

Practical Incremental View Maintenance

CPS 296.1
Topics in Database Systems

Roadmap

- Zhuge et al. "View Maintenance in a Warehousing Environment." *SIGMOD*, 1995
 - Identified the problem of changing base table states
 - Proposed the idea of compensation
- Salem et al. "How to Roll a Join: Asynchronous Incremental View Maintenance." *SIGMOD*, 2000
 - Proposed the idea of asynchronous change propagation based on compensation
 - Prototyped in a commercial DBMS

2

Data warehousing

- The ETL process
 - Extract data from operational data sources
 - Transform (cleanse and integrate) data
 - Load data into a central warehouse
- Data warehouse data = materialized views over source data
 - Supports fast OLAP (On-Line Analytical Processing)
 - Needs to be kept up-to-date w.r.t. source data
 - The view maintenance problem!

3

Correct view maintenance

Source:		Warehouse:
$R: \begin{array}{ c c } \hline W & X \\ \hline 1 & 2 \\ \hline \end{array}$	$S: \begin{array}{ c c } \hline X & Y \\ \hline 2 & 4 \\ \hline 2 & 3 \\ \hline \end{array}$	$\pi_W(R \triangleright \triangleleft S): \begin{array}{ c } \hline W \\ \hline 1 \\ \hline 1 \\ \hline \end{array}$

- Update $U_1 = \text{insert}(S, [2, 3])$ occurs at the source and is reported to the warehouse
- Warehouse sends $Q_1 = \pi_W(R \triangleright \triangleleft [2, 3])$ to the source
 - Recall change propagation equations

$$R \triangleright \triangleleft (S \oplus \Delta S) = (R \triangleright \triangleleft S) \oplus (R \triangleright \triangleleft \Delta S)$$

$$\pi_W(T \oplus \Delta T) = \pi_W(T) \oplus \pi_W(\Delta T)$$
- Source evaluates Q_1 and returns answer $A_1 = [1]$
- Warehouse receives A_1 and adds it to the view

4

Observations

- To maintain a warehouse view, we may need to send queries back to the sources
 - For a join view, we need to join a delta with the other base tables
 - Source queries are not needed for selection and projection views (assuming minimal deltas, i.e., no over-delete)
- ... Unless we store enough information at the warehouse to make it self-maintainable
 - Thursday

5

A maintenance anomaly

Source:		Warehouse:
$R: \begin{array}{ c c } \hline W & X \\ \hline 1 & 2 \\ \hline 4 & 2 \\ \hline \end{array}$	$S: \begin{array}{ c c } \hline X & Y \\ \hline 2 & 3 \\ \hline \end{array}$	$\pi_W(R \triangleright \triangleleft S): \begin{array}{ c } \hline W \\ \hline 1 \\ \hline 4 \\ \hline 4 \\ \hline \end{array}$ Wrong!

- Source executes and sends $U_1 = \text{insert}(S, [2, 3])$
- Warehouse receives U_1 and sends $Q_1 = \pi_W(R \triangleright \triangleleft [2, 3])$
- Source executes and sends $U_2 = \text{insert}(R, [4, 2])$
- Warehouse receives U_2 and sends $Q_2 = \pi_W([4, 2] \triangleright \triangleleft S)$
- Source receives and evaluates Q_1 , returns $A_1 = [1], [4]$
- Warehouse receives A_1 and adds $[1], [4]$ to the view
- Source receives and evaluates Q_2 , returns $A_2 = [4]$
- Warehouse receives A_2 and adds $[4]$ to the view

6

Another maintenance anomaly

Source:

R:

W	X
1	1

S:

X	Y
1	1

Warehouse:

$\pi_W(R \triangleright \triangleleft S)$:

W
1

 Wrong! ①

- Source executes and sends $U_1 = \text{delete}(R, [1, 2])$
- Warehouse receives U_1 and sends $Q_1 = \pi_W([1, 2] \triangleright \triangleleft S)$
- Source executes and sends $U_2 = \text{delete}(S, [2, 3])$
- Warehouse receives U_2 and sends $Q_2 = \pi_W(R \triangleright \triangleleft [2, 3])$
- Source receives and evaluates Q_1 , returns $A_1 = \emptyset$
- Warehouse receives A_1 and does nothing
- Source receives and evaluates Q_2 , returns $A_2 = \emptyset$
- Warehouse receives A_2 and does nothing

7

What went wrong?

- Change propagation equations should be evaluated over the original state of the base tables
 - Example: $R \triangleright \triangleleft (S \oplus \Delta S) = (R \triangleright \triangleleft S) \oplus (R \triangleright \triangleleft \Delta S)$, where $(R \triangleright \triangleleft \Delta S)$ should read the state of R at the time when ΔS occurs
- But when source receives the maintenance query, base tables might have changed already
 - Example: R changes after the warehouse receives ΔS and before the source receives $(R \triangleright \triangleleft \Delta S)$

8

Solution: compensation

- Augment maintenance queries with compensating queries to offset the effect of concurrent updates
- Example
 - Warehouse receives ΔS
 - Warehouse sends $Q_1 = (R \triangleright \triangleleft \Delta S)$ to source
 - Warehouse receives ΔR before receiving answer to Q_1
 - Instead of just sending $Q_2 = (\Delta R \triangleright \triangleleft S)$, warehouse sends $Q_2 = (\Delta R \triangleright \triangleleft S) \ominus (\Delta R \triangleright \triangleleft \Delta S)$, where the term $(\Delta R \triangleright \triangleleft \Delta S)$ compensates for Q_1

9

ECA (Eager Compensating Algorithm)

- Warehouse maintains UQS (Unanswered Query Set)
 - That is, the set of queries that were sent by the warehouse, but whose answers have not been received
- When warehouse receives source update U_i
 - Formulate a maintenance query Q_i based on U_i
 - For each query in UQS , formulate a compensating query Q' based on U_i , and augment Q_i with $\ominus Q'$
 - Send the augmented Q_i to the source
- Assumption: If A_i is received after U_i , then A_i has seen the effect of U_i
 - Send the message in the same transaction
 - Assume in-order message delivery

10

A note on negative deltas

- If tuples are allowed to have negative counts, then we can capture all changes in a single ΔR rather than the pair ∇R and ΔR
- Everything continues to work
 - \oplus adds counts of matching tuples
 - \ominus subtracts counts of matching tuples
 - \times multiplies tuple counts
 - Yes, negative times negative is positive

11

ECA example

Source:

R:

W	X
1	2
4	2

 S:

X	Y
2	5

 T:

Y	Z
5	3

Warehouse:

$\pi_W(R \triangleright \triangleleft S \triangleright \triangleleft T)$:

W
4
1

- $U_1 = \text{insert}(R, [4, 2])$
- $Q_1 = \pi_W([4, 2] \triangleright \triangleleft S \triangleright \triangleleft T)$
- $U_2 = \text{insert}(T, [5, 3])$
- $Q_2 = \pi_W(R \triangleright \triangleleft S \triangleright \triangleleft [5, 3]) \ominus \pi_W([4, 2] \triangleright \triangleleft S \triangleright \triangleleft [5, 3])$ for Q_1
- $U_3 = \text{insert}(S, [2, 5])$
- $Q_3 = \pi_W(R \triangleright \triangleleft [2, 5] \triangleright \triangleleft T) \ominus \pi_W([4, 2] \triangleright \triangleleft [2, 5] \triangleright \triangleleft T)$ for Q_1
 $\ominus (\pi_W(R \triangleright \triangleleft [2, 5] \triangleright \triangleleft [5, 3]) \ominus \pi_W([4, 2] \triangleright \triangleleft [2, 5] \triangleright \triangleleft [5, 3]))$ for Q_2
- $A_1 = [4]; A_2 = [1]; A_3 = \emptyset$

12

Summary

- Problem: changing base table states
- Solution: compensation
- Trick: negative counts
- Lots, lots of follow-on work
 - More efficient algorithms
 - Multi-source version
 - Parallel version

13

Traditional database view maintenance


- Incremental maintenance is executed as an atomic transaction
 - Blocks updates to base tables
 - Blocks reads of views
 - Could be broken into separate propagation and apply phases
- The maintenance transaction is synchronous and needs to see particular states of the base tables around the time of the refresh

14

Synchronous propagation using pre-states

Griffin and Libkin

- $V = R \triangleright \triangleleft S$
- a : last refresh time; b : current refresh time
- $V_{a,b} = R_{a,b} \triangleright \triangleleft S_a \oplus R_a \triangleright \triangleleft S_{a,b} \oplus R_{a,b} \triangleright \triangleleft S_{a,b}$




15

Synchronous propagation using after-states

Oracle (?)


- $V = R \triangleright \triangleleft S$
- a : last refresh time; b : current refresh time
- $V_{a,b} = R_{a,b} \triangleright \triangleleft S_b \oplus R_b \triangleright \triangleleft S_{a,b} \ominus R_{a,b} \triangleright \triangleleft S_{a,b}$



16

Synchronous propagation using mixed states

- $V = R \triangleright \triangleleft S$
- a : last refresh time; b : current refresh time
- $V_{a,b} = R_{a,b} \triangleright \triangleleft S_a \oplus R_b \triangleright \triangleleft S_{a,b}$




17

Asynchronous propagation

- $V = R \triangleright \triangleleft S$
- a : last refresh time
- b : target refresh time
- c, d : some later points in time

- $V_{a,b} = R_{a,b} \triangleright \triangleleft S_c \ominus R_{a,b} \triangleright \triangleleft S_{b,c} \oplus R_d \triangleright \triangleleft S_{a,b} \ominus R_{a,d} \triangleright \triangleleft S_{a,b}$



18

Advantages of asynchronous propagation


$$\underbrace{R_{a,b} \triangleright \triangleleft S_c}_{\text{forward query}} \ominus \underbrace{R_{a,b} \triangleright \triangleleft S_{b,c}}_{\text{compensation query}} \oplus \underbrace{R_{d,c} \triangleright \triangleleft S_{a,b}}_{\text{forward query}} \ominus \underbrace{R_{a,d} \triangleright \triangleleft S_{a,b}}_{\text{compensation query}}$$

- Flexibility: Reading of base tables can happen any later time (independent of a and b)
 - Although every read of a base table must be properly compensated
- More concurrency: Each term can be evaluated in a different transaction

19

Continuous propagation process

- Choose a propagation interval length δ
- Asynchronously compute $V_{t, t+\delta}$
 - May be executed after $t + \delta$
- $t \leftarrow t + \delta$
- Repeat
- δ is tunable



20

Apply process

- Delta timestamps
 - Assume that base table delta tuples are timestamped by transactions that update them
 - Compute timestamps for view delta tuples
 - Join returns the smallest timestamp (may sound counterintuitive, but works with compensation)
- An independent apply process can refresh the view to any time t before the current high-water mark
- Without these timestamps, the apply process can refresh the view only to a high-water mark

21


Example of timestamp computation

- $V = R \triangleright \triangleleft S$
 - Last refreshed at time t ; need to refresh to time t'
- At time a , insert(R, x); at time b , insert(S, y)
 - $t < a < b < t'$
- At time $c > t'$, calculate forward query $R_{t,t'} \triangleright \triangleleft S_c$
 - Adds xy to view delta with timestamp a
- Compensate by $R_{t,t'} \triangleright \triangleleft S_{t',c}$
 - Empty
- At time $d > t'$, calculate forward query $R_d \triangleright \triangleleft S_{t,t'}$
 - Adds xy to view delta with timestamp b ← Correct effect
- Compensate by $R_{t,d} \triangleright \triangleleft S_{t,t'}$
 - Subtracts xy from view delta with timestamp a

22

Rolling propagation

- Flexibility: Different base tables may be updated at different rates, so allow each base table's propagation interval to be tuned independently
- Efficiency: Do not set target refresh time for a view in advance; less compensation work



23

Complications

- Regions that require compensation may not be rectangular
 - A query always returns a rectangular region
 - So multiple compensation queries may be needed
- High-water marks are no longer determined in advance
 - The current high-water mark must be calculated as the beginning of the oldest query that has not been completely compensated
 - Prior to this point, all forward queries have been completely compensated

24

Implementation issues

- Detecting and timestamping base table deltas
 - Log-based approach
 - Trigger-based approach
- Determining the evaluation time of a query (or the base table state that it reads)

25

Log-based approach

- Used by the paper on DB2
- A tool continuously examines the database transaction log and populates base table deltas
 - Transaction ID
 - Commit sequence number (unique “timestamp”)
 - Commit timestamp (not necessarily unique)
- Advantage: does not disrupt normal database operations
- Disadvantage: needs to scan through many unnecessary log entries if we are only interested in a few base tables

26

Trigger-based approach

- Define a trigger on the base table that fires whenever the table is updated and populates the delta table
 - What is the timestamp then?
 - A regular trigger has no access to the commit sequence number because it is not known until commit time
 - A commit trigger (fired at commit time) is required but is not a standard DBMS feature
- Disadvantage: interferes with normal database operations

27

Determining query evaluation time

- We need the commit sequence number of the transaction in which a propagation query is evaluated
- But it is difficult to tell which log entries belong to this particular transaction
 - Hack: make this transaction write a unique value into a special table
 - These solutions are very system-dependent!

28

Next time

- All the continuous changing base table states give me headaches!
- Self-maintainable views—do not rely on base tables for view maintenance!

29