



Introduction to SQL

Main Textbook



 Database Systems: The Complete Book
 Hector Garcia-Molina
 Jeffrey D. Ullman
 Jennifer Widom



Alternative Textbook



- Database Management
 Systems
 - Raghu Ramakrishnan
 - Johannes Gehrke





Goals of Course

- To obtain a firm background in database systems, e.g.,
 - how to talk to database systems in a standard language
 - how to improve the efficiency of database systems
 - the theories behind database design and some algorithms behind database implementation
- Mostly "basic stuff" about databases



What will NOT be taught

- Advanced database technologies
 - Geographical information systems
 - Data mining
 - ••••
 - (This is an introductory course only)
- Specific instructions on how to install and use a specific database system on a specific platform
 - Try the user manual or Google

Teaching Style



- There will be a lot of in-lecture exercises
- Questions will be welcomed
- Lecture notes will be released on the Drive (at least several days before lectures)



Course Overview

- What is a database?
- A large collection of data organized especially for rapid search and retrieval (as by a computer)
- What is a database system? (more formally, a database management system, i.e., DBMS)
- A management system that helps us retrieve information from databases

Database and DBMS





Tables, Relations, Relational Model



Tables, Relations, Relational Model











Database Schema Design





Database Schema Design

Taxpayer_ID	Annual_Income
51248297	100000
33891634	50000
67904777	70000

 Assume that we want to capture parent-child relationships

SULEYMAN DEMIREL UN IV FRSITY

Database Schema Design

Taxpayer_ID	Annual_Income	Child_I	D
51248297	100000		
33891634	50000		
67904777	70000		

Is one column enough?



Database Schema Design

Taxpayer_ID	Annual_Income	Child_ID1	Child_ID2
51248297	100000		
33891634	50000		
67904777	70000		

- Are two columns enough?
- Assume that two columns are enough
- Does everyone have two children?

Schema designs based on the Entity-Relationship model

Course Content

- SQL
- Constraints and Triggers
- Conceptual Design
- Indices
- Relation Algebra
- Query Processing/Optimization
- Concurrency Control
- Recovery
- Current trend (e.g., NOSQL)

Database Design

Database Implementation





What do you want from a DBMS?



Why do we need it?

- Keep data around (persistent)
- Answer queries (questions) about data
- Update data

Requirements from high-end applications

 Massive amounts of data (terabytes ~ petabytes)
 High throughput (thousands ~ millions
 transactions/min)

The Relational Revolution

The Relational Revolution (1970's)

- •IBM and Univ of Berkeley
- •A simple data model: Data is stored in relations (tables)
- A declarative query language: SQL
 - Programmer specifies what answers a query should return, but not how the query is executed
 - DBMS picks the best execution strategy
- Hide the physical organization of the database from applications

Provided only logical view of the data

Turing Award!

Edgar C Codd

Relational model is the dominating technology today
Graphs/Streams/Arrays are hot wanna-be!

revolu





"Relational databases are the foundation of western civilization."



Bruce Lindsay IBM Fellow IBM Almaden Research Center



Structured Query Language (SQL)



Structured Query Language (SQL)

- A declarative (computer) language for managing data in a relational database management system
- Two parts
 - Data Definition Language (DDL)
 - Create/Alter/Delete tables
 - Will be discussed in the next week
 - Data Manipulation Language (DML)
 - Query one or more tables
 - Insert/Delete/Modify tuples in tables
 - Will be discussed in the following





Data Types in SQL

Character strings

- CHAR(20)
- VARCHAR(50)

••••

Numbers

- INT
- FLOAT
- ••••

Others

- BOOLEAN
- DATETIME

••••

Product		
<u>PName</u>	Price	Category
iPhone 4	888	Phone
iPad 2	668	Tablet
Milestone	798	Phone
EOS 550D	1199	Camera



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon
SELECT * FROM Product WHERE Category = 'Phone'		Ĺ	"selection"	
	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	Milestone	798	Phone	Motorola



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon
			_	

SELECT	*
FROM	Product
WHERE	Category <> 'Phone'

<u>PName</u>	Price	Category	Manufacturer
iPad 2	668	Tablet	Apple
EOS 550D	1199	Camera	Canon



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT * FROM Product WHERE Category = 'Phone' AND Price > 800

<u>PName</u>	Price	Category	Manufacturer
iPhone 4	888	Phone	Apple



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT * FROM Product WHERE Category = 'Tablet' OR Price > 1000

<u>PName</u>	Price	Category	Manufacturer
iPad 2	668	Tablet	Apple
EOS 550D	1199	Camera	Canon



Simple SQL Query (cont.)

Product	<u>PName</u>	Price	Category	Manufacturer		
	iPhone 4	888	Phone	Apple		
	iPad 2	668	Tablet	Apple		
	Milestone	798	Phone	Motorola		
	EOS 550D	1199	Camera	Canon		
SELECT PName, Price, Manufacturer Selection						

FROM Product WHERE Price > 800

<u>PName</u>	Price	Manufacturer
iPhone 4	888	Apple
EOS 550D	1199	Canon

projection"

Details



- SQL is NOT case sensitive (when it comes to keywords and names)
 - SELECT = Select = select

Product = product

- Constants must use single quotes
 - 🗅 'abc' OK
 - "abc" NOT OK



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT	*	П
FROM	Product	\checkmark
WHERE	PName LIKE 'iPh%'	

	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
% stands for	"any string			



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT FROM WHERE	* Product PName LIKE '%Ph%'	Ĺ

	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
% stands for	"any string	,11		



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT	*	
FROM	Product	7
WHERE	PName LIKE '%P%e%'	

	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
% stands for	"any string	11		



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

FROM Product	
	7
WHERE PName LIKE '_Phone 4'	

	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
_stands for				



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT	*	
FROM	Product	マ
WHERE	PName LIKE '_Phone'	

	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
_stands for				



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon


Eliminating Duplicates

Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon
SELECT Cat			Category	





Eliminating Duplicates (cont.)

Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product





Ordering the Results

Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT	PName, Price
FROM	Product
WHERE	Price < 800
ORDER B	37 PName

<u>PName</u>	Price
Milestone	798
iPad 2	668



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT	PName, Price
FROM	Product
WHERE	Price < 800
ORDER E	BY PName DESC

<u>PName</u>	Price
iPad 2	668
Milestone	798



Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT PName, Category FROM Product WHERE Price < 1000 ORDER BY Category, PName





Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT PName, Category FROM Product WHERE Price < 1000 ORDER BY Category DESC, PName





Product	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

SELECT PName, Category FROM Product WHERE Price < 1000 ORDER BY Category DESC, PName DESC



Exercise		<u>PName</u>	Price	Category	Manufacturer
		iPhone 4	888	Phone	Apple
Product		iPad 2	668	Tablet	Apple
		Milestone	798	Phone	Motorola
		EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product ORDER BY Category



SULEYMAN DEMIREL



Exercise		<u>PName</u>	Price	Category	Manufacturer
		iPhone 4	888	Phone	Apple
		iPad 2	668	Tablet	Apple
	Product	Milestone	798	Phone	Motorola
		EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product ORDER BY Category WHERE Price < 1000



UNIVERSITY

Exercise		<u>PName</u>	Price	Category	Manufacturer
		iPhone 4	888	Phone	Apple
		iPad 2	668	Tablet	Apple
	Product	Milestone	798	Phone	Motorola
		EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product ORDER BY Category WHERE Price < 1000

Error!

 "WHERE" should always proceed "ORDER BY"

Exercise		<u>PName</u>	Price	Category	Manufacturer
		iPhone 4	888	Phone	Apple
		iPad 2	668	Tablet	Apple
	Product	Milestone	798	Phone	Motorola
		EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product ORDER BY PName



SULEYMAN DEMIREL

Exercise		<u>PName</u>	Price	Category	Manufacturer
		iPhone 4	888	Phone	Apple
		iPad 2	668	Tablet	Apple
	Product	Milestone	798	Phone	Motorola
		EOS 550D	1199	Camera	Canon

SELECT DISTINCT Category FROM Product ORDER BY PName

Error!

 "ORDER BY" items must appear in the select list if "SELECT DISTINCT" is specified

Joins				<u>CN</u>	<u>ame</u>	StockPric	e	Country
		Company		Ca	non	45		Japan
		compe	~~~~	Mot	orola	40		USA
	Product			Ар	ple	374		USA
	<u>PName</u>	Price	Cate	gory	Manu	facturer		
	iPhone 4	888	Pho	one	A	pple		
	iPad 2	668	Tab	olet	A	pple		
	Milestone	798	Pho	one	Mo	torola		
	EOS 550D	1199	Carr	nera	С	anon		

A user wants to know the names and prices of all products by Japan companies. How?

Taina			<u>CN</u>	<u>ame</u>	StockPric	e	Country
JOINS	Comp	Company		non	45		Japan
	Comp			orola	40		USA
Produc	:†		Ap	ple	374		USA
PName	e Price	Cate	gory	Manu	facturer		
iPhone	4 888	Pho	one	A	pple		
iPad 2	668	Tab	olet	A	pple		
Milesto	ne 798	Pho	one	Мо	torola		
EOS 550	DD 1199	Carr	nera	С	anon		

 SELECT PName, Price FROM Product, Company WHERE Country = 'Japan' AND Manufacturer = CName





 Find the names of the persons who work for companies in USA
 SELECT PName FROM Person, Company WHERE Country = 'USA' AND WorksFor = CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT PName, Address FROM Person, Company
 WHERE Country = 'USA' AND WorksFor = CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT PName, Company.Address FROM Person, Company WHERE Country = 'USA' AND WorksFor = CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT PName, Company.Address
 FROM Person, Company
 WHERE Country = 'USA'
 AND CName = CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT PName, Company.Address
 FROM Person, Company
 WHERE Country = 'USA'
 AND Person.CName = Company.CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT X.PName, Y.Address
 FROM Person AS X, Company AS Y
 WHERE Y.Country = 'USA'
 AND X.CName = Y.CName





- Find the names the persons who work for companies in USA, as well as their company addresses
- SELECT X.PName, Y.Address
 FROM Person X, Company Y
 WHERE Y.Country = 'USA'
 AND X.CName = Y.CName





- Exercise: Find the names of the companies in China that produce products in the 'tablet' category
- SELECT DISTINCT CName
 EPOM
 Company
 Product
 - FROM Company, Product WHERE Manufacturer = CName AND Country = 'China'
 - AND Country = 'China' AND Category = 'Tablet'







- Exercise: Find the names of the companies in China that produce products in the 'tablet' or 'phone' category
- SELECT DISTINCT CName

 - FROM Company, Product WHERE Manufacturer = CName AND Country = 'China' AND (Category = 'Tablet' OR Category = 'Phone')



	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
Product	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

- Exercise: Find the manufacturers that produce products in both the 'tablet' and 'phone' categories
 SELECT DISTINCT Manufacturer FROM Product
 - WHERE Category = 'Tablet' AND Category = 'Phone'

Error!



	<u>PName</u>	Price	Category	Manufacturer
	iPhone 4	888	Phone	Apple
Product	iPad 2	668	Tablet	Apple
	Milestone	798	Phone	Motorola
	EOS 550D	1199	Camera	Canon

- Exercise: Find the manufacturers that produce products in both the 'tablet' and 'phone' categories
- SELECT DISTINCT X.Manufacturer
 FROM Product AS X, Product AS Y
 WHERE X.Manufacturer = Y.Manufacturer
 AND X.Category = 'Tablet'
 AND Y.Category = 'Phone'

Subqueries



- A subquery is a SQL query nested inside a larger query
- Queries with subqueries are referred to as nested queries
- A subquery may occur in
 - SELECT
 - FROM
 - WHERE

SQL subquery
SQL subquery

A special subquery: Scalar Subquery



Scalar Subquery

- return a single value which is then used in a comparison.
- If query is written so that it expects a subquery to return a single value, and it returns multiple values or no values, a run-time error occurs.

Example Query

From Sells(bar, beer, price), find the bars that serve Heineken for the same price Ku De Ta charges for Bud.



Find the bars that serve Heineken at that price.





Example Scalar Subquery

```
SELECT bar
FROM Sells
WHERE beer = 'Heineken' AND
price = (SELECT price
FROM Sells
WHERE bar = 'Ku De Ta'
AND beer = 'Bud');
```



Subqueries in FROM



- Find all products in the 'phone' category with prices under 1000
- SELECT X.PName FROM (SELECT * FROM Product WHERE category = 'Phone') AS X WHERE X.Price < 1000



Subqueries in FROM (cont.)



- Find all products in the 'phone' category with prices under 1000
- SELECT PName FROM Product
 WHERE Category = 'Phone' AND Price < 1000
- This is a much more efficient solution





- Find all companies that make some products with price < 100
- SELECT DISTINCT CName FROM Company AS X WHERE X.CName IN (SELECT Y.CName FROM Product AS Y WHERE Y.Price < 100)





- Find all companies that make some products with price < 100
- SELECT DISTINCT CName FROM Company AS X WHERE EXISTS (SELECT * FROM Product AS Y WHERE X.CName = Y.Cname AND Y.Price < 100)
- •A nested query is **correlated** with the outer query if it contains a reference to an attribute in the outer query.

•A nested query is *correlated* with the outside query if it must be re-computed for every tuple produced by the outside query.





- Find all companies that make some products with price < 100
- SELECT DISTINCT CName FROM Company AS X WHERE X.CName IN (SELECT * FROM Product AS Y WHERE Y.Price < 100)

•The number of attributes in the SELECT clause in the subquery must match the number of attributes compared to with the comparison operator.





- Find all companies that make some products with price < 100
- SELECT DISTINCT CName EDOM Company AS X
 - FROM Company AS X WHERE 100 > ANY

(SELECT Price FROM Product AS Y WHERE X.CName = Y.Cname)




- Find all companies that make some products with price < 100
- SELECT DISTINCT CName
 FROM Product
 WHERE Price < 100
- This is more efficient than the previous solutions

Operators in Subqueries



IN

<tuple> IN <relation> is true if and only if the tuple is a member of the relation.

ANY

x = ANY(<relation>) is a
boolean cond. meaning
that x equals at least one
tuple in the relation.

EXISTS

- EXISTS(<relation>) is true if and only if the <relation> is not empty.
- Returns true if the nested query has 1 or more tuples.

ALL

x <> ALL(<relation>) is true if and only if for every tuple t in the relation, x is not equal to t.

Note

The keyword NOT can proceed any of the operators (s NOT IN R)



Avoiding Nested Queries

- In general, nested queries tend to be more inefficient than un-nested queries
 - query optimizers of DBMS do not generally do a good job at optimizing queries containing subqueries
- Therefore, they should be avoided whenever possible
- But there are cases where avoiding nested queries is hard...





- Find all companies that do not make any product with price < 100
- SELECT DISTINCT CName FROM Company AS X WHERE NOT EXISTS (SELECT * FROM Product AS Y WHERE X.CName = Y.Cname AND Y.Price < 100)





- Find all companies that do not make any product with price < 100</p>
- SELECT DISTINCT CName FROM Company AS X WHERE 100 <= ALL

(SELECT Price FROM Product AS Y WHERE X.CName = Y.Cname)





- Find all companies that does not make any products with price < 100</p>
- SELECT DISTINCT CName FROM Company AS X WHERE 100 <= ALL

(SELECT Price FROM Product AS Y WHERE X.CName = Y.Cname)



Beer

...

Drinker

...



Find all drinkers that frequent some bar that serves some beer they like
 SELECT DINSTINT F.Drinker FROM Likes AS L, Frequent AS F, Serve AS S
 WHERE L.Drinker = F.Drinker AND F.Bar = S.Bar AND L.Beer = S.Beer



Find all drinkers that frequent some bar that does not serve any beer they like
 SELECT DISTINCT F.Drinker
 FROM Frequent AS F, Serves AS S
 WHERE F.Bar = S.Bar AND NOT EXIST
 (SELECT *
 FROM Likes as L
 WHERE L.Beer = S.Beer
 AND L.Drinker = F.Drinker)





SULEYMAN DEMIREL UINIVERSITY

Roadmap --SQL

- Table
- SELECT FROM WHERE
- ORDER BY
- Joins
- Subqueries
- Aggregations
- UNION, INTERSECT, EXCEPT
- NULL
- Outerjoin
- Insert/Delete tuples
- Create/Alter/Delete tables
- View





Find all drinkers that do not frequent any bar that serve some beer they like
 SELECT DISTINCT F.Drinker
 FROM Frequent AS F
 WHERE NOT EXIST
 (SELECT *
 FROM Likes AS L, Serves AS S
 WHERE L.Beer = S.Beer
 AND L.Drinker = F.Drinker
 AND S.Bar = F.Bar)