#### Introduction to Database Systems CSE 414

Lecture 21: Views

CSE 414 - Spring 2013

#### Announcements

- Homework 6 is due Wednesday
- Today:
  - Views (just enough to know what they are)
  - Start transactions
- Next: transactions and serializibility

#### Views

- A view in SQL =
  - A table computed from other tables, s.t., whenever the base tables are updated, the view is also updated
- More generally:
  - A view is derived data that keeps track of changes in the original data
- Compare:
  - A function computes a value from other values, but does not keep track of changes to the inputs

StorePrice(store, price)

#### A Simple View

Create a view that returns for each store the prices of products purchased at that store

> CREATE VIEW StorePrice AS SELECT DISTINCT x.store, y.price FROM Purchase x, Product y WHERE x.product = y.pname

> > This is like a new table StorePrice(store,price)

### We Use a View Like Any Table

- A "high end" store is a store that sell some products over 1000.
- For each customer, return all the high end stores that they visit.

```
SELECT DISTINCT u.name, u.store
FROM Purchase u, StorePrice v
WHERE u.store = v.store
AND v.price > 1000
```

# Types of Views

#### <u>Virtual</u> views

- Used in databases
- Computed only on-demand slow at runtime
- Always up to date
- <u>Materialized</u> views
  - Used in data warehouses
  - Pre-computed offline fast at runtime
  - May have stale data (must recompute or update)
  - Indexes are materialized views

### **Query Modification**

For each customer, find all the high end stores that they visit.

CREATE VIEW StorePrice AS SELECT DISTINCT x.store, y.price FROM Purchase x, Product y WHERE x.product = y.pname

SELECT DISTINCT u.name, u.store FROM Purchase u, StorePrice v WHERE u.store = v.store AND v.price > 1000

#### StorePrice(store, price)



**Query Modification** 

CSE 414 - Spring 2013

#### **Query Modification**

For each customer, find all the high end stores that they visit.



### Further Virtual View Optimization

Retrieve all stores whose name contains ACME

CREATE VIEW StorePrice AS SELECT DISTINCT x.store, y.price FROM Purchase x, Product y WHERE x.product = y.pname

SELECT DISTINCT v.store FROM StorePrice v WHERE v.store like '%ACME%'

## Further Virtual View Optimization

Retrieve all stores whose name contains ACME

CREATE VIEW StorePrice AS SELECT DISTINCT x.store, y.price FROM Purchase x, Product y WHERE x.product = y.pname

SELECT DISTINCT v.store FROM StorePrice v WHERE v.store like '%ACME%' Modified query:

SELECT DISTINCT v.store FROM (SELECT DISTINCT x.store, y.price FROM Purchase x, Product y WHERE x.product = y.pname) v WHERE v.store like '%ACME%'

### Further Virtual View Optimization

Retrieve all stores whose name contains ACME



## Further Virtual View Optimization

Retrieve all stores whose name contains ACME

SELECT DISTINCT x.store FROM Purchase x, Product y WHERE x.product = y.pname —AND-x.store like '%ACME%'

Modified and unnested query:

Assuming Product.pname is a key <u>and</u> Purchase.product is a foreign key



**Final Query** 

SELECT DISTINCT x.store FROM Purchase x WHERE x.store like '%ACME%'

# **Applications of Virtual Views**

- Increased physical data independence. E.g.
  - Vertical data partitioning
  - Horizontal data partitioning
- Logical data independence. E.g.
  - Change schemas of base relations (i.e., stored tables)
- Security
  - View reveals only what the users are allowed to know

#### **Vertical Partitioning**

Resumes	<u>SSN</u>		Name		A	Address			Resume			Picture	
	234234		Mary		Hı	Huston		Clo	Clob1		Blob1		
	345345		Sue		Se	Seattle		Clob2		Blob2			
	345343		Joan		Se	Seattle (		Clo	Clob3		Blob3		
	432432		Ann		Po	Portland		Clob4		Blob4			
T1				7	<b>[2</b> —					Т3			
<u>SSN</u>	Name	Add	Address		<u>SSN</u>		Resume			SSN Pic		Picture	
234234	Mary	Huston			23423	34	Clob1			234234		Blob1	
345345	Sue	Seattle		4	345345		Clob2			34534	5	Blob2	

**T2**.SSN is a key <u>and</u> a foreign key to **T1**.SSN. Same for **T3**.SSN <sup>15</sup>

T1(<u>ssn</u>,name,address) T2(<u>ssn</u>,resume) T3(<u>ssn</u>,picture)

#### **Vertical Partitioning**

CREATE VIEW Resumes AS SELECT T1.ssn, T1.name, T1.address, T2.resume, T3.picture FROM T1,T2,T3 WHERE T1.ssn=T2.ssn AND T1.ssn=T3.ssn

T1(<u>ssn</u>,name,address) T2(<u>ssn</u>,resume) T3(<u>ssn</u>,picture)

#### **Vertical Partitioning**

CREATE VIEW Resumes AS SELECT T1.ssn, T1.name, T1.address, T2.resume, T3.picture FROM T1,T2,T3 WHERE T1.ssn=T2.ssn AND T1.ssn=T3.ssn

SELECT address FROM Resumes WHERE name = 'Sue'

T1(<u>ssn</u>,name,address) T2(<u>ssn</u>,resume) T3(<u>ssn</u>,picture)

#### **Vertical Partitioning**

CREATE VIEW Resumes AS SELECT T1.ssn, T1.name, T1.address, T2.resume, T3.picture FROM T1,T2,T3 WHERE T1.ssn=T2.ssn AND T1.ssn=T3.ssn



SELECT address FROM Resumes WHERE name = 'Sue'

Modified query:

SELECT T1.address FROM T1, T2, T3 WHERE T1.name = 'Sue' AND T1.SSN=T2.SSN AND T1.SSN = T3.SSN

T1(<u>ssn</u>,name,address) T2(<u>ssn</u>,resume) T3(<u>ssn</u>,picture)

**SELECT** address

#### Vertical Partitioning



Modified query:



### **Vertical Partitioning Applications**

#### 1. Advantages

- Speeds up queries that touch only a small fraction of columns
- Single column can be compressed effectively, reducing disk I/O

#### 2. Disadvantages

- Updates are expensive!
- Need many joins to access many columns
- Repeated key columns add overhead

Hot trend today for data analytics: e.g., Vertica startup acquired by HP They use a highly-tuned column-oriented data store AND engine

#### Horizontal Partitioning

k

#### Customers

SSN	Name	City			
234234	Mary	Houston			
345345	Sue	Seattle			
345343	Joan	Seattle			
234234	Ann	Portland			
	Frank	Calgary			
	Jean	Montreal			

#### **CustomersInHouston**

SSN	Name	City
234234	Mary	Houston
	-	

#### **CustomersInSeattle**

	SSN	Name	City	
_ /	345345	Sue	Seattle	
V	345343	Joan	Seattle	

. . . . .

. . . .

Customers(<u>ssn</u>,name,city)

### Horizontal Partitioning

CREATE VIEW Customers AS CustomersInHouston UNION ALL CustomersInSeattle UNION ALL

. . . . .

Customers(<u>ssn</u>,name,city)

### **Horizontal Partitioning**



Which tables are inspected by the system ?

. . . . .

Customers(<u>ssn</u>,name,city)

### Horizontal Partitioning



Which tables are inspected by the system ?

All tables! The systems doesn't know that CustomersInSeattle.city = 'Seattle'

CSE 414 - Spring 2013

Customers(<u>ssn</u>,name,city)

# Horizontal Partitioning

Better: remove CustomerInHuston.city etc

CREATE VIEW Customers AS (SELECT SSN, name, 'Houston' as city FROM CustomersInHouston) UNION ALL (SELECT SSN, name, 'Seattle' as city FROM CustomersInSeattle) UNION ALL

. . . .

Customers(<u>ssn</u>,name,city)

#### **Horizontal Partitioning**



# Horizontal Partitioning Applications

- Performance optimization
  - Especially for data warehousing
  - E.g. one partition per month
  - E.g. archived applications and active applications
- Distributed and parallel databases
- Data integration