Extending the Optimizer

Ralf Hartmut Güting

Fernuniversität Hagen, Germany

Secondo Extensions

1. New algebras for attribute types:
   - Write a display function for a new type constructor to show corresponding values.
   - Define the syntax of a new operator to be used within SQL and in SECONDO.
   - Write optimization rules to perform selections and joins involving a new operator, including rules to use appropriate indexes.
   - Define a cost function or constant for using the operator.
2. New query processing operators in the relational algebra:
   - Define the operator syntax to be used in SECONDO.
   - Write optimization rules using the new operator.
   - Write a cost function (predicate) for the new operator.
3. New types of indexes:
   - Define the operator syntax to be used in SECONDO for search operations on the index.
   - Write optimization rules using the index.
   - Write cost functions for access operations.

Extension Tasks

- Writing a display function for a type constructor
- Defining operator syntax for SQL
- Defining operator syntax for SECONDO
- Writing optimization rules
- Writing cost functions

Writing a Display Predicate for a Type Constructor

```
display(Type, Value) :-
```

Display the `Value` according to its `Type` description.

```
display(int, N) :-
  !,
  write(N).


display([rel, [tuple, Attrs]], Tuples) :-
  !,
  nl,
  max_attr_length(Attrs, AttrLength),
  displayTuples(Attrs, Tuples, AttrLength).


displayTuples(_, [], _).

displayTuples(Attrs, [Tuple | Rest], AttrLength) :-
  displayTuple(Attrs, Tuple, AttrLength),
  nl,
  displayTuples(Attrs, Rest, AttrLength).
```

```
displayTuple([], _, _).

displayTuple([[Name, Type] | Attrs], [Value | Values], AttrNameLength) :-
  atom_length(Name, NLength),
  PadLength is AttrNameLength - NLength,
  write_spaces(PadLength),
  write(Name),
  write(' : '),
  display(Type, Value),
  nl,
  displayTuple(Attrs, Values, AttrNameLength).
```

Defining Operator Syntax for SQL

Atomic operators working on attribute types like

```
+, <, mod, inside, starts, distance, ...
```

can be used directly in SQL.

- explain PROLOG syntax
- explain translation to Secondo syntax

Several cases:

- syntax is predefined in PROLOG, e.g. for

  ```
  +, -, *, /, <, >
  ```

- write in prefix syntax (always possible)
- define infix syntax:

  ```
  :- op(800, xfx, adjacent).
  ```

Defining Operator Syntax for SECONDO

Query language terms written in prefix notation in the optimizer.

```
x y product filter[cond] consume
```

written as

```
consume(filter(product(x, y), cond))
```

For each operator, optimizer needs to know translation to SECONDO.

1. By default
   - 1 or 3 arguments: prefix syntax
   - 2 arguments: infix syntax

```
length(x), x adjacent y, translate(x, y, z)
```

2. Syntax specification via predicate `secondoOp` in file `opsyntax.pl`

```
secondoOp(distance, prefix, 2).

secondoOp(feed, postfix, 1).
secondoOp(consume, postfix, 1).
```

3. Programming a `plan_to_atom` rule.

```
plan_to_atom(X, Y) :-
```

`Y` is the SECONDO expression corresponding to term `X`.

```
plan_to_atom(sortmergejoin(X, Y, A, B), Result) :-
 plan_to_atom(X, XAtom),
 plan_to_atom(Y, YAtom),
 plan_to_atom(A, AAtom),
 plan_to_atom(B, BAtom),
 concat_atom([XAtom, YAtom, 'sortmergejoin[',
   AAtom, ', ', BAtom, '] '], '', Result),
 !.
```

Writing Optimization Rules

See what the optimizer does. Consider the query

```
select count(*)
from [staedte as s, plz as p]
where [s:sname = p:ort, p:plz > 40000]
```

See the predicate order graph:

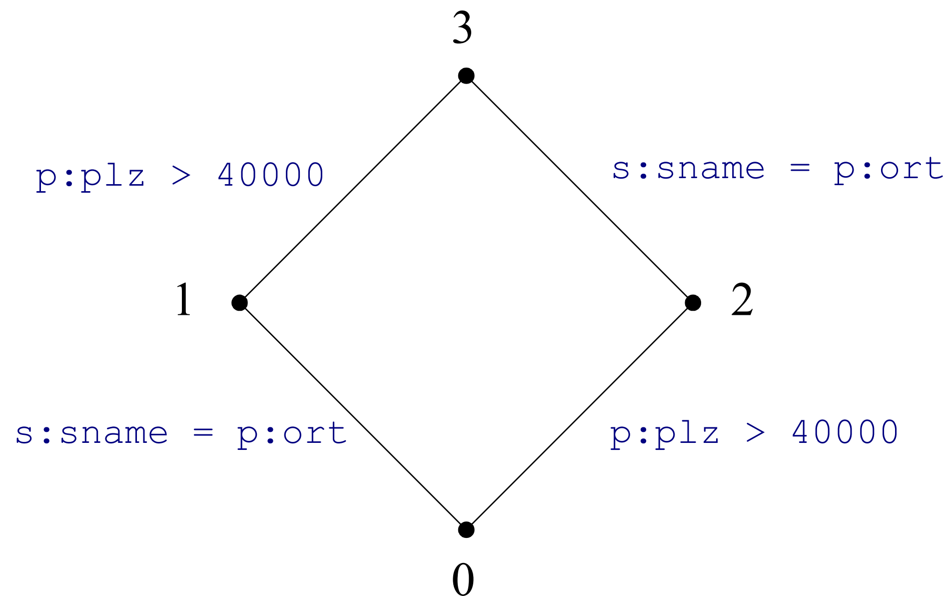

```
16 ?- writeNodes.
Node: 0
Preds: []
Partition: [arp(arg(2), [rel(plz, p, l)], []), arp(arg(1), [rel(staedte, s, u)],
[])]

Node: 1
Preds: [pr(attr(s:sName, 1, u)=attr(p:ort, 2, u), rel(staedte, s, u), rel(plz, p,
l))]
Partition: [arp(res(1), [rel(staedte, s, u), rel(plz, p, l)], [attr(s:sName, 1,
u)=attr(p:ort, 2, u)])]

Node: 2
Preds: [pr(attr(p:pLZ, 1, u)>40000, rel(plz, p, l))]
Partition: [arp(res(2), [rel(plz, p, l)], [attr(p:pLZ, 1, u)>40000]),
arp(arg(1), [rel(staedte, s, u)], [])]

Node: 3
Preds: [pr(attr(p:pLZ, 1, u)>40000, rel(plz, p, l)), pr(attr(s:sName, 1,
u)=attr(p:ort, 2, u), rel(staedte, s, u), rel(plz, p, l))]
Partition: [arp(res(3), [rel(staedte, s, u), rel(plz, p, l)], [attr(p:pLZ, 1,
u)>40000, attr(s:sName, 1, u)=attr(p:ort, 2, u)])]

Yes
17 ?-
```

```
17 ?- writeEdges.
Source: 0
Target: 1
Term: join(arg(1), arg(2), pr(attr(s:sName, 1, u)=attr(p:ort, 2, u),
rel(staedte, s, u), rel(plz, p, l)))
Result: 1

Source: 0
Target: 2
Term: select(arg(2), pr(attr(p:pLZ, 1, u)>40000, rel(plz, p, l)))
Result: 2

Source: 1
Target: 3
Term: select(res(1), pr(attr(p:pLZ, 1, u)>40000, rel(plz, p, l)))
Result: 3

Source: 2
Target: 3
Term: join(arg(1), res(2), pr(attr(s:sName, 1, u)=attr(p:ort, 2, u),
rel(staedte, s, u), rel(plz, p, l)))
Result: 3

Yes
18 ?-
```

Translating the Arguments

```
res(N) => res(N).

arg(N) => feed(rel(Name, *, Case)) :-
  argument(N, rel(Name, *, Case)), !.

arg(N) => rename(feed(rel(Name, Var, Case)), Var) :-
  argument(N, rel(Name, Var, Case)), !.
```

See what happens:

```
18 ?- sql select count(*)
from [staedte as s, plz as p]
where [s:sname = p:ort, p:plz > 40000].
...

19 ?- arg(1) => X.

X = rename(feed(rel(staedte, s, u)), s)

Yes
20 ?-
```

Translating Selection Predicates

Rule for filtering a stream:

```
select(Arg, pr(Pred, _)) => filter(ArgS, Pred) :-
  Arg => ArgS.
```

Rules for using a B-tree index:

```
select(arg(N), Y) => X :-
  indexselect(arg(N), Y) => X.

indexselect(arg(N), pr(attr(AttrName, Arg, Case) = Y, Rel)) => X :-
  indexselect(arg(N), pr(Y = attr(AttrName, Arg, Case), Rel)) => X.

indexselect(arg(N), pr(Y = attr(AttrName, Arg, AttrCase), _)) =>
  exactmatch(IndexName, rel(Name, *, Case), Y)
  :-
  argument(N, rel(Name, *, Case)),
  hasIndex(rel(Name, *, Case), attr(AttrName, Arg, AttrCase), IndexName).

indexselect(arg(N), pr(Y = attr(AttrName, Arg, AttrCase), _)) =>
  rename(exactmatch(IndexName, rel(Name, Var, Case), Y), Var)
  :-
  argument(N, rel(Name, Var, Case)), Var \= * ,
  hasIndex(rel(Name,Var,Case), attr(AttrName, Arg, AttrCase), IndexName).
```

See translations:

```
3 ?- writePlanEdges.
Source: 0
Target: 1
Plan: Staedte  feed {s}  plz  feed {p}  product  filter[(.SName_s = .Ort_p)]
Result: 1

Source: 0
Target: 1
Plan: Staedte  feed {s}  loopjoin[plz_Ort plz  exactmatch[.SName_s] {p} ]
Result: 1

Source: 0
Target: 1
Plan: Staedte  feed {s} plz  feed {p} sortmergejoin[SName_s, Ort_p]
Result: 1

Source: 0
Target: 1
Plan: Staedte  feed {s} plz  feed {p} hashjoin[SName_s, Ort_p, 999997]
Result: 1

Source: 0
Target: 1
Plan: plz  feed {p} Staedte  feed {s} hashjoin[Ort_p, SName_s, 999997]
```

```
Result: 1

Source: 0
Target: 2
Plan: plz  feed {p}  filter[(.PLZ_p > 40000)]
Result: 2

Source: 1
Target: 3
Plan: res(1)  filter[(.PLZ_p > 40000)]
Result: 3

Source: 2
Target: 3
Plan: Staedte  feed {s}  res(2)  product  filter[(.SName_s = .Ort_p)]
Result: 3

Source: 2
Target: 3
Plan: Staedte  feed {s} res(2) sortmergejoin[SName_s, Ort_p]
Result: 3

Source: 2
Target: 3
Plan: Staedte  feed {s} res(2) hashjoin[SName_s, Ort_p, 999997]
Result: 3
```

```
Source: 2
Target: 3
Plan: res(2) Staedte  feed {s} hashjoin[Ort_p, SName_s, 999997]
Result: 3

Yes
4 ?-
```

Writing Cost Functions

First assign sizes to the nodes and selectivities to the edges of the predicate order graph:

```
7 ?- assignSizes.

Yes
```

```
8 ?- writeSizes.
Node: 1
Size: 7419.81

Node: 2
Size: 22696.9

Node: 3
Size: 4080.89

Source: 0
Target: 1
Selectivity: 0.0031

Source: 0
Target: 2
Selectivity: 0.55

Source: 1
Target: 3
Selectivity: 0.55

Source: 2
Target: 3
Selectivity: 0.0031
Yes
```

Create cost edges:

```
14 ?- createCostEdges.

Yes
15 ?- writeCostEdges.
Source: 0
Target: 1
Plan: Staedte  feed {s}  plz  feed {p}  product  filter[(.SName_s = .Ort_p)]
Result: 1
Size: 7419.81
Cost: 7.07402e+06

Source: 0
Target: 1
Plan: Staedte  feed {s}  loopjoin[plz_Ort plz  exactmatch[.SName_s] {p} ]
Result: 1
Size: 7419.81
Cost: 75027.0

Source: 0
Target: 1
Plan: Staedte  feed {s} plz  feed {p} sortmergejoin[SName_s, Ort_p]
Result: 1
Size: 7419.81
Cost: 157724.0
```

```
Source: 0
Target: 1
Plan: Staedte  feed {s} plz  feed {p} hashjoin[SName_s, Ort_p, 999997]
Result: 1
Size: 7419.81
Cost: 238243.0

Source: 0
Target: 1
Plan: plz  feed {p} Staedte  feed {s} hashjoin[Ort_p, SName_s, 999997]
Result: 1
Size: 7419.81
Cost: 114616.0

Source: 0
Target: 2
Plan: plz  feed {p}  filter[(.PLZ_p > 40000)]
Result: 2
Size: 22696.9
Cost: 89962.1

Source: 1
Target: 3
Plan: res(1)  filter[(.PLZ_p > 40000)]
Result: 3
```

```
Size: 4080.89
Cost: 12465.3

Source: 2
Target: 3
Plan: Staedte  feed {s}  res(2)  product  filter[(.SName_s = .Ort_p)]
Result: 3
Size: 4080.89
Cost: 3.87937e+06

Source: 2
Target: 3
Plan: Staedte  feed {s} res(2) sortmergejoin[SName_s, Ort_p]
Result: 3
Size: 4080.89
Cost: 71374.0

Source: 2
Target: 3
Plan: Staedte  feed {s} res(2) hashjoin[SName_s, Ort_p, 999997]
Result: 3
Size: 4080.89
Cost: 119751.0

Source: 2
Target: 3
```

```
Plan: res(2) Staedte  feed {s} hashjoin[Ort_p, SName_s, 999997]
Result: 3
Size: 4080.89
Cost: 51834.0

true.

16 ?-
```

The Cost Function

```
cost(Term, Sel, Size, Cost) :-
```

```
The cost of an executable Term representing a predicate with selectivity Sel is
Cost and the size of the result is Size. Here Term and Sel have to be instanti-
ated, and Size and Cost are returned.
```

Terms look like this:

```
17 ?- planEdge(Source, Target, Term, Result).
```

One of the solutions listed (for `example5`) is

```
Source = 0,
Target = 1,
Term = hashjoin(rename(feed(rel(plz, p, l)), p), rename(feed(rel(staedte, s,
u)), s), attrname(attr(p:ort, 2, u)), attrname(attr(s:sName, 1, u)), 999997),
Result = 1 ;
```

Some cost predicates:

```
cost(rel(Rel, _, _), _, Size, 0) :-
  card(Rel, Size).

cost(res(N), _, Size, 0) :-
  resultSize(N, Size).

cost(feed(X), Sel, S, C) :-
  cost(X, Sel, S, C1),
  feedTC(A),
  C is C1 + A * S.

cost(filter(X, _), Sel, S, C) :-
  cost(X, 1, SizeX, CostX),
  filterTC(A),
  S is SizeX * Sel,
  C is CostX + A * SizeX.
```

```
cost(hashjoin(X, Y, _, _, 999997), Sel, S, C) :-
  cost(X, 1, SizeX, CostX),
  cost(Y, 1, SizeY, CostY),
  hashjoinTC(A, B, D),
  S is SizeX * SizeY * Sel,
  C is CostX + CostY +                    % producing the arguments
    A * SizeY +                           % A - time [microsecond] per build
    B * SizeX +                           % B - time per probe
    D * S.                                % D - time per result tuple
                                          % table fits in memory assumed
```

Determining cost constants

Create large test relations:

```
let plz100Even = plz feed thousand feed filter[.No <=100] product
  extend[R: randint(1000000)] filter[(.R mod 2) = 0] consume

let plz100Odd = plz feed thousand feed filter[.No <=100] product
  extend[R: randint(1000000)] filter[(.R mod 2) = 1] consume
```

Cardinalities: plz100Even 2063237, plz100Odd 2064214

Determine time for build:

```
query plz100Even feed count

4.63 seconds, 4.55

query plz100Even feed {a} plz100Even feed {b} hashjoin[R_a, R_b, 999997] head[1]
count

11.21, 11.06 seconds
```

Time for build phase: 11 - 4.5 seconds = 6.5 seconds needed for 2063237 tuples. Hence time per tuple is 6500000 / 2063237 = 3.15 microseconds.

Determine time for probe:

```
query plz100Even feed {a} plz100Odd feed {b} hashjoin[R_a, R_b, 999997] count

Result: 0
17.3199, 17.3145 seconds
```

Time for probe phase:

- Scanning first argument: 4.5 seconds
- Scanning second argument: 4.5 seconds
- Build phase: 6.5 seconds
- Remaining: 17.31 - 15.5 = 1.81 seconds
- per tuple: 1810000 / 2064214 = 0.876 microseconds

Determine time for producing result tuples:

```
query plz100Even feed {a} plz100Even feed {b} hashjoin[R_a, R_b, 999997] count

Result: 10589309
33.78, 33.67 seconds
```

Subtract time for previous join which did not have any result tuples. Remaining time is 33.7 - 17.3 seconds = 16.4 seconds. Per tuple 16400000 / 10589309 = 1.548 microseconds.

Enter in file `operators.pl`:

```
hashjoinTC(3.15, 0.876, 1.55).
```