

Towards a linear algebra semantics for SQL

J.N. Oliveira

INFOBLENDER SEMINAR SERIES 2016

February 17th, 2016



INESC TEC & University of Minho

Grant FP7-ICT 619606

Context



About project **LeanBigData**:

*" (...) **queries** [identifying] facts of interest take hours, days, or weeks, whereas business processes demand today shorter cycles.*



FP7-ICT 619606

Project motto: bend big data!

However — **what** are we actually **leaning**?

What is, after all, a **query**?

Back to basics



There are **jobs**:

```
create table jobs (  
  j_code char (15) not null,  
  j_desc char (50),  
  j_salary decimal (15,2) not null);
```

<i>j_code</i>	<i>j_desc</i>	<i>j_salary</i>
<i>Pr</i>	<i>Programmer</i>	<i>1000</i>
<i>SA</i>	<i>System Analyst</i>	<i>1100</i>
<i>GL</i>	<i>Group Leader</i>	<i>1333</i>

Back to basics



There are **jobs**:

```
create table jobs (  
  j_code char (15) not null,  
  j_desc char (50),  
  j_salary decimal (15, 2) not null);
```

j_code	j_desc	j_salary
<i>Pr</i>	<i>Programmer</i>	<i>1000</i>
<i>SA</i>	<i>System Analyst</i>	<i>1100</i>
<i>GL</i>	<i>Group Leader</i>	<i>1333</i>

Back to basics



There are **employees**:

```
create table empl (  
  e_id      integer not null,  
  e_job     char (15) not null,  
  e_name    char (15),  
  e_branch  char (15) not null,  
  e_country char (15) not null);
```

<i>e_id</i>	<i>e_job</i>	<i>e_name</i>	<i>e_branch</i>	<i>e_country</i>
1	Pr	Mary	Mobile	UK
2	Pr	John	Web	UK
3	GL	Charles	Mobile	UK
4	SA	Ana	Web	PT
5	Pr	Manuel	Web	PT

Back to basics



There are **employees**:

```
create table empl (  
  e_id      integer not null,  
  e_job     char (15) not null,  
  e_name    char (15),  
  e_branch  char (15) not null,  
  e_country char (15) not null);
```

e_id	e_job	e_name	e_branch	e_country
1	Pr	Mary	Mobile	UK
2	Pr	John	Web	UK
3	GL	Charles	Mobile	UK
4	SA	Ana	Web	PT
5	Pr	Manuel	Web	PT

Query



Monthly salary total per country / branch:

```
select e_country, e_branch, sum (j_salary)
from empl, jobs
where j_code = e_job
group by e_country, e_branch
order by e_country;
```

sqlite3:

```
PT|Web|2100
UK|Mobile|2333
UK|Web|1000
```

Query



Impact of

```
insert into "jobs" values ('SA', 'System Admin', 1000);
```

that is, *j_code* no longer a key.

sqlite3:

```
PT|Web|3100
UK|Mobile|2333
UK|Web|1000
```

Fine — so *SA* is taken as a kind of “multi-job”.

But — where are these **semantics** specified?

Standard semantics



Given in English:

“The result of evaluating a query-specification can be explained in terms of a multi-step algorithm. The order of [the 7] steps in this algorithm follows the mandatory order of the clauses (FROM, WHERE, and so on) of the SELECT statement”

Cf. pages 71-73 of

*X/Open CAE Specification Data Management:
Structured Query Language (SQL) Version 2 March
1996, X/Open Company Limited*

7 steps



1. *For each table-reference that is a joined-table, conceptually join the tables (...) to form a single table*
2. *Form a Cartesian product of all the table-references (...)*
3. *Eliminate all rows that do not satisfy the search-condition in the WHERE clause.*
4. *Arrange the resulting rows into groups (...)*
 - *If there is a GROUP BY clause specifying grouping columns, then form groups so that all rows within each group have equal values for the grouping columns (...)*
5. *If there is a HAVING clause, eliminate all groups that do not satisfy its search- condition (...)*
6. *Generate result rows based on the result columns specified by the select-list (...)*
7. *In the case of SELECT DISTINCT, eliminate duplicate rows from the result (...)*

Background



Join operator — ok, well defined in Codd's relation algebra.

However,

[...] relational DBMS were never intended to provide the very powerful functions for data synthesis, analysis and consolidation that is being defined as multi-dimensional data analysis.

*E.F.Codd*¹

[...] expressing roll-up, and cross-tab queries with conventional SQL is daunting. [...] GROUP BY is an unusual relational operator [...]

*J. Gray et al*²

¹Providing OLAP to User-Analysts: An IT Mandate (1998)

²Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals (1997)

Background



December 4,
2014

In sql

12 Comments

Do You Really Understand SQL's GROUP BY and HAVING clauses?

★★★★☆ 27 Votes

There are some things in SQL that we simply take for granted without thinking about them properly.

One of these things are the **GROUP BY** and the **less popular HAVING** clauses.

[<http://blog.jooq.org/2014/12/04/do-you-really-understand-sqls-group-by-and-having-clauses/>]


Background



Why these shortcomings / questions ?

*While **relation algebra** "à la Codd" [works] well for qualitative data science [it is] rather clumsy in handling the quantitative side [...] we propose to solve this problem by suggesting **linear algebra** (LA) as an alternative suiting both sides [...]*

*H. Macedo, J. Oliveira*³



Linear algebra ...

³A linear algebra approach to OLAP (2015)

Formalizing SQL data aggregation



VLDB'87, among other research:

SQL Query	Calculus Expression
SELECT f_1, \dots, f_l FROM $r_1(v_1), \dots, r_n(v_n)$ WHERE P_w	$(f'_1, \dots, f'_l) (* : r_1(v_1), \dots, r_n(v_n) : P_w)$
SELECT $t_1, \dots, t_l (\neq f)$ FROM $r_1(v_1), \dots, r_n(v_n)$ WHERE P_w	$(t_1, \dots, t_l) : r_1(v_1), \dots, r_n(v_n) : P_w$
SELECT t_1, \dots, t_l FROM $r_1(v_1), \dots, r_n(v_n)$ WHERE P_w GROUP BY $v_{i_1}[A_{i_1}], \dots, v_{i_k}[A_{i_k}]$ HAVING P_h	$(t'_1, \dots, t'_l) : \alpha(v) : P'_h$ $\alpha = (\phi_{\langle (A_{i_1}, \dots, A_{i_k}), (J'_1, \dots, J'_m) \rangle} (* : r_1(v_1), \dots, r_n(v_n) : P_w));$ $(t'_1, \dots, t'_l, P'_h) = (t_1, \dots, t_l, P_h)[f_i/v[k+i], v_{i_j}[A_{i_j}]/v[j]];$ $(f_1, \dots, f_m \text{ aggregate functions in } t_1, \dots, t_l, P_h)$


G. Bultzingsloewen⁴

⁴Translating and optimizing SQL queries having aggregates (1987)



“Star” diagrams

Entities (cf. tables) surrounded (placed at the center of) by their **attributes**:



Entities marked in bold.

Attribute types made explicit, linking entities to each other.

“Star” diagrams



What is the (formal) meaning of the **arrows** in the diagram?

There is one arrow per **attribute** — **column** in the database table.


Assigning meanings to the arrows amounts to formalizing a **columnar** approach to SQL.⁵

Let us do so using the **linear algebra of programming** (LAoP).⁶

⁵D. Abadi et al, *The Design and Implementation of Modern Column-Oriented Database Systems* (2012).

⁶J. Oliveira, *Towards a Linear Algebra of Programming* (2012).

Formal star-diagram in (**typed**) LAoP



Legend:

- **Types:**

K — Job code

C — Country

B — Branch

$\#e$ — *empl* record nrs

$\#j$ — *jobs* record nrs

- **Dimensions:**

- *branch*

- *code*

- *country*

- *job*

- **Measures:**

- *salary*



Dimensions

Dimension attribute columns are captured by **bitmap** matrices:

e_{branch}	1	2	3	4	5
Mobile	1	0	1	0	0
Web	0	1	0	1	1

e_{job}	1	2	3	4	5
GL	0	0	1	0	0
Pr	1	1	0	0	1
SA	0	0	0	1	0

$e_{country}$	1	2	3	4	5
PT	0	0	0	1	1
UK	1	1	1	0	0

j_{desc}	1	2	3
Group Leader	0	0	1
Programmer	1	0	0
System Analyst	0	1	0

j_{code}	1	2	3
GL	0	0	1
Pr	1	0	0
SA	0	1	0

Meaning of bitmap **matrix** t_d , for d a dimension of table t :

$$v \ t_d \ i = 1 \Leftrightarrow t[i].d = v \quad (1)$$



Measures

However — main difference wrt. **relation algebra** — we won't build

j^{salary}	1	2	3
1000	1	0	0
1100	0	1	0
1333	0	0	1

but rather the **row vector** $j^{\text{salary}} : \#j \rightarrow 1$ which “internalizes” the quantitative information:

j^{salary}	1	2	3
1	1000	1100	1333

Summary:

Measures are *vectors*, **dimensions** are *matrices*.



Linear algebra

Matrix **multiplication**, given matrices $B \xleftarrow{M} C \xleftarrow{N} A$:

$$b (M \cdot N) a = \langle \sum c :: (b M c) \times (c N a) \rangle \quad (2)$$

Matrix **converse**:

$$c M^\circ b = b M c \quad (3)$$

Functions are (special cases of Boolean) matrices:

$$y f x = \begin{cases} 1 & \text{if } y = f x \\ 0 & \text{otherwise} \end{cases} \quad (4)$$



Examples

$1 \xleftarrow{j^{salary}} \#j \xleftarrow{j^{code}} K$	<i>Pr</i>	<i>SA</i>	<i>GL</i>
1	1000	1100	1333

Calculation:

$$1 (j^{salary} \cdot j^{code}) k$$

$$\Leftrightarrow \{ \text{multiplication (2)} \}$$

$$\langle \sum y :: (1 j^{salary} y) \times (y j^{code} k) \rangle$$

$$\Leftrightarrow \{ \text{converse (3)} ; \text{vector } j^{salary} \}$$

$$\langle \sum y :: (k j^{code} y) \times (j[y].salary) \rangle$$

$$\Leftrightarrow \{ \text{functions (4)} ; \text{quantifier notation (details soon)} \}$$

$$\langle \sum y : k = j[y].code : j[y].salary \rangle$$

□



Examples

In case of the addition of

insert into "jobs" values ('SA', 'System Admin', 1000);

we get non-injective bitmap

j_{code}	1	2	3	4
GL	0	0	1	0
Pr	1	0	0	0
SA	0	1	0	1

and

j_{salary}	1	2	3	4
1	1000	1100	1333	1000

Therefore:

1	$\xleftarrow{j_{salary}} \#j$	$\xleftarrow{j_{code}} K$	Pr	SA	GL
1		1	1000	2100	1333



Pointwise LAoP calculus

Quantifier notation follows the Eindhoven style,

$$\langle \sum x : R : T \rangle$$

where R is a predicate (**range**) and T is a numeric term.

In case $T = B \times M$ where Boolean $B = \llbracket P \rrbracket$ encodes predicate P , we have the **trading rule**:

$$\langle \sum x : R : \llbracket P \rrbracket \times M \rangle = \langle \sum x : R \wedge P : M \rangle \quad (5)$$

Thus

$$y(f \cdot N)x = \langle \sum z : y = f z : z N x \rangle \quad (6)$$

$$y(g^\circ \cdot N \cdot f)x = (g y) N (f x) \quad (7)$$

hold.

Pointwise LAoP calculus



Given a binary predicate $p : B \times A \rightarrow Bool$, we denote by $\llbracket p \rrbracket : B \leftarrow A$ the Boolean matrix which encodes p , that is,

$$b \llbracket p \rrbracket a = \mathbf{if } p(b, a) \mathbf{ then } 1 \mathbf{ else } 0 \quad (8)$$

In case of a unary predicate $q : A \rightarrow Bool$, $\llbracket q \rrbracket : 1 \leftarrow A$ is the Boolean vector such that:

$$1 \llbracket q \rrbracket a = \llbracket q \rrbracket [a] = \mathbf{if } q a \mathbf{ then } 1 \mathbf{ else } 0 \quad (9)$$




Joins and tabulations

SQL querying amounts to **following paths** in star diagrams.

The **meaning of a path** is obtained by composing (multiplying) the matrices involved.

Two particular such compositions deserve special reference, as they correspond to well-known operations in data processing:



- **Join:** $X = t_B^\circ \cdot M \cdot p_B$
- **Tabulation:** $Y = p_B \cdot N \cdot p_A^\circ$

M and N are whatever matrices of their type.



Simple Examples

Equi-join ($M = id$):

$j_{code} \cdot e_{job}$	1	2	3	4	5
1	1	1	0	0	1
2	0	0	0	1	0
3	0	0	1	0	0

Pointwise meaning:

$$j[y].code = e[x].job$$

recall (7).

Counting tabulation ($N = id$):

$e_{country} \cdot e_{branch}$	Mobile	Web
PT	0	2
UK	2	1

Pointwise meaning:

$$\langle \sum k : y = e[k].country \wedge x = e[k].branch : 1 \rangle$$

recall (6), for y a country, x a branch.



Columnar joins

Excerpt from Abadi et al⁷

For example, the figure below shows the results of a join of a column of size 5 with a column of size 4:

$$\begin{array}{|c|} \hline 42 \\ \hline 36 \\ \hline 42 \\ \hline 44 \\ \hline 38 \\ \hline \end{array} \bowtie \begin{array}{|c|} \hline 38 \\ \hline 42 \\ \hline 46 \\ \hline 36 \\ \hline \end{array} = \begin{array}{|c|} \hline 1 \\ \hline 2 \\ \hline 3 \\ \hline 5 \\ \hline \end{array} \begin{array}{|c|} \hline 2 \\ \hline 4 \\ \hline 2 \\ \hline 1 \\ \hline \end{array}$$

shows **columnar-join** “isomorphic” to our matrix joins:

	1	2	3	4	5
1	0	0	0	0	1
2	1	0	1	0	0
3	0	0	0	0	0
4	0	1	0	0	0

⁷The Design (..) of Modern Column-Oriented Database Systems (2012).

Back to the starting SQL query



Minimal diagram accommodating query:

select

e_branch,

e_country,

sum (j_salary)

from *empl, jobs*

where *j_code = e_job*


group by

e_country,

e_branch

order by

e_country;



Clearly,

group by \Rightarrow *tabulation Q*

where \Rightarrow *join J*

Back to the starting SQL query




select

```

e_branch,
e_country,
sum (j_salary)
from empl, jobs
  where j_code = e_job
group by
  e_country,
  e_branch
order by
  e_country;

```

How do **salaries** get involved? We need a direct path from employees to (their) salaries,



involving the **where**-clause join:

$$v = j^{\text{salary}} \cdot j_{\text{code}}^{\circ} \cdot e_{\text{job}} \quad (10)$$



Query = Group by + Join

The **group by** clause calls for a tabulation — but, how does vector

$j^{\text{salary}} \cdot j^{\text{code}} \cdot e_{\text{job}}$	1	2	3	4	5
1	1000	1000	1333	1100	1000

get into the place of N in the generic scheme?

Easy: every vector v can be turned into a **diagonal** matrix, e.g.

$v \nabla id$	1	2	3	4	5
1	1000	0	0	0	0
2	0	1000	0	0	0
3	0	0	1333	0	0
4	0	0	0	1100	0
5	0	0	0	0	1000

and vice versa.



Khatri-Rao product

This diagonalization resorts to another LA operator, termed Khatri-Rao product ($M \nabla N$) defined by

$$(b, c) (M \nabla N) a = (b M a) \times (c N a) \quad (11)$$

Then:

$$b (v \nabla id) c = v [c] \times (b id c)$$

$$\Leftrightarrow \{ \text{Khatri-Rao (11) ; function } id \}$$

$$b (v \nabla id) c = v [c] \times (b = c)$$

$$\Leftrightarrow \{ \text{pointwise LAoP (8)} \}$$

$$b (v \nabla id) c = \mathbf{if } b = c \mathbf{ then } v [c] \mathbf{ else } 0$$

i.e. non-zeros can only be found in the **diagonal**.



Linear algebra

Property of diagonal matrices:

$$(v \nabla id) \cdot (u \nabla id) = (v \times u) \nabla id \quad (12)$$

where $M \times N$ is the Hadamard product:

$$b (M \times N) a = (b M a) \times (b N a) \quad (13)$$

Moreover, for f a function, rule

$$f \nabla v = f \cdot (v \nabla id) \quad (14)$$

is easy to derive:

$$\begin{aligned} & b (f \cdot (v \nabla id)) a \\ \Leftrightarrow & \quad \{ \text{composition ; Khatri-Rao} \} \\ & \langle \sum c :: (b f c) \times (v [a] \times (c id a)) \rangle \\ \Leftrightarrow & \quad \{ \text{trading (5) ; cancel } \sum \text{ cf. } c = a \} \\ & (b f a) \times v [a] \\ \Leftrightarrow & \quad \{ \text{Khatri-Rao} \} \\ & b (f \nabla v) a \end{aligned}$$

□




Query = Group by + Join

Query:

```

select
  e_branch,
  e_country,
  sum (j_salary)
from empl, jobs
  where j_code = e_job
group by
  e_country,
  e_branch
order by
  e_country;
  
```

Diagram:



LA semantics:

$$Q = e_{country} \cdot (v \nabla id) \cdot e_{branch}^{\circ} \quad (15)$$

where $v = j^{salary} \cdot j_{code}^{\circ} \cdot e_{job}$



Pointwise semantics

Of vector v first:

$$\begin{aligned}
 & v[k] \\
 = & \quad \{ \text{definition (10)} \} \\
 & 1 (j^{\text{salary}} \cdot j_{\text{code}}^{\circ} \cdot e_{\text{job}}) k \\
 = & \quad \{ \text{matrix multiplication (2)} \} \\
 & \langle \sum_i i :: (1 j^{\text{salary}} i) \times (i (j_{\text{code}}^{\circ} \cdot e_{\text{job}}) k) \rangle \\
 = & \quad \{ \text{trading rules (7) and (5)} \} \\
 & \langle \sum_i i : j_{\text{code}} i = e_{\text{job}} k : (1 j^{\text{salary}} i) \rangle \\
 = & \quad \{ \text{pointwise notation conventions} \} \\
 & \langle \sum_i i : j[i].\text{code} = e[k].\text{job} : j[i].\text{salary} \rangle
 \end{aligned}$$

□



Pointwise semantics

Of the whole query:

$$\begin{aligned}
 & c \ Q \ b \\
 = & \quad \{ \text{definition (15)} ; \text{diagonal } v \triangleright id \} \\
 & \langle \sum k :: (c \ e_{country} \ k) \times (k \ (v \triangleright id) \ k) \times (k \ e_{branch}^{\circ} \ b) \rangle \\
 \Leftrightarrow & \quad \{ \text{trading rule (5)} \} \\
 c \ Q \ b = & \langle \sum k : c = e_{country} \ k \wedge b = e_{branch} \ k : v \ [k] \rangle
 \end{aligned}$$

Putting both together:


$$\begin{aligned}
 query \ (c, b) = & \sum k, i : \\
 & c = e[k].country \wedge b = e[k].branch \wedge j[i].code = e[k].job : \\
 & j[i].salary
 \end{aligned}$$

Rest point :-)



Clearly:

- SQL is a **path-language**
- SQL is **pointfree** — see how the surface language **hides** the double-cursor k, i pointwise **for-loop**.



SQL tries to be as **pointfree** as **natural** language is so, compare

“dogs are mammals”

to the (boring!)

$\langle \forall d : d \in Dog : d \in Mammal \rangle$

We don't **speak** using *cursors*!

Simplification



LA script (15)

$$Q = e_{country} \cdot (v \triangleright id) \cdot e_{branch}^{\circ} \quad \text{where } v = j^{salary} \cdot j_{code}^{\circ} \cdot e_{job}$$

can be simplified into

$$Q = (e_{country} \triangleright v) \cdot e_{branch}^{\circ}$$

thanks to Khatri-Rao law (14). Note how matrix

$e_{country} \triangleright v$	1	2	3	4	5
PT	0	0	0	1100	1000
UK	1000	1000	1333	0	0

nicely combines **qualitative** (functional) with **quantitative** information.

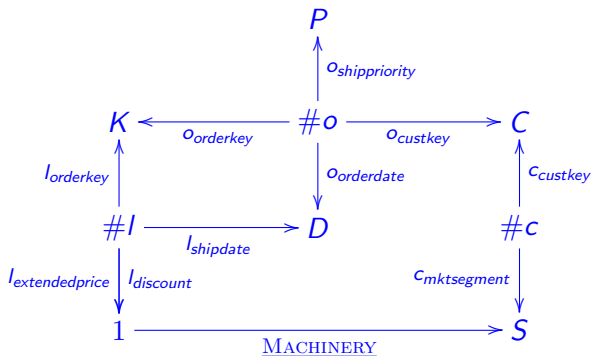
LA script for TPC-H query3



```
query3 =  
select  
  l_orderkey, o_orderdate, o_shippriority;  
  sum (l_extendedprice * (1 - l_discount)) as revenue  
from  
  orders, customer, lineitem  
where  
  c_mktsegment = 'MACHINERY'  
  and c_custkey = o_custkey  
  and l_orderkey = o_orderkey  
  and o_orderdate < date '1995-03-10'  
  and l_shipdate > date '1995-03-10'  
group by  
  l_orderkey, o_orderdate, o_shippriority  
order by  
  revenue desc, o_orderdate;
```



Diagram for TPC-H query3



“Big-plan” **tabulation** again dictated by the **group by** clause:

$$Q = K \xleftarrow{I_{orderkey}} \#I \xleftarrow{X} \#O \xleftarrow{(O_{shippriority} \nabla O_{shipdate})^\circ} P \times D$$

LA semantics for TPC-H query3




Data aggregation is performed over a derived vector

$$revenue = l_{extendedprice} \times (! - l_{discount}) \quad (16)$$

where $! : \#I \rightarrow 1$ is the unique (constant) function of its type — a row vector wholly filled with ones.


We move on:



LA semantics for TPC-H query3



As expected, the link Y between the two tables is the join in the **where** clause:



LA semantics for TPC-H query3



Moving on, clauses

```
o_orderdate < date '1995-03-10'  
and l_shipdate > date '1995-03-10'
```

convert to vectors $v : \#o \rightarrow 1$ and $u : \#l \rightarrow 1$ defined by


```
 $v[i] = \llbracket o[i].orderdate < '1995-03-10' \rrbracket$   
 $u[k] = \llbracket l[k].shipdate > '1995-03-10' \rrbracket$ 
```

recall (9).



LA semantics for TPC-H query3

Altogether, thus far:



where $v[i] = \llbracket o[i].orderdate < '1995-03-10' \rrbracket$

and $u[k] = \llbracket l[k].shipdate > '1995-03-10' \rrbracket$

LA semantics for TPC-H query3



Finally, clauses

$c_mktsegment = 'MACHINERY'$ and $c_custkey = o_custkey$

amount to Boolean path (vector)


$$z = 1 \xleftarrow{MACHINERY^\circ} S \xleftarrow{c_mktsegment} \#c \xleftarrow{c_custkey^\circ} C \xleftarrow{o_custkey} \#o$$

which **counts** how many customers exhibit the specified market segment:

$$z[k] = \langle \sum i : c[i].custkey = o[k].custkey \wedge c[i].mktsegment = MACHINERY : 1 \rangle$$



Query final path



Simplification of (“water fall”) path



Thanks to LA laws:

$$\begin{array}{c}
 Q3 = \quad \#o \xleftarrow{(o_{shippriority} \nabla o_{shipdate})^\circ} P \times D \\
 \quad \quad \quad \downarrow o_{orderkey} \nabla (v \times z) \\
 \quad \quad \quad \#l \xleftarrow{(l_{orderkey})^\circ} K \\
 \quad \quad \quad \downarrow (l_{orderkey}) \nabla (revenue \times u) \\
 \quad \quad \quad K
 \end{array}$$

Notice the same overall pattern: a **join** inside a **tabulation**.

Other simplifications possible, likely changing **performance** — in what sense ?

Divide and conquer



Block linear algebra enables **distributed** evaluation of query paths by “divide & conquer” laws for **all** operators involved, cf.

$$[A|B] \cdot \begin{bmatrix} C \\ D \end{bmatrix} = A \cdot C + B \cdot D \quad (17)$$

$$\begin{bmatrix} A \\ B \end{bmatrix}^\circ = [A^\circ | B^\circ] \quad (18)$$

and

$$[A|B] \triangleright [C|D] = [A \triangleright C | B \triangleright D] \quad (19)$$

$$[A|B] \times [C|D] = [A \times C | B \times D] \quad (20)$$

which generalize to **any finite** number of blocks.

Map-reduce



Overall path splits in two parts,

- Workload over table $\#o$:

$$\begin{array}{c} \#o \xleftarrow{(o_{shippriority} \nabla o_{shipdate})^\circ} P \times D \\ \downarrow o_{orderkey} \nabla (v \times z) \\ K \end{array}$$

- Workload over table $\#l$:

$$\begin{array}{c} \#l \xleftarrow{(l_{orderkey})^\circ} K \\ \downarrow (l_{orderkey}) \nabla (revenue \times u) \\ K \end{array}$$

With n **machines**, each table is divided into n **slices**, each slice residing into its machine.

Map runs the two workloads on each machine, in parallel.

Reduce joins all machine-contributions together, then performing the final composition of the 2 paths.

Summary



Recall the X/Open CAE Specification:

“The result of evaluating a query-specification can be explained in terms of a multi-step algorithm. The order of [the 7] steps in this algorithm follows the mandatory order of the clauses (FROM, WHERE, and so on) of the SELECT statement”

Our **evaluation order** is clearly different !

It is “demand driven” by the **group by** clause.

In theory, everything is **embarrassingly parallel**... but read Rogério’s MSc dissertation ⁸ before getting too excited...

⁸R. Pontes, Benchmarking a Linear Algebra Approach to OLAP (2015)

Practical side of all this



Future (practical) work:

- Define a **DSL** for the LA **path** language
- Mount a **map-reduce** interpreter for such a DSL running on a data-distributed environment
- Write a **compiler** mapping (a subset of) **SQL** to the DSL
- Enjoy experimenting with the overall toy :-)

In particular,

- Compare LA paths with TPC-H query plans
- Complete the benchmark already carried out.⁹

⁹R.Pontes, Benchmarking a Linear Algebra Approach to OLAP (2015).

Theory side of all this



A lot!

- Compare with related work on **columnar** DB systems
- Parametrize DSL on appropriate **semirings** for non arithmetic aggregations (*min*, *max* etc)
- Extend semantic coverage as much as possible, keeping the LA encoding such as e.g. in

$$t_B^\circ \cdot t_B = id$$

expressing **UNIQUE** constraints, or **integrity constraints** such as in e.g.

$$p_F \leq t_K \cdot t_K^\circ \cdot p_F$$

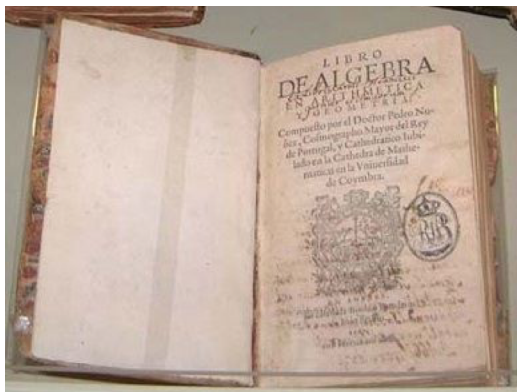
(*K* primary key, *F* foreign key.)

- **Null values** ? ...

Today, as in 1567...



... *quien sabe por Algebra, sabe científicamente*¹⁰



¹⁰(...) *who knows by Algebra knows scientifically* — Pedro Nunes, Libro de Algebra (1567).



Appendix

What about queries without **group by**?




Query:¹¹

```

select
  sum (r_a)
from r, s
  where r_c = s_b and
    5 < r_a < 20 and
    40 < r_b < 50 and
    30 < s_a < 40;
  
```

Star diagram:




Define

```

u i = 5 < r[i].a < 20
v i = 40 < r[i].b < 50
x j = 30 < s[j].a < 40
  
```

in the reduction:




¹¹Example taken from D. Abadi et al, *The Design (...) Systems* (2012).



Faster, this time

Vector $\#s \xrightarrow{!} 1$ models the implicit '**group by all**' clause:



Thanks to (LA)

$$(M \nabla N)^\circ \cdot (P \nabla Q) = (M^\circ \cdot P) \times (N^\circ \cdot Q) \quad (22)$$

$$b (v^\circ \cdot u) a = v[b] \times u[a] \quad (23)$$

$$1 (! \cdot M) a = \langle \sum b :: b M a \rangle \quad (24)$$


we get the expected output scalar:

$$\rho = \langle \sum j, i : u i \wedge v i \wedge r[i].c = s[j].b \wedge x j : r[i].a \rangle$$

Details



Details about the “hidden” tabulation in (21):



$$t = ! \cdot (v \nabla id) \cdot !^\circ$$

$$\Leftrightarrow \{ (14) \}$$

$$t = (v \nabla !) \cdot !^\circ$$

$$\Leftrightarrow \{ ! \text{ is the unit of Khatri-Rao} \}$$

$$t = v \cdot !^\circ$$

$$\Leftrightarrow \{ \text{definition of } \rho \}$$

$$t = \rho$$

□