

Query Processing for Non-traditional Applications


CS848 Spring 2013

Cheriton School of CS

Query Plans

Physical Design and Query Compilation: Overview

$$\Sigma = (\Sigma \cup \Sigma' \cup \Sigma'')$$



Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

$$\forall x, y_1, y_2, z_1, z_2. \text{employee}(x, y_1, z_1) \wedge \text{employee}(x, y_2, z_2) \\ \rightarrow (y_1 = y_2 \wedge z_1 = z_2)$$

Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

$$\forall x, y_1, y_2, z_1, z_2. \text{employee}(x, y_1, z_1) \wedge \text{employee}(x, y_2, z_2) \\ \rightarrow (y_1 = y_2 \wedge z_1 = z_2)$$

Physical Design: $S_P = S_A = \{ \text{emp-rid-eid}/2/1, \text{emp-eid-rid}/2/0(/1) \\ \text{emp-rid-name}/2/1, \text{emp-name-rid}/2/1 \\ \text{emp-rid-slry}/2/1, \text{emp-slry-rid}/2/1 \}$

Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

$$\forall x, y_1, y_2, z_1, z_2. \text{employee}(x, y_1, z_1) \wedge \text{employee}(x, y_2, z_2) \\ \rightarrow (y_1 = y_2 \wedge z_1 = z_2)$$

Physical Design: $S_P = S_A = \{ \text{emp-rid-eid}/2/1, \text{emp-eid-rid}/2/0(/1) \\ \text{emp-rid-name}/2/1, \text{emp-name-rid}/2/1 \\ \text{emp-rid-slry}/2/1, \text{emp-slry-rid}/2/1 \}$

$$\forall x, y, z. \text{employee}(x, y, z) \rightarrow \\ \exists a. \text{emp-rid-eid}(a, x) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$$

Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

$$\forall x, y_1, y_2, z_1, z_2. \text{employee}(x, y_1, z_1) \wedge \text{employee}(x, y_2, z_2) \\ \rightarrow (y_1 = y_2 \wedge z_1 = z_2)$$

Physical Design: $S_P = S_A = \{ \text{emp-rid-eid}/2/1, \text{emp-eid-rid}/2/0(/1) \\ \text{emp-rid-name}/2/1, \text{emp-name-rid}/2/1 \\ \text{emp-rid-slry}/2/1, \text{emp-slry-rid}/2/1 \}$

$$\forall x, y, z. \text{employee}(x, y, z) \rightarrow \\ \exists a. \text{emp-rid-eid}(a, x) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z) \\ \forall a, x_1, x_2. \text{emp-rid-eid}(a, x_1) \wedge \text{emp-rid-eid}(a, x_2) \rightarrow x_1 = x_2 \\ \forall a, x. \text{emp-rid-eid}(a, x) \rightarrow \exists y, z. \text{employee}(x, y, z) \\ \forall a, y_1, y_2. \text{emp-rid-name}(a, y_1) \wedge \text{emp-rid-name}(a, y_2) \rightarrow y_1 = y_2 \\ \forall a, y. \text{emp-rid-name}(a, y) \rightarrow \exists x. \text{emp-rid-eid}(a, x) \quad (\text{same for -slry})$$

Example Physical Design: Column Store

Logical Design: $S_L = \{\text{employee}/3\}$

$$\forall x, y_1, y_2, z_1, z_2. \text{employee}(x, y_1, z_1) \wedge \text{employee}(x, y_2, z_2) \\ \rightarrow (y_1 = y_2 \wedge z_1 = z_2)$$

Physical Design: $S_P = S_A = \{ \text{emp-rid-eid}/2/1, \text{emp-eid-rid}/2/0(/1) \\ \text{emp-rid-name}/2/1, \text{emp-name-rid}/2/1 \\ \text{emp-rid-slry}/2/1, \text{emp-slry-rid}/2/1 \}$

$$\forall x, y, z. \text{employee}(x, y, z) \rightarrow \\ \exists a. \text{emp-rid-eid}(a, x) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z) \\ \forall a, x_1, x_2. \text{emp-rid-eid}(a, x_1) \wedge \text{emp-rid-eid}(a, x_2) \rightarrow x_1 = x_2 \\ \forall a, x. \text{emp-rid-eid}(a, x) \rightarrow \exists y, z. \text{employee}(x, y, z) \\ \forall a, y_1, y_2. \text{emp-rid-name}(a, y_1) \wedge \text{emp-rid-name}(a, y_2) \rightarrow y_1 = y_2 \\ \forall a, y. \text{emp-rid-name}(a, y) \rightarrow \exists x. \text{emp-rid-eid}(a, x) \quad (\text{same for -slry}) \\ \forall a, x. \text{emp-rid-eid}(a, x) \leftrightarrow \text{emp-eid-rid}(x, a) \quad (\text{same for -name, -slry})$$

Queries over Column Store Physical Design

1 `employee(x, y, z)`

$\exists a. \text{emp-rid-rid}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

Queries over Column Store Physical Design

1 $\text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

2 $\exists y, z. \text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a)$

Queries over Column Store Physical Design

1 $\text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

2 $\exists y, z. \text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a)$

3 $\exists z. \text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a) \wedge \text{emp-rid-name}(a, y)$

Queries over Column Store Physical Design

1 $\text{employee}(x, y, z)$

$\exists a. \text{emp-rid-name}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

2 $\exists y, z. \text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a)$

3 $\exists z. \text{employee}(x, y, z)$

$\exists a. \text{emp-rid-rid}(x, a) \wedge \text{emp-rid-name}(a, y)$

4 $\exists x. \text{employee}(x, y, z)\{y\}$

$\{\exists a. \text{emp-name-rid}(y, a) \wedge \text{emp-rid-slry}(a, z)\}$

Queries over Column Store Physical Design

1 $\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

2 $\exists y, z.\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a)$

3 $\exists z.\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a) \wedge \text{emp-rid-name}(a, y)$

4 $\exists x.\text{employee}(x, y, z)\{y\}$


$\{\exists a.\text{emp-name-rid}(y, a) \wedge \text{emp-rid-slry}(a, z)\}$

Issues to resolve (today)


- Why are the above *plans* implement the *user queries*?
- What “formulas” do qualify as *plans*?
 \Rightarrow how do we interpret *logical connectives* as programs?
- Are all (desired) *plans* captured by *appropriate formulas*?

Physical Design and Query Compilation: Plans

$$\Sigma = (\Sigma \cup \Sigma' \cup \Sigma'')$$



Physical Design and Query Compilation: Plans



We consider the structure of Q' :

- what formulas can be interpreted as *plans*
⇒ how do we deal with *SETS* in programs?
- what additional *non-logical* operations can be used.

Outline

- 1 Iterator Protocols to communicate Sets (review)
- 2 Atomic Plan Operations: Access Paths (review)
- 3 Logical Connectives/Quantifiers as Plan Operators
- 4 Beyond Logical Operators: Dealing with Duplicates

ACME Case: Access Path Code Templates

Pseudo-code templates realizing a *first/next* protocol for `emp-array0` might be given as follows (variables would be renamed for each occurrence of `emp-array0` in a query plan).

```
function emp-array0-first
   $i := 0$ 
  return emp-array0-next
```

```
function emp-array0-next
   $i := i + 1$ 
  if ( $i > n$ ) return false
   $x_1 := \text{emp-array}[i].\text{emp-salary}$ 
   $x_2 := \text{emp-array}[i].\text{emp-num}$ 
   $x_3 := \text{emp-array}[i].\text{emp-name}$ 
  return true
```

ACME Case: Access Path Code Templates

Pseudo-code templates realizing a *first/next* protocol for `emp-array0` might be given as follows (variables would be renamed for each occurrence of `emp-array0` in a query plan).

```
function emp-array0-first
   $i := 0$ 
  return emp-array0-next
```

```
function emp-array0-next
   $i := i + 1$ 
  if ( $i > n$ ) return false
   $x_1 := emp-array[i].emp-salary$ 
   $x_2 := emp-array[i].emp-num$ 
   $x_3 := emp-array[i].emp-name$ 
  return true
```

Assumes a global state recording bindings of (possible copies of) variables.

- 1 x_1 , x_2 and x_3 to communicate the contents of `emp-array`.
- 2 i and n to record scanning status and size of `emp-array`.

Note: Code templates for access paths must be provided by ACME's DBA department.

Iterator Execution

Hereon, assume **C** denotes the following code that *prints a line* for each result computed by a query plan Q , where $(\text{In}(Q) \cup \text{Out}(Q)) = \{x_1, \dots, x_m\}$.

```
if Q-first
  repeat
    printline("x1" = x1, ..., "xm" = xm)
  until not Q-next
```

Iterator Execution

Hereon, assume **C** denotes the following code that *prints a line* for each result computed by a query plan Q , where $(\text{In}(Q) \cup \text{Out}(Q)) = \{x_1, \dots, x_m\}$.

```
if Q-first
  repeat
    printline("x1" = x1, ..., "xm" = xm)
  until not Q-next
```

Let Q be a query plan that scans `emp-array`.

```
emp-array0(x3, x1, x2)
```

Iterator Execution

Hereon, assume **C** denotes the following code that *prints a line* for each result computed by a query plan Q , where $(\text{In}(Q) \cup \text{Out}(Q)) = \{x_1, \dots, x_m\}$.

```
if Q-first
  repeat
    printline("x1" = x1, ..., "xm" = xm)
  until not Q-next
```

Let Q be a query plan that scans `emp-array`.

`emp-array0(x3, x1, x2)`

Running **C** for a database (given by interpretation) \mathcal{I} produces the following output.

$$\begin{aligned} x_1 &= e_{1,1}, & x_2 &= e_{1,2}, & x_3 &= e_{1,3} \\ x_1 &= e_{2,1}, & x_2 &= e_{2,2}, & x_3 &= e_{2,3} \\ & \vdots \\ x_1 &= e_{n,1}, & x_2 &= e_{n,2}, & x_3 &= e_{n,3} \end{aligned}$$

Access Paths

The *access paths* $S_A \subseteq S_P$ are predicate symbols *associated with a physical capability* realized by an iterator implementation.

⇒ some attributes can be designated as *parameters*
(by convention the left-most ones)

Access Paths

The *access paths* $S_A \subseteq S_P$ are predicate symbols associated with a *physical capability* realized by an iterator implementation.

⇒ some attributes can be designated as *parameters*
(by convention the left-most ones)

Requirements (for access path $AP/k + I/I$):

- given a *fixed values for parameters* there are only finitely many answers (bindings) to the remaining variables, i.e., the set

$$\{a_1, \dots, a_k \mid \mathcal{I}, \mathcal{V} \models \exists x_1, \dots, x_l. AP(x_1, \dots, x_l, y_1, \dots, y_k) \wedge (\bigwedge x_i = p_i)\}$$

(where $\mathcal{V} = [y_1 = a_1, \dots, y_k = a_k]$) is finite.

- the invocation of the iterator protocol *outputs* all and only the valuations \mathcal{V} that satisfy the condition above.

Access Paths

The *access paths* $S_A \subseteq S_P$ are predicate symbols associated with a *physical capability* realized by an iterator implementation.

⇒ some attributes can be designated as *parameters*
(by convention the left-most ones)

Requirements (for access path $AP/k + I/I$):

- given a *fixed values for parameters* there are only finitely many answers (bindings) to the remaining variables, i.e., the set

$$\{a_1, \dots, a_k \mid \mathcal{I}, \mathcal{V} \models \exists x_1, \dots, x_l. AP(x_1, \dots, x_l, y_1, \dots, y_k) \wedge (\bigwedge x_i = p_i)\}$$

(where $\mathcal{V} = [y_1 = a_1, \dots, y_k = a_k]$) is finite.

- the invocation of the iterator protocol *outputs* all and only the valuations \mathcal{V} that satisfy the condition above.

A *plan interpretation* \mathcal{I} satisfies the above for every access path in S_A .

(More Esoteric) Access Paths

- 1 Built-in “operations”:
 - arithmetic (`plus/3/2`, `times/3/2`, etc.)
 - string manipulation (`concat/3/2`, `substr/4/3`, etc.)
 - ...
- 2 data type tests (`is-integer/1/1`)
- 3 pointer dereference and field extraction from records
- 4 (page) reads from external storage
- 5 ...

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

$$Q ::= (Q_1 \wedge Q_2) \quad (\textit{nested loop join})$$

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

$$\begin{array}{l} Q ::= (Q_1 \wedge Q_2) \quad (\textit{nested loop join}) \\ \quad | \exists x. Q, \text{ where } x \in V \quad (\textit{duplicate preserving projection}) \end{array}$$

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

Q	$::=$	$(Q_1 \wedge Q_2)$	<i>(nested loop join)</i>
		$\exists x.Q$, where $x \in V$	<i>(duplicate preserving projection)</i>
		$\{Q\}$	<i>(duplicate elimination)</i>

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

Q	$::=$	$(Q_1 \wedge Q_2)$	<i>(nested loop join)</i>
		$\exists x.Q, \text{ where } x \in V$	<i>(duplicate preserving projection)</i>
		$\{Q\}$	<i>(duplicate elimination)</i>
		$[Q]_i$	<i>(cut introduction)</i>

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

Q	$::=$	$(Q_1 \wedge Q_2)$	<i>(nested loop join)</i>
		$\exists x.Q$, where $x \in V$	<i>(duplicate preserving projection)</i>
		$\{Q\}$	<i>(duplicate elimination)</i>
		$[Q]_i$	<i>(cut introduction)</i>
		$!_i$	<i>(named cut)</i>

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

$Q ::=$	$(Q_1 \wedge Q_2)$	<i>(nested loop join)</i>
	$\exists x. Q$, where $x \in V$	<i>(duplicate preserving projection)</i>
	$\{Q\}$	<i>(duplicate elimination)</i>
	$[Q]_i$	<i>(cut introduction)</i>
	$!_i$	<i>(named cut)</i>

Also require any query plan to satisfy two conditions.

- 1 If $Q = "(Q_1 \wedge Q_2)"$ then $(Fv(Q_1) \cap Out(Q_2)) = \emptyset$.

Conjunctive Query Plans: Syntax

The *conjunctive query plans* induced by S are as follows:

$Q ::=$	$(Q_1 \wedge Q_2)$	<i>(nested loop join)</i>
	$\exists x.Q$, where $x \in V$	<i>(duplicate preserving projection)</i>
	$\{Q\}$	<i>(duplicate elimination)</i>
	$[Q]_i$	<i>(cut introduction)</i>
	$!_i$	<i>(named cut)</i>

Also require any query plan to satisfy two conditions.

- 1 If $Q = "(Q_1 \wedge Q_2)"$ then $(Fv(Q_1) \cap Out(Q_2)) = \emptyset$.
- 2 If $Q = "\exists x.Q_1"$ then $x \notin In(Q_1)$.

Parameters and User Query Embeddings

Input variables, output variables and user query mapping are as follows, where $\text{Param}(\text{Uq}(Q)) = \text{In}(Q)$ always holds.

$$\text{In}(Q) = \begin{cases} \text{In}(Q_1) \cup (\text{In}(Q_2) - \text{Out}(Q_1)) & \text{if } Q = "(Q_1 \wedge Q_2)", \\ \text{In}(Q_1) & \text{if } Q = "\exists x.Q_1", "\{Q_1\}", \text{ or } "[Q_1]_i", \text{ and} \\ \emptyset & \text{if } Q = "!_j". \end{cases}$$

$$\text{Out}(Q) = \begin{cases} \text{Out}(Q_1) \cup \text{Out}(Q_2) & \text{if } Q = "(Q_1 \wedge Q_2)", \\ \text{Out}(Q_1) \setminus \{x\} & \text{if } Q = "\exists x.Q_1", \\ \text{Out}(Q_1) & \text{if } Q = "\{Q_1\}" \text{ or } "[Q_1]_i", \text{ and} \\ \emptyset & \text{if } Q = "!_j". \end{cases}$$

$$\text{Uq}(Q) = \begin{cases} (\text{Uq}(Q_1) \wedge \text{Uq}(Q_2)) & \text{if } Q = "(Q_1 \wedge Q_2)", \\ \exists x. \text{Uq}(Q_1) & \text{if } Q = "\exists x.Q_1", \\ \text{Uq}(Q_1) & \text{if } Q = "\{Q_1\}" \text{ or } "[Q_1]_i", \text{ and} \\ \text{true} & \text{if } Q = "!_j". \end{cases}$$

Conjunctive Query Plans: Semantics

```
function  $(Q_1 \wedge Q_2)$ -first
  if not  $Q_1$ -first return false
  while not  $Q_2$ -first do
    if not  $Q_1$ -next return false
  return true
```

```
function  $(Q_1 \wedge Q_2)$ -next
  if  $Q_2$ -next return true
  while  $Q_1$ -next do
    if  $Q_2$ -first return true
  return false
```

```
function  $(\exists x.Q_1)$ -first
  return  $Q_1$ -first
```

```
function  $(\exists x.Q_1)$ -next
  return  $Q_1$ -next
```

```
function  $\{Q_1\}$ -first
  if not exists store  $S$ 
    create S
  if  $Q_1$ -first
    empty S
    add  $\langle x_1, \dots, x_n \rangle$  to S
  return true
return false
```

```
function  $\{Q_1\}$ -next
  while  $Q_1$ -next do
    if not  $\langle x_1, \dots, x_n \rangle \in S$ 
      add  $\langle x_1, \dots, x_n \rangle$  to S
  return true
return false
```

Conjunctive Query Plans: Semantics

```
function ( $[Q_1]_i$ )-first  
  cut $i$  := false  
  return  $Q_1$ -first
```

```
function ( $[Q_1]_i$ )-next  
  if cut $i$  return false  
  return  $Q_1$ -next
```

```
function (! $i$ )-first  
  cut $i$  := true  
  return true
```

```
function (! $i$ )-next  
  return false
```

Comparing and Assigning

Hereon, we assume a given physical signature includes two additional access paths.

$$\{\text{compare}/2/2, \text{assign}/2/1\} \subseteq S_A$$

Comparing and Assigning

Hereon, we assume a given physical signature includes two additional access paths.

$$\{\text{compare}/2/2, \text{assign}/2/1\} \subseteq S_A$$

Semantics of $\text{compare}(x_1, x_2)$ and $\text{assign}(x_1, x_2)$ is given as follows.

```
function compare-first
  if  $x_1 = x_2$  return true
  return false
```

```
function compare-next
  return false
```

```
function assign-first
   $x_2 := x_1$ 
  return true
```

```
function assign-next
  return false
```

Comparing and Assigning

Hereon, we assume a given physical signature includes two additional access paths.

$$\{\text{compare}/2/2, \text{assign}/2/1\} \subseteq S_A$$

Semantics of $\text{compare}(x_1, x_2)$ and $\text{assign}(x_1, x_2)$ is given as follows.

```
function compare-first
  if  $x_1 = x_2$  return true
  return false
```

```
function compare-next
  return false
```

```
function assign-first
   $x_2 := x_1$ 
  return true
```

```
function assign-next
  return false
```

Also assume any theory Σ includes the following.

$$\forall x_1, x_2. \text{compare}(x_1, x_2) \equiv \text{assign}(x_1, x_2) \equiv (x_1 \approx x_2)$$

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z .*

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

Execution proceeds as follows.

- 1 Use access path `emp-array0` to scan `emp-array` (an atomic query subplan).

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

Execution proceeds as follows.

- 1 Use access path `emp-array0` to scan `emp-array` (an atomic query subplan).
- 2 For each element of `emp-array` returned by this scan, compare the `salary` field with the supplied parameter value z (by using an operator for nested cross product coupled with another atomic query subplan using access path `compare`).

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

Execution proceeds as follows.

- 1 Use access path `emp-array0` to scan `emp-array` (an atomic query subplan).
- 2 For each element of `emp-array` returned by this scan, compare the `salary` field with the supplied parameter value z (by using an operator for nested cross product coupled with another atomic query subplan using access path `compare`).
- 3 If this comparison evaluates to `true` then add the contents of this element to the result.

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

Execution proceeds as follows.

- 1 Use access path `emp-array0` to scan `emp-array` (an atomic query subplan).
- 2 For each element of `emp-array` returned by this scan, compare the `salary` field with the supplied parameter value z (by using an operator for nested cross product coupled with another atomic query subplan using access path `compare`).
- 3 If this comparison evaluates to `true` then add the contents of this element to the result.

Note: (a) Duplicate preserving projection has no effect on execution

ACME Case: Scanning and Selection

Consider where the APS department needs a plan that will *report the employee-number x and name y of each employee that has a given salary z* . A query plan using access path `emp-array0` can now be formulated.

$$\exists u.(\text{emp-array0}(u, x, y) \wedge \text{compare}(z, u))$$

Input and output variables: $\{z\}$ and $\{x, y\}$, respectively.

Execution proceeds as follows.

- 1 Use access path `emp-array0` to scan `emp-array` (an atomic query subplan).
- 2 For each element of `emp-array` returned by this scan, compare the `salary` field with the supplied parameter value z (by using an operator for nested cross product coupled with another atomic query subplan using access path `compare`).
- 3 If this comparison evaluates to `true` then add the contents of this element to the result.

Note: (a) Duplicate preserving projection has no effect on execution, and (b) duplicate elimination is not required since employees must have unique employee numbers.

ACME Case: Scanning and Cutting

Consider where the APS department needs a plan that will *report the* name y *of any* employee *that has a given* salary z *and* employee-number x .

ACME Case: Scanning and Cutting

Consider where the APS department needs a plan that will *report the name y of any employee that has a given salary z and employee-number x* . A query plan Q using access path `emp-array0` can also be formulated.

$$\exists u, v. (\text{emp-array0}(u, v, y) \wedge \text{compare}(x, v) \wedge \text{compare}(z, u))$$

Input and output variables: $\{y\}$ and $\{x, z\}$, respectively.

ACME Case: Scanning and Cutting

Consider where the APS department needs a plan that will *report the* name y of any employee that has a given salary z and employee-number x . A query plan Q using access path `emp-array0` can also be formulated.

$$\exists u, v. (\text{emp-array0}(u, v, y) \wedge \text{compare}(x, v) \wedge \text{compare}(z, u))$$

Input and output variables: $\{y\}$ and $\{x, z\}$, respectively. Pseudo-code templates realizing a *first/next* protocol for Q is then defined as follows.

```
function Q-first
  i := 0
  while i < n do
    i := i + 1
    u := emp-array[i].emp-salary
    v := emp-array[i].emp-num
    y := emp-array[i].emp-name
    if x = v
      if z = u return true
  return false
```

```
function Q-next
  while i < n do
    i := i + 1
    u := emp-array[i].emp-salary
    v := emp-array[i].emp-num
    y := emp-array[i].emp-name
    if x = v
      if z = u return true
  return false
```


ACME Case: Scanning and Cutting

Recall that employees have unique employee numbers.

ACME Case: Scanning and Cutting

Recall that employees have unique employee numbers. Can therefore add cut introduction and named cut operators to Q to improve performance.

$$\exists u, v. ([\text{emp-array0}(u, v, y) \wedge \text{compare}(x, v)]_1 \wedge !_1 \wedge \text{compare}(z, u))$$

ACME Case: Scanning and Cutting

Recall that employees have unique employee numbers. Can therefore add cut introduction and named cut operators to Q to improve performance.

$$\exists u, v. ([\text{emp-array}0(u, v, y) \wedge \text{compare}(x, v)]_1 \wedge !_1 \wedge \text{compare}(z, u))$$

Pseudo-code templates realizing a *first/next* protocol for Q are then modified as follows.

```
function  $Q$ -first
   $i := 0$ 
  while  $i < n$  do
     $i := i + 1$ 
     $u := \text{emp-array}[i].\text{emp-salary}$ 
     $v := \text{emp-array}[i].\text{emp-num}$ 
     $y := \text{emp-array}[i].\text{emp-name}$ 
    if  $x = v$ 
      if  $z = u$  return true
    return false
return false
```

```
function  $Q$ -next
  return false
```

ACME Case: Eliminating Duplicates

Consider where the APS department needs a plan that will *find integers z that occur as the salary value for some employee.*

ACME Case: Eliminating Duplicates

Consider where the APS department needs a plan that will *find integers z that occur as the salary value for some `employee`*. A user query that specifies this requirement for OPTION 1 is as follows.

$$\exists x, y. \text{employee}(x, y, z)$$

ACME Case: Eliminating Duplicates

Consider where the APS department needs a plan that will *find integers z that occur as the salary value for some `employee`*. A user query that specifies this requirement for OPTION 1 is as follows.

$$\exists x, y. \text{employee}(x, y, z)$$

A query plan using access path `emp-array0` implementing this query is as follows.

$$\{\exists x, y. \text{emp-array0}(z, x, y)\}$$

ACME Case: Eliminating Duplicates

Consider where the APS department needs a plan that will *find integers z that occur as the salary value for some `employee`*. A user query that specifies this requirement for OPTION 1 is as follows.

$$\exists x, y. \text{employee}(x, y, z)$$

A query plan using access path `emp-array0` implementing this query is as follows.

$$\{\exists x, y. \text{emp-array0}(z, x, y)\}$$

The plan uses a top-level duplicate elimination operation that will use temporary store to accumulate new salary values as they occur in a scan of `emp-array` with access path `emp-array0`.

General Query Plans: Syntax

The *query plans* induced by S add two final productions:

$$Q ::= (Q_1 \vee Q_2) \quad (\textit{concatenation})$$

General Query Plans: Syntax

The *query plans* induced by S add two final productions:

$$\begin{array}{l} Q ::= (Q_1 \vee Q_2) \quad (\textit{concatenation}) \\ \quad | \neg Q \quad \quad \quad (\textit{simple complement}) \end{array}$$

General Query Plans: Syntax

The *query plans* induced by S add two final productions:

$$Q ::= (Q_1 \vee Q_2) \quad (\textit{concatenation}) \\ | \neg Q \quad (\textit{simple complement})$$

Also require any query plan to satisfy two additional conditions.

- 1 If $Q = "(Q_1 \vee Q_2)"$ then $\text{Out}(Q_1) = \text{Out}(Q_2)$.

General Query Plans: Syntax

The *query plans* induced by S add two final productions:

$$Q ::= (Q_1 \vee Q_2) \quad (\textit{concatenation}) \\ | \neg Q \quad (\textit{simple complement})$$

Also require any query plan to satisfy two additional conditions.

- 1 If $Q = "(Q_1 \vee Q_2)"$ then $\text{Out}(Q_1) = \text{Out}(Q_2)$.
- 2 If $Q = "\neg Q_1"$ then $\text{Out}(Q_1) = \emptyset$.

Parameters and User Query Embeddings

Input variables, output variables and user query mapping are extended as follows, where $\text{Param}(\text{Uq}(Q)) = \text{In}(Q)$ always holds.

$$\text{In}(Q) = \begin{cases} \text{In}(Q_1) \cup \text{In}(Q_2) & \text{if } Q = "(Q_1 \vee Q_2)", \text{ and} \\ \text{In}(Q_1) & \text{if } Q = "\neg Q_1". \end{cases}$$

$$\text{Out}(Q) = \begin{cases} \text{Out}(Q_1) \cap \text{Out}(Q_2) & \text{if } Q = "(Q_1 \vee Q_2)", \text{ and} \\ \emptyset & \text{if } Q = "\neg Q_1". \end{cases}$$

$$\text{Uq}(Q) = \begin{cases} (\text{Uq}(Q_1) \vee \text{Uq}(Q_2)) & \text{if } Q = "(Q_1 \vee Q_2)", \text{ and} \\ \neg \text{Uq}(Q_1) & \text{if } Q = "\neg Q_1". \end{cases}$$

General Query Plans: Semantics

```
function  $(Q_1 \vee Q_2)$ -first
   $(Q_1 \vee Q_2)$ -flag := true
  if  $Q_1$ -first return true
   $(Q_1 \vee Q_2)$ -flag := false
  return  $Q_2$ -first
```

```
function  $(Q_1 \vee Q_2)$ -next
  if  $(Q_1 \vee Q_2)$ -flag
    if  $Q_1$ -next return true
   $(Q_1 \vee Q_2)$ -flag := false
  return  $Q_2$ -next
```

```
function  $(\neg Q_1)$ -first
  if  $Q_1$ -first return false
  return true
```

```
function  $(\neg Q_1)$ -next
  return false
```

ACME Case: Concatenation

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a salary matching either the parameter p_1 or the parameter p_2 .*

ACME Case: Concatenation

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a salary matching either the parameter p_1 or the parameter p_2 .*

A query plan using access path `emp-array0` implementing this query is as follows.

$$\left\{ \begin{array}{l} \exists y, z. (\text{emp-array0}(z, x, y) \wedge \text{compare}(p_1, z)) \\ \vee \exists u, v. (\text{emp-array0}(v, x, u) \wedge \text{compare}(p_2, v)) \end{array} \right\}$$

ACME Case: Concatenation

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a salary matching either the parameter p_1 or the parameter p_2 .*

A query plan using access path `emp-array0` implementing this query is as follows.

$$\left\{ \begin{array}{l} \exists y, z. (\text{emp-array0}(z, x, y) \wedge \text{compare}(p_1, z)) \\ \vee \exists u, v. (\text{emp-array0}(v, x, u) \wedge \text{compare}(p_2, v)) \end{array} \right\}$$

An execution proceeds as follows.

- 1 Scan `emp-array` and add the employee number of any employee with a salary given by input parameter p_1 to a temporary store S if not already there.

ACME Case: Concatenation

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a salary matching either the parameter p_1 or the parameter p_2 .*

A query plan using access path `emp-array0` implementing this query is as follows.

$$\{(\exists y, z.(\text{emp-array0}(z, x, y) \wedge \text{compare}(p_1, z)) \vee \exists u, v.(\text{emp-array0}(v, x, u) \wedge \text{compare}(p_2, v)))\}$$

An execution proceeds as follows.

- 1 Scan `emp-array` and add the employee number of any employee with a salary given by input parameter p_1 to a temporary store S if not already there.
- 2 Scan `emp-array` for a second time and add the employee number of any employee with a salary given by input parameter p_2 to S if not already there.

ACME Case: Concatenation

Note: The user query mapping of the plan is a *positive query* (since it is a union of two conjunctive queries).

ACME Case: Concatenation

Note: The user query mapping of the plan is a *positive query* (since it is a union of two conjunctive queries).

An alternative plan can be formulated that avoids the need for two scans of `emp-array`.

$$\{\exists y, z. (\text{emp-array}_0(z, x, y) \wedge (\text{compare}(p_1, z) \vee \text{compare}(p_2, z)))\}$$

The plan illustrates a common plan idiom for determining if a given value occurs in a given small fixed set of values.

ACME Case: Concatenation

Note: The user query mapping of the plan is a *positive query* (since it is a union of two conjunctive queries).

An alternative plan can be formulated that avoids the need for two scans of `emp-array`.

$$\{\exists y, z. (\text{emp-array0}(z, x, y) \wedge (\text{compare}(p_1, z) \vee \text{compare}(p_2, z)))\}$$

The plan illustrates a common plan idiom for determining if a given value occurs in a given small fixed set of values.

Two possible reasons for requiring top-level duplicate elimination.

- 1 Individual employee numbers may be related to more than one salary or employee name.

ACME Case: Concatenation

Note: The user query mapping of the plan is a *positive query* (since it is a union of two conjunctive queries).

An alternative plan can be formulated that avoids the need for two scans of `emp-array`.

$$\{\exists y, z. (\text{emp-array}_0(z, x, y) \wedge (\text{compare}(p_1, z) \vee \text{compare}(p_2, z)))\}$$

The plan illustrates a common plan idiom for determining if a given value occurs in a given small fixed set of values.

Two possible reasons for requiring top-level duplicate elimination.

- 1 Individual employee numbers may be related to more than one salary or employee name.
- 2 Parameters p_1 and p_2 may not be distinct.

ACME Case: Concatenation

The first reason is ruled out by the logical design of `payroll`.

ACME Case: Concatenation

The first reason is ruled out by the logical design of `payroll`.

The second reason can be ruled out by modifying the plan to ensure the second subplan for the concatenation operator will only return additional results when p_1 and p_2 are distinct.

$$\exists y, z. (\text{emp-array0}(z, x, y) \wedge (\text{compare}(p_1, z) \vee (\neg \text{compare}(p_1, p_2) \wedge \text{compare}(p_2, z))))$$

ACME Case: Concatenation

The first reason is ruled out by the logical design of `payroll`.

The second reason can be ruled out by modifying the plan to ensure the second subplan for the concatenation operator will only return additional results when p_1 and p_2 are distinct.

$$\exists y, z. (\text{emp-array0}(z, x, y) \wedge (\text{compare}(p_1, z) \vee (\neg \text{compare}(p_1, p_2) \wedge \text{compare}(p_2, z)))))$$

Note that the query mapping for this plan is no longer a positive query.

ACME Case: Concatenation

The first reason is ruled out by the logical design of `payroll`.

The second reason can be ruled out by modifying the plan to ensure the second subplan for the concatenation operator will only return additional results when p_1 and p_2 are distinct.

$$\exists y, z. (\text{emp-array0}(z, x, y) \wedge (\text{compare}(p_1, z) \vee (\neg \text{compare}(p_1, p_2) \wedge \text{compare}(p_2, z))))$$

Note that the query mapping for this plan is no longer a positive query.

An alternative plan avoids simple complement by using a cut.

$$\exists y, z. (\text{emp-array0}(z, x, y) \wedge [(\text{compare}(p_1, z) \vee \text{compare}(p_2, z))]_1 \wedge !_1)$$

ACME Case: Simple Complement

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a name that is also unique.*

ACME Case: Simple Complement

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a name that is also unique.*

Two plans using access path `emp-array0` and a simple complement operator are as follows.

$$\begin{aligned} & \exists y, z. (\text{emp-array0}(z, x, y) \\ & \quad \wedge \neg \exists u, v, w. (\text{emp-array0}(u, v, w) \\ & \quad \quad \wedge \text{compare}(y, w) \\ & \quad \quad \wedge \neg \text{compare}(x, v))) \end{aligned}$$

$$\begin{aligned} & \exists y, z. (\text{emp-array0}(z, x, y) \\ & \quad \wedge \neg \exists u, v, w. ([\text{emp-array0}(u, v, w)]_1 \\ & \quad \quad \wedge \text{compare}(y, w) \\ & \quad \quad \wedge \neg \text{compare}(x, v)) \\ & \quad \wedge !_1) \end{aligned}$$

ACME Case: Simple Complement

Consider where the APS department needs a plan that will *find the employee-number x for any employee that has a name that is also unique.*

Two plans using access path `emp-array0` and a simple complement operator are as follows.

$$\begin{aligned} & \exists y, z. (\text{emp-array0}(z, x, y) \\ & \quad \wedge \neg \exists u, v, w. (\text{emp-array0}(u, v, w) \\ & \quad \quad \wedge \text{compare}(y, w) \\ & \quad \quad \wedge \neg \text{compare}(x, v))) \end{aligned}$$

$$\begin{aligned} & \exists y, z. (\text{emp-array0}(z, x, y) \\ & \quad \wedge \neg \exists u, v, w. ([\text{emp-array0}(u, v, w)]_1 \\ & \quad \quad \wedge \text{compare}(y, w) \\ & \quad \quad \wedge \neg \text{compare}(x, v)) \\ & \quad \wedge !_1) \end{aligned}$$

Question: Which is more efficient?

When do Plans Implement User Queries?

Requirements

Given a *plan interpretation* \mathcal{I} , the plan Q' (driven by \mathbf{C}) outputs exactly the valuations that make the *user query* Q true in \mathcal{I} .

When do Plans Implement User Queries?

Requirements

Given a *plan interpretation* \mathcal{I} , the plan Q' (driven by \mathbf{C}) outputs exactly the valuations that make the *user query* Q true in \mathcal{I} .

What can we do with this?

- we show how to *model* many (if not most) features of standard SQL implementations using the operators introduced today with the help of creative physical design and selection of access paths (next time).

When do Plans Implement User Queries?

Requirements

Given a *plan interpretation* \mathcal{I} , the plan Q' (driven by \mathbf{C}) outputs exactly the valuations that make the *user query* Q true in \mathcal{I} .

What can we do with this?

- we show how to *model* many (if not most) features of standard SQL implementations using the operators introduced today with the help of creative physical design and selection of access paths (next time).

How do we *actually* find plans?

- 1 we search for Q' that (as a formula) *is logically equivalent* to Q under the logical and physical schema constraints Σ , and
- 2 we *improve* Q' by eliminating *duplicate elimination operations* and by inserting *cuts*.

Queries over Column Store Physical Design Revisited

1 $\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a) \wedge \text{emp-rid-name}(a, y) \wedge \text{emp-rid-slry}(a, z)$

2 $\exists y, z.\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a)$

3 $\exists z.\text{employee}(x, y, z)$

$\exists a.\text{emp-oid-rid}(x, a) \wedge \text{emp-rid-name}(a, y)$

4 $\exists x.\text{employee}(x, y, z)\{y\}$

$\{\exists a.\text{emp-name-rid}(y, a) \wedge \text{emp-rid-slry}(a, z)\}$

Results:

- Plans for queries 1-4 are logically equivalent to the given user queries.
- Plans 1-3 can avoid duplicate elimination operator.

Related Issues

- Relational Algebra
- Domain Independence and Range restricted Queries
- Temporary storage
- Ordered properties of iterated semantics (merge joins?)
- Streaming queries