

Distributed Databases Data Warehousing

CPS 216
Advanced Database Systems

Review

Distributed DBMS

- Top-down approach
 - Data partitioning
 - Query processing
 - Query optimization
 - Concurrency control and recovery
- Bottom-up approach
 - Query processing and optimization

2

Optimizing distributed queries

What is different from optimizing centralized queries?


- New strategies: parallel joins, semijoins, ...
- Plans have a new property: “interesting sites”
- Communication cost is a big factor besides I/O
 - Per-message cost, per-byte cost, CPU cost to pack/unpack data
- Parallelism: response time versus total resource consumption



3

Example: two-step optimization


- Step 1 (compile time): decide the join order, join methods, and access paths
 - Same complexity as in a centralized DBMS
- Step 2 (run time): decide where to execute each operator
 - Can cope with changing load and network characteristics
 - Can use data that has been dynamically allocated to a site (caching or replication)



4

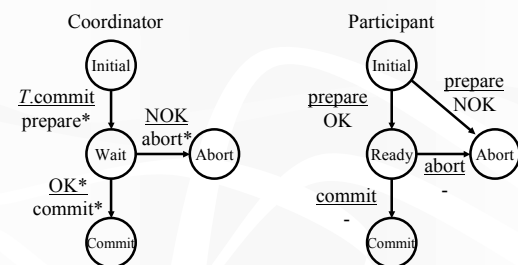
Concurrency control & recovery in Distributed DBMS

- Rich and interesting field
- We will just sample the field by looking at the problem of distributed transaction commit



5

Two-phase commit



Notation: Incoming message * = everyone
Outgoing message

6

Key points of 2PC

- By sending OK a participant promises the coordinator to commit
 - But it can only commit when instructed to do so by the coordinator
 - The coordinator could tell it to abort instead
- After sending NOK a participant can abort unilaterally
- Coordinator can decide to commit only if all participants have responded OK
- Logging of all messages are required at each site


7

Bottom-up approach to Building a distributed DBMS

- Data already in various sources
- Build a distributed DBMS to provide global, uniform access to all data
 - How to integrate data?
 - How to deal with heterogeneous and autonomous sources?
- » Mediation approach

8


Wrapper/mediator architecture



9

Mediator


- Accept queries from clients
- Rewrite and optimize queries
- Send subplans to be executed by wrappers
- Combine results from wrappers and perform any additional local processing necessary
- Mediator catalog stores global schema and external schema of sources as exported by wrappers
- » No source-specific code in a mediator!



10

Wrapper

- Hide heterogeneity away from mediator
- Translate mediator requests so that they are understood by sources
 - Example: `SELECT * FROM Books WHERE title LIKE '%Databases'`; → a form-based search request for books with title matching “*Databases”
- Translate results returned by a source so that they are compliant with its external schema
 - Example: result HTML page → Books tuples



11

Query optimization with wrappers

Basic questions

- Capability: What types of subplans can be handled by a wrapper?
 - How do we enumerate valid plans?
- Cost: What is the cost of executing a subplan by a wrapper?
 - How do we pick the optimal plan?

12

Example: Garlic query optimization

- Haas et al., VLDB 1997
- Incorporated in DB2
- Rules for generating valid plans
 - Supplied by wrappers and mediator
 - Plugged into the optimizer
- Plans have “interesting properties”
 - Order (as in Selinger)
 - Site (where the output is produced)
 - Columns (in the output)
 - Predicates (that have been applied)
 - Cost, etc.

13

Example rules for a DBMS source

- $\text{wrap_access}(table, columns, predicates) = \text{SCAN}_{\text{DBMS}}(table, columns, predicates)$
 - Condition: *table* is at my site
 - I can handle any projection and selection (by converting them to a single-table SELECT-FROM-WHERE SQL statement)
- $\text{wrap_join}(subplan_1, subplan_2, predicates) = \text{JOIN}_{\text{DBMS}}(subplan_1, subplan_2, predicates)$
 - Condition: $subplan_1.site = subplan_2.site = \text{my site}$
 - I can handle any local join (by converting it to a multi-table SELECT-FROM-WHERE SQL statement)

14

Example rules for a Web source

- $\text{wrap_access}(table, columns, predicates) = \text{FETCH}_{\text{Web}}(\text{Books}, \text{title LIKE } string)$
 - Condition: $table = \text{Books}, (\text{title LIKE } string) \in predicates$
 - I can search books by title (with wildcards); no projection
- $\text{wrap_access}(table, columns, predicates) = \text{FETCH}_{\text{Web}}(\text{Books}, \text{author} = string)$
 - Condition: $table = \text{Books}, (\text{author} = string) \in predicates$
 - I can search books by exact author names; no projection
 - I cannot search books by title and author at the same time
- No wrap_join rule
 - I cannot handle process joins

15

Example rules for the mediator

- $\text{med_pushdown}(subplan) = \text{RECEIVE}(\text{SEND}(subplan))$
 - Condition: $subplan.site \neq \text{mediator}$
- $\text{med_pushdown}(subplan) = subplan$
 - Condition: $subplan.site = \text{mediator}$
- $\text{med_access}(table, columns, predicates) = \forall plan \in \text{wrap_access}(table, columns, predicates): \text{FILTER}_{\text{med}}(\text{med_pushdown}(plan), predicates - plan.predicates)$
 - I can get the result of a single-table scan from a wrapper and then evaluate remaining selection predicates

16

More rules for the mediator

- $\text{med_join}(subplan_1, subplan_2, predicates) = \forall plan \in \text{wrap_join}(subplan_1, subplan_2, predicates): \text{med_pushdown}(plan)$
 - Condition: $subplan_1.site = subplan_2.site \neq \text{mediator}$
 - I can push down a join to a wrapper
- $\text{med_join}(subplan_1, subplan_2, predicates) = \text{JOIN}_{\text{med}}(\text{med_pushdown}(subplan_1), \text{med_pushdown}(subplan_2), predicates)$
 - I also can handle a join locally
- And more...

17

Plan enumeration

- Call all wrap_access and med_access rules to generate single-table access plans
- Repeatedly call all wrap_join and med_join rules to generate multi-table join plans
- Example final plans
 - $\text{FILTER}_{\text{med}}(\text{RECEIVE}(\text{SEND}(\text{FETCH}_{\text{Web}}(\text{Books}, \text{title LIKE } string))), \text{author} = string)$, versus
 - $\text{FILTER}_{\text{med}}(\text{RECEIVE}(\text{SEND}(\text{FETCH}_{\text{Web}}(\text{Books}, \text{author} = string))), \text{title LIKE } string)$
 - $\text{RECEIVE}(\text{SEND}(\text{JOIN}_{\text{DBMS}}(R, S)))$, versus
 - $\text{JOIN}_{\text{med}}(\text{RECEIVE}(\text{SEND}(R)), \text{RECEIVE}(\text{SEND}(S)))$

18

Costing

- Wrapper-supplied cost model
 - Lots of work for wrapper developers
- Calibration
 - Define a generic cost model with parameters for all wrappers
 - Example: $\text{cost} = c \cdot (\# \text{ of tuples})$
 - Run test queries to measure the parameters for each wrapper
- Learning curve
 - Use recent statistics to adjust cost estimates
 - Example: $\text{cost} = \text{average over last three runs}$

19

Summary of wrapper/mediator

Not all sources are created equal!

- What's in a source?
 - Wrapper: source schema \leftrightarrow external schema
 - Mediator: external schema \leftrightarrow global schema
- What can it do?
 - Wrappers and mediators supply rules describing their query processing capabilities
- How much does it cost?
 - Wrappers supply cost model, or
 - Mediator calibrates or learns the cost model

20

Data warehousing

- Data resides in many distributed, heterogeneous OLTP (On-Line Transaction Processing) sources
 - Sales, inventory, customer, ...
 - NC branch, NY branch, CA branch, ...
 - Need to support OLAP (On-Line Analytical Processing) over an integrated view of the data
- » Store the integrated data at a central repository called the data warehouse

21

OLTP versus OLAP

OLTP

- Mostly updates
- Short, simple transactions
- Clerical users
- Goal: ACID, transaction throughput

OLAP

- Mostly reads
- Long, complex queries
- Analysts, decision makers
- Goal: fast queries

22

Warehousing versus mediation

Warehousing

- Eager "integration"
 - In advance: before queries
 - Answer could be stale
- Copy data from sources
 - Need to maintain consistency
 - Query processing is local to the warehouse
 - Faster
 - Can operate when sources are unavailable

Mediation

- Lazy "integration"
 - On demand: at query time
 - Answer is more up-to-date
- Leave data at sources
 - No need to maintain consistency
 - Sources participate in query processing

23

Maintaining a data warehouse

Buzz word: the "ETL" process

- Extraction: extract relevant data and/or changes from sources
- Transformation: transform data to match the warehouse schema
- Loading: integrate data/changes into the warehouse

» Can still use a wrapper/mediator architecture

24

Warehouse data = materialized views

- If the transformation process can be captured by SQL, then warehouse data can be seen as a view
 - CREATE VIEW warehouse_table AS
SELECT ...
FROM source_table1, source_table2, ...
WHERE ...;
- Except the view is materialized
 - That is, the result is stored
 - And needs to be maintained when source data changes

25

Maintaining materialized views

$$V_{old} = Q(R_{old}, \dots)$$

$$\text{Change detected: } R_{new} \leftarrow R_{old} - \nabla R \cup \Delta R$$

What is V_{new} ?

- Recomputation: $V_{new} \leftarrow Q(R_{new}, \dots)$
 - Done periodically, e.g., every “night”
 - What if there is no “night,” e.g., an international organization?
 - What if recomputation takes longer than a day?
- Incremental maintenance
 - Compute only the changes to V : ∇V and ΔV
 - Apply the changes to V_{old} : $V_{new} \leftarrow V_{old} - \nabla V \cup \Delta V$
 - » Potentially much faster if changes are small

26

Incremental maintenance

Example: $V = \sigma_p R$

- Change: $R_{new} \leftarrow R_{old} - \nabla R$
 - $V_{new} = \sigma_p R_{new} = \sigma_p (R_{old} - \nabla R) = \sigma_p R_{old} - \sigma_p \nabla R$
 $= V_{old} - \nabla V$
 - Change: $R_{new} \leftarrow R_{old} \cup \Delta R$
 - $V_{new} = \sigma_p R_{new} = \sigma_p (R_{old} \cup \Delta R) = \sigma_p R_{old} \cup \sigma_p \Delta R$
 $= V_{old} \cup \Delta V$
- Change propagation equations

27

Change propagation

- More change propagation equations
 - $(R \cup \Delta R) \triangleright \triangleleft S = (R \triangleright \triangleleft S) \cup (\Delta R \triangleright \triangleleft S)$
 - $(R - \nabla R) \triangleright \triangleleft S = (R \triangleright \triangleleft S) - (\nabla R \triangleright \triangleleft S)$
- Repeatedly apply change propagation equations to “factor out” changes
 - $(\sigma_{pr} (R \cup \Delta R)) \triangleright \triangleleft_{prs} \sigma_{ps} S = (\sigma_{pr} R \cup \sigma_{pr} \Delta R) \triangleright \triangleleft_{prs} \sigma_{ps} S = (\sigma_{pr} R \triangleright \triangleleft_{prs} \sigma_{ps} S) \cup (\sigma_{pr} \Delta R \triangleright \triangleleft_{prs} \sigma_{ps} S)$

28

Self-maintainability

- A warehouse is self-maintainable if it can be maintained without accessing the sources
- Self-maintainable: $V = \sigma_p R$
- Not self-maintainable: $V = R \triangleright \triangleleft S$
 - ΔR and ∇R need to be joined with S
 - ΔS and ∇S need to be joined with R
 - Problem: need to query the source for maintenance
 - What if the source/network is slow?
 - What if the source/network is down?
 - What if the source has been updated again?

29

Making warehouse self-maintainable

- Add auxiliary views
- Example: Order $\triangleright \triangleleft_{O.OID = L.OID \text{ AND } O.month = 'nov' \text{ AND } L.product = 'book'}$ Lineitem
- Naïve approach: add base tables O and L
- Better approach: push selections down and then add selection views $\sigma_{month = 'nov'} O$ and $\sigma_{product = 'book'} L$
- Use constraints
 - The join is a foreign-key join ($L.OID$ references $O.OID$), so only $\sigma_{month = 'nov'} O$ is needed
 - If we only insert matching orders and lineitems together, then no auxiliary view is needed

30

Next time

- Warehouse design
- Data cube
- ROLAP versus MOLAP

31