The number of tuples deleted depends on the number of tuples in the relation that satisfy the WHERE-clause

\* Referential integrity should be enforced

### DELETE (cont.)

Examples:
 U4A: DELETE FROM EMPLOYEE
 WHERE LNAME='Brown'

U4B:DELETE FROM WHERE EMPLOYEE SSN='123456789'

U4C:DELETE FROM WHERE EMPLOYEE DNO IN (SELECT FROM WHERE

DNUMBER DEPARTMENT DNAME='Research')

U4D:DELETE FROM EMPLOYEE

#### UPDATE

\* Used to modify attribute values of one or more selected tuples

\* A WHERE-clause selects the tuples to be modified

\* An additional SET-clause specifies the attributes to be modified and their new values

\* Each command modifies tuples in the same relation

\* Referential integrity should be enforced

## UPDATE (cont.)

\* Example: Change the location and controlling department number of project number 10 to 'Bellaire' and 5, respectively.

U5: UPDATE PROJECT SET PLOCATION = 'Bellaire', DNUM = 5 WHERE PNUMBER=10

# UPDATE (cont.)

Example: Give all employees in the 'Research' department a 10% raise in salary.

U6: UPDATE EMPLOYEE SET SALARY = SALARY \*1.1 WHERE DNO IN (SELECT DNUMBER FROM DEPARTMENT WHERE DNAME='Research')

- In this request, the modified SALARY value depends on the original SALARY value in each tuple
- \* The reference to the SALARY attribute on the right of = refers to the old SALARY value before modification
- \* The reference to the SALARY attribute on the left of = refers to the new SALARY value after modification

#### **DATABASE IN PRACTICE**

#### Examples

\* A suite for distributed execution of general purpose applications on the Grid for small- and mid-size VOs

\* Usage and performance analysis of Dirac databases (just a few comments)

#### Submission system

#### Aim

- \* Develop a flexible and dynamic model to manage (simulation) data production on the Grid for smalland medium VOs
- \* Allow a easy, quick and customizable access to the Grid Infrastructure to research groups/organizations
- Successfully tested with SuperB Full- and Fast Simulation Production

### Distributed Architecture and Infrastructure

- \* The LHC Computing Grid (LCG) architecture was adopted to provide the minimum set of services and applications upon which the distributed production system has been built.
- \* Authentication and authorization is provided by VOMS service, LFC is the file catalog, WMS is used for brokering purpose and for Grid flavor interoperability features, transfers are done via Lcg-Utility, GANGA is the submitting interface.



# System Design

- \* The simulation production system heavily relies on a bookkeeping database, storing both applicationspecific and infrastructure metadata, which is tightly coupled with a Web-based user-interface (WebUI).
- \* The first makes available to the users information on the execution status of jobs and their specific meaning and parameters, and contributes in orchestrating the submission mechanism.
- \* The latter provides job submission management for several simulation applications, bookkeeping database interactions and basic monitoring functionalities.



# System Design

\* The bookkeeping database is implemented with PostgreSQL rDBMS in a centralized way, the WebUI in PHP and JQuery.

\* The database interactions with the submission portal and the job in execution on the WNs are managed by a direct interface to PostgreSQL or a RESTful interface (with X509 proxy-certificate cipher-encryption auths), respectively.

#### Job Workflow

- \* The structure of services and job workflow follow a semi-centralized design: job management service, bookkeeping database and default storage repository are hosted in a central site.
- \* Jobs executed into remote sites update the bookkeeping database with status, logging and timing information and transfer their output back to central repository or to a predefined site, discriminating on execution metadata.
- \* The system requires a proper configuration of the remote Grid sites.
- Bookkeeping metadata are integrated with Grid Logging & Bookkeeping service (LB) information provided by the infrastructure.
- In addition, the submission mechanism takes into account sites availability data from Nagios monitoring service.
- \* Simulation jobs are also exposed to the Grid dashboard for monitoring.

#### Job Workflow



207

#### WebUl

- \* The WebUI provides separate management for several type of simulation productions. Both sections are divided in configuration, submission and a monitor subsections. Their content is dynamically generated from the bookkeeping database schema and state in order to include the simulationspecific fields.
- \* A production cycle consists of several requests (defined by a specific set of job parameters values and events), which in turn are divided in several submissions, each consisting of several jobs.
- \* A configuration interface for requests definition per production cycle and simulation type, is provided
- Multi-site submissions based on requests and fine grain parametric submission interfaces complete the set of available services permitting a shift based scheduled session and a debugging specific console, respectively.

#### Sessions and status



#### Database schema (simplified ER/Relational schema)



#### Database structure



R INFO:	YOUR SESSIONS' STATIST	105:							
H Enrice Vianello!	session job launch	hed job-graph			events generated	testal wet (+)			
	FastSim 4100				5454079	4412912.80			
to-sub: @ enabled	Puttin 253				253	120790.90			
are member of \$									
up(s): webui									_
webuiFastProdMgr	YOUR LAST 10 10 AC	INTERS:							
webuilfuliPredHgr	Show 10 4 entries						Search		
WebUilbeveloper	time v	type	session	info					
or privileges for	2011-11-27 22:26:15	3085U6M25530N	FastSim			6 jobs (6 events) or	2010_July	25	-
create productions	2011-11-24 11:52:49	JOBSUEMISSION	FastSim			10 jobs (300000 e	vents) on 2010_July_test	25	-
create requests shifter ui	2011-11-22 18:02:40	JOBSUBHISSION	FastSim			3 jobs (3 events) or	2010_July_best	10	
expert ui	2011-11-22 16:35:15	JOBSUEMISSION	FulSim			7 jobs (7 events) or	2011_Pul_grid_test_test	1	-
monitoring jobs or privileges for	2011-11-22 12:01:11	JORGUEMISSION	FastSim			15 jobs (750000 e	vents) on 2010_July_test	23	
tiles:	2011-11-18 09:32:01	X085UBHISSION	FastSim			50 jobs (50 events)	on 2010_July	10	-
create requests	2011-11-16 12:27:21	X085UBMISSION	FullSim			1 jobs (1 events) or	2011_Full_prid_test_test	15	
shifter ui expert ui	2011-11-16 12:24:23	JORSURMISSION	FastSim			1 jobs (1 events) or	2010_July_test	25	-
monitoring jobs	2011-11-16 11:52:31	2085UBMISSION	FulSim			Ljobs (Levents) or	2011_Full_grid_test_test	10	
	3011-11-10 13-03-33	VORT RATE TON	F.Min			2 into (2 metals) or	2011 Full arid test test	10	



	_	
	_	

ew 10 🗘	entries		hide/sh	now columns				Searchs	
			4030 jobs for	und matching the	e search criteria.				
winnum v	events	Analysis	Generator	Geometry	Background	Status	bit D	Site	wet (s
t0002949 ovianelio	1	HadRecolCocktail	B0B0bar_Btag-HD_Cocktail	DG_4	MixSuperbillikg_NoPair	done	gjid	CYFROMET-LCG2	289.40
t0002948 evianelio	1	HadRecolCocktail	8080bar_8tag-HD_Cocktail	DG_4	MixSuperbBkg_NoPair	done	gjid	CYFROMET-LCG2	736.30
t0002947 evianello	1	HadRecolCocktall	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFROMET-LOG2	997.80
tevianelie	1	HadRecolCocktall	8080bar_SL	DG_4	PacProduction	done	gjid	CYFROMET-LOG2	589.90
10002945 evianelio	1	HadRecolCocktal	8080bar_SL	DG_4	PacProduction	done	gjid	CYFROMET-LOG2	655.00
LOOO2944 evianelio	1	HadRecolCocktal	8080bar_SL	DG_4	PacProduction	done	gjid	CYFROMET-LOG2	737.00
L0002943 evianello	1	HadRecolCocktal	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYPROMET-LOG2	1206.3
LOOD2942 evianelis	1	HadRecolCocktal	8080bar_SL	DG_4	PacProduction	done	gjid	CYFROMET-LOG2	897.00
10002941 palvani	1	BloKNuNu	8+8Btag-HD_Cocktal	DG_0	MixSuperbillikg	failed	gjid	INTR-BARL	5.00
10002940	1	BISKNuNu	B+BBlag-HD_Cocktal	DG_0	MixSuperbilling	done	stid	INFIN-BARD	89.10

table generated in '0000:00:02.65' (destroyed in: 0000:00:00.00)

#### Productions

#### [FastSim]

show test productions				-		
now 10 entries				Search:		
2010 July	Status	V0.2.5 311	Job Stats	2 Requests (0 open, 0 valid, 2 running, 0 closed)	close	note
2010 July test	0	V0.2.5 311		7 Requests		
2010 September 307	•	V0.2.5 307		56 Requests (0 open, 0 valid, 0 running, 56 close	d) <u>open</u>	
2010 September 307 test	0	V0.2.5 307		2 Requests		
2010 September 311	•	V0.2.5 311		43 Requests (0 open, 0 valid, 0 running, 43 close	d) open	
2010 September 311 test	0	V0.2.5 311	and the second se	18 Requests		
TSIM CREATE PRODUC	TION:			GUIDE		
orm fields are required.				<ul> <li>Type in the new Prod</li> <li>Specify the release v</li> </ul>	uction Series. ersion.	
Production Series:	2011_N	ovember		<ul> <li>The system will check series. On success, it</li> </ul>	of for uniquenes	s of the
Software Release:	V0.2.5	311	•	into the bookkeeping	database.	

refres	sh all every 10 6 seconds						SUBRIT ON A	IL SEUCTED	SHE
8	Nama	Last Submission	Last Submission Peni - But - Durs - Pal	[107]	Site Load Prog - Sub - But	[101]	Threahold Peni - Fai - Bate	Status	**
0	017_H07_01 X			(8)	0 - 0 - 0	- 10	20-12-300	(100/100)	-
9	CYTRONET-LOBE #*	2012-13-15 12:00:58 extendio 115.16-07-58		[14]		(8)	20-9-92	(10,10)	-
e.	967 y	2011-11-18 09-32-01 evianelle 112:08:14-7	0 × 0 × N × 0	[19]		- 10	20-12-511	(35/30)	- •
0	GROBU, UNDAN X			10		(1)	2-10-30	0000	- 4
e.	382#3-0E **	2012-12-15 12:00:18 eventsis 110:10:47:50		(194)			10-10-700	0	-
8	18299-0888 y'	2011-11-15 11:00:18 eviewile [10:16:47:57]		(5)		10	1-3-8	e level	-
e.	INFR-BARL V	2011-11-15 11-00-18 evieneto 115.16-47.00		[14]		10	30-10-300	(Revise)	-
b	INFN-FERRARA X			10		10	13-18-300	Pag	-
8	Dere-Col-3 y	2011-11-15 11:00:18 evaneta 113:16:47:20		[10]		.10	20-12-300	(10/10)	-
e	IMPR-RELANO-RTLABE #*	2013-13-12 13:43:53 existencio (38:36:4:38)	8 - 0 - 100 - 0	[100]		(4)	20-10-200	9 [100/100]	-
D	INTR-RAPOLI-ATLAS ¥			10		- 10	10-10-310	(20,00)	-
•	INTR-PERIODA **	2013-11-27 22-26-15 evidentite		14		.00	3-3-6	9	
ø	INTROPILIE of	2012-11-15 12-00-18 evianelle (15-16-47-56)	9 - 0 - 130 - 1	(1991)		10	10-10-1008	(100/100)	-
e	19879-71 V	2011-11-13 11:00-18 evientite (15:16:47:00)		[100]	1 - 2 - 4	10	100-10-900	(200-200)	-
e	INTR-TORINO #*	2011-11-15 11:00:18 evianeta (10:10:47:00)		[14]		- 191	30-10-300	-	-
	8AL-L003 ¥			11		.01	10-10-800	().00/10/0	-
	UK2-LT2-QMUL //	2011-11-19 11-00-18 evianeto (15-16-47-36)	0 - 0 -100 - 0	[100]	8 - 100 - 0	(100)	10-10-910	(1)-100	
e	010-501715833-0X-88# y/	2011-11-15 11:00:18 events (18:18-47-98)		[100]		(8)	10-10-300	0.00/1000	-
e	UKD-BOUTHERED-RALPP y	2012-10-15 10:00:00 evaluation	1 - 0 - 10 - 0	[10]		10	30-20-530	(10:101)	-
-							1.24		

#### FASTSIM CREATE NEW REQUEST

roduction: 2010_July	reques	t: [HadRecoilCocktail, B0B0b	ar_Btag-HD_Cockt	tail, DG_4, MixSuperbBkg_NoPair]	import
ob Parameters:					
nalysis:	Generator:	Geometry:		Background:	
HadRecoilCocktail	B0B0bar_Btag-HD_Cock	tail 🔹 DG_0	•	All Bkg, No Pair (MixSuperb)	
A Production Request with	the selected parameters doesn	't exists and will be created			
Request settings:					
riority: 10	minrunnum:	required events:	evt/job:	number of jobs:	
	10000000	10	1	10	
xpected event time (sec):	max wct (sec):				
2	57600				
Input:					
nut moder 0					
NONE					
Output & Log:					
utput directory (Ifn):					0
fn:/grid/superbvo.org/produ	ction/FastSim/2010_July/output	/HadRecoilCocktail/DG_0/B0B	Obar_Btag-HD_Cod	cktail/%RUNNUM%/	
utput directory (absolute):					0
/production/FastSim/2010_Ju	ily/output/HadRecoilCocktail/DC	_0/B0B0bar_Btag-HD_Cocktail	/%RUNNUM%/		
g file:					0
/production/FastSim/2010_Ju	ly/output/HadRecoilCocktail/DC	_0/B0B0bar_Btag-HD_Cocktail	/log/run%RUNNU!	M%.log	
Edit output directories for est request's output director	the associated test request y (Ifn):				0
fn:/grid/superbvo.org/produ	ction/FastSim/2010_July_test/ou	tput/HadRecoilCocktail/DG_0	/8080bar_8tag-HD	_Cocktail/%RUNNUM%/	
est request's output director	y (absolute):				0
/production/FastSim/2010_Ju	ily_test/output/HadRecoilCockta	I/DG_0/8080bar_Btag-HD_Co	ktail/%RUNNUM%	78	
est log file:					0
/production/FastSim/2010_Ju	ily_test/output/HadRecoilCockta	I/DG_0/8080bar_Btag-HD_Co	ktail/log/run%RU	NNUM%.log	
utputs' sites:	CIT_HEP_CE	CYFRONET-LCG2	GRIF	GRISU_UNINA	
utputs' sites: CIT_CMS_T2B		IN2P3-LPSC	INFN-BARI	INFN-CAGLIARI	
utputs' sites: CIT_CMS_T28 IN2P3-CC	IN2P3-IRES		INFN-MILAN	IO-ATLASC INFN-NAPOLI-AT	LAS
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA	IN2P3-IRES	INFN-LNL-2	Charles		
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERUGIA	IN2P3-IRES INFN-FRASCATI INFN-PISA	□ INFN-LNL-2 ✓ INFN-T1	INFN-TORIN	NO RAL-LCG2	
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERUGIA UKI-LT2-QMUL	IN2P3-IRES INFN-FRASCATI INFN-PISA UKI-SOUTHGRID-OX-HEP	INFN-LNL-2 INFN-T1 UKI-SOUTHGRID-RALPP	UNINA-EGE	e CICTORIA-LCG2	
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERUGIA UKI-LT2-QMUL VICTORIA-LCG2-SL4	IN2P3-IRES INFN-FRASCATI INFN-PISA UKI-SOUTHGRID-OX-HEP WT2	INFN-LNL-2 INFN-T1 UKI-SOUTHGRID-RALPP	UNINA-EGE	IO RAL-LCG2 E VICTORIA-LCG2	
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERUGIA UKI-LT2-QMUL VICTORIA-LCG2-SL4 Pther:	IN2P3-IRES INFN-FRASCATI INFN-PISA UKI-SOUTHGRID-OX-HEP WT2	□ INFN-LNL-2 ☑ INFN-T1 □ UKI-SOUTHGRID-RALPP	UNFN-TORIN	IO RAL-LCG2 E VICTORIA-LCG2	
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERUGIA UKI-LT2-QMUL VICTORIA-LCG2-SL4 Otheri	IN2P3-IRES INFN-FRASCATI INFN-PISA UKI-SOUTHGRID-OX-HEP WT2	□ INFN-LNL-2 ☑ INFN-T1 □ UKI-SOUTHGRID-RALPP	UNINA-EGE	NO CRAL-LCG2 E VICTORIA-LCG2	
utputs' sites: CIT_CMS_T2B IN2P3-CC INFN-FERRARA INFN-PERLIGIA UKI-LT2-QMUL VICTORIA-LCG2-SL4 Otheri ote: (max. 255 characters)	IN2P3-IRES INFN-FRASCATI INFN-PISA UKI-SOUTHGRID-OX-HEP WT2	□ INFN-LNL-2 ☑ INFN-T1 □ UKI-SOUTHGRID-RALPP	UNINA-EGE	NO CRAL-LCG2 E VICTORIA-LCG2	

215

#### JOB DETAILS (REQUESTS)

w 1	0 📫 entries				Search:	
s	Job Parameters	Priority 🔻	Graph Pending – Running – Failed – Done [TOT]	events per job	jobs d/nf/r	events d/nf/i
Θ	PRODSCRIPT : HadRecoilCocktail GENERATOR : BOBObar_SL DG : DG_4 TCL : PacProduction	70	0 - 0 - 65 - 2381 [2446]	1	2 381 (6.33%) 2 381 (6.33%) 37 600	2 381 (6.33%) 2 381 (6.33%) 37 60(
0	PRODSCRIPT : HadRecoilCocktail GENERATOR : B0B0bar_Btag-HD_Cocktail DG : DG_4 TCL : MixSuperbBkg_NoPair	50	0 - 0 - 0 - 5 [5]	80000	5 (0.11%) 5 (0.11%) 4 700	400 000 (0.11% 400 000 (0.11% 376 000 00
# in use...

refres	sh all every 10 Seconds						SUBMIT ON A	LL SELECTED	-1814	3
1	Name	Last Submission	Last Submission		Site Load		Threshold	Status		eti
-			Pend - Run - Done - Fail	[TOT]	Prep - Sub - Run	[101]	Pend - Feil - Runn			1
	CTT_HEP_CE X		0 - 0 - 0 - 0	(0)	0 - 0 - 0	[0]	20-10-300	[100/100]	-+	
3	CYFRONET-LCG2	2011-11-15 11:00:18 evianello [15:16:47:50]	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	20-5-50	(\$0/\$0]	-+	
ø	GRIF 🖌	2011-11-18 09:32:01 evianello {12:18:16:7}	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	30-10-551	(\$0/\$0]	-•	
	GRISU_UNINA X		0 - 0 - 0 - 0	(0)	0 - 0 - 0	[0]	2-10-30	(10/10)	-+	
9	1м2р3-сс 🏑	2011-11-15 11:00:18 evianello (15:16:47:50)	0 - 0 -100 - 0	(100)	0 - 0 - 0	[0]	50-10-700	(100/100)	-+	
9	IN2P3-IRES W	2011-11-15 11:00:18 evianello (15:16:47:50)	0 - 0 - 5 - 0	[5]	0 - 0 - 0	[0]	5-3-5	(\$/5]	-+	
3	INFN-BARI 🖌	2011-11-15 11:00:18 evianello [15:16:47:50]	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	30-10-300	(so/so)	-	
	INFN-FERRARA X		0 - 0 - 0 - 0	(0)	0 - 0 - 0	[0]	12-10-300	(3/3]	-+	
J	INFN-LNL-2 🖌	2011-11-15 11:00:18 evianello [15:16:47:50]	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	30-10-500	(\$0/\$0]	-•	
ø	INFN-MILANO-ATLASC 💅	2011-11-12 11:43:55 evianello (18:16:4:13)	0 - 0 -100 - 0	(100)	0 - 0 - 0	[0]	20-10-300	(100/100)	-	
9	INFN-NAPOLI-ATLAS X		0 - 0 - 0 - 0	(0)	0 - 0 - 0	[0]	30-10-300	(so/so)	-+	
	INFN-PERUGIA 🖌	2011-11-27 22:26:15 evianello (3:5:22:2)	0 - 0 - 0 - 6	(6)	0 - 0 - 0	[0]	3-3-6	(6/6)		
J	INFN-PISA 🗸	2011-11-15 11:00:18 evianello (15:16:47:59)	0 - 0 -100 - 0	(100)	0 - 0 - 0	[0]	50-10-1000	(100/100)	-+	
J	INFN-T1 🖌	2011-11-15 11:00:18 evianello [15:16:47:59]	0 - 0 -100 - 0	(100)	0 - 3 - 0	[3]	100-10-900	(200/200)	-+	
a	INFN-TORINO 🖌	2011-11-15 11:00:18 evianello [15:16:47:59]	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	30-10-300	(\$0/\$0]	-•	
	RAL-LOG2 X		0 - 0 - 0 - 0	(0)	0 - 0 - 0	[0]	30-10-800	[100/100]	-+	
0	UKI-LT2-QMUL	2011-11-15 11:00:18 evianello (15:16:47:59)	0 - 0 -100 - 0	(100)	0 - 100 - 0	[100]	50-10-950	(0/100)		
ø	UKI-SOUTHGRID-OX-HEP y	2011-11-15 11:00:18 evianello (15:16:47:59)	0 - 0 - 99 - 1	[100]	0 - 0 - 0	[0]	20-10-300	(100/100)	-+	
3	UKI-SOUTHGRED-RALPP	2011-11-15 11:00:18 evianello (15:16:47:59)	0 - 0 - 50 - 0	[50]	0 - 0 - 0	[0]	30-10-550	(so/so)	-+	
							1.1.4			

217

#### in use...

#### Automatic submission ou

#### Submission status per site:

/data1/script_webul/FastSim/submissionscripts/2010_July/CYFRONET-LCG2_10001643_10001692.php	[32.89 s] log
/data1/script_webui/FastSim/submissionscripts/2010_July/IN2P3-CC_10001693_10001792.php	 [52.508 s] log
/data1/script_webui/FastSim/submissionscripts/2010_July/IN2P3-IRES_10001793_10001797.php	C
/data1/script_webui/FastSim/submissionscripts/2010_July/INFN-BARI_10001798_10001847.php	0
/data1/script_webui/FastSim/submissionscripts/2010_July/INFN-LNL-2_10001848_10001897.php	0
/data1/script_webui/FastSim/submissionscripts/2010_July/INFN-PI5A_10001898_10001997.php	O
/data1/script_webui/FastSim/submissionscripts/2010_July/INFN-T1_10001998_10002097.php	0
/data1/script_webui/FastSim/submissionscripts/2010_July/INFN-TORINO_10002098_10002147.php	0
/data1/script_webui/FastSim/submissionscripts/2010_July/UKI-LT2-QHUL_10002148_10002247.php	0
/data1/script_webui/FastSim/submissionscripts/2010_July/UKI-SOUTHGRID-OX-HEP_10002248_10002347.php	O
/data1/script_webui/FastSim/submissionscripts/2010_July/UKI-SOUTHGRID-RALPP_10002348_10002397.php	O

2010\_July

PR	ODUCTION DATA	- SUMMARY	
	Production Series:	2010_July	
	Software:	V0.2.5 311	
	Test release:	lfn:/grid/superbvo.org/production/FastSim/2010_September/test_release/V0.2.5_SL- 5.3_x86_64_rev311.tjz	

Request Para	meters:					
PRO	DISCRIPT		GENERATOR	DG	TCL	
HadRecoilCocktail		В	0B0bar_Btag-HD_Cocktail	DG_4	MixSuperbBkg_NoPair	
Jobs' Type	Events per job	Log on db	Target site(s)			
normal	80000	no	INFN-T1			
#Jobs	Site		First RUNNUM	Last RUNNUM	#Events	
50	CYFRONET-LCG2		10002454	10002503	4000000	
50	GRIF		10002504	10002553	4000000	
100	IN2P3-CC		10002554	10002653	8000000	
5	IN2P3-IRES		10002654	10002658	400000	
50	INFN-BARI		10002659	10002708	4000000	
50	INFN-LNL-2		10002709	10002758	4000000	
100	INFN-MILANO-ATLA	SC	10002759	10002858	8000000	
100	INFN-PISA		10002859	10002958	8000000	
200	INFN-T1		10002959	10003158	16000000	
50	INFN-TORINO		10003159	10003208	4000000	
100	UKI-SOUTHGRID-O	X-HEP	10003209	10003308	8000000	
50	UKI-SOUTHGRID-R	ALPP	10003309	10003358	4000000	
905					72400000	

Production Request existal nput mode: DIR input path: Ifn:/grid/superbyo.org/production/FastSim/2010\_September/input output directory's name: /production/FastSim/2010\_July/output/HadRecollCocktail/DG\_4/B0B0bar\_Btag-HD\_Cocktail/%RUNNUH% output directory's lifn: Ifn:/grid/superbyo.org/production/FastSim/2010\_July/output/HadRecollCocktail/DG\_4/B0B0bar\_Btag-HD\_Cocktail/%RUNNUH% output directory's lifn: Ifn:/grid/superbyo.org/production/FastSim/2010\_July/output/HadRecollCocktail/DG\_4/B0B0bar\_Btag-HD\_Cocktail/%RUNNUH%

production/FestSim/2010\_July/output/HadRecollCocktail/DG\_4/B0B0bar\_Btag-HD\_Cocktail/log/run%RUNNUM%.it

(Production Initialization and Script Generation ) M Automatic Submission

#### in use...

#### FASTSIM JOB LIST

w 10	entries		hide/sh	now columns =			5	Search:	
4030 jobs found matching the search criteria.									
unnum 🔫	events	Analysis	Generator	Geometry	Background	Status	GJid	Site	wct (s)
vianello	1	HadRecollCocktail	B0B0bar_Btag-HD_Cocktail	DG_4	MixSuperbBkg_NoPair	done	gjid	CYFRONET-LCG2	289.40
0002948 vianello	1	HadRecoilCocktail	B0B0bar_Btag-HD_Cocktail	DG_4	MixSuperbBkg_NoPair	done	gjid	CYFRONET-LCG2	736.30
0002947 vianello	1	HadRecollCocktall	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	997.80
vianello	1	HadRecollCocktall	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	589.90
vianelio	1	HadRecollCocktall	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	655.00
vianello	1	HadRecollCocktail	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	737.00
0002943 vianello	1	HadRecollCocktail	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	1206.30
0002942 vianello	1	HadRecollCocktail	B0B0bar_SL	DG_4	PacProduction	done	gjid	CYFRONET-LCG2	897.00
0002941 alvani	1	BtoKNuNu	B+BBtag-HD_Cocktail	DG_0	MixSuperbBkg	failed	gjid	INFN-BARI	5.00
0002940 anzali	1	BtoKNuNu	B+BBtag-HD_Cocktail	DG_0	MixSuperbBkg	done	gjid	INFN-BARI	89.10

table generated in '0000:00:02.65' (destroyed in: 0000:00:00.00)

🗹 runnum production 📃 uid 🗹 events 🗹 Analysis Generator 🗹 Geometry 🗹 Background 🗹 Status 📃 Status Reason ErrorMsg 🗹 GJid 🗹 Site 📃 ts prepared 📃 ts submitted 📃 ts running 📃 time 🗹 wct (s)

#### Results



# Issues with "production software"

## MySQL

MySQL is one of the most used 'open-source' rDBMS

You may already know it has two main DB engines:

#### **\* MyISAM**

These tables have a small footprint. <u>Table-level locking</u> limits the performance in read/write workloads, so it is often used in read-only or read-mostly workloads in Web and data warehousing configurations.

#### # InnoDB

A transaction-safe (ACID compliant) storage engine for MySQL that has commit, rollback, and crash-recovery capabilities to protect user data. InnoDB row-level locking (without escalation to coarser granularity locks) and Oraclestyle consistent nonlocking reads increase multi-user concurrency and performance. InnoDB stores user data in clustered indexes to reduce I/O for common queries based on primary keys. To maintain data integrity, InnoDB also supports FOREIGN KEY referential-integrity constraints.

### MySQL

- To be able to use transactions + take advantage of referential integrity, InnoDB must be used
- If we use DBMS, we should trust them!(I understand, sometimes we are paranoic):

 If we defined a table as CREATE TABLE T (ID INT PRIMARY KEY, ...); why should we continuously do SELECT ID FROM T WHERE ID=<some value>; (to check for existence) before doing UPDATE T SET ... WHERE ID=<some value>;? or INSERT INTO T2 (..., IDfk, ...) VALUES (..., <some value>,...); where IDfk is defined as Foreign Key referencing ID in T?

## DIRAC Transformation System



Name: Transformations Engine: InnoDB Version: 10 Row format: Compact Rows: 102398 Avg row length: 6303 Data length: 645447680 Max data length: 0 Index length: 4767744 Data free: 4194304 Auto increment: 24618 Create time: 2010-05-18 13:24:05 Update time: NULL Check time: NULL Collation: latin1 swedish ci Checksum: NULL Create options: Comment:

Name: TransformationTasks Engine: MyISAM Version: 10 Row format: Fixed Rows: 3244338 Avg row length: 316 Data length: 1218232456 Max data length: 88946092640567295 Index length: 109328384 Data free: 193021648 Auto increment: 1 Create time: 2010-05-27 12:23:07 Update time: 2013-06-04 18:34:44 Check time: 2010-05-27 12:23:09 Collation: latin1 swedish ci Checksum: NULL Create options: Comment:

Name: TransformationFiles Engine: MyISAM Version: 10 Row format: Dynamic Rows: 16451371 Avg row length: 91 Data length: 1511429528 Max data length: 281474976710655 Index length: 549894144 Data free: 0 Auto increment: NULL Create time: 2010-05-27 12:23:03 Update time: 2013-06-04 18:35:33 Check time: 2010-05-27 12:23:07 Collation: latin1 swedish ci Checksum: NULL Create options: Comment:

Name: TransformationFileTasks Engine: MyISAM Version: 10 Row format: Fixed Rows: 63564637 Avg row length: 13 Data length: 826340281 Max data length: 3659174697238527 Index length: 1405467648 Data free: 0 Auto increment: NULL Create time: 2010-08-09 20:53:22 Update time: 2013-06-04 18:35:50 Check time: NULL Collation: latin1 swedish ci Checksum: NULL Create options: Comment:

Name: DataFiles Engine: MyISAM Version: 10 Row format: Dynamic Rows: 50722343 Avg row length: 90 Data length: 4581560000 Max data length: 281474976710655 Index length: 7016133632 Data free: 0 Auto increment: 50722344 Create time: 2010-02-15 22:55:18 Update time: 2013-06-04 18:35:10 Check time: NULL Collation: latin1 swedish ci Checksum: NULL Create options: Comment:

#### Clean-up

- \* Assuming all table InnoDB
- \* All FK defined
- \* To delete finished tasks (+files+...)
  - \* is there some order in the DELETE operations?
  - \* may we use the CASCADE OPTION?



#### Slow queries

Some query may require long time to execute...needs for indexes

\* modify db design or application logic?

\* Log analysis is crucial to figure out the problems

## Slow queries

# # # #	Files Overa Time Attri	s: 22_slow.1 all: 2.57M t range: 2013 ibute	log total, 17 3-04-16 1 total	6 unique, 8:31:40 to 	206.73 5 21:59: max	QPS, 0. :14 avg	98x conc 95%	currenc	cy lev mec	 lian
# # #	Exec Lock	time time	12224s 6342s	1us 0	114s 96s	5ms 2ms	 316us 44us	s 386 s 345	oms 1	.6us 0
# # #	Rows Rows Query	sent examine y size	11.49M 754.03M 159.38M	1 0 1 0 1 6	12.46k 10.82M 79.15k	4.68 307.09 64.91	0.99 0.99 212.52	9 94. 9 20.9 2 385.	10 90k 80 26	0 0 5.08
# # # # <b>#</b> # # #	Profi Rank ==== 1 2 3 4 5 6	ile Query ID ====================================	2FB888A8 2A574963 2DCA4EA4 092EB3A8 3E26019C .76F3EDA	Response 6400.6257 878.2083 804.9874 642.1562 479.3516 423.3719	time ( 52.4% 7.2% 6.6% 5.3% 3.9% 3.5%	Calls 11330 30592 9 9741 3632 29471	R/Call 0.5649 0.0287 89.4430 0.0659 0.1320 0.0144	V/M ===== 27.07 40.62 <b>8.12</b> 27.56 42.20 0.06	Item SELECT UPDATE SELECT SELECT SELECT SELECT	TransformationFiles TransformationFiles <b>TransformationFiles</b> TransformationFiles TransformationFiles TransformationFiles
# # #	7 8 9	0x12A6600A7 0x2B2651709 0x7E3A04D71	7E70CB2E 93495FE9 EF9BF1B	312.4006 292.0841 255.9953	2.6% 2.4% 2.1%	1411 1304 217	0.2214 0.2240 1.1797	0.00 0.01 75.80	SELECT SELECT SELECT	Transformations Transformations TransformationFiles

. . .

#### Slow queries

Filter with WHERE

+- Bookmark lookup
+- Table
| table TransformationFiles
| possible\_keys Status
+- Index range scan
 key TransformationFiles->Status
 possible\_keys Status
 key\_len 35
 rows 11334650

# Query 3: 0.03 QPS, 2.55x concurrency, ID 0xAFA91D2D2DCA4EA4 at byte 313302282 # This item is included in the report because it matches --limit. # Scores: V/M = 8.12 # Time range: 2013-04-16 20:26:12 to 20:31:28 # Attribute pct total 95% stddev median min max avq # =========== ==== =========== \_\_\_\_\_ \_\_\_\_ # Count. 0 9 # Exec time 6 805s 51s 114s 89s 113s 27s 102s 2 157s # Lock time 44us 53s 17s 52s 24s 134ms # Rows sent 0 56 1 6.24 3.89 18 6.22 17.65 # Rows examine 12 97.38M 10.82M 10.82M 10.82M 10.76M 0 10.76M # Query size 0 2.53k 225 429 287.67 420.77 72.04 258.32 # String: # Databases ProductionDB volhcb19.c... (4/44%), volhcb29.c... (4/44%), volhcb24.c... (1/11%) # Hosts # Users Dirac # Query time distribution 10ms # 100ms 15 # Tables SHOW TABLE STATUS FROM `ProductionDB` LIKE 'TransformationFiles'\G SHOW CREATE TABLE `ProductionDB`. `TransformationFiles`\G # EXPLAIN /\*!50100 PARTITIONS\*/ SELECT TransformationID, FileID, Status, TaskID, TargetSE, UsedSE, ErrorCount, LastUpdate, InsertedTime, RunNumber FROM TransformationFiles WHERE `Status` IN ( "Processed", "Problematic" ) AND `FileID` IN ("29643136", "29643865", "35351500", "29826552", "35351597") LIMIT 1000\G

#### Conclusions

\* Software for Physics is generally written by physicists (also Experiment's Core Software)

- \* so... if you need to interface your software to or base your software on databases: you must 'learn' databases, especially database design
- Interaction with experts is recommended for both db design and programming techniques

#### Credits

\* Elmasri and Navathe Fundamentals of Database Systems, 5/6th Ed. Pearson, Addison Wesley

- \* E. Vianello, A. Fella, E. Luppi
- \* DIRAC and LHCbDIRAC developers

#### Notice 2

# I left out:

\* Relational Calculus

\* Functional Dependencies and Normalization

\* Transactions, Concurrency, Recovery, Physical level, ...

\* all the emerging non-relational databases (NoSQL), such as MongoDB, CouchDB, Hadoop-based, ...