

Database Languages and their Compilers

Prof. Dr. Torsten Grust

Database Systems Research Group
U Tübingen

Winter 2010



Copyright © 1995 United Feature Syndicate, Inc.
Redistribution in whole or in part prohibited

5 Query Plan Generation

- ▶ If we inspect the algorithms provided by a relational DBMS query engine, we typically will **not** find *fold*.
- ▶ Instead, query engines provide a small set of basic **query operators**, commonly referred to as the **relational algebra**.
- ▶ The operators of the relational algebra are chosen so that
 - ① (**Expressivity**) any SQL query can be computed if we build an appropriate tree of relational algebra operators, and
 - ② (**Efficiency**) each operator may be implemented efficiently on top of the buffer manager, index structures, and file operations provided by the DBMS.
- ▶ This chapter will discuss how we can **generate relational algebra operator trees from MCC expressions**.

5.1 Basic Relational Algebra Operators

- ▶ For our discussion, we will assume the presence of the following query algebra.
- ▶ Since SQL is a query language operating on **bags** of tuples (notable exceptions are `DISTINCT` and `ORDER BY`), we will introduce relational algebra as an algebra over bags.

5.1.1 Access - α

- ▶ Accessing a relation on secondary storage is always the first step before query execution can begin. This is what operator α provides.

- Synopsis:

$$\alpha (id)$$

Argument *id* is the name of a relation registered in the DBMS catalog. Note that bogus values for *id* would have been caught by the type checker (*name-unknown-4*, see Section 3.3).

- Example (see Chapter 0):

α (Mail) =		sender	subject	domain
		'Alice'	'I love you'	' .nl '
		'Bob'	'Re: You are fired'	' .uk '
		'Cathy'	'You are fired'	' .uk '
		'Dave'	'I love you'	' .de '
		⋮	⋮	⋮

5.1.2 Projection – π

- ▶ Operator π projects its input relation onto a list of given attributes or computes new attributes from old ones. Duplicates are retained.

- **Synopsis:**

$$\pi_L (r)$$

The **projection list** L is p_1, p_2, \dots, p_n , where each p_i is of the form $e_i \rightarrow l_i$, l_i being the name of the new result attribute and e_i being an expression involving attribute names (of input r), constants, and arithmetic operators.

- **Example:**

$$\begin{array}{c|c|c|c} r & a & b & c \\ \hline & 0 & 1 & 2 \\ & 0 & 1 & 2 \\ & 3 & 4 & 5 \\ & 0 & 0 & 3 \end{array}, \quad \pi_{a \rightarrow a, b+c \rightarrow x} (r) = \begin{array}{c|c|c} & a & x \\ \hline & 0 & 3 \\ & 0 & 3 \\ & 3 & 9 \\ & 0 & 3 \end{array}.$$

- Projection items $a \rightarrow a$ may be abbreviated by a .

5.1.3 Selection – σ

- ▶ Selection σ operates as a filter that rejects all incoming tuples who fail to satisfy a given predicate (expression of type \mathbb{B}). Duplicates are retained.

- **Synopsis:**

$$\sigma_C (r)$$

The **selection condition** $C :: \mathbb{B}$ may be built from attribute names (present in r), constants, arithmetic and relational ($<$, \leq , $=$, ...) operators, and boolean connectives \neg , \vee , \wedge).

- **Example:**

r	a	b	c
	0	1	2
	0	1	2
	3	4	5
	0	0	3

,
$$\sigma_{a+b>c \vee b=1} (r) =$$

	a	b	c
	0	1	2
	0	1	2
	3	4	5

.

5.1.4 Cross Product – \times

- ▶ Operator \times forms all combinations of tuples – via tuple merging (\parallel) – of its two inputs (cartesian product). Duplicates are retained.

- Synopsis:

$$r \times s$$

Note that, since result tuples are built via \parallel , \times is commutative as well as associative.

- Example:

r	a	b		s	c	d			a	b	c	d
	0	1			1	4			0	1	1	4
	2	3			1	4			0	1	2	5
	2	3			2	5			2	3	1	4
									2	3	1	4
									2	3	2	5
									2	3	1	4
									2	3	1	4
									2	3	2	5

5.1.5 Duplicate Elimination – δ

- ▶ The duplicate eliminator δ removes duplicate tuples – decided via tuple equality (=) – in its input, thus producing a set of tuples.

- Synopsis:

$$\delta (r)$$

- Example:

r	a	b		s	c	d		$\delta (r \times s)$	=	a	b	c	d
	0	1			1	4				0	1	1	4
	2	3			1	4	,			0	1	2	5
	2	3			2	5				2	3	1	4
										2	3	2	5

 **Implementing δ ?**

How would you try to efficiently implement the δ operator?

5.1.6 Grouping – γ

- ▶ γ partitions its input relation using a list of grouping attributes and optionally applies aggregates to each partition.

- **Synopsis:**

$$\gamma_{G|A}(r)$$

$G = g_1, \dots, g_r$ is the list of grouping attributes (of input r). The optional $A = agg_1(a_1) \rightarrow l_1, \dots, agg_n(a_n) \rightarrow l_n$ describes the list of aggregations to perform on each partition ($agg_i \in \{sum, count, max, min\}$). $\{g_i\} \cap \{a_j\} = \emptyset$.

- **Example:**

r	a	b	c
0	1	4	
0	1	2	
2	3	0	
2	3	6	
0	1	1	

,

$$\gamma_{a,b|sum(c) \rightarrow s, max(c) \rightarrow m}(r) =$$

	a	b	s	m
0	1	3	4	
2	3	6	6	

.

5.1.7 Bag Operations – \uplus , \setminus^+ , \uplus

- These binary operators respect the multiplicity of elements in their inputs (\uplus adds multiplicities, \setminus^+ subtracts, \uplus preserves the minimum multiplicity).

■ Synopses:

$$\begin{array}{l} r \uplus s \\ r \setminus^+ s \\ r \uplus s \end{array}$$

The tuple types of the elements of r, s have to be identical.

■ Examples:

r	a	b
0	1	
0	1	
2	3	

s	a	b
0	1	
2	3	
2	3	
4	5	

$r \uplus s$	a	b
0	1	
0	1	
0	1	
2	3	
2	3	
2	3	
2	3	
4	5	

$r \setminus^+ s$	a	b
0	1	

$r \uplus s$	a	b
0	1	
2	3	

5.1.8 Sorting – \mathcal{T}

► Operator \mathcal{T} imposes an ordering on the elements of a collection, thus producing a list of tuples.

■ Synopsis:

$$\mathcal{T}_L (r)$$

$L = s_1, \dots, s_n$, the **sort criteria**, define a lexicographic ordering on the tuples of r ; $s_i = (a, \leq)$ or $s_i = (a, \geq)$ to specify ascending/descending sort order for attribute a .

■ Example:

$$\begin{array}{c|c|c} r & a & b \\ \hline & 1 & 4 \\ & 0 & 1 \\ & 0 & 2 \\ & 1 & 3 \end{array}, \quad \mathcal{T}_{(a,\leq),(b,\geq)} (r) = \begin{array}{c|c|c} & a & b \\ \hline & 0 & 2 \\ & 0 & 1 \\ & 1 & 4 \\ & 1 & 3 \end{array}.$$

5.2 The Construction of Algebraic Operator Trees

- ▶ As input to the plan generation phase we assume a normalized and simplified (see Chapter 4) expression of MCC.
- ▶ Before we present a systematic approach to generate relational algebra operator trees from MCC expressions, let us try to develop an intuition for the problem.

5.2.1 Step 1: Algebraic Operators and MCC

- ▶ In the previous section, we could have introduced the operators of relational algebra by defining them through MCC expressions, e.g., like this:

$$\pi_{e_1 \rightarrow l_1, \dots, e_n \rightarrow l_n} (r) \stackrel{\text{def}}{=} [(l_1 : e_1, \dots, l_n : e_n) \mid v \leftarrow r]^{bag}$$

$$\sigma_p (r) \stackrel{\text{def}}{=} [v \mid v \leftarrow r, p]^{bag}$$

$$r \times s \stackrel{\text{def}}{=} [v_1 \parallel v_2 \mid v_1 \leftarrow r, v_2 \leftarrow s]^{bag}$$


 **Operator δ 's definition in terms of MCC?**

How would you express $\delta(r)$ in the monoid comprehension calculus?

- The same can be done for γ (refer to case (j) in the definition of \mathbb{T}).
- Operators $\uplus, \uplus, \setminus^+$ as well as τ ($= \textit{sort}$) have already been introduced by \mathbb{T} so that these need no further translation.

5.2.2 Step 2: Algebraic Operator Trees and MCC

- Now that we have seen that each operator on its own has a counterpart in MCC, let us try to map more complex MCC expressions to algebraic operator trees.


 **Which tree of algebraic operators is equivalent to the comprehension below $(R :: \textit{bag} (a : \mathbb{S}, b : \mathbb{N}), S :: \textit{bag} (c : \mathbb{N}, d : \mathbb{B}))$?**

$$[(a : x.a, d : y.d) \mid x \leftarrow R, y \leftarrow S, x.b = y.c]^{bag}$$


- ▶ To generalize, remember that **comprehension normalization** generates MCC terms of the following form:

$$[(l_1 : e_1, \dots, l_n : e_n) \mid v_1 \leftarrow r_1, v_2 \leftarrow r_2, \dots, v_r \leftarrow r_r, p]^{bag}$$

- Nested comprehensions have been unnested as far as possible, i.e., most r_i will simply be relation names;
- occurring predicates p_i have been pushed to the far right, with $p = \bigwedge_i p_i$.
- (Left-deep) **Algebraic operator tree**:



- ▶ To verify that this translation from MCC into algebra is indeed correct, simply replace the algebraic operators with their MCC definitions, then normalize. We should get back the original MCC expression.
- ▶ Example (using the query shown at the beginning of this section):



=
MCC definitions

$$\left[(a : a, d : d) \mid v \leftarrow \left[v' \mid v' \leftarrow \underbrace{\left[x \parallel y \mid x \leftarrow R, y \leftarrow S \right]^{bag}}_{\text{unnest via rule (4)}}, b = c \right]^{bag} \right]^{bag}$$

=

$$\left[(a : a, d : d) \mid v \leftarrow \left[v' \mid x \leftarrow R, y \leftarrow S, v' \equiv x \parallel y, b = c \right]^{bag} \right]^{bag}$$

=
remove \equiv

$$\left[(a : a, d : d) \mid v \leftarrow \underbrace{\left[x \parallel y \mid x \leftarrow R, y \leftarrow S, b = c \right]^{bag}}_{\text{unnest via rule (4)}} \right]^{bag}$$

=

$$\left[(a : a, d : d) \mid x \leftarrow R, y \leftarrow S, b = c, v \equiv x \parallel y \right]^{bag}$$

=
remove \equiv

$$\left[(a : a, d : d) \mid x \leftarrow R, y \leftarrow S, b = c \right]^{bag}$$

=
explicit use of variables


$$\left[(a : x.a, d : y.d) \mid x \leftarrow R, y \leftarrow S, x.b = y.c \right]^{bag} .$$

- Complex predicates (in an SQL WHERE clause) may lead to **existential and universal quantifiers** in MCC expressions. Operator σ cannot deal with these directly:

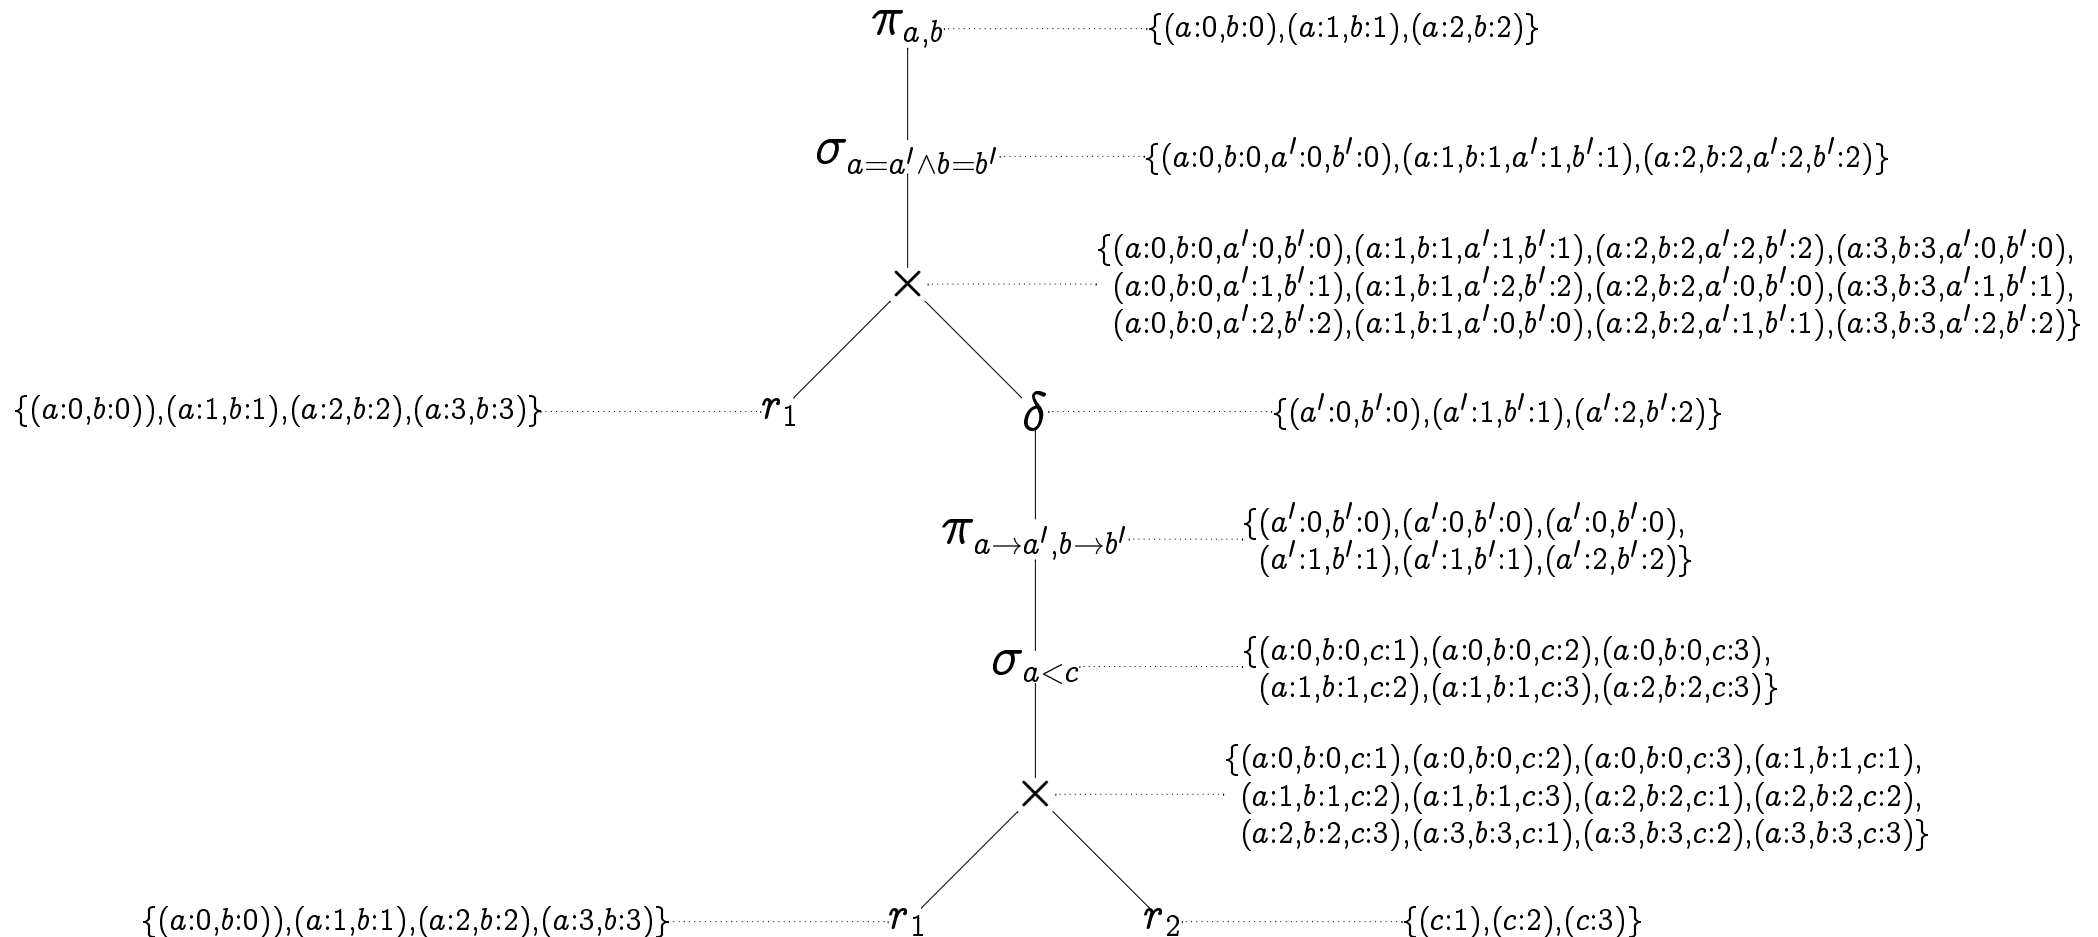
- **Existential quantification:**

$$[(l_1 : e_1, \dots, l_n : e_n) \mid v_1 \leftarrow r_1, [p \mid v_2 \leftarrow r_2]^{some}]^{bag}$$

- **Semantics:** a tuple v_1 of r_1 contributes to the result if there exists **any** v_2 in r_2 so that p is satisfied (if we find two or more v_2 this makes no difference).
- **Algebraic operator tree** ($r_1 :: bag(a_1 : \tau_1, \dots, a_r : \tau_r)$):



$$\mathbb{T} \left(\begin{array}{l} \text{SELECT } x.a, x.b \\ \text{FROM } r_1 \text{ AS } x \\ \text{WHERE EXISTS } y \text{ IN } r_2 : x.a < y.c \end{array} \right) = [(a : x.a, b : x.b) \mid x \leftarrow r_1, [x.a < y.c \mid y \leftarrow r_2]^{some}]^{bag}$$



► **N.B.:**

- ① Adding δ on top of the root of the operator trees makes the previous discussion **applicable to set comprehensions** ($[\mid]^{set}$);
- ② comprehension normalization (5) unnests existential quantifiers ($[\mid]^{some}$) inside *set* comprehensions.

✍ **Universal quantification and relational algebra?**


Try to extend the mapping from MCC to algebraic operator trees so that it correctly generates plans for universal quantifiers:

$$[(l_1 : e_1, \dots, l_n : e_n) \mid v_1 \leftarrow r_1, [p \mid v_2 \leftarrow r_2]^{all}]^{bag} .$$

Hint: $[p \mid qs]^{all} = \neg[\neg p \mid qs]^{some}$.

5.2.3 Mapping MCC to Query Plans

- ▶ Let us complete this chapter by putting our findings about the generation of algebraic operator trees into a systematic framework.



- The missing piece is Plan , a function that maps an expression of MCC to a query plan ready for execution (or further optimization).
- To design Plan , we will need to
 - ① consider the different expression forms that may occur in MCC;
 - ② find an equivalent algebraic operator tree for each of these forms.

► Mapping Plan

Input: MCC expression

Output: Query plan (algebraic operator tree)

$$\text{Plan } (e \uplus e') = \begin{array}{c} \uplus \\ \swarrow \quad \searrow \\ \text{Plan } (e) \quad \text{Plan } (e') \end{array} \quad (8)$$

$$\text{Plan } (e \uparrow e') = \begin{array}{c} \uparrow \\ \swarrow \quad \searrow \\ \text{Plan } (e) \quad \text{Plan } (e') \end{array} \quad (9)$$

$$\text{Plan } (e \uplus e') = \begin{array}{c} \oplus \\ \swarrow \quad \searrow \\ \text{Plan } (e) \quad \text{Plan } (e') \end{array} \quad (10)$$

$$\text{Plan } (\text{sort } os \ e) = \begin{array}{c} \tau_{os} \\ | \\ \text{Plan } (e) \end{array} \quad (11)$$

$$\text{Plan } ([e \mid qs]^{bag}) = \mathbb{P} (\{()\}, [e \mid qs]^{bag}) \quad (12)$$

$$\text{Plan } ([e \mid qs]^{set}) = \begin{array}{c} \delta \\ | \\ \text{Plan } ([e \mid qs]^{bag}) \end{array} \quad (13)$$

$$\text{Plan } (R) = \begin{array}{c} \alpha \\ | \\ R \end{array} \quad (14)$$

► Auxiliary Mapping \mathbb{P}

Input: Partial query plan \times Monoid comprehension

Output: Query plan

$$\mathbb{P}(e, [|]^{bag}) = e \quad (15)$$

$$\mathbb{P}(e, [(l_1 : e_1, \dots, l_n : e_n) |]^{bag}) = \pi_{e_1 \rightarrow l_1, \dots, e_n \rightarrow l_n} \Big|_e \quad (16)$$

$$\mathbb{P}(e, [e' | p, qs]^{bag}) = \mathbb{P}\left(\sigma_p \Big|_e, [e' | qs]^{bag}\right) \quad (17)$$

$$\mathbb{P}(e, [e' | v \leftarrow e'', qs]^{bag}) = \mathbb{P}\left(e \times \text{Plan}(e''), [e' | qs]^{bag}\right) \quad (18)$$

■ Predicate $p :: \mathbb{B}$ above is **simple**, i.e., p contains no quantifiers.

■ **N.B.** $\{()\} \times r = r$ for any relation r .

\mathbb{P} continued on next slide ...

► \mathbb{P} is extended to cover existential and universal quantifiers:

$$\mathbb{P} (e, [e' \mid [p \mid qs]^{some}, qs']^{bag}) = \mathbb{P} \left(\begin{array}{c} \pi_{a_i} \\ \sigma_{\wedge_i a_i = a'_i} \\ \times \\ e \quad \delta \\ \pi_{a_i \rightarrow a'_i} \\ \sigma_p \\ \times \\ e \quad \text{Plan} ([\mid qs]^{bag}) \end{array} \right), [e' \mid qs']^{bag} \quad (19)$$

$$\mathbb{P} (e, [e' \mid [p \mid qs]^{all}, qs']^{bag}) = \mathbb{P} \left(\begin{array}{c} \vee \\ e \quad \pi_{a_i} \\ \sigma_p \\ \times \\ e \quad \text{Plan} ([\mid qs]^{bag}) \end{array} \right), [e' \mid qs']^{bag} \quad (20)$$

■ Let $e :: \text{bag} (a_1 : \tau_1, \dots, a_r : \tau_r)$.

► **Example.** Plan compilation (assume $R :: \text{bag } (a : \mathbb{S}, b : \mathbb{N}), S :: \text{bag } (c : \mathbb{N}, d : \mathbb{B})$):

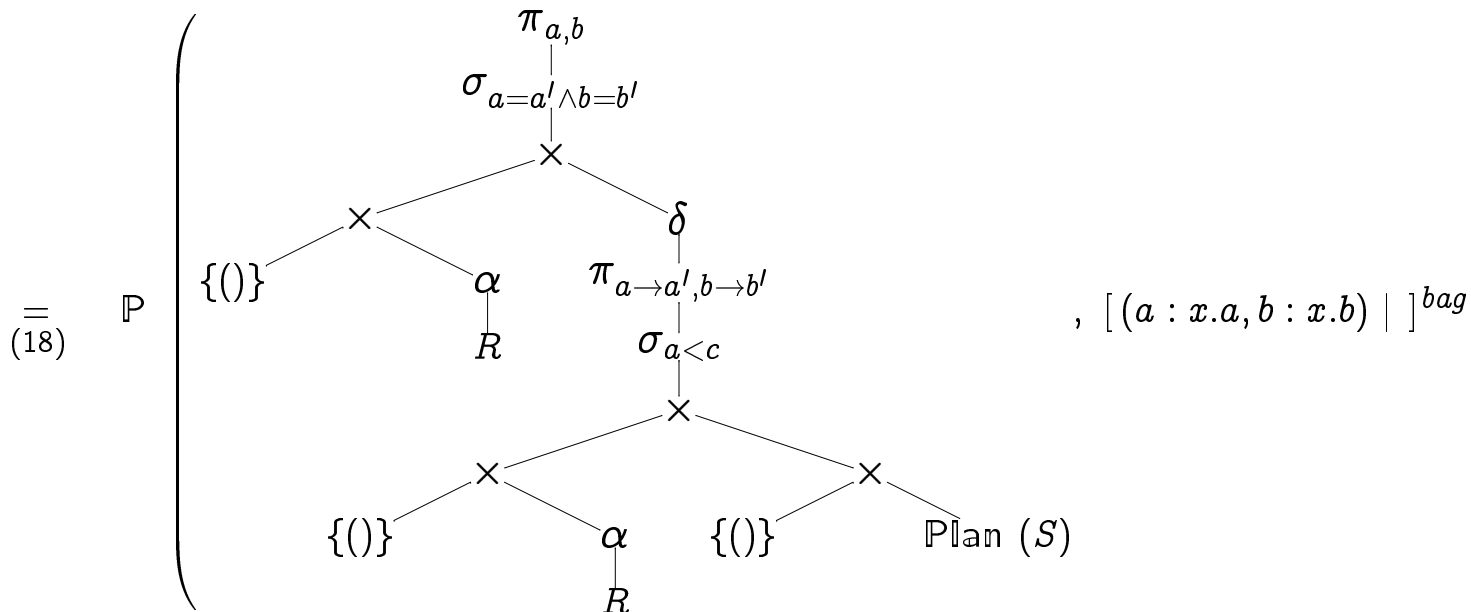
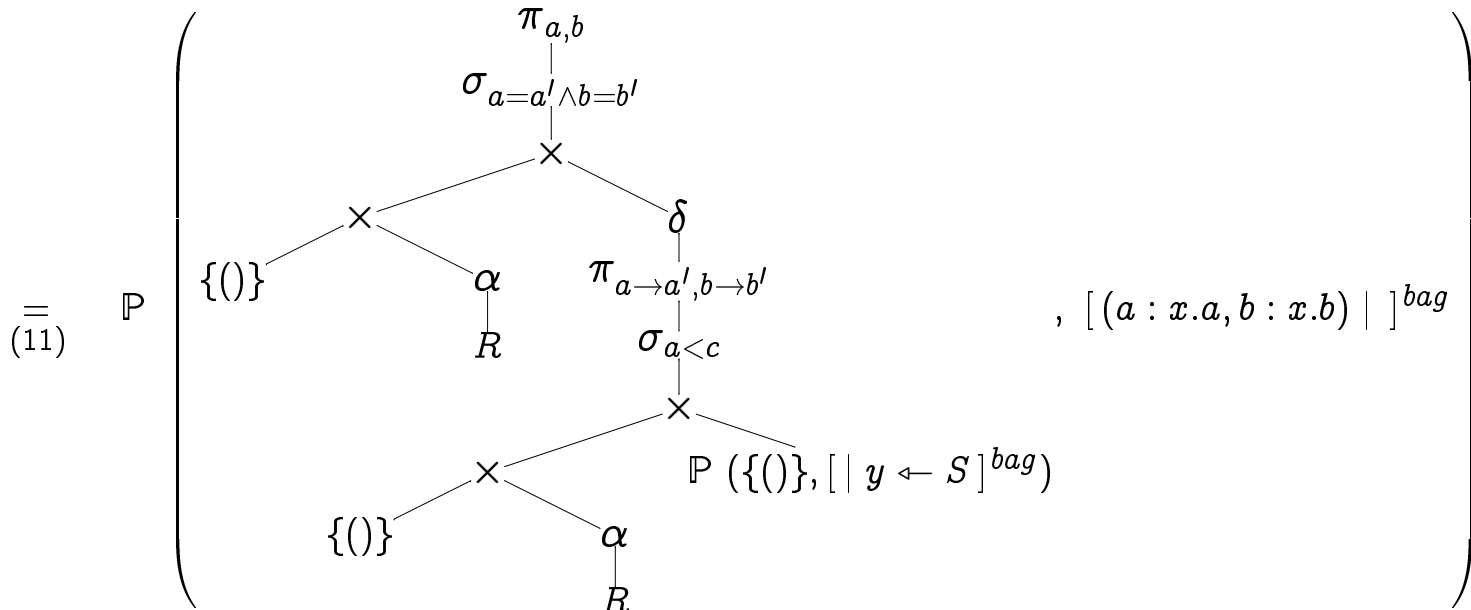
$$\text{Plan } \left([(a : x.a, b : x.b) \mid x \leftarrow R, [x.a < y.c \mid y \leftarrow S]^{some}]^{bag} \right)$$


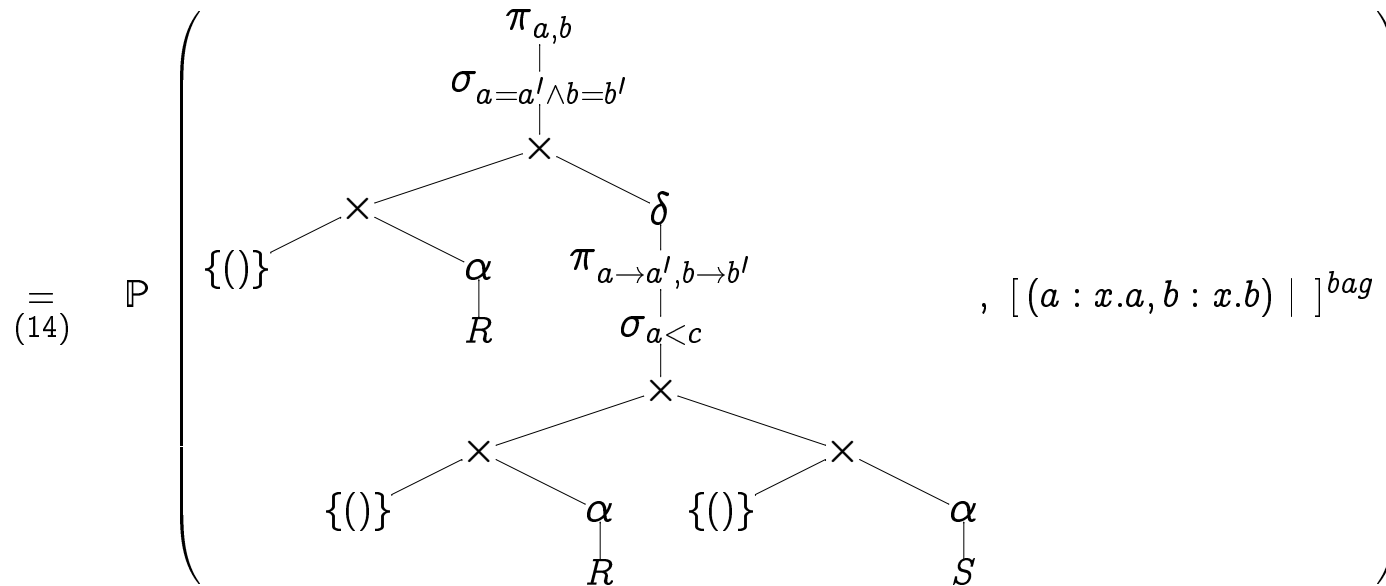
$$\stackrel{=}{(12)} \mathbb{P} \left(\{()\}, [(a : x.a, b : x.b) \mid x \leftarrow R, [x.a < y.c \mid y \leftarrow S]^{some}]^{bag} \right)$$

$$\stackrel{=}{(11)} \mathbb{P} \left(\{()\} \times \text{Plan } (R), [(a : x.a, b : x.b) \mid [x.a < y.c \mid y \leftarrow S]^{some}]^{bag} \right)$$

$$\stackrel{=}{(14)} \mathbb{P} \left(\{()\} \times \begin{array}{c} \alpha \\ R \end{array}, [(a : x.a, b : x.b) \mid [x.a < y.c \mid y \leftarrow S]^{some}]^{bag} \right)$$


$$\stackrel{=}{(19)} \mathbb{P} \left(\{()\} \times \begin{array}{c} \pi_{a,b} \\ \sigma_{a=a' \wedge b=b'} \\ \times \\ \begin{array}{c} \alpha \\ R \end{array} \end{array}, \begin{array}{c} \delta \\ \pi_{a \rightarrow a', b \rightarrow b'} \\ \sigma_{a < c} \\ \times \\ \text{Plan } ([y \leftarrow S]^{bag}) \end{array} \right)$$





5.2.4 Optimizer Desperately Needed!

- ▶ As you can see, the plans generated by $\mathbb{P}lan$ provide obvious opportunities for simplification and optimization in general.
- To facilitate modular design in a query processor, this optimization phase is carried out in a separate unit, the **query optimizer**.
- The optimizer comes with a **rule base** of optimization opportunities, e.g.



- The optimizer selects the next rule to be **fired** using a **heuristic** that characterizes useful or “promising” rules.

 The list goes on ...

Which other **algebraic optimization rules** can you think of?

5.3 Recommended Reading

- [1] H. GARCIA-MOLINA, J. D. ULLMAN, AND J. WIDOM, *Database System Implementation*, Prentice Hall, New Jersey, 2000. ISBN 0-13-040264-8.

Contents

0	Introduction to Query Compilation	1
0.1	Welcome!	1
0.2	Administrativa	3
0.3	Some Remarks on these Slides	4
0.4	Relational Databases and SQL	6
0.5	A Guided Tour through an SQL Query Processor	14
0.5.1	Character Streams and Tokens	15
0.5.2	Identify Valid SQL Syntax	16
0.5.3	Resolve the Meaning of Variables and Identifiers	17
0.5.4	Check Types and Schemas	18

0.5.5	How Could Query Evaluation Look Like?	19
0.5.6	One Query, Many Programs	20
0.5.7	Query Operators	21
0.5.8	Query Operator Trees and Rewriting	22
0.5.9	The Cheaper, the Better: Query Cost Models	23
0.5.10	Executing Operator Trees on a Machine	24
0.6	Recommended Reading	27
1	Query Parsing	28
1.1	The Scanner (Lexer)	28
1.1.1	Regular Expressions	33
1.1.2	Regular Expressions for SQL	35
1.1.3	lex – A Scanner Generator	39

1.2	Syntax Analysis (Parsing)	42
1.2.1	Context-Free Grammars	44
1.2.2	Derivations and Parse Trees	46
1.2.3	Ambiguous Grammars	50
1.2.4	Predictive Parsing and Recursive Descent	54
1.2.5	yacc – A Parser Generator	63
1.3	Recommended Reading	67
2	A Query Calculus for SQL	68
2.1	SQL's Nested Loops Semantics	69
2.2	<i>fold</i>	71
2.3	The Monoid Comprehension Calculus	75
2.4	Mapping SQL to MCC	81

2.4.1	Applying \mathbb{T} to Translate SQL Queries	91
2.5	Recommended Reading	95
3	Variable Scopes and Type Inference for SQL	96
3.1	Variables in SQL	97
3.1.1	Variable Environments	99
3.2	The Type of an SQL Query	102
3.2.1	Type Inference	103
3.3	Semantic Attributes	105
4	Query Normalization	115
4.1	Comprehension Normalization	117
4.2	Further Rewriting Laws	124
4.2.1	Constant Folding	124

4.2.2	Reordering the Qualifier List	126
-------	---	-----

5	Query Plan Generation	129
----------	------------------------------	------------

5.1	Basic Relational Algebra Operators	130
-----	--	-----

5.1.1	Access - α	131
-------	-----------------------------	-----

5.1.2	Projection - π	132
-------	------------------------------	-----

5.1.3	Selection - σ	133
-------	--------------------------------	-----

5.1.4	Cross Product - \times	134
-------	------------------------------------	-----

5.1.5	Duplicate Elimination - δ	135
-------	--	-----

5.1.6	Grouping - γ	136
-------	-------------------------------	-----

5.1.7	Bag Operations - $\uplus, \setminus^+, \oplus$	137
-------	--	-----

5.1.8	Sorting - τ	138
-------	----------------------------	-----

<u>5.2</u>	<u>The Construction of Algebraic Operator Trees</u>	<u>139</u>
------------	---	------------

5.2.1	Step 1: Algebraic Operators and MCC	139
5.2.2	Step 2: Algebraic Operator Trees and MCC	140
5.2.3	Mapping MCC to Query Plans	146
5.2.4	Optimizer Desperately Needed!	153
5.3	Recommended Reading	154