

Implementation Alternatives for Bottom-Up Evaluation

STEFAN BRASS

*Martin-Luther-Universität Halle-Wittenberg, Institut für Informatik,
Von-Seckendorff-Platz 1, D-06099 Halle (Saale), Germany
(e-mail: brass@informatik.uni-halle.de)*

submitted —; revised —; accepted —

Abstract

Bottom-up evaluation is a central part of query evaluation / program execution in deductive databases. It is used after a source code optimization like magic sets or SLDmagic that ensures that only facts relevant for the query can be derived. Then bottom-up evaluation simply performs the iteration of the standard TP -operator to compute the minimal model. However, there are different ways to implement bottom-up evaluation efficiently. Since this is most critical for the performance of a deductive database system, and since performance is critical for the acceptance of deductive database technology, this question deserves a thorough analysis. In this paper we start this work by discussing several different implementation alternatives. Especially, we propose a new implementation of bottom-up evaluation called “Push-Method”. First runtime comparisons are given.

KEYWORDS: deductive databases, bottom-up evaluation, implementation

1 Introduction

Deductive databases (Minker 1988; Ullman and Zaniolo 1990; Freitag et al. 1991; Ramakrishnan et al. 1994; Vaghani et al. 1994; Friesen et al. 1995; Ramakrishnan and Ullman 1995) have not yet been very successful in practice (at least in terms of market share). However, their basic idea is practically very important: Deductive databases aim at an integrated system of database and programming language that is based on the declarative paradigm which was so successful in database languages. Currently, database programming is typically done in languages like PHP or Java. The programs construct SQL statements, send them to the database server, fetch the results, and process them. The interface is not very smooth, and although the situation can be improved with specific database languages like PL/SQL and server-side procedures / user-defined functions within the DBMS, the language paradigms remain different. Object-oriented databases were one approach to develop an integrated system based on a single paradigm, but there the declarativity of the database query part was sacrificed, and they did not get a significant market share, either. Nevertheless, there is an obvious demand for integrated database/programming systems, and this demand has even grown

because of object-relational features that need programming inside the database server, and because of web and XML applications.

One of the reasons why deductive databases were not yet successful is the non-satisfying performance of many prototypes. If somebody moves from a standard relational database to a deductive database, he or she would expect that the performance of query evaluation is not significantly worse in the deductive database.

Another, and even more important reason, is that deductive database prototypes have not really tried to be useful also as a programming platform — they were concentrated only on recursive query evaluation. For instance, whereas a large fraction of a Prolog system is usually written in Prolog, deductive database prototypes are usually completely written in C++.

There has been a lot of progress over the years on recursive query evaluation (Bancilhon and Ramakrishnan 1986), which is certainly an important task, because many applications use tree-structured or graph-structured data. A large part of this work was about source-level optimizations, like the well-known magic-set method (Bancilhon et al. 1986; Beeri and Ramakrishnan 1991) and its many optimizations, including the SLDMagic-method of the author (Brass 2000).

However, this all depends on an efficient implementation of bottom-up evaluation. If one wants to build a new deductive database system which a real chance for acceptance in practice, one needs to clarify first how bottom-up evaluation should be done. This is not obvious, and several alternatives will be discussed in this paper.

This is also related to the second problem mentioned above: For using a Datalog-variant as a real programming language efficiently, at least sometimes something must be known about the basic evaluation method. Of course, one must be careful so that the meaning of programs can be understood declaratively. But at least to some degree programmers might help the system to efficiently execute a program if a typical basic execution method is known.

Note that programming in deductive databases is not the same as programming in Pure Prolog. In deductive databases one thinks in the direction of the arrow, because they are based on bottom-up evaluation. For instance left recursion is very natural in this way, whereas in Prolog it must be avoided.

The XSB system must be mentioned here (Sagonas et al. 1994), because it is still actively developed and used, in contrast to many deductive database prototypes, which are no longer supported. It is also much faster than many deductive database prototypes. However, it is not based on bottom-up evaluation, it uses to a significant extent Prolog technology (however, it offers tabling to avoid infinite loops).

Our belief is that the bottom-up approach has still room for improvement in order to deliver competitive performance. In (Brass 2000) we proposed a source-level transformation that is for tail-recursions asymptotically faster than the standard magic set method (and also than the tabling method underlying the XSB system). It is also interesting because it unifies many improvements which were proposed for the magic set method over time.

After a source-level transformation like SLDmagic, which solves the problem of goal-direction, one needs an efficient implementation of bottom-up evaluation. The research reported in this paper is a step in this direction.

2 Basic Framework

The approach we want to follow is to translate Datalog into C++, which can then be compiled to machine code. A first prototype was developed in a diploma thesis and is currently being evaluated.

There are three types of predicates:

- EDB predicates (“extensional database”), the given database relations,
- IDB predicates (“intensional database”), which are defined by means of rules,
- built-in predicates like $<$, which usually have an infinite extension, and are defined by means of program code inside the system.

The purpose of bottom-up evaluation is to compute the extensions of the IDB relations. Actually, only one of them is the “answer predicate”, the extension of which must be printed, or otherwise made available to the user.

Bottom-up evaluation works by applying the rules from right to left, so basically it computes the minimal model by iterating the standard T_P -operator. Of course, an important goal is to apply every applicable rule instance only once via rule-ordering and managing deltas for recursive rules (“seminaive evaluation”).

Because of the infinite extension, built-in predicates can only be called when certain arguments are bound (i.e. input arguments, known values). In contrast, a free argument position permits a variable (output argument). The restrictions for the predicates are described by binding patterns (modes, adornments), e.g. $<$ can be called for the binding pattern **bb** only (every letter in a binding pattern corresponds to an argument position, **b** means bound, **f** means free).

A basic interface for relations is that it is possible to open a cursor (scan, iterator) over the relation, which permits to loop over all tuples. We assume that for every normal predicate p , there is a class p_cursor with the following methods:

- `void open()`: Open a scan over the relation, i.e. place the cursor before the first tuple.
- `bool fetch()`: Move the cursor to the next tuple. This function must also be called to access the first tuple. It returns `true` if there is a first/next tuple, or `false`, if the cursor is at the end of the relation.
- `T col_i()`: Get the value of the i -th column (attribute). Here T is the type of the i -th column.
- `close()`: Close the cursor.

For the push method, we will also need

- `push()`: Save the state of the cursor on a global stack.
- `pop()`: Restore the state of the cursor.

A relation may have special access structures (e.g. it might be stored in a B-tree or an array). Then not only a full scan (corresponding to binding pattern `ff...f`) is possible, but also scans only over tuples with a given value for certain arguments. We assume that in such cases there are additional cursor classes called p_cursor_β , with a binding pattern β . These classes have the same methods as the other cursor

classes, only the `open`-method has parameters for the bound arguments. E.g. if `p` is a predicate of arity 3 which permits especially fast access to tuples with a given value of the first argument, and if this argument has type `int`, the class `p_cursor_bff` would have the method `open(int x)`.

Actually, some access structures can efficiently evaluate small conjunctions with parameters, e.g. a B-tree over the first argument of `p` would also support a query of the form `p(X,Y,Z) ∧ X ≥ c1 ∧ X < c2`, where `c1` and `c2` are integer constants or bound variables. This is not in the focus of the current paper, but a realistic system must be able to make use of such possibilities.

In addition, we need a possibility to create new tuples for predicates defined by rules (IDB predicates). We assume that for each predicate `p` there is a class `p` with a class method `insert` that creates a new tuple in `p`. Of course, it is possible that the objects of class `p` correspond to individual tuples, but since we only use the cursor interface, this is only one possible implementation.

3 Materializing Derived Predicates

The first, most basic method for implementing bottom-up evaluation is to explicitly create a stored relation for every IDB-predicate. As an example, let us consider

$$p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z).$$

We assume that all columns in the examples have type `int`.

Of course, one option is to use a standard relational database, create a table for every IDB predicate, and send SQL statements to the database to execute the rules. But using a separate system for the management of facts causes performance penalties. Furthermore, for the seminaive evaluation of recursive rules, there is no good and efficient way to manage the deltas with SQL (the set of tuples newly derived in an iteration). Another interesting problem is that a good sideways information passing (SIP) strategy in the magic set method or selection function in the SLD-Magic method needs already knowledge about existing indexes and relation sizes. Therefore it is not a good idea to do query optimization in two completely separate systems: The chosen SIP strategy/selection function more or less prescribes the evaluation of the resulting rules. For instance, if one wants to use a merge join, less “sideways information passing” is possible than with a nested loop/index join.

Therefore, one would do basic rule evaluation in the deductive database system itself, although it might be possible to use parts of a standard relational system (e.g. the storage manager). E.g. with a nested loop/index join, the implementation of the above rule would look as shown in Figure 1.

Of course, the materialization method causes a lot of copying. E.g., disjunctions must be expressed in standard Datalog with a derived predicate:

$$\begin{aligned} p(X, Y) &\leftarrow q(X, Y). \\ p(X, Y) &\leftarrow r(X, Y). \end{aligned}$$

The materialization method would copy all `q`- and `r`-facts. This is especially expensive if the data values `X` and `Y` are large (e.g. longer strings). The problem can

```

q_cursor q1;
q1.open();
while(q1.fetch()) {
    int X = q1.col_1();
    int Y = q1.col_2();
    r_cursor_bff r1; // if there is an index on the first argument
                    // (and none for the binding pattern bbf)

    r1.open(Y);
    while(r1.fetch()) {
        if(r1.col_2() == 5) {
            int Z = r1.col_3();
            p::insert(X, Z, 2);
        }
    }
}

```

Fig. 1. Materializing an IDB-Predicate: $p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z)$.

be reduced by working only with pointers to the real data values. (but that might not make optimal use of the memory cache in the CPU, because values are more scattered around in memory).

4 Pull-Method

Of course, explicitly materializing every intermediate predicate needs a lot of memory. Therefore, it is a standard technique in databases to compute tuples only on demand, or actually not even compute the entire tuple, but permit access to its columns (in this way, possibly large data values do not have to be copied). In order to get such “lazy” evaluation, one only needs to support the cursor interface for each predicate. Let us consider again the rule

$$p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z).$$

If p is nonrecursive, and this is the only rule, and no duplicate elimination is needed, the code would look as shown in Figure 2.

If duplicate elimination is needed, storing all tuples is necessary (unless the tuples for the body literals are generated in a fitting sort order). One can then apply the materialization method, or extend the pull-method by building a hash table of all previously returned tuples, and adding a check that the derived tuple is new.

An important disadvantage of the pull-method is that it causes recomputation if multiple scans over a predicate are performed. This does not only happen when there are several body literals with the same predicate p , but also when a single p -literal appears in the inner loop of a nested loop join. However, recomputation is not necessarily something evil that must be avoided at any price. If the recomputation is not expensive, as in the example with the predicate describing a disjunction, it is a possible alternative.

```

class p_cursor {
public:
    void open()
    {
        q1.open();
        q1_more = q1.fetch();
        if(q1_more)
            r1.open(q1.col_2()); // Assuming again index on first argument
    }
    bool fetch()
    {
        while(q1_more) {
            while(r1.fetch()) {
                if(r1.col_2() == 5)
                    return true;
            }
            r1.close();
            q1_more = q1.fetch();
            if(q1_more)
                r1.open(q1.col_2());
        }
        return false;
    }
    int col_1() { return q1.col_1(); }
    int col_2() { return r1.col_3(); }
    int col_3() { return 2; }
private:
    q_cursor q1;
    r_cursor_bff r1;
    bool q1_more;
};

```

Fig. 2. Pull-Method (Lazy Evaluation): $p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z)$.

Recursion with the pure pull-method is not possible: If one tries to implement recursion with the recursive opening of cursors, this leads to an infinite recursion even for acyclic relations since no bindings are passed to the recursive call: Each call does the same work again. Of course, it is possible to integrate standard seminaive evaluation, but this simply means to use the materialization method at least once in each recursive cycle.

5 Push-Method

It is also possible to apply the rules strictly from right to left, and move generated facts immediately to the place where they are needed. In contrast to the materialization method, a rule is not applied to produce all consequences in the current state, but only a single fact is derived each time. This reduces the need for interme-

mediate storage and copying, which was also the main motivation for the pull-method. But here the producer of facts is in control, not the consumer as in the pull-method.

Of course, usually several facts can be derived with a rule. Therefore, once a fact is derived, one must store the current state of rule application on a stack for later backtracking. Then control jumps to a rule where this fact matches a body literal. In general there can be several rules that might use the produced fact, in which case again a backtrack point is generated.

This method basically works only with rules that have at most one body literal with IDB-predicate, because then matching facts for the other (EDB) body literals are readily available when a fact for the IDB body literal arrives. The SLDMagic method (Brass 2000) produces such rules as output of the program transformation, therefore this case is practically interesting. Furthermore, when there are several body literals with IDB-predicates, it is often possible to use the materialization or pull method for the predicates of all but one body literal.

The push method is applicable to linear recursion, and that is in fact one of its strengths (it can be very efficient in this case).

Let us explain how it works. First one creates a variable for every column of an IDB predicate. Consider again the example rule:

$$p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z).$$

If q is an IDB predicate, and r is an EDB-predicate, and all arguments have type `int`, we get:

```
int q_1, q_2;
int p_1, p_2, p_3;
```

In addition, there is a code piece for every body literal with IDB-predicate. Control jumps to this code piece when a new fact for this predicate was derived. The argument values of the fact are stored in the above variables. The purpose of the code is to check whether new facts can be derived with this rule with the given instantiation of the IDB body literal, and if yes, to store the arguments of the derived fact in the corresponding variables and to jump to every place where the newly derived fact is used. For the example rule, the code looks as shown in Figure 3 (there are in fact many optimization possibilities, discussed below). Rules with only EDB-predicates in the body act as starting points. For instance, consider

$$q(X, Y) \leftarrow s(X, Y) \wedge Y \geq 0.$$

Then for every s -fact, we would fill the variables q_1 and q_2 and jump to the place where q -facts are used (label q :). Again, this loop is implemented with backtracking.

There are a number of optimization possibilities:

1. If q is non-recursive, pushing and restoring q_1 and q_2 is not necessary.
2. If the rule is non-recursive (i.e. not body literal depends on the predicate in the head), pushing and restoring the open cursor $r1$ is not necessary.
3. It is also an optimization parameter how far one follows possible alternatives before saving the computation state. In the example, it is checked that the

```

q: r_cursor_bff r1;
  r1.open(q_2); // Assuming index on first arg
  while(r1.fetch()) {
    if(r1.col_2() == 5) {
      p_1 = q_1;
      p_2 = r1.col_3();
      p_3 = 2;
      if(r1.fetch()) {
        push_int(q_1);
        push_int(q_2);
        r1.push();
        push_cont(CONT_q);
      }
      goto p;
    }
  }
  goto backtrack; // if this is the last place where q is used
cont_q:
  r1.pop();
  q_2 = pop_int();
  q_1 = pop_int();
  do {
    if(r1.col_2() == 5) {
      p_1 = q_1;
      p_2 = r1.col_3();
      p_3 = 2;
      if(r1.fetch()) {
        push_int(q_1);
        push_int(q_2);
        r1.push();
        push_cont(CONT_q);
      }
      goto p;
    }
  } while(r1.fetch());
  goto backtrack; // if this is the last place where q is used
backtrack:
  if(stack_empty()) return false;
  switch(pop_task()) {
  case CONT_q; goto cont_q;
  ...
  }

```

Fig. 3. Push-Method: $p(X, Z, 2) \leftarrow q(X, Y) \wedge r(Y, 5, Z)$.

cursor is not at the end before the variables are pushed. It would also be possible to find first a tuple with the value 5 in the second column. If there is none, it is not necessary to save the computation state.

```

init:
  s_cursor s1;
  s1.open();
  while(s1.fetch()) {
    if(s1.col_2() >= 0) {
      q_1 = s1.col_1();
      q_2 = s1.col_2();
      if(s1.fetch()) {
        s1.push();
        push_cont(CONT_init);
      }
      goto q;
    }
  }
  return false; // if this is the last/only initialization rule
cont_init:
  s1.pop();
  do {
    if(s1.col_2() >= 0) {
      q_1 = s1.col_1();
      q_2 = s1.col_2();
      if(s1.fetch()) {
        s1.push();
        push_cont(CONT_init);
      }
      goto q;
    }
  } while(s1.fetch());
  return false; // if this is the last/only initialization rule

```

Fig. 4. Push-Method: Initialization with $q(X, Y) \leftarrow s(X, Y) \wedge Y \geq 0$.

4. Of course, if we would know that the first two columns are a key of the relation r , it would not be necessary to search for another solution at all.
5. This technique to check first whether there really is another solution before creating a backtrack point seems also important if q appears in the bodies of several rules: Not seldom such rules are in fact mutually exclusive and treat different cases. Of course, the best case is if we can detect this by static analysis of the rules. If this is not possible, one could also follow the alternative until the point where it is clear that it really gives another solution before pushing the computation state.
6. The assignment of value 2 to p_3 could be done before the loop, if it is clear that no other rule generating p -tuples can be executed before backtracking returns to this rule. This will be the case if p is non-recursive. In this way, the assignment is executed only once instead of for every generated tuple. The same is possible for $p_1 = q_1$, because the first argument of the generated tuple does not depend on the second body literal over which the loop runs.

7. It is also not given that the i -th argument of predicate p must be stored in a variable p_i . In the example, if one chooses the same variable for the first arguments of p and q , the assignment $p_1 = q_1$ becomes unnecessary.

Note also that the push method can be made to fit into the cursor interface: If for instance p needs to be queried with a cursor, the above code is inside the `fetch` method. The code jumps to label `p`: when a new p -fact is derived, the `fetch`-method returns `true` to the caller. Each call to the `fetch`-method starts at the label `backtrack`. In the `open`-method the backtrack-stack is initialized with a value that causes a jump to the initialization (`init`).

6 Pull-Method with Passing of Bindings

In deductive databases, one normally uses first a program transformation like magic sets, which is responsible for passing bindings from the caller to the callee, so that only relevant facts are computed when the transformed rules are evaluated strictly bottom-up (i.e. the entire minimal model of the transformed program is computed). The “magic predicates” contain values for the input/bound arguments of a predicate. Calls to these predicates are added as conditions to the rule body, so that the rule can only “fire” when the result is needed.

With the magic set transformation, all calls to a predicate are put together in one set. This is good, if there are several calls to a complex predicate with the same input values: Then the answer is computed only once. But it is also bad, because the results for different calls to a predicate are all in one set, from which one has to select the result matching the current input arguments.

The Pull Method works already not completely bottom-up, but is controlled from the caller (top-down) who requests the next tuple. Therefore it is very natural that the caller passes all information he has about the required tuples. This would replace the magic set transformation, but it is not exactly the same, because now there is not one big magic set for a predicate (and a binding pattern), but for each call there are given values for certain arguments. This has positive and negative effects: When the call returns, one gets a tuple with the required values in the given positions, so no further check/selection is necessary. This is especially important since the pull method repeats computations, so non-matching tuples would simply be wasted. On the negative side, if the same call appears more than once, the result is computed repeatedly.

Passing bindings to called predicates fits nicely into the cursor interface, because for EDB-predicates with special access structures, it is already possible. In this way, this is also possible for IDB predicates.

For the materialization method and the push method, magic sets (or one of its variants) works well. However, if one wants to combine the different methods in one program (which is advisable, since each has its strengths and weaknesses), it would be possible to treat the magic set specially and to initialize it each time with only a single tuple.

Example code for the pull method with passing of bindings is available at

<http://www.informatik.uni-halle.de/~brass/botup/test2bind.cpp>

7 Performance Evaluation

Test 1: Basic Comparison, Push Implementation Variants

We use the following example program:

$$\begin{aligned} p3(X, Z) &\leftarrow p2(X, Y) \wedge q(Y, Z). \\ p2(X, Z) &\leftarrow p1(X, Y) \wedge q(Y, Z). \\ p1(X, Z) &\leftarrow p0(X, Y) \wedge q(Y, Z). \\ p0(X, Y) &\leftarrow r(X, Y). \end{aligned}$$

The EDB-predicates are defined as

$$\begin{aligned} q &:= \{(i, i+1) \mid 1 \leq i \leq 5000\} \\ r &:= \{(i, i) \mid 1 \leq i \leq 5000\} \end{aligned}$$

The EDB-predicates are implemented by computation, so no time is needed for loading facts from disk. The code assumes no index for the EDB-predicates, thus always a full table scan is needed. The main program opens a scan over $p3$, reads all tuples and adds the column values (so that no time is lost for printing). The relations for the materialization method were implemented as a linked list.

The times listed under “Windows PC” were measured on a PC with Intel Core2 Duo CPU E8200 @ 2.66GHz with 4 GB main memory, running Windows XP Professional, the programs were compiled with Microsoft Visual C++ 2008 Express Edition with the standard settings for the “Release” configuration.

The times listed under “Linux Server” were measured on a system with four Opteron 852 CPUs (2.6 GHz) and 16 GB memory running Ubuntu 8.04 64bit. The used compiler is gcc 4.2.4 with -O2 code optimization.

The times listed under “Solaris Server” were measured on a SunFire V490 (4 dual core UltraSparc IV processors, 16 GB memory) running SunOS 5.10. The gcc version was 3.4.6, -O2 optimization was used.

Method	Windows PC	Linux Server	Solaris Server
Method 1: Materialization	515ms	380ms	430ms
Method 2: Pull	937ms	320ms	930ms
Method 3: Push (no opt.)	562ms	250ms	950ms
Method 3: Push (with opt.)	546ms	250ms	950ms
Method 3: Push (static)	578ms	220ms	1100ms
Method 3: Push (local)	515ms	250ms	490ms
Method 3: Push (main)	515ms	240ms	330ms

It is interesting that the optimal method in this case depends a lot on hardware and compiler. It is also interesting that the materialization method is in fact very

good here, whereas we assumed that the copying of values makes it slow (well, the values were integers in this case, which are efficient to copy).

The “Push Method with optimizations” behaves not better than the version without optimizations. The optimizations done were that the “push” and “pop” of the cursors and argument variables was removed (since all rules are non-recursive), and the variables `p0_1`, `p1_1`, `p2_1`, `p3_1` were merged to one variable `X` (which also saves some assignments, since the first argument of all predicates is the same). However, merging the variables creates additional data dependencies, which might hinder pipelined execution in the processor. In one test (on the Linux machine), the “optimized version” even ran slower.

So the accessing of these variables (arguments of predicates and open cursors) is especially critical. In the versions marked “no opt” and “with opt” these variables are attributes of the class `p3_cursor`, of which the method `fetch` contains the main code of the push method. In the line “Push (static)” these variables were declared as “static”, i.e. class attributes. In the line “Push (local)” there were declared as local variables (then it is important that they are restored from the stack before using, so the optimization cannot be applied). The version “Push (main)” does not support the cursor-interface, but has the code of the push-method in the main procedure (however, the EDB-predicates `q` and `r` are accessed via cursors). All three, static, local, and main version are without the optimizations of the push-method.

In this example, the passing of bindings does not occur: When `p3` is called with binding pattern `ff`, also the other IDB predicates `p2`, `p1`, `p0` are called with both arguments free. Therefore, Method 4 is the same as Method 2 in this case.

Just for curiosity, we tried this logic program with SWI-Prolog 5.6.64 on the PC. The predicates `q` and `r` were defined as

```
q(I, J) :- between(1, 5000, X), X=I, J is I+1.
r(I, I) :- between(1, 5000, X), X=I.
```

Since the C++ programs had to do a loop over all `q`- and `r`-facts, it was necessary to call the built-in predicate `between` with the last argument free. Since the C++ program accesses all `p3`-tuples, and adds the column values, the following `main`-predicate seems comparable (the last condition always fails):

```
main :- p3(X, Y), Z is X+Y, Z < 0.
main.
```

The runtime in SWI-Prolog was 13.19s CPU time. B-Prolog 7.4 needs 2.55s CPU time with `between` defined as

```
between(I, J, J) :- J >= I.
between(I, J, K) :- J > I, J1 is J - 1, between(I, J1, K).
```

(In SWI-Prolog, `between` is a built-in predicate.) XSB 3.1 needs 12.9s for the same program. While this comparison is probably unfair for many reasons, it shows at least that the method we investigate has a certain potential.

If tabling is switched on for all predicates in B-Prolog, it needs 0.86s, which is comparable to our runtimes. However, since it presumably uses a hash table or

other index to access the q - and r -facts, it has a big advantage on the join. Our runtimes would also improve if q and r supported the binding pattern bf . XSB with tabling needs 4.36s.

Test 2: Loop Order, Passing Bindings

Next, we use a very similar program, but with the inverse order of the predicates in the body:

$$\begin{aligned} p3(X, Z) &\leftarrow q(X, Y) \wedge p2(Y, Z). \\ p2(X, Z) &\leftarrow q(X, Y) \wedge p1(Y, Z). \\ p1(X, Z) &\leftarrow q(X, Y) \wedge p0(Y, Z). \\ p0(X, Y) &\leftarrow r(X, Y). \end{aligned}$$

The performance results are as follows:

Method	Windows PC	Linux Server	Solaris Server
Method 1: Materialization	781ms	830ms	1950ms
Method 2: Pull	>1h		
Method 2: Pull (size 200)	24843ms	3400ms	19660ms
Method 3: Push (local)	531ms	290ms	340ms
Method 4: Pull+Bindings	609ms	260ms	330ms

As expected, the pull method behaves very badly in this case, because a lot of recomputations are done (the predicates $p0$, $p1$, $p2$ are in the inner loop of a nested loop join). For the problem size 200 (instead of 5000) it already needs more than 24s on the PC.

However, if one passes bindings to the predicates $p0$, $p1$, $p2$ according to the binding pattern bf , the recomputation is avoided and the pull method has a good performance.

The reason that the materialization method needs more time than for the first program is that now the inner loop is over the predicate implemented with a linked list, whereas the predicates q and r are implemented by computation (faster). If the iteration is not equally fast, it is better if the faster one is in the inner loop.

The push method works very similar as for the first program. This is not a surprise since it is controlled by the generated IDB facts, so not much changes here.

Test 3: Using Keys

We use the same program as in Test 2, but make use of knowledge of keys: If for q and r , the first argument is declared as a key, a program analysis shows that this also holds for $p0$, $p1$, $p2$, $p3$. That means that once a matching tuple is found in the inner loop, one can stop that loop. So on average one can save about half of the time. Of course, if one had special access structures (trees, hash tables), one could

save much more, but that is a different topic. For the pull method with passing of bindings, we tried two variants: One that kept the cursor interface, and one that uses functions with result parameters.

The push method cannot make use of the keys mentioned above, therefore, one would have the same runtime as in Test 2. However, it can use the fact that the second argument of q functionally determines the first, and we have reported this runtime in the table. Actually, the materialization method cannot use the key on q and r , it uses only the knowledge about the keys of p_0 , p_1 , p_2 . The pull method with passing of bindings uses the keys on the first arguments of q , r , p_0 , p_1 , p_2 . It is interesting that the three methods differ in the keys that they can utilize.

Method	Windows PC	Linux Server	Solaris Server
Method 1: Materialization	390ms	360ms	740ms
Method 3: Push (other key)	250ms	110ms	230ms
Method 4: Pull+Bindings	281ms	110ms	180ms
Method 4: Pull+Bind+Fun	265ms	110ms	190ms

Test 4: Using Ordering, Merge Join

The generators for q and r facts produce facts with a strictly ascending order of values in both arguments (this also implies that both columns are keys). From this, one can conclude that p_0 , p_1 , p_2 , and p_3 -facts are also produced in sorted order (with all three methods). If one uses this knowledge, one can do a merge join (without explicit sorting), which is extremely efficient. Already on the PC, all three methods show 0ms time. Therefore, we increased the number of q - and r -facts to 1 million.

Method	Windows PC	Linux Server	Solaris Server
Method 1: Materialization	765ms	710ms	1710ms
Method 2: Pull	93ms	10ms	90ms
Method 2: Push	78ms	20ms	80ms

Test 5: Disjunctions, Larger Tuples

The following program uses disjunctions and creates tuples with up to 7 columns:

```

p1(A) ← q(A).
p1(A) ← q(A).
p2(A,1) ← p1(A).
p2(A,2) ← p1(A).
p3(A,B,1) ← p2(A,B).
p3(A,B,2) ← p2(A,B).
...
p7(A,B,C,D,E,F,1) ← p6(A,B,C,D,E,F).
p7(A,B,C,D,E,F,2) ← p6(A,B,C,D,E,F).

```

The (computed) predicate q returned 50000 tuples.

Method	Windows PC	Linux Server	Solaris Server
Method 1: Materialization	1562ms	1220ms	4150ms
Method 2: Pull	390ms	20ms	100ms
Method 2: Push	453ms	300ms	400ms
Method 2: Push (with opt)	171ms	100ms	150ms

As expected, the large number of columns is bad for the materialization method. The pull method performs very well, although there is a certain amount of duplicated work: We implemented it in the straightforward way, which treats each rule for a predicate in isolation. A more clever way would be to note that the body of both rules of a predicate is identical and therefore a single scan over the body literal suffices. The push method inherently does this, because it is controlled by the IDB body literal. Therefore, the result tuples are also produced in a different order. In deductive databases, the order of derived facts is not important, but as we have seen, certain optimization possibilities depend on a specific order.

8 Related Work

There are still more variants of bottom-up evaluation proposed in the literature, which we intend to include in our comparison in a future version of this article:

- In (Liu and Stoller 2003), an extreme form of materialization is proposed: Not only facts about the derived predicates are explicitly stored, but also intermediate results during rule evaluation, e.g. partial joins. This is combined with a clever selection of data structures.
- In (Wunderwald 1995), a method similar to the push method is used, but with a materialization of the derived predicates (our push method avoids this materialization). Similar methods are also used to propagate changes from base relations to materialized views.

In summary, implementations of bottom-up evaluation differ in what they materialize (store in memory for a longer time), what they recompute, and the order in which applicable rule instances are considered (and the duration: the pull method has a rule instance “open” for a longer time).

9 Conclusions

I am convinced that deductive databases can still become a serious competitor to standard relational and object-relational databases for many practical projects. Declarative programming has many advantages, and for single queries this is already standard in the database field (SQL is a declarative language). Deductive databases would lift the declarativity to the level of programs, but this is not as easy as it was expected in the times when deductive databases were a hype topic. More research is still needed.

Our long-term goal is to develop a deductive database system that supports step-wise migration from classical SQL. Of course, the system will use our SLDmagic method (Brass 2000) for goal-direction, but it also needs an efficient bottom-up engine to run the transformed program. In this paper, we investigated several implementation variants based on a translation to C++. It turned out that the optimal method depends on the input program and also on the compiler and the hardware, as well as keys and access structures for the relations. But the push method performed constantly quite well. It has the restriction that it can work directly only with rules having only a single IDB-literal in the body, but it can be combined with other methods, or applied in several steps with different components of the program. Also, the SLDmagic method produces rules with only one IDB-predicate in the body.

The push method is also especially interesting because it seems very simple to execute on parallel hardware: The stack of tasks is in fact a set (if one does not use the optimizations). Since each task is small, the load balancing should be easy.

The next steps will be to continue the tests, also with parallel execution, with other data structures for the relations, and with access to external memory. Then we will also develop algorithms to find for a given logic program a good selection of the above methods and of access structures for the relations (similar to cost based optimization in standard relational databases).

The source code and performance results for the tests are available at

<http://www.informatik.uni-halle.de/~brass/botup/>

Results of future tests will also be posted on this page.

References

- BANCILHON, F., MAIER, D., SAGIV, Y., AND ULLMAN, J. D. 1986. Magic sets and other strange ways to implement logic programs. In *Proc. of the 5th ACM Symp. on Principles of Database Systems (PODS'86)*. ACM Press, 1–15.
- BANCILHON, F. AND RAMAKRISHNAN, R. 1986. An amateur's introduction to recursive query processing. In *Proceedings of the 1986 ACM SIGMOD international conference on Management of Data (SIGMOD'86)*, C. Zaniolo, Ed. 16–52.
- BEERI, C. AND RAMAKRISHNAN, R. 1991. On the power of magic. *The Journal of Logic Programming* 10, 255–299.
- BRASS, S. 2000. SLDMagic — the real magic (with applications to web queries). In *First International Conference on Computational Logic (CL'2000/DOOD'2000)*, W. Lloyd et al., Eds. Number 1861 in LNCS. Springer, Heidelberg, Berlin, 1063–1077.

- FREITAG, B., SCHÜTZ, H., AND SPECHT, G. 1991. LOLA — A logic language for deductive databases and its implementation. In *Proc. of 2nd Int. Symp. on Database Systems for Advanced Applications (DASFAA'91)*. World Scientific, 216–225.
- FRIESEN, O., GAUTHIER-VILLARS, G., LEFEBVRE, A., AND VIEILLE, L. 1995. Applications of deductive object-oriented databases using DEL. In *Applications of Logic Databases*, R. Ramakrishnan, Ed. Kluwer, 1–22.
- LIU, Y. A. AND STOLLER, S. D. 2003. From datalog rules to efficient programs with time and space guarantees. In *Proceedings of the 5th ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming (PPDP'03)*. ACM, 172–183.
- MINKER, J. 1988. Perspectives in deductive databases. *The Journal of Logic Programming* 5, 33–60.
- RAMAKRISHNAN, R., SRIVASTAVA, D., AND SUDARSHAN, S. 1994. Rule ordering in bottom-up fixpoint evaluation of logic programs. *IEEE Transactions on Knowledge and Data Engineering* 6, 4, 501–517.
- RAMAKRISHNAN, R. AND ULLMAN, J. D. 1995. A survey of deductive database systems. *The Journal of Logic Programming* 23, 125–149.
- SAGONAS, K., SWIFT, T., AND WARREN, D. S. 1994. XSB as an efficient deductive database engine. In *Proc. of the 1994 ACM SIGMOD Int. Conf. on Management of Data (SIGMOD'94)*, R. T. Snodgrass and M. Winslett, Eds. 442–453.
- ULLMAN, J. D. AND ZANIOLO, C. 1990. Deductive databases: Achievements and future directions. *ACM SIGMOD Record* 19, 75–82.
- VAGHANI, J., RAMAMOCHANARAO, K., KEMP, D., SOMOGYI, Z., STUCKEY, P. J., LEASK, T. S., AND HARLAND, J. 1994. The Aditi deductive database system. *The VLDB Journal* 3, 245–288.
- WUNDERWALD, J. E. 1995. Memoing evaluation by source-to-source transformation. In *Logic Programming Synthesis and Transformation, 5th International Workshop (LOPSTR'95)*, M. Proietti, Ed. Number 1048 in LNCS. Springer-Verlag, 17–32.