Query Execution

CS157B

Chris Pollett

Mar. 9, 2005.

Outline

- An Algebra for Queries
- Physical Query Plan Operators
- One Pass Algorithms for Database Operations

An Algebra for Queries

- One of the first steps in evaluating a query is to turn the original SQL into an internal representation, based on relational algebra, that we can then manipulate.
- We will view our algebra as operating on bags rather than sets.

Relational Algebra Expressions

- Relations R1, R2 … are expressions:
- If Q1, Q2 are an expression so are:
  1. $Q1 \cup Q2$, $Q1 \cap Q2$, $Q1/Q2$
  2. $\sigma_{expr}(1)$ where expr is a valid selection expression. (Boolean formulas built out of things like R.A=S.B or R.A=value, R.A>= value, etc.)
  3. $\pi_{A1,\ldots An}(Q1)$. (Projection) Also, allow A_j-> A_k to mean renaming
  4. Q1 x Q2 (Cartesian Product)
  5. $Q1 \rhd\!\lhd_{expr} Q2$ (Natural Join is a special case)
  6. $\delta(Q1)$ -- duplicate elimination on Q1
  7. $\gamma_L(Q1)$ -- Grouping and Aggregate Operators applied to Q1 (To handle COUNT, GROUP BY, HAVING)
  8. $\tau_L(Q1)$ -- sort Q1 (ORDER BY)

Expression Trees

- These are the intermediate representations of SQL queries. (There can be many equivalent)
- For example:

```
SELECT title, birthdate
FROM  MovieStar, StarsIn
WHERE year =1996 AND
gender='F' AND
starName = 'Starlet'
```

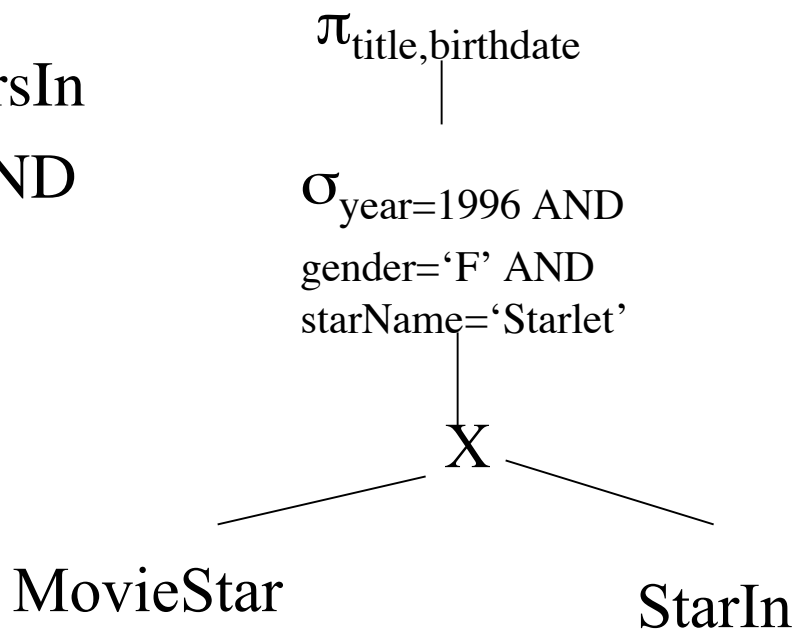

Physical Query Plan Operators

- Physical Query Plans are similarly built out of operators each one of which would enable us to implement a relational operator operation (or at least part of one).
- Some simple example operators:
  - Table-scan
  - Index Scan (can be more specific to index type)
  - Sort-Scan (tree-sort, main memory, merge)

Model of Computation for Physical Operators

- Basic model will just count the number of disk I/Os.
- Later might add network communication costs.
- Assume arguments for any query come from disk, but results are left in main memory. (Reasonable assuming pipelining)

Parameters for Measuring Cost

- M - number of main memory buffers available
- B(R) - number of blocks in R. (Just B if R understood.)
- T(R) - number of tuples in R. (Just T if R understood.)
- V(R,a) - number of distinct values of attribute a.

I/O Cost for Scan Operators

- Cost(Table-Scan(R)) = B(R)
- Cost(Sort-Scan(R)) = B -- provided R fits in main memory. 3B -- if use multiway merge sort and only need one pass in Phase II.
  - If table is not *clustered* (that is, all blocks nearly filled with tuples) then sort scan will depend on T and be more like T+2B to do.

Iterators for Implementing Physical Operators

- Many physical operators are implemented using *Iterators*. An iterator is an object which supports:
  - Open
  - GetNext
  - Close
- Example: A one pass algorithm for R union S might first open R use GetNext in a loop to get each tuple in R till R done. Then Open S and use GetNext in a loop to get each tuple of S. Finally, Close both R and S.