

# Views

*April 6*

```
SELECT l.partkey
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
ORDER BY l.shipdate DESC
LIMIT 10;
```

```
SELECT l.partkey
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
ORDER BY l.shipdate DESC
LIMIT 10;
```

```
SELECT l.partkey, COUNT(*)
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
GROUP BY l.partkey;
```

```
SELECT l.supkey, COUNT(*)
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
GROUP BY l.supkey;
```

```
SELECT l.partkey
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
ORDER BY l.shipdate DESC
LIMIT 10;
```

**“orders since last month”**

```
SELECT l.partkey, COUNT(*)
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
GROUP BY l.partkey;
```

```
SELECT l.supkey, COUNT(*)
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
GROUP BY l.supkey;
```

```
CREATE VIEW salesSinceLastMonth AS
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')
```

```
CREATE VIEW salesSinceLastMonth AS
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')

  SELECT partkey FROM salesSinceLastMonth
  ORDER BY shipdate DESC LIMIT 10;
```

```
CREATE VIEW salesSinceLastMonth AS
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')
```

```
  SELECT partkey FROM salesSinceLastMonth
  ORDER BY shipdate DESC LIMIT 10;
```

```
    SELECT suppkey, COUNT(*)
    FROM salesSinceLastMonth
    GROUP BY suppkey;
```

```
CREATE VIEW salesSinceLastMonth AS
SELECT l.*
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
```

```
SELECT partkey FROM salesSinceLastMonth
ORDER BY shipdate DESC LIMIT 10;
```

```
SELECT suppkey, COUNT(*)
FROM salesSinceLastMonth
GROUP BY suppkey;
```

```
SELECT partkey, COUNT(*)
FROM salesSinceLastMonth
GROUP BY partkey;
```



```
CREATE VIEW salesSinceLastMonth AS
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')

  SELECT partkey FROM ordersSinceLastMonth
  ORDER BY shipdate DESC LIMIT 10;
```

```
CREATE VIEW salesSinceLastMonth AS
SELECT l.*
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
```

```
SELECT partkey FROM ordersSinceLastMonth
ORDER BY shipdate DESC LIMIT 10;
```

```
SELECT partkey FROM
(
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')
) AS salesSinceLastMonth
ORDER BY shipdate DESC LIMIT 10;
```

# Views

- ... contain and abstract complex concepts.
- Complex query patterns can be given a shorthand.
- It's easier to change view logic “in the background”
- ... act as normal relations.
- View references can be expanded inline into nested subqueries.
- Updates are trickier.....

# View Updates

What happens when we Insert Into/Update a view?

# View Updates

```
UPDATE salesSinceLastMonth  
    SET statusCode = 'q';  
WHERE orderkey = 22;
```

# View Updates

```
UPDATE salesSinceLastMonth
  SET statusCode = 'q';
WHERE orderkey = 22;
```

Rows in `salesSinceLastMonth` correspond 1-1 with rows in `lineitem`. Update `lineitem`!

# View Updates

```
INSERT INTO salesSinceLastMonth
  (orderkey, partkey, suppkey, ...)
VALUES
  (22, 99, 42, ...);
```

# View Updates

```
INSERT INTO salesSinceLastMonth
  (orderkey, partkey, suppkey, ...)
VALUES
  (22, 99, 42, ...);
```

Lots of problems...

- What if order # 22 doesn't exist?
- How does the insertion interact with sequences (e.g., `lineitem.lineno`)



# View Updates

# View Updates

**Solution 1:** Data Integration

# View Updates

**Solution 1:** Data Integration  
(CSE 636)

# View Updates

**Solution 1:** Data Integration  
(CSE 636)

**Solution 2:** INSTEAD OF triggers

# View Updates

```
CREATE TRIGGER salesSinceLastMonthInsert
INSTEAD OF INSERT ON salesSinceLastMonth
REFERENCING NEW ROW AS newRow
FOR EACH ROW
  IF NOT EXISTS (
    SELECT * FROM ORDERS
    WHERE ORDERS.orderkey = newRow.orderKey)
  ) THEN
    INSERT INTO ORDERS(orderkey)
      VALUES (orderkey)
  END IF;
  INSERT INTO LINEITEM VALUES newRow;
END FOR;
```

Can we use views for anything else?

# Materialization

Views exist to be queried frequently

Pre-compute and save the view's contents!  
(like an index)

# Materialization Challenges

- How do we maintain the views as data changes?
- What if the view is not explicitly referenced?
- What views should be materialized?



# Delta Queries

- If  $D$  is your Database and  $Q$  is your Query:
  - $Q(D)$  is the result of your query on the database.
- Let's say you make a change  $\Delta D$  (Insert tuple)
  - $Q(D+\Delta D)$  is the new result
- If we have  $Q(D)$ , can we get  $Q(D+\Delta D)$  faster?
  - Analogy to Sum:  $\{34, 29, 10, 15\} + \{12\}$  ( $88+12$ )

# Query Rewriting

```
CREATE MATERIALIZED VIEW salesSinceLastMonth AS
  SELECT l.*
  FROM lineitem l, orders o
  WHERE l.orderkey = o.orderkey
        AND o.orderdate > DATE('2015-03-31')
```

# Query Rewriting

```
CREATE MATERIALIZED VIEW salesSinceLastMonth AS
SELECT l.*
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')

SELECT l.partkey
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
ORDER BY l.shipdate DESC
LIMIT 10;
```

# Query Rewriting

```
CREATE MATERIALIZED VIEW salesSinceLastMonth AS
SELECT l.*
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')

SELECT l.partkey
FROM lineitem l, orders o
WHERE l.orderkey = o.orderkey
      AND o.orderdate > DATE('2015-03-31')
ORDER BY l.shipdate DESC
LIMIT 10;
```

**We can use a materialized view to speed the query up**

# Query Rewriting

View Query

```
SELECT Lv  
FROM Rv  
WHERE Cv
```

User Query

```
SELECT Lq  
FROM Rq  
WHERE Cq
```

When are we allowed to rewrite this query?

# Query Rewriting

View Query

```
SELECT LV  
FROM RV  
WHERE CV
```

User Query

```
SELECT LQ  
FROM RQ  
WHERE CQ
```

$$R_V \subseteq R_Q$$

All relations in the view are part of the query join

$$C_Q = C_V \wedge C'$$

The view condition is weaker than the query condition

$$\text{attrs}(C') \cap \text{attrs}(R_V) \subseteq L_V \quad L_Q \cap \text{attrs}(R_V) \subseteq L_V$$

The view doesn't project away needed attributes

# Query Rewriting

View Query

```
SELECT Lv  
FROM Rv  
WHERE Cv
```

User Query

```
SELECT Lq  
FROM Rq  
WHERE Cq
```

What does the query rewrite to?

# Query Rewriting

View Query

```
SELECT Lv  
FROM Rv  
WHERE Cv
```

User Query

```
SELECT Lq  
FROM Rq  
WHERE Cq
```


```
SELECT Lq  
FROM (Rq-Rv), VIEW  
WHERE Cq
```



# Materialized Views



# Materialized Views



# Materialized Views



# Materialized Views



**When the base data changes, the view needs to be updated**

# View Maintenance

**VIEW** ← **Q(D)**

# View Maintenance

```
WHEN D ← D + ΔD DO:  
    VIEW ← Q(D + ΔD)
```

# View Maintenance

```
WHEN D ← D+ΔD DO:  
  VIEW ← Q(D+ΔD)
```

**Re-evaluating the query from scratch is expensive!**

# View Maintenance

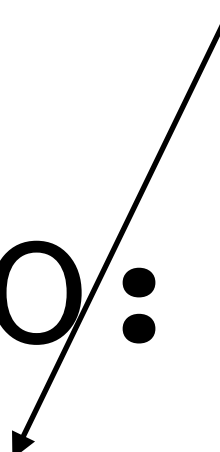
```
WHEN D ← D + ΔD DO:  
    VIEW ← VIEW + ΔQ̄(D, ΔD)
```



# View Maintenance

(ideally) Smaller & Faster Query

WHEN  $D \leftarrow D + \Delta D$  DO:  
VIEW  $\leftarrow \text{VIEW} + \Delta Q(D, \Delta D)$



# View Maintenance

(ideally) Smaller & Faster Query

WHEN  $D \leftarrow D + \Delta D$  DO:

VIEW  $\leftarrow \text{VIEW} + \Delta Q(D, \Delta D)$

(ideally) Fast “merge” operation.

# Intuition

$$D = \{1, 2, 3, 4\} \quad \Delta D = \{5\}$$

$$Q(D) = \text{SUM}(D)$$

# Intuition

$$D = \{1, 2, 3, 4\} \quad \Delta D = \{5\}$$

$$\underline{Q}(D) = \text{SUM}(D)$$

$$\underline{Q}(D + \Delta D) \sim O(|D| + |\Delta D|)$$

# Intuition

$$D = \{1, 2, 3, 4\} \quad \Delta D = \{5\}$$

$$\underline{Q}(D) = \text{SUM}(D)$$

$$\underline{Q}(D+\Delta D) \sim O(|D| + |\Delta D|)$$

$$\text{VIEW} + \text{SUM}(\Delta D) \sim O(|\Delta D|)$$

# Intuition

$$R = \{1, 2, 3\}, S = \{5, 6\} \quad \Delta R = \{4\}$$

$$Q(R, S) = \text{COUNT}(R \times S)$$

# Intuition

$$R = \{1, 2, 3\}, S = \{5, 6\} \quad \Delta R = \{4\}$$

$$Q(R, S) = \text{COUNT}(R \times S)$$

$$Q(R + \Delta R, S) \sim O( (|R| + |\Delta R|) * |S| )$$

# Intuition

$$R = \{1, 2, 3\}, S = \{5, 6\} \quad \Delta R = \{4\}$$

$$Q(R, S) = \text{COUNT}(R \times S)$$

$$Q(R + \Delta R, S) \sim O( (|R| + |\Delta R|) * |S| )$$

$$\text{VIEW} + \text{COUNT}(|\Delta R| * |S|) \sim O(|\Delta R| * |S|)$$



# Intuition

**+**   **~**   **U**

**\***   **~**   **X**

# Intuition

**+ ~ U**

**\* ~ X**

**Are these kinds of patterns common?**

# Rings/Semirings

This kind of pattern occurs frequently.

**Semiring** :  $\langle \mathbf{S}, +, \times, \mathbf{S}_0, \mathbf{S}_1 \rangle$

Any set of ‘things’  $\mathbf{S}$  such that...

Closed

$$\mathbf{S}_i + \mathbf{S}_j = \mathbf{S}_k$$

$$\mathbf{S}_i \times \mathbf{S}_j = \mathbf{S}_k$$

$$\mathbf{S}_i + \mathbf{S}_0 = \mathbf{S}_i$$

$$\mathbf{S}_i \times \mathbf{S}_1 = \mathbf{S}_i$$

$$\mathbf{S}_i \times \mathbf{S}_0 = \mathbf{S}_0$$

Additive &  
Multiplicative  
“zeroes”

$$\mathbf{S}_i \times (\mathbf{S}_j + \mathbf{S}_k) = (\mathbf{S}_i \times \mathbf{S}_j) + (\mathbf{S}_i \times \mathbf{S}_k)$$

Distributive

# Rings/Semirings

**Ring :  $\langle S, +, \times, S_0, S_1, - \rangle$**

Any semiring where every element has an additive inverse...

$$S_i + (-S_i) = S_0$$



**THE TANGENT ENDS NOW**

# Incremental View Maintenance

WHEN  $D \leftarrow D + \Delta D$  DO:

$VIEW \leftarrow VIEW + \Delta Q(D, \Delta D)$

# Incremental View Maintenance

WHEN  $D \leftarrow D + \Delta D$  DO:

$VIEW \leftarrow VIEW + \Delta Q(D, \Delta D)$

Basic Challenges of IVM

What does  $\Delta R$  represent?

# Incremental View Maintenance

**WHEN**  $D \leftarrow D + \Delta D$  **DO:**

**VIEW**  $\leftarrow \text{VIEW} + \Delta Q(D, \Delta D)$

Basic Challenges of IVM

What does  $\Delta R$  represent?

How to interpret  $R \pm \Delta R$ ?



# Incremental View Maintenance

**WHEN  $D \leftarrow D + \Delta D$  DO:**

**$VIEW \leftarrow VIEW + \Delta Q(D, \Delta D)$**

Basic Challenges of IVM

What does  $\Delta R$  represent?

How to interpret  $R \pm \Delta R$ ?

How to compute  $\Delta Q$ ?

What is  $\Delta R$ ?

# What is $\Delta R$ ?

What does it need to represent?

# What is $\Delta R$ ?

What does it need to represent?

Insertions

Deletions

Updates

# What is $\Delta R$ ?

What does it need to represent?



Insertions

Deletions

Updates

(Delete Old Record & Insert Updated Record)

What is  $\Delta R$ ?

# What is $\Delta R$ ?

A Set/Bag of Insertions

A Set/Bag of Deletions

# What is $+$ ?

**R**

A Set/Bag

**$\Delta R$**

A Set/Bag of Insertions

A Set/Bag of Deletions



# What is $+$ ?

$$\mathbf{R} \quad + \quad \mathbf{\Delta R}$$

A Set/Bag

$+$

A Set/Bag of Insertions

A Set/Bag of Deletions

# What is $+$ ?

**R**      **+**       **$\Delta R$**

A Set/Bag

**+**

A Set/Bag of Insertions

A Set/Bag of Deletions

R

U

$\Delta R_{\text{inserted}}$

-

$\Delta R_{\text{deleted}}$

# What is $+$ ?

**R**      **+**       **$\Delta R$**

A Set/Bag

A Set/Bag of Insertions

**+**

A Set/Bag of Deletions

R

U

$\Delta R_{\text{inserted}}$

-

$\Delta R_{\text{deleted}}$

**But this breaks closure of ' $+$ '!**

# Incremental View Maintenance

$$\text{VIEW} \leftarrow \text{VIEW} + \Delta Q (D, \Delta D)$$

# Incremental View Maintenance

$VIEW \leftarrow VIEW - \Delta Q(D, \Delta D)$

# Incremental View Maintenance

$VIEW \leftarrow VIEW - \Delta Q(D, \Delta D)$

Given  $Q(R, S, \dots)$

Construct  $\Delta Q(R, \Delta R, S, \Delta S, \dots)$

# Delta Queries

$$\Delta(\sigma(R))$$

$\sigma$

|

$R$

# Delta Queries

$$\Delta(\sigma(R))$$

$\sigma$

|

R

R

$\Delta R$

Original R

Inserted  
Tuples of R



# Delta Queries

$$\Delta(\sigma(R))$$

$\sigma$

|

R

Original R

$\sigma$

|

$\Delta R$

Inserted  
Tuples of R

# Delta Queries

$$\Delta(\sigma(R)) = \sigma(\Delta R)$$

$\sigma$   
|  
R

R

Original R

$\sigma$   
|  
 $\Delta R$

Inserted  
Tuples of R

# Delta Queries

$$\Delta(\sigma(R)) = \sigma(\Delta R)$$

$\sigma$   
|  
R

R

$\sigma$   
|  
 $\Delta R$

Original R

Inserted  
Tuples of R

**Does this work for deleted tuples?**

# Delta Queries

$$\Delta(\pi(R)) = \pi(\Delta R)$$

$\pi$   
|  
R

R

$\pi$   
|  
 $\Delta R$

# Delta Queries

$$\Delta(\pi(R)) = \pi(\Delta R)$$


$\pi$   
|  
R

$\pi$   
|  
R       $\Delta R$

**Does this work (completely) under set semantics?**


# Delta Queries

$$\Delta(R_1 \cup R_2)$$




# Delta Queries

$$\Delta(R_1 \cup R_2) = \Delta R_1 \cup \Delta R_2$$




# Delta Queries

$$\Delta(R_1 \cup R_2) = \Delta R_1 \cup \Delta R_2$$





# Delta Queries



# Delta Queries

$R : \{ 1, 2, 3 \}$        $S : \{ 5, 6 \}$

# Delta Queries

$$R : \{ 1, 2, 3 \} \quad S : \{ 5, 6 \}$$

$$R \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle 2, 5 \rangle, \langle 2, 6 \rangle, \langle 3, 5 \rangle, \langle 3, 6 \rangle \}$$

# Delta Queries

$$R : \{ 1, 2, 3 \} \quad S : \{ 5, 6 \}$$

$$R \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle 2, 5 \rangle, \langle 2, 6 \rangle, \langle 3, 5 \rangle, \langle 3, 6 \rangle \}$$

$$\Delta R_{\text{inserted}} = \{ 4 \}$$

$$\Delta R_{\text{deleted}} = \{ 3, 2 \}$$

# Delta Queries

$$R : \{ 1, 2, 3 \} \quad S : \{ 5, 6 \}$$

$$R \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle 2, 5 \rangle, \langle 2, 6 \rangle, \langle 3, 5 \rangle, \langle 3, 6 \rangle \}$$

$$\Delta R_{\text{inserted}} = \{ 4 \}$$

$$\Delta R_{\text{deleted}} = \{ 3, 2 \}$$

$$(R + \Delta R) \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle \mathbf{4}, 5 \rangle, \langle \mathbf{4}, 6 \rangle \}$$

# Delta Queries

$$R : \{ 1, 2, 3 \} \quad S : \{ 5, 6 \}$$

$$R \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle 2, 5 \rangle, \langle 2, 6 \rangle, \langle 3, 5 \rangle, \langle 3, 6 \rangle \}$$

$$\Delta R_{\text{inserted}} = \{ 4 \}$$

$$\Delta R_{\text{deleted}} = \{ 3, 2 \}$$

$$(R + \Delta R) \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle \mathbf{4}, 5 \rangle, \langle \mathbf{4}, 6 \rangle \}$$

$$\Delta_{\text{inserted}}(R \times S) = \Delta R_{\text{inserted}} \times S$$

$$\Delta_{\text{deleted}}(R \times S) = \Delta R_{\text{deleted}} \times S$$

# Delta Queries

$$R : \{ 1, 2, 3 \} \quad S : \{ 5, 6 \}$$

$$R \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle 2, 5 \rangle, \langle 2, 6 \rangle, \langle 3, 5 \rangle, \langle 3, 6 \rangle \}$$

$$\Delta R_{\text{inserted}} = \{ 4 \}$$

$$\Delta R_{\text{deleted}} = \{ 3, 2 \}$$

$$(R + \Delta R) \times S = \{ \langle 1, 5 \rangle, \langle 1, 6 \rangle, \langle \mathbf{4}, 5 \rangle, \langle \mathbf{4}, 6 \rangle \}$$

$$\Delta_{\text{inserted}}(R \times S) = \Delta R_{\text{inserted}} \times S$$

$$\Delta_{\text{deleted}}(R \times S) = \Delta R_{\text{deleted}} \times S$$

**What if R and S both change?**

# Delta Queries

Computing a Delta Query

$$\Delta(\sigma(R)) = \sigma(\Delta R)$$

$$\Delta(\pi(R)) = \pi(\Delta R)$$

$$\Delta(R_1 \cup R_2) = \Delta R_1 \cup \Delta R_2$$

$$\Delta(R_1 \times R_2) = ??$$



# Delta Queries

$$(R_1 \cup \Delta R_1) \times (R_2 \cup \Delta R_2)$$

# Delta Queries

$$(R_1 \cup \Delta R_1) \times (R_2 \cup \Delta R_2)$$

$$(R_1 \times R_2) \cup (R_1 \times \Delta R_2) \cup (\Delta R_1 \times R_2) \cup (\Delta R_1 \times \Delta R_2)$$

# Delta Queries

$$(R_1 \cup \Delta R_1) \times (R_2 \cup \Delta R_2)$$

$$(R_1 \times R_2) \cup (R_1 \times \Delta R_2) \cup (\Delta R_1 \times R_2) \cup (\Delta R_1 \times \Delta R_2)$$

**The original  
query**

# Delta Queries

$$(R_1 \cup \Delta R_1) \times (R_2 \cup \Delta R_2)$$


$$(R_1 \times R_2) \cup (R_1 \times \Delta R_2) \cup (\Delta R_1 \times R_2) \cup (\Delta R_1 \times \Delta R_2)$$

**The original  
query**

**The delta query**


How about an example...

# Delta Queries



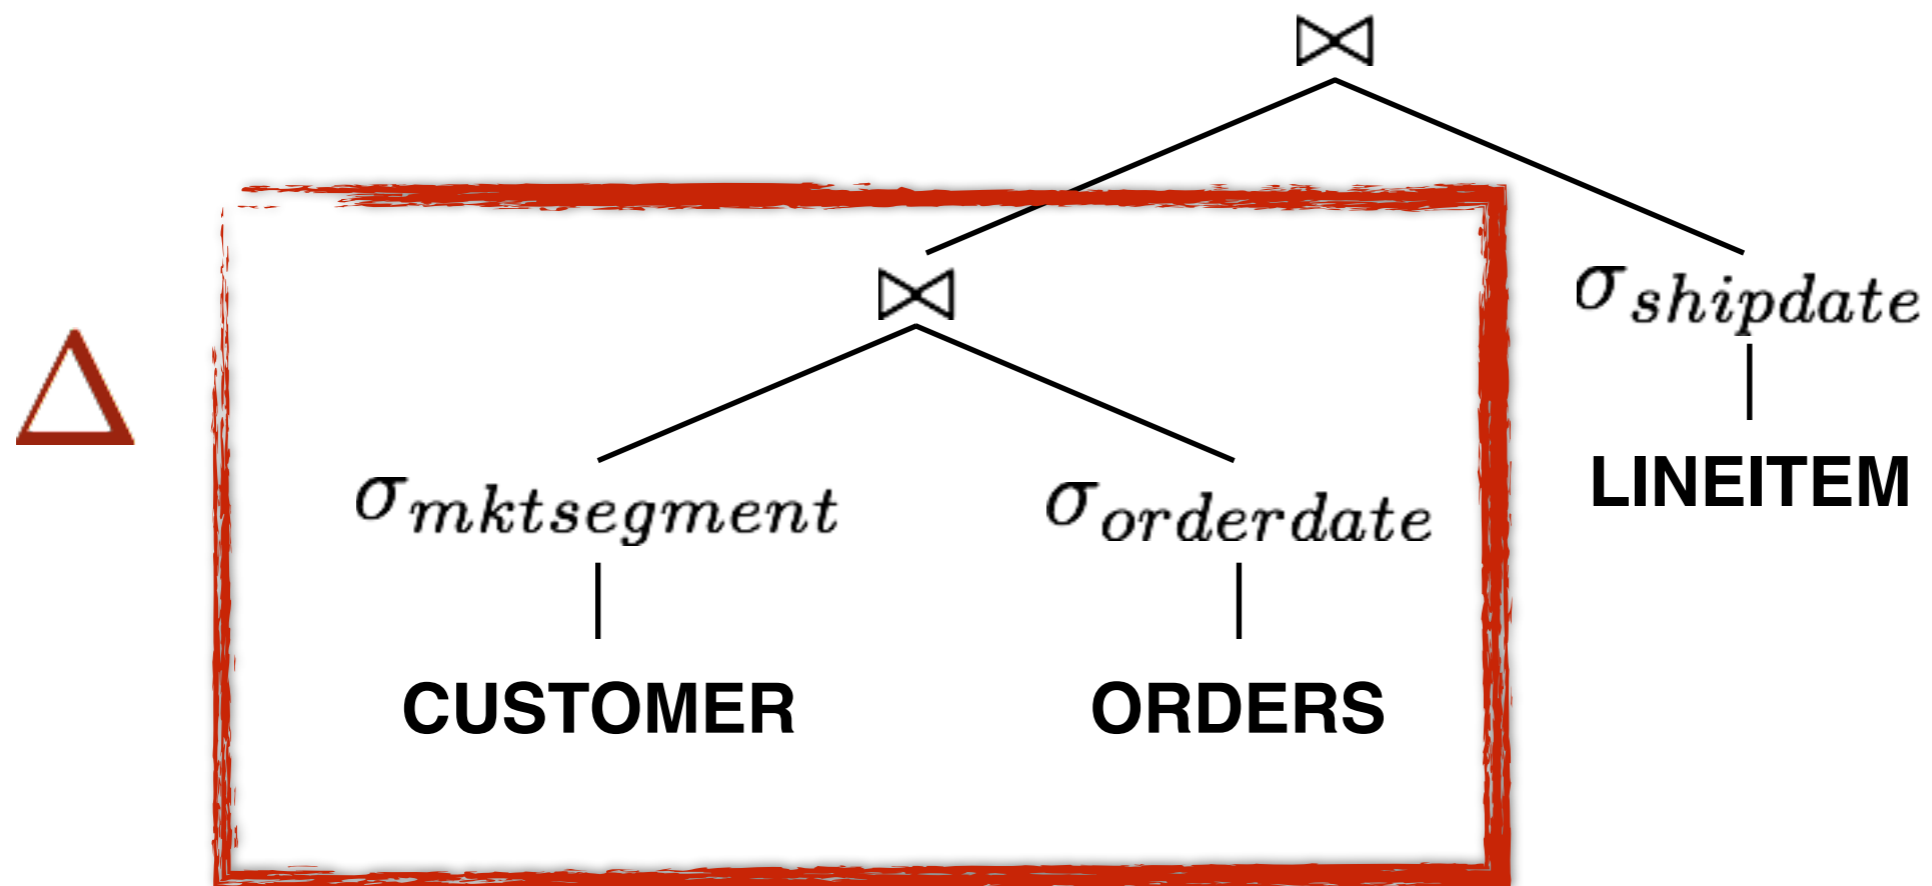
Let's say you have an insertion into LINEITEM

# Delta Queries



$$\Delta((\sigma(C) \bowtie \sigma(O)) \bowtie (\sigma(L)))$$

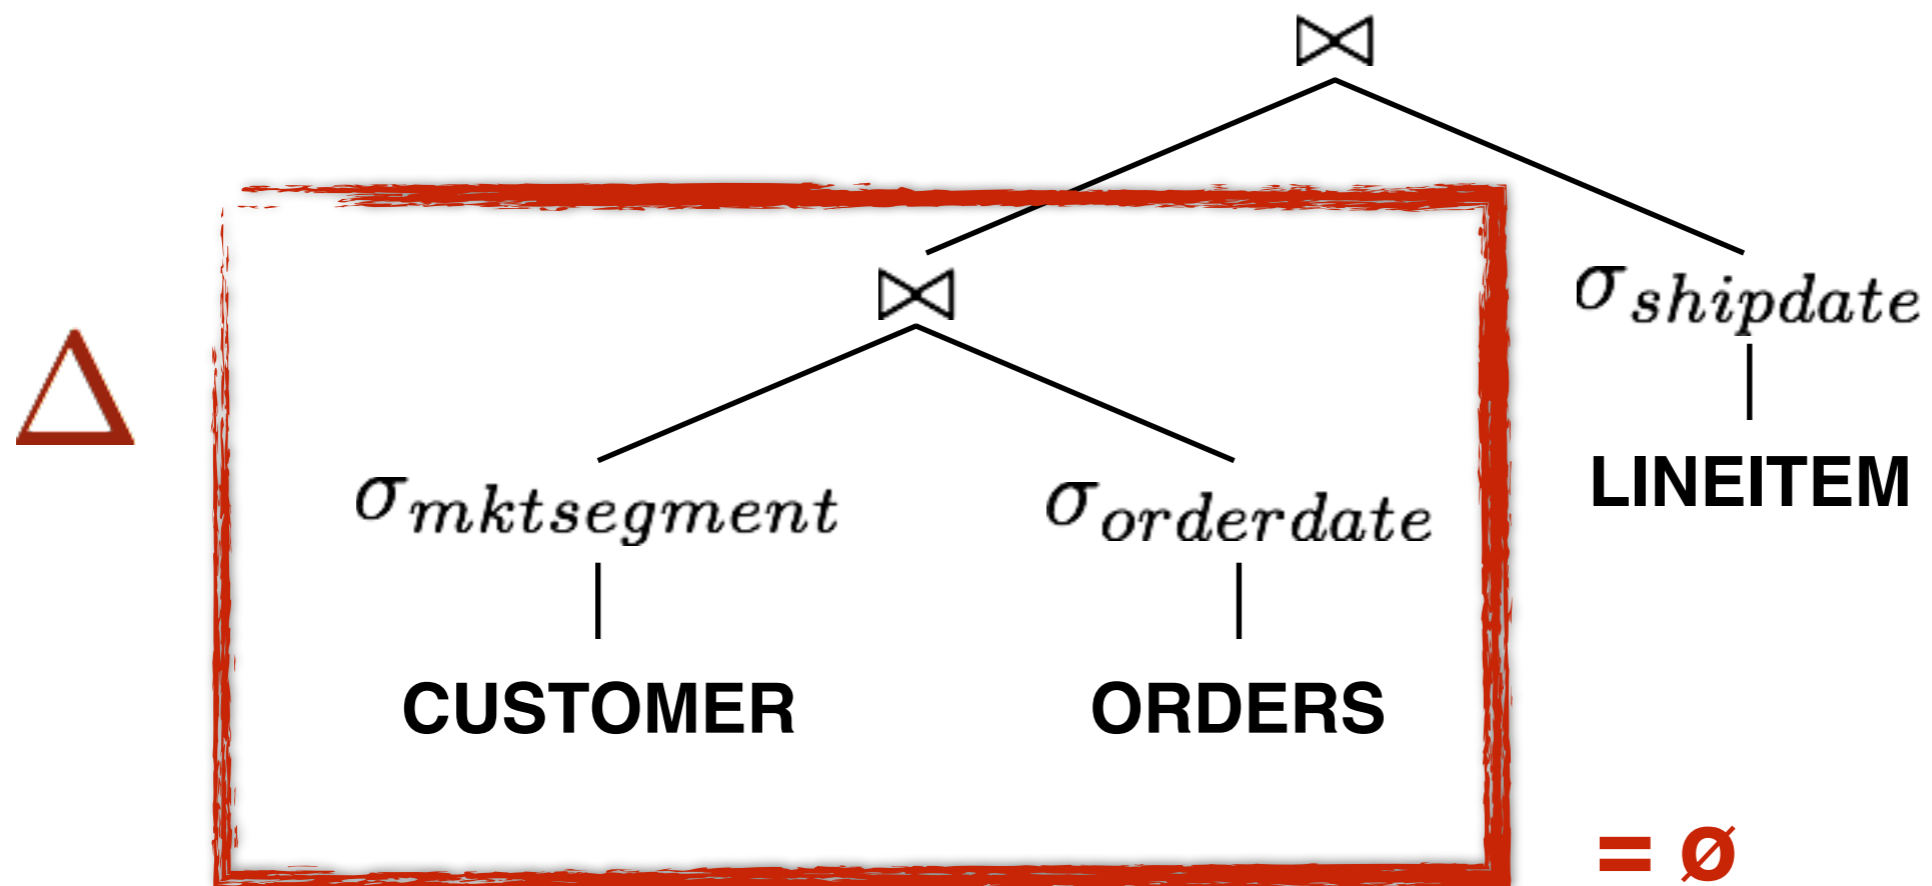
# Delta Queries



$$\Delta((\sigma(C) \bowtie \sigma(O)) \bowtie (\sigma(L)))$$




# Delta Queries




$$\Delta((\sigma(C) \bowtie \sigma(O)) \bowtie (\sigma(L)))$$

# Delta Queries



$$((\sigma(C) \bowtie \sigma(O)) \bowtie \Delta(\sigma(L)))$$

# Delta Queries



# Delta Queries

```
SELECT *
FROM CUSTOMER C, ORDERS O, DELTA_LINEITEM DL
WHERE C.custkey = O.custkey
      AND DL.orderkey = O.orderkey
      AND C.mktsegment = ...
      AND O.orderdate = ...
      AND DL.shipdate = ...
```

# Multisets

**{ 1, 1, 1, 2, 2, 2, 2, 2, 3, 3, 4, 4, 4, 4, 4, 4, 5 }**  
(not compact)

**{ 1 → x3, 2 → x5, 3 → x2, 4 → x6, 5 → x1 }**

Multiset representation: Tuple → # of occurrences

# Multisets

**{ 1, 1, 1, 2, 2, 2, 2, 2, 3, 3, 4, 4, 4, 4, 4, 4, 5 }**  
(not compact)

**{ 1 → x3, 2 → x5, 3 → x2, 4 → x6, 5 → x1 }**

Multiset representation: Tuple → ~~# of occurrences~~  
**multiplicity**

# Multiset Deltas

Insertions = Positive Multiplicity

Deletions = Negative Multiplicity

+ = Bag/Multiset Union