

SQL (Advanced)

Database Management
COP4540, SCS, FIU

Advanced SQL

- View Definition
- SQL Transactions
- NULL Search Condition
- Sub-query with an Aggregate Function
- Use of ANY/SOME, ALL with Sub-query
- UNION, INTERSECT, MINUS Operations

Database Management
COP4540, SCS, FIU

View Definition

- A view is a derived relation and is defined based on base relations or other views.
- A view is used to store a query and to allow other queries to be defined on the view.
- Views can also help to provide data security and improve programming productivity.

Database Management
COP4540, SCS, FIU

View Definition

Employee(SSN, Name, Salary, DeptNo)
Dept(DeptNo, DeptName)

```
CREATE VIEW DeptInfo(DeptName, No_of_Emp, Total_Salary) AS  
SELECT DeptName, COUNT(*), SUM(Salary)  
FROM Employee AS E, Dept AS D  
WHERE E.DeptNo = D.DeptNo  
GROUP BY DeptName;
```

Query: List all the departments with more than ten employees

```
SELECT DeptName  
FROM DeptInfo  
WHERE No_of_Emp > 10;
```

Database Management
COP4540, SCS, FIU

SQL Transactions

- A SQL Transaction is a collection of SQL commands that are committed or aborted as a unit.
- Transaction management is an important concept for concurrency control in multi-user database systems.
- Transaction integrity means to maintain correctness of data in the database after executing each transaction.

Database Management
COP4540, SCS, FIU

SQL Transactions

```
BEGIN Transaction  
    INSERT ....  
    INSERT ...  
    ...  
END Transaction
```

Database Management
COP4540, SCS, FIU

Sub-query with an Aggregate Function

Publication(ISBN, Title, Price, Author, PubDate)

Query: Find all books cost more than the average price.

```
SELECT AVG(Price) FROM Publication;
```

```
SELECT Title, Price
FROM Publication
WHERE Price < (SELECT AVG(Price)
               FROM Publication);
```

Database Management
COP4540, SCS, FIU

Use of ANY/SOME with Sub-query

- **ANY/SOME** behaves as existential quantifiers.

Publication(ISBN, Title, Price, Author, PubDate)

Query: Find title and price of the books that cost more than at least one other book published after June 15, 1990.

```
SELECT Title, Price
FROM Publication
WHERE Price > ANY ( SELECT Price
                   FROM Publication
                   WHERE PubDate > '15-jun-1990');
```

Database Management
COP4540, SCS, FIU

Use of ALL with Sub-query

- **ALL** behaves as universal quantifiers

Publication(ISBN, Title, Price, Author, PubDate)

Query: Find title and price of the books that cost more than all the books published before June 15, 1985.

```
SELECT Title, Price
FROM Publication
WHERE Price > ALL ( SELECT Price
                  FROM Publication
                  WHERE PubDate < '15-jun-1985');
```

Database Management
COP4540, SCS, FIU

Union, intersection, and difference of Queries

- The keywords used are UNION, INTERSECT, and MINUS for \cup , \cap , and - respectively.
- Keywords like UNION are used between two queries.
- The the results of the two queries should be union compatible.
- All the three operations are SET operations by default.

Database Management
COP4540, SCS, FIU

Use of UNION Operation

Student(SSN, Name, DeptName)

Faculty(SSN, Name, Salary, DeptName)

Query: Find all the people with their name who either work or study in the CS department.

```
SELECT DISTINCT Name
FROM Student
WHERE DeptName = 'CS'
UNION
SELECT DISTINCT Name
FROM Faculty
WHERE DeptName = 'CS';
```

Database Management
COP4540, SCS, FIU

Nulls

- In place of a value in a tuple's component.
 - Regarding any arithmetic operator, like \times or $+$, the result is NULL.
 - When comparing a NULL value and any value, including another NULL, the result is UNKNOWN. The value UNKNOWN is another truth-value, like TRUE and FALSE.
 - The correct way to ask if x has the value NULL is with the expression x is NULL. Similarly, x IS NOT NULL has the value true unless the value of x is NULL

Database Management
COP4540, SCS, FIU

Three-valued logic

- In addition to *true* and *false*, there is a third value *unknown*.
- Think of *true* = 1; *false* = 0, and *unknown* = 1/2.
- Then consider logic operators:
 - AND = min.
 - OR = max.
 - NOT(x) = 1 - x

Database Management
COP4540, SCS, FIU

Applying three logical operators

x	y	x AND y	x OR y	NOT x
TRUE	TRUE	1	1	0
TRUE	UNKNOWN	1/2	1	0
TRUE	FALSE	0	1	0
UNKNOWN	TRUE	1/2	1	1/2
UNKNOWN	UNKNOWN	1/2	1/2	1/2
UNKNOWN	FALSE	0	1/2	1/2
FALSE	TRUE	0	1	1
FALSE	UNKNOWN	0	1/2	1
FALSE	FALSE	0	0	1

TRUE = 1; FALSE = 0; UNKNOWN = 1/2

Database Management
COP4540, SCS, FIU

SQLqueries with three-valued logic

- In the WHERE clauses of select-from-where statements, only the tuples for which the condition has the value TRUE become part of the answer; tuples with either UNKNOWN or FALSE as value are excluded from the answer.
- It is the same case regarding DELETE statements.

Database Management
COP4540, SCS, FIU

Query example

Attri ₁	Attri ₂
1	NULL
2	1
3	NULL
4	3
5	4
6	5

SELECT *
FROM R
WHERE Attri₂ >=2 **OR** Attri₂ < 2;

SELECT *
FROM R
WHERE Attri₂ **IS NOT NULL**;

Database Management
COP4540, SCS, FIU