

An Amateur's Introduction to Recursive Query Processing Strategies

Francois Bancilhon (1)
Raghu Ramakrishnan (1,2)

1 MCC
9430 Research Blvd
Austin, Texas 78759

2 University of Texas at Austin
Austin, Texas 78712

ABSTRACT

This paper surveys and compares various strategies for processing logic queries in relational databases. The survey and comparison is limited to the case of Horn Clauses with evaluable predicates but without function symbols. The paper is organized in three parts. In the first part, we introduce the main concepts and definitions. In the second, we describe the various strategies. For each strategy, we give its main characteristics, its application range and a detailed description. We also give an example of a query evaluation. The third part of the paper compares the strategies on performance grounds. We first present a set of sample rules and queries which are used for the performance comparisons, and then we characterize the data. Finally, we give an analytical solution for each query/rule system. Cost curves are plotted for specific configurations of the data.

1. Introduction

The database community has recently manifested a strong interest in the problem of evaluating "logic queries" against relational databases. This interest is motivated by two converging trends: (i) the desire to integrate database technology and artificial intelligence technology, i.e., to extend database systems, to provide them with the functionality of expert systems, thus creating "knowledge base systems" and (ii) the desire to

integrate logic programming technology and database technology, i.e., to extend the power of the interface to the database system to that of a general purpose language. The second goal is of a somewhat different nature and has found in its ranks proponents of object oriented, functional and imperative as well as logic based programming languages. The logic programming camp is relying on the fact that logic programming and relational calculus have the same underlying mathematical model, namely first order logic.

Of course, database researchers already know how to evaluate logic queries: the view mechanism, as offered by most relational systems, is a form of support of a restricted set of logic queries. But those logic queries are restricted to be non-recursive and the problem of efficiently supporting recursive queries is still open.

In the past five years, following the pioneering work by Chang, Shapiro and McKay, and Henschen and Naqvi, numerous strategies have been

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1986 ACM 0-89791-191-1/86/0500/0016 \$00.75

proposed to deal with recursion in logic queries. The positive side of this work is that there are a lot of algorithms offered to solve *the* problem. The negative side is that we do not know how to make a choice of an algorithm. It seems reasonable to say that all these strategies can only be compared on three grounds: functionality (i.e., application domain), performance and ease of implementation. However, each of these algorithms is described at a different level of detail, and it is sometimes difficult to understand their differences. In fact, we shall claim later in this paper that some of them are indeed identical. Each comes with little or no performance analysis, and the application domain is not always easy to identify. We try in this paper to evaluate these algorithms with respect to these three criteria. We describe all the algorithms at the same level of detail and demonstrate their behavior on common examples. This is not always easy to do since some of them are fairly well formalized while others are merely sketched as an idea.

For each one of them, we state in simple terms the application domain. Finally, we give a first simple comparison of the performance of these algorithms. Choosing a simple set of typical queries, a simple characterization of the data and a simple cost function, we give an analytical evaluation of the cost of each strategy. The results give a first insight into the respective value of all the proposed strategies.

The rest of the paper is organized as follows. In section 2 we present our definitions and notations, and introduce the main ideas. In section 3 we present the main features of the strategies, and describe each one individually, and finally, in section 4, we present the performance evaluation methodology and results.

2. Logic Databases

2.1. An Example

Let us start by discussing informally an example. Here is what we call a "logic database"

```
parent(cain,adam)
parent(abel,adam)
parent(cain,eve)
parent(abel,eve)
parent(sem,abel)
ancestor(X,Y) - ancestor(X,Z),ancestor(Z,Y)
ancestor(X,Y) - parent(X,Y)
generation(adam,1)
generation(X,I) - generation(Y,J),
    parent(X,Y),J=I-1
generation(X,I) - generation(Y,J),
```

parent(Y,X),J=I+1

Note that this is a purely syntactic object. In this database, we have a set of predicate or relation names (parent, ancestor and generation), a set of arithmetic predicates ($I=J+1$, $I=J-1$) and a set of constants (adam, eve, cain, sem and abel). Finally, we have a set of variables (X, Y and Z). The database consists of a set of sentences ending with a period. "parent(cain,adam)" is a fact, and "ancestor(X,Y) - parent(X,Y)" is a rule.

Let us now associate a meaning with the database. We first associate with each constant an object from the real world: thus, with "adam" we associate the individual whose name is "adam". Then, we associate with each arithmetic predicate name the corresponding arithmetic operator. Then we can interpret intuitively each fact and each rule. For instance we interpret "parent(cain,adam)" by saying that the predicate parent is true for the couple (cain,adam), and we interpret the rule

ancestor(X,Y) - ancestor(X,Z), ancestor(Z,Y)

by saying that if there are three objects X, Y and Z such that ancestor(X,Z) is true and ancestor(Z,Y) is true then ancestor(X,Y) is true.

This leads to an interpretation which associates with each predicate a set of tuples. For instance with the predicate ancestor we associate the interpretation $\{(cain,adam), (abel,adam), (cain,eve), (abel,eve), (sem,abel), (sem,adam), (sem,eve)\}$, and with the predicate generation we associate the interpretation $\{(adam,1), (eve,1), (cain,2), (abel,2), (sem,3)\}$.

The problem is to answer queries, given the logic database. For instance given a query of the form generation(sem,?) or ancestor(?,adam), how do we find the answer generation(sem,3) and $\{ancestor(cain,adam), ancestor(abel,adam), ancestor(sem,adam)\}$?

Let us now formalize all the notions encountered in this example and define a logic database. We first define it syntactically, then we attach an interpretation to this syntax.

2.2. Syntax of a Logic Database

We first define four sets of names: *variable* names, *constant* names, *predicate* or *relation* names and *evaluable predicate* names.

We adopt the Prolog convention of denoting variables by strings of characters starting with an upper case letter and constants by strings of characters starting with a lower case letter or integers. For instance X1, Father and Y are variables, while

john, salary and 345 are constants

We use identifiers starting with lower case letters for predicate names and relation names (evaluable and non-evaluable)

We use the term relation (from database terminology) and predicate (from logic terminology) indifferently to represent the same object. We shall however interpret them differently: a relation will be interpreted by a set of tuples and a predicate by a true/false function. There is a fixed arity associated with each relation/predicate.

The set of evaluable predicate names is a subset of the set of predicate names. We will not be concerned with their syntactic recognition, in the examples it will be clear from the name we use. The main examples of evaluable predicate names are arithmetic predicates. For instance, sum, difference and greater-than are examples of evaluable predicates of arity 3, 3 and 2 respectively, while parent and ancestor are non-evaluable predicates of arity 2.

A *literal* is of the form $p(t_1, t_2, \dots, t_n)$ where p is a predicate name of arity n and each t_i is a constant or a variable. For instance $\text{father}(\text{john}, X)$, $\text{ancestor}(Y, Z)$, $\text{id}(\text{john}, 25, \text{austin})$ and $\text{sum}(X, Y, Z)$ are literals. An *instantiated literal* is one which does not contain any variables. For instance $\text{id}(\text{john}, \text{doe}, 25, \text{austin})$ is an instantiated literal, while $\text{father}(\text{john}, \text{Father})$ is not.

We allow ourselves to write evaluable literals using functions and equality for the purpose of clarity. For instance, $Z = X + Y$ denotes $\text{sum}(X, Y, Z)$, $I = J + 1$ denotes $\text{sum}(J, 1, I)$, and $X > 0$ denotes $\text{greater-than}(X, 0)$.

If $p(t_1, t_2, \dots, t_n)$ is a literal, we call (t_1, t_2, \dots, t_n) a *tuple*.

A *rule* is a statement of the form

$$p \text{ - } q_1, q_2, \dots, q_n$$

where p and the q_i 's are literals such that the predicate name in p is a non-evaluable predicate. p is called the *head* of the rule, and each of the q_i 's is called a *goal*. The conjunction of the q_i 's is the *body* of the rule. We have adopted the Prolog notation of representing implication by $'\text{-}'$ and conjunction by $'\text{'}$. For instance

$\text{uncle}(\text{john}, X) \text{ - } \text{brother}(X, Y), \text{parent}(\text{john}, Y)$

is a rule with head $\text{"uncle}(\text{john}, X)\text{"}$ and body $\text{"brother}(X, Y), \text{parent}(\text{john}, Y)\text{"}$.

A *ground clause* is a rule in which the body is empty. A *fact* is a ground clause which contains no variables. For instance

$\text{loves}(X, \text{john})$
 $\text{loves}(\text{mary}, \text{susan})$

are ground clauses, but only the second of these is a fact.

A *database* is a set of rules, note that this set is not ordered. Given a database, we can partition it into a set of facts and the set of all other rules. The set of facts is called the *extensional* database, and the set of all other rules is called the *intensional* database.

2.3. Interpretation of a Logic Database

Up to now our definitions have been purely syntactical. Let us now give an interpretation of a database. This will be done by associating with each relation name in the database a set of instantiated tuples. We first assume that with each evaluable predicate p is associated a set $\text{natural}(p)$ of instantiated tuples which we call its *natural interpretation*. For instance, with the predicate sum is associated an infinite set of all the 3-tuples (x, y, z) of integers such that the sum of x and y is z . In general the natural interpretation of an evaluable predicate is infinite.

Given a database, an *interpretation* of this database is a mapping which associates with each relation name a set of instantiated tuples.

A *model* of a database is an interpretation I such that

- (1) for each evaluable predicate p ,
 $I(p) = \text{natural}(p)$, and,
- (2) for any rule,
 $p(t) \text{ - } q_1(t_1), q_2(t_2), \dots, q_n(t_n)$

for any instantiation σ of the variables of the rule such that $\sigma(t_i)$ is in the interpretation of q_i for all i then $\sigma(t)$ is in the interpretation of p .

This is simply a way of saying that, in a model, if the right hand side is true then the left hand side is also true. This implies that for every fact $p(x)$ of the database the tuple x belongs to the interpretation of p .

Of course, for a given database there are many models. The nice property of Horn Clauses is that among all these models there is a *minimal* one (minimal in the sense of set inclusion), which is the one we choose as *the* model of the database [Van Emden and Kowalski 76]. Therefore from now on, when we talk about the model or the interpretation of a database, we mean its minimal model.

Notice that because of the presence of evaluable arithmetic predicates the minimal model is, in gen-

eral, not finite

Let p be an n -ary predicate. An *adornment* of p is a sequence a of length n of b 's and f 's [Ullman 85]. For instance bbf is an adornment of a ternary predicate, and $fbff$ is an adornment of predicate of arity 4. An adornment is to be interpreted intuitively as follows: the i -th variable of p is bound (respectively free) if the i -th element of a is b (respectively f). Let $p(x_1, x_2, \dots, x_n)$ be a literal, an adornment $a_1 a_2 \dots a_n$ of that literal is an adornment of p such that

- (i) if x_i is a constant then a_i is b ,
- (ii) if $x_i = x_j$ then $a_i = a_j$

We denote adornments by superscripts. A *query form* is an adorned predicate. Examples of query forms are father^{bf} , id^{bfff}

A *query* is a query form and an instantiation of the bound variables. We denote it by an adorned literal where all the bound positions are filled with the corresponding constants and the free positions are filled by distinct free variables. Therefore $\text{father}^{bf}(\text{john}, X)$ and $\text{id}^{bfff}(\text{john}, X, Y, 25)$ are queries. The distinction between queries and query forms is that query forms are actually compiled, and at run-time their parameters will be instantiated. Notice that $\text{father}(X, X)$ is not a query form in this formalism.

The *answer* to a query $q(t)$ is the set $\{q(\sigma(t)) \mid \sigma \text{ is an instantiation of } t, \text{ and } \sigma(t) \text{ is in the interpretation of } q\}$

2.4. Structuring and Representing the Database

A predicate which only appears in the intensional database is a *derived* predicate. A predicate which appears only in the extensional database or in the body of a rule is a *base* predicate.

For performance reasons, it is good to decompose the database into a set of pure base predicates (which can then be stored using a standard DBMS) and a set of pure derived predicates. Fortunately, such a decomposition is always possible, because every database can be rewritten as an "equivalent" database containing only base and derived predicates. By equivalent, we mean that all the predicate names of the original database appear in the modified database and have the same interpretation.

We obtain this equivalent database in the following way: consider any predicate p that is neither base nor derived. By definition, we have a set of facts for p , and p appears on the left of some rules. So we simply introduce a new predicate p_ext and do

the following

- (1) replace p by p_ext in each fact of p ,
- (2) add a new rule of the form $p(X_1, X_2, \dots, X_n) \text{ - } p_ext(X_1, X_2, \dots, X_n)$ where n is the arity of p

Example

$\text{father}(a, b)$
 $\text{parent}(b, c)$
 $\text{grandfather}(b, d)$
 $\text{grandfather}(X, Y) \text{ - } \text{father}(X, Z), \text{parent}(Z, Y)$

becomes

$\text{father}(a, b)$
 $\text{parent}(b, c)$
 $\text{grandfather_ext}(b, d)$
 $\text{grandfather}(X, Y) \text{ - } \text{father}(X, Z), \text{parent}(Z, Y)$
 $\text{grandfather}(X, Y) \text{ - } \text{grandfather_ext}(X, Y)$

Most authors have chosen to describe a set of rules through some kind of graph formalism. Predicate Connection Graphs, as presented in [McKay and Shapiro 81], represent the relationship between rules and predicates. Rule/goal graphs, as presented in [Ullman 85], carry more information because predicates and rules are adorned by their variable bindings. We have chosen here to keep the rule/goal graph terminology while using unadorned predicates.

The *rule/goal* graph has two sets of nodes: square nodes which are associated with predicates, and oval nodes which are associated with rules. If there is a rule of the form


$r \text{ - } p_1, p_2, \dots, p_n$

in the intensional database, then there is an arc going from node r to node p , and for each predicate p_i there is an arc from node p_i to node r .

Here is an example of an intensional database. For the sake of simplicity, we have omitted the variables in the rules.

$r_1 \quad p_1 \text{ :- } p_3, p_4$
 $r_2 \quad p_2 \text{ - } p_4, p_5$
 $r_3 \quad p_3 \text{ - } p_6, p_4, p_3$
 $r_4 \quad p_4 \text{ - } p_5, p_3$
 $r_5 \quad p_3 \text{ - } p_6$
 $r_6 \quad p_5 \text{ :- } p_5, p_7$
 $r_7 \quad p_5 \text{ :- } p_6$
 $r_8 \quad p_7 \text{ :- } p_8, p_9$

The rule/goal graph is:



2.5. Recursion

Recursion is often discussed in the single rule context. For the purpose of clarity and simplicity, let us first give some temporary definitions in this context. We say that a rule is recursive if it is of the form

$$p(t) - \dots, p(t')$$

For instance the rule

$$\text{ancestor}(X,Y) - \text{ancestor}(X,Z), \text{parent}(Z,Y)$$

is recursive

An interesting subcase is that of linear rules. Linear rules play an important role because (i) there is a belief that most "real life" recursive rules are indeed linear, and (ii) algorithms have been developed to handle them efficiently.

We say that a rule is linear if it is recursive, and the recursive predicate appears once and only once on the right. This property is sometime referred to as regularity [Chang 81]. We believe the term linear to be more appropriate, and we think that regularity should be kept for another concept (which is not defined here).

For instance the rule

$$\text{sg}(X,Y) - p(X,XP), p(Y,YP), \text{sg}(XP,YP)$$

is linear, while the rule

$$\text{ancestor}(X,Y) - \text{ancestor}(X,Z), \text{ancestor}(Z,Y)$$

is not

These definitions are fairly simple in the single rule context. They are a little more involved in the context of a set of rules where properties have to be attached to predicates instead of rules. Consider the following database

$$\begin{aligned} p(X,Y) &- b1(X,Z), q(Z,Y) \\ q(X,Y) &- p(X,Z), b2(Z,Y) \end{aligned}$$

Neither of the rules are recursive according to the above definition, while clearly both predicates p and q are recursive.

We now come to the general definitions of recursion in the multirule context. Let p and q be two predicates. We say that p derives q (denoted $p \rightarrow q$) if it occurs in the body of a rule whose head predicate is q . We define $\rightarrow+$ to be the transitive closure (not the reflexive transitive closure) of \rightarrow . A predicate p is said to be recursive if $p \rightarrow+ p$. Two predicates p and q are mutually recursive if $p \rightarrow+ q$ and $q \rightarrow+ p$. It can be easily shown that mutual recursion is an equivalence relation on the set of recursive predicates. Therefore the set of recursive predicates can be decomposed into disjoint blocs of mutually recursive predicates.

Given a set of rules, we say that the rule

$$p - p_1, p_2, \dots, p_n$$

is recursive iff there exists p_i in the body of the rule which is mutually recursive to p .

A recursive rule $p - p_1, p_2, \dots, p_n$ is linear if there is one and only one p_i in the body of the rule which is mutually recursive to p . A set of rules is linear if every recursive rule in it is linear. For instance, the following system is linear.

$$\begin{aligned} r1 & p(X,Y) - p1(X,Z), q(Z,Y) \\ r2 & q(X,Y) - p(X,Z), p2(Z,Y) \\ r3 & p(X,Y) - b3(X,Y) \\ r4 & p1(X,Y) - b1(X,Z), p1(Z,Y) \\ r5 & p1(X,Y) - b4(X,Y) \\ r6 & p2(X,Y) - b2(X,Z), p2(Z,Y) \\ r7 & p2(X,Y) - b5(X,Y) \end{aligned}$$

The set of recursive predicates is $\{p, q, p1, p2\}$, the set of base predicates is $\{b1, b2, b3, b4, b5\}$. The blocks of mutually recursive predicates are $\{p, q, p1, p2\}$. The recursive rules are $r1, r2, r4$ and $r6$, and the system is linear even though rules $r1$ and $r2$ both have two recursive predicates on their right.

We say that two recursive rules are mutually recursive iff the predicates in their heads are mutu-


ally recursive. This defines an equivalence relation among the recursive rules.

Thus mutual recursion defines an equivalence class among recursive predicates and among the recursive rules, [Bancilhon 85]. Therefore, it groups together all predicates which are mutually recursive to one another, i.e. which must be evaluated as a whole. It also groups together all the rules which participate in evaluating those blocks of predicates. Let us now see how this can be represented in the rule/goal graph. We define the *reduced rule/goal graph* as follows.

Square nodes are associated with non-recursive predicates or with blocks of mutually recursive predicates and, oval nodes are associated with non-recursive rules or with blocks of mutually recursive rules. The graph essentially describes the non-recursive part of the database by grouping together all the predicates which are mutually recursive to one another and isolating the recursive parts. For every non-recursive rule of the form $r: p - p_1, p_2, \dots, p_n$, there is an arc going from node r to node p (if p is non-recursive), or to node $[p]$, which is the node representing the set of predicates mutually recursive to p (if p is recursive). For each non-recursive predicate p_i , there is an arc from the node p_i to the node r , and for each recursive predicate p_j there is an arc going from $[p_j]$ (the node representing the set of predicates mutually recursive to p_j)

Finally, each bloc of recursive rules $[r]$ is uniquely associated to a set of mutually recursive predicates $[p]$, and we draw an arc from $[p]$ to $[r]$ and an arc from $[r]$ to $[p]$. We also draw an arc from q (if q is non-recursive) or from $[q]$ (if q is recursive) to $[r]$ if there is a rule in $[r]$ which has q in its body. This grouping of recursive predicates in blocks of strongly connected components is presented in [Morris et al 86].

Here is the representation of the previous database



2.6. Safety of Queries

Given a query q in a database D , we say that q is *safe* in D if the answer to q is finite. Obviously unsafe queries are highly undesirable.

Sources of unsafeness are of two kinds:

(i) the evaluable arithmetic predicates are interpreted by infinite tables. Therefore they are unsafe by definition. For instance the query $\text{greater-than}(27, X)$ is unsafe.

(ii) rules with free variables in the head which do not appear in the body are a source of unsafeness in the presence of evaluable arithmetic predicates (the arithmetic predicates provide an infinite underlying domain, and the variable from the head of the rule which does not appear on the right ranges over that domain). Thus for instance, in the system

```
good-salary(X) - X > 100000
like(X, Y) - nice(X)
nice(john)
```

the query $\text{like}(\text{john}, X)?$ is unsafe because, in the minimal model of the database $\text{like}(\text{john}, x)$ is true for every integer x . Note that if the first rule was not there, $\text{like}(\text{john}, X)?$ would be safe and have answer $\text{like}(\text{john}, \text{john})$.

The problem of safety has received a lot of attention recently [Afrati et al 86, Ullman 85, Zaniolo

86] We shall not survey those results here but merely present some simple sufficient syntactic conditions to guarantee safety. A rule is *range restricted* if every variable of the head appears somewhere in the body. Thus in this system

r1 loves(X,Y) - nice(X)
 r2 loves(X,Y) - nice(X),human(Y)

r1, which corresponds to "nice people love everything", is not range restricted while r2, which corresponds to "nice people love all humans", is. Obviously, every ground rule which is not a fact is not range restricted. For instance

loves(john,X)

is not range restricted

A set of rules is range restricted if every rule in this set is range restricted.

It is known [Reiter 78] that if each evaluable predicate has a finite natural interpretation, and if the set of rules is range restricted, then every query defined over this set of rules is safe. This applies obviously to the case where there are no evaluable predicates. However, if there are evaluable predicates with infinite natural interpretations, safety is no longer assured. We now present a simple sufficient condition for safety in the presence of such predicates.

A rule is *strongly safe* iff (1) it is range restricted, and (2) every variable in an evaluable predicate term also appears in at least one base predicate.

For example, the rule

well-paid(X) - has-salary(X,Y), Y > 100K

is strongly safe, whereas

great-salary(X) - X > 100K

is not strongly safe.

A set of rules is strongly safe if every rule in this set is strongly safe.

Any query defined over a set of strongly safe rules is safe. However, while this is a sufficient condition, it is not necessary. We can develop better conditions for testing safety, or leave it to the user to ensure that his queries are safe.

2.7. Effective Computability.

Safety, in general, does not guarantee that the query can be effectively computed. Consider for instance

p1(1,X,Y) - X ≥ Y
 p2(X,Y,2) - X ≤ Y
 p(X,Y) - p1(X,Z,Z),p2(Z,Z,Y)

The query p(X,Y) is safe (the answer is {p(1,2)}), but there is no safe computation for it.

However, strongly safe rules are guaranteed to be safe and safely computable.

In fact, while we might often be willing to let the user ensure that his queries are safe, it is desirable to ensure that the query can be computed without materializing "infinite" intermediate results. We now present a sufficient condition for ensuring this.

We first need some information about the way arithmetic predicates can propagate bindings. So we characterize each arithmetic predicate by a set of *safety dependencies* [Zaniolo 86]. A safety dependency is a couple (X → Y) where X is a set of attributes and Y is a set of attributes. It is to be interpreted intuitively as "if the values of the X attributes are fixed then there is a finite number of values of the Y attributes associated with them". Therefore, while their semantics is different from that of functional dependencies, they behave in the same fashion (and have the same axiomatization). Of course, we assume that the natural interpretation of the evaluable predicate satisfies the set of safety dependencies.

For instance, the ternary arithmetic predicate "sum" has the safety dependencies

{1,2} → {3}
 {1,3} → {2}
 {2,3} → {1}

while the arithmetic predicate "greater than" has only trivial safety dependencies.

Now consider a rule, and define each variable in the body to be *secure* if it appears in a non-evaluable predicate in the body or if it appears in position 1 in an evaluable predicate p and there is a subset I of the variables of p which are secure and I → {1}. Note that the definition is recursive.

A rule is *bottom-up evaluable* if

- 1 it is range restricted, and
- 2 every variable in the body is secure

For instance

p(X,Y) - Y=X+1, X=Y1+Y2, p(Y1,Y2)

is bottom-up evaluable because (i) Y1 and Y2 are secure (they appear in p which is non-evaluable), (ii) in X=Y1+Y2, the safety dependency {Y1,Y2} → {X} holds, therefore X is secure, and (iii) in Y=X+1, the safety dependency {X} → {Y} holds,

therefore Y is secure

On the contrary

$$p(X,Y) - X > Y1, q(Y1,Y)$$

is not bottom-up evaluable because X is not secure

A set of rules is bottom-up evaluable if every rule in this set is bottom-up evaluable

Any computation using only a set of bottom-up evaluable rules can be carried out without materializing infinite intermediate results. The computation proceeds in a strictly bottom-up manner, using values for the body variables to produce values for the head variables. The bottom-up evaluability criterion ensures that the set of values for body variables is finite at each step. However, there may be an infinite number of steps. For example, if we repeatedly apply the bottom-up evaluable rule given above, at each step we have a finite number of values (in this case, a unique value) for Y1 and Y2, and hence for X and Y. However, we can apply the rule an infinite number of times, producing new values for X and Y at each step.

3. The Strategies

In the past five years, a large number of strategies to deal with Horn rules have been presented in the literature. A strategy is defined by (i) an application domain (i.e., a class of rules for which it applies) and (ii) an algorithm for replying to queries given such a set of rules.

In studying the strategies, we found that the methods were described at different levels of detail and using different formalisms, that they were sometimes very difficult to understand (and sometimes were understood differently by subsequent authors), that the application domain was not always very clearly defined, and that no performance evaluation was given for any of the strategies, which left the choice of a given strategy completely open when the application domain was the same. Finally, we found that some of the strategies were in fact the same.

We think that the strategies should be compared according to the following criteria: (i) size of the application domain, (the larger the better), (ii) performance of the strategy, (the faster the better) and (iii) ease of implementation (the simpler the better). While the last criterion is somehow subjective, the first two should be quantifiable. In this section, we give a complete description of our understanding of the strategies and of their application domains, and we demonstrate each one of them through an example. As much as possible, we have tried to use the same example, except for

some "specialized" strategies where we have picked a specific example which exhibits its typical behavior.

3.1. Characteristics of the Strategies

3.1.1. Query Evaluation vs. Query Optimization

Let us first distinguish between two approaches. One first class of strategies consists of an actual query evaluation algorithm, i.e. a program which, given a query and a database, will produce the answer to the query. We will call these *methods*. Representatives of this class are *Henschen-Nagvi*, *Query/Subquery (QSQ)* or *Extension Table*, *APEX*, *Prolog*, *Naive Evaluation* and *Semi-Naive Evaluation*.

The strategies in the second class assume an underlying simple strategy (which is in fact naive or semi-naive evaluation) and optimize the rules to make their evaluation more efficient. They can all be described as *term rewriting systems*. These include *Aho-Ullman*, *Counting and Reverse Counting*, *Magic Sets* and *Kifer-Lozinski*.

Note that this distinction is somehow arbitrary. Each of the optimization strategies could be described as a method (when adding to it naive or semi-naive evaluation). However, this decomposition has two advantages: (i) it *might* make sense from an implementation point of view to realize the optimization strategies as term rewriting systems on top of an underlying simpler method such as naive evaluation, and (ii) from a pedagogical standpoint, they are much easier to understand this way, because presenting them as term rewriting systems indeed captures their essence.

The subsequent characteristics only relate to pure methods.

3.1.2. Interpretation vs. Compilation

A method can be *interpreted* or *compiled*. The notion is somehow fuzzy, and difficult to characterize formally. We say that the strategy is compiled if it consists of two phases: (i) a compilation phase, which accesses only the intensional database, and which generates an "object program" of some form, and (ii) an execution phase, which executes the object program against the facts only. A second characteristic of compiled methods is that all the database query forms (i.e., the query forms on base relations which are directly sent to the DBMS) are generated during the compilation phase. This condition is very important, because it allows the DBMS to precompile the the query

forms. Otherwise the database query forms are repetitively compiled by the DBMS during the execution of the query, which is a time consuming operation. If these two conditions do not hold, we say that the strategy is interpreted. In this case, no object code is produced and there is a fixed program, the "interpreter", which runs against the query, the set of rules and the set of facts.

3.1.3. Recursion vs. Iteration

A rule processing strategy can be *recursive* or *iterative*. It is iterative if the "target program" (in case of a compiled approach) or the "interpreter" (in case of the interpreted approach) is iterative. It is recursive if this program is recursive, i.e., uses a stack as a control mechanism. Note that in the iterative methods, the data we deal with is statically determined. For instance, if we use temporary relations to store intermediate results, there are a finite number of such temporary relations. On the contrary, in recursive methods the number of temporary relations maintained by the system is unbounded.

3.1.4. Potentially Relevant Facts

Let D be a database and q be a query. A fact $p(a)$ is *relevant* to the query iff there exists a derivation $p(a) \rightarrow^* q(b)$ for some b in the answer set. The notion of relevant fact was introduced in [Lozinski 85], we use it here with a somewhat different meaning. If we know all the relevant facts in advance, instead of using the database to reply to the query, we can use the relevant part of the database only, thus cutting down on the set of facts to be processed. A *sufficient set of relevant facts* is a set of facts such that replacing the database by this set of facts gives the same answer to the query. Unfortunately, in general there does not exist a unique minimal set of facts as the following example shows.

```

suspect(X) - long-hair(X)
suspect(X) - alien(X)
long-hair(antoine)
alien(antoine)

```

Minimal sets of facts with respect to the query $\text{suspect}(X)?$ are $\{\text{long-hair}(\text{antoine})\}$ and $\{\text{alien}(\text{antoine})\}$. The second unfortunate thing about relevant facts is that it is in general impossible to find all the relevant facts in advance without spending as much effort as in replying to the query. Thus, all methods have a way of finding a super-set of relevant facts. We call this set the *set of potentially relevant facts*. A set of potentially relevant facts is *valid* if it contains a sufficient set

of relevant facts. An obvious but not very interesting valid set is the set of all facts of the database.

3.1.5. Top Down vs. Bottom Up

Consider the following set of rules and the query

```

ancestor(X,Y) - parent(X,Z), ancestor(Z,Y)
ancestor(X,Y) - parent(X,Y)
query(X) - ancestor(john,X)

```

We can view each of these rules as productions in a grammar. In this context, the database predicates (parent in this example) appear as terminal symbols, and the derived predicates (ancestor in this example) appear as the non-terminal symbols. Finally, to pursue the analogy, we shall take the distinguished symbol to be $\text{query}(X)$. Of course, we know that the analogy does not hold totally, for two reasons: (i) the presence of variables and constants in the literals and (ii) the lack of order between the literals of a rule (for instance "parent(X,Z), ancestor(Z,Y)" and "ancestor(Z,Y), parent(X,Z)" have the same meaning). But we shall ignore these differences, and use the analogy informally.

Let us now consider the language generated by this "grammar". It consists of

```

{parent(john,X),
parent(john,X),parent(X,X1),
parent(john,X),parent(X,X1),parent(X1,X2),
}

```

This language has two interesting properties: (i) it consists of first order sentences involving only base predicates, i.e., each word of this language can be directly evaluated against the database, and (ii) if we evaluate each word of this language against the database and take the union of all these results, we get the answer to the query.

There is a minor problem here: the language is not finite, and we would have to evaluate an infinite number of first order sentences. To get out of this difficulty, we use termination conditions which tell us when to stop. An example of such a termination condition is: if one word of the language evaluates to the empty set, then all the subsequent words will also evaluate to the empty set, so we can stop generating new words. Another example of a termination condition is: if a word evaluates to a set of tuples, and all these tuples are already in the evaluation of the words preceding it, then no new tuple will ever be produced by the evaluation of any subsequent word, thus we can stop at this point.

All query evaluation methods in fact do the following

(i) generate the language, (ii) while the language is generated, evaluate all its sentences and (iii) at each step, check for the termination condition

Therefore, there are essentially two classes of methods those which generate the language bottom up, and those which generate the language top-down. The bottom-up strategies start from the terminals (i.e., the base relations) and keep assembling them to produce non-terminals (i.e. derived relations) until they generate the distinguished symbol (i.e., the query). The top-down strategies start from the distinguished symbol (the query) and keep expanding it by applying the rules to the non-terminals (derived relations). As we shall see, top-down strategies are often more efficient because they "know" which query is being solved, but they are more complex. Bottom up strategies are simpler, but they compute a lot of useless results because they do not know what query they are evaluating.

3.2. The Methods

We shall use the same example for most of the methods. The intensional database and query are

```
r1 ancestor(X,Y) - parent(X,Z),ancestor(Z,Y)
r2 ancestor(X,Y) - parent(X,Y)
r3 query(X) - ancestor(aa,X)
```

The extensional database is

```
parent(a,aa)
parent(a,ab)
parent(aa,aaa)
parent(aa,aab)
parent(aaa,aaaa)
parent(c,ca)
```

3.2.1. Naive Evaluation

Naive Evaluation is a bottom-up, compiled, iterative strategy

Its application domain is the set of bottom-up evaluable rules

In a first phase, the rules which derive the query are compiled into an iterative program. The compilation process uses the reduced rule/goal graph. It first selects all the rules which derive the query. A temporary relation is assigned to each derived predicate in this set of rules. A statement which computes the value of the output predicate from the value of the input predicates is associated with each rule node in the graph. With each set of mutually recursive rules, there is associated a loop

which applies the rules in that set until no new tuple is generated. Each temporary relation is initialized to the empty set. Then computation proceeds from the base predicates capturing the nodes of the graph.

In this example, the rules which derive the query are {r1, r2, r3}, and there are two temporary relations ancestor and query. The method consists in applying r2 to parent, producing a new value for ancestor, then applying r1 to ancestor until no new tuple is generated, then applying r3.

The object program is

```
begin
initialize ancestor to the empty set,
evaluate (ancestor(X,Y) - parent(X,Y)),
insert the result in ancestor,
while "new tuples are generated" do
begin
evaluate (ancestor(X,Y) - parent(X,Z),
ancestor(Z,Y))
using the current value of ancestor,
insert the result in ancestor
end,
evaluate (query(X) - ancestor(aa,X)),
insert the result in query
end
```

The execution of the program against the data goes as follows

Step 1 Apply r1

The resulting state is

```
ancestor = {(a,aa), (a,ab), (aa,aaa), (aa,aab),
            (aaa,aaaa), (c,ca)}
query = {}
```

Step 2 Apply r2

The following new tuples are generated in ancestor {(a,aaa), (a,aab), (aa,aaaa)}

And the resulting state is

```
ancestor = {(a,aa), (a,ab), (aa,aaa), (aa,aab),
            (aaa,aaaa), (c,ca), (a,aaa), (a,aab), (aa,aaaa)}
query = {}
```

New tuples have been generated so we continue

Step 3 Apply r2

The following tuples are generated

```
{(a,aaa), (a,aab), (aa,aaaa), (a,aaaa)}
```

The new state is

```
ancestor = {(a,aa), (a,ab), (aa,aaa), (aa,aab),
            (aaa,aaaa), (c,ca), (a,aaa), (a,aab),
            (aa,aaaa), (a,aaaa)}
query = {}
```

Because (a,aaaa) is new, we continue

Step 4 Apply r2

The following tuples are generated
 $\{(a,aaa), (a,aab), (aa,aaaa), (a,aaaa)\}$
 Because there are no new tuples, the state
 does not change and we move to r3

Step 5 Apply r3

The following tuples are produced in query
 $\{(aa,aaa), (aa,aaaa)\}$ and the new state is
 $\text{ancestor} = \{(a,aa), (a,ab), (aa,aaa), (aa,aab),$
 $(aaa,aaaa), (c,ca), (a,aaa), (a,aab),$
 $(aa,aaaa), (a,aaaa)\}$
 $\text{query} = \{(aa,aaa), (aa,aaaa)\}$

The algorithm terminates

In this example, we note the following problems (i)
 the entire relation is evaluated, i.e., the set of
 potentially relevant facts is the set of facts of the
 base predicates which derive the query, and (ii)
 step 3 completely duplicates step 2

Naive evaluation is the most widely described
 method in the literature. It has been presented in a
 number of papers under different forms. The infer-
 ence engine of SNIP, presented in [Shapiro and
 McKay 80, Shapiro et al 82, McKay and Shapiro
 81], is in fact an interpreted version of naive
 evaluation. The method described in [Chang 81],
 while based on a very interesting language para-
 digm and restricted to linear systems, is a compiled
 version of naive evaluation based on relational
 algebra. The method in [Marque-Pucheu 83,
 Marque-Pucheu et al 84] is a compiled version of
 naive evaluation using a different algebra of rela-
 tions. The method in [Bayer 85] is another descrip-
 tion of naive evaluation. The framework presented
 in [Delobel 86] also uses naive evaluation as its
 inference strategy. SNIP is, to our knowledge, the
 only existing implementation in the general case.

3.2.2. Semi-Naive Evaluation

Semi-naive evaluation is a bottom-up, compiled
 and iterative strategy.

Its application range is the set of bottom-up evalu-
 able rules.

This method uses the same approach as naive
 evaluation, but tries to cut down on the number of
 duplications. It behaves exactly as naive evalua-
 tion, except for the loop mechanism where it tries
 to be smarter.

Let us first try to give an idea of the method as an
 extension of naive evaluation. Let p be a recursive
 predicate, consider a recursive rule having p as a
 head predicate and let us write this rule

$$p = \phi(p_1, p_2, \dots, p_n, q_1, q_2, \dots, q_m)$$

where ϕ is a first order formula, p_1, p_2, \dots, p_n are

mutually recursive to p , and q_1, q_2, \dots, q_m are base or
 derived predicates, which are not mutually recur-
 sive to p .

In the naive evaluation strategy, all the q_i 's are
 fully evaluated when we start computing p and the
 p_i 's. On the other hand p and the p_i 's are all
 evaluated inside the same loop (together with the
 rest of predicates mutually recursive to p).

Let $p_j(i)$ be the value of the predicate p_j at the i -th
 iteration of the loop. At this iteration, we compute

$$\phi(p_1(i), p_2(i), \dots, p_n(i), q_1, q_2, \dots, q_m)$$

During that same iteration each p_j receives a set of
 new tuples. Let us call this new set $dp_j(i)$. Thus the
 value of p_j at the beginning of step $(i+1)$ is $p_j(i) +$
 $dp_j(i)$ (where $+$ denotes union).

At step $(i+1)$ we evaluate

$$\phi((p_1(i)+dp_1(i)), \dots, (p_n(i)+dp_n(i)), q_1, \dots, q_m),$$

which, of course, recomputes the previous expres-
 sion (because ϕ is monotonic).

The ideal however, is to compute only the *new*
 tuples i.e. the expression

$$d\phi(p_1(i), dp_1(i), \dots, p_n(i), dp_n(i), q_1, \dots, q_m) =$$

$$\phi((p_1(i)+dp_1(i)), \dots, (p_n(i)+dp_n(i)), q_1, \dots, q_m)$$

$$- \phi(p_1(i), \dots, p_n(i), q_1, \dots, q_m)$$

The basic principle of the semi-naive method is the
 evaluation of the differential of ϕ instead of the
 entire ϕ at each step. The problem is to come up
 with a first order expression for $d\phi$, which does not
 contain any difference operator. Let us assume
 there is such an expression, and describe the algo-
 rithm. With each recursive predicate p are associ-
 ated four temporary relations p before, p after,
 dp before and dp after. The object program for a
 loop is as follows

```

while "the state changes" do
  begin
    for all mutually recursive predicates p do
      begin
        initialize dp after to the empty set,
        initialize p after to p before,
      end
    for each mutually recursive rule do
      begin
        evaluate  $d\phi(p_1, dp_1, \dots, p_n, dp_n, q_1, \dots, q_n)$ 
        using the current values of
         $p_i$  before for  $p_i$  and of  $dp_i$  before for  $dp_i$ ,
        add the resulting tuples to dp after,
        add the resulting tuples to p after
      end
    end
  end

```

All we have to do now is provide a way to generate

$d\phi$ from ϕ . The problem is not solved in its entirety and only a number of transformations are known. In [Bancilhon 85], some of them are given in terms of relational algebra.

It should be noted however, that for the method to work, the only property we have to guarantee is that

$$\phi(p1+dp1,) - \phi(p1,) \subseteq d\phi(p1,dp1,) \\ \subseteq \phi(p1+dp1,)$$

Clearly, the closer $d\phi(p1,dp1,)$ is to $(\phi(p1+dp1,) - \phi(p1,))$, the better the optimization is. In the worse case, where we use ϕ for $d\phi$, semi-naive evaluation behaves as naive evaluation. Here are some simple examples of rewrite rules.

if $\phi(p,q) = p(X,Y),q(Y,Z)$,
then $d\phi(p,dp,q) = dp(X,Y),q(Y,Z)$

More generally when ϕ is linear in p , the expression for $d\phi$ is obtained by replacing p by dp .

if $\phi(p1,p2) = p1(X,Y),p2(Y,Z)$,
then $d\phi(p,dp) = p1(X,Y),dp2(Y,Z)$
 $+dp1(X,Y),p2(Y,Z)+dp1(X,Y),dp2(Y,Z)$

Note that this is not an exact differential but a reasonable approximation.

The idea of semi-naive evaluation underlies many papers. A complete description of the method based on relational algebra is given in [Bancilhon 85]. The idea is also present in [Bayer 85].

It should also be pointed out that, in the particular case of linear rules, because the differential of $\phi(p)$ is simply $\phi(dp)$, it is sufficient to have an inference engine which only uses the new tuples. Therefore many methods which are restricted to linear rules do indeed use semi-naive evaluation. Note also that when the rules are not linear, applying naive evaluation only to the "new tuples" is an incorrect method (in the sense that it does not produce the whole answer to the query). This can be easily checked on the recursive rule

$ancestor(X,Y) - ancestor(X,Z),ancestor(Z,Y)$

In this case, if we only feed the new tuples at the next stage, the relation which we compute consists of the ancestors whose distance to one another is a power of two.

To our knowledge, outside of the special case of linear rules, the method as a whole has not been implemented.

3.2.3. Iterative Query/Subquery

Iterative Query/Subquery (QSQI) is an interpreted, top-down strategy.

Its application domain is the set of range restricted rules without evaluable predicates.

The method associates a temporary relation with every relation which derives the query, but the computation of the predicates deriving the query is done at run time. QSQI also stores a set of queries which are currently being evaluated. When several queries correspond to the same query form, QSQI stores and executes them as a single object. For instance, if we have the queries $p(a,X)$ and query $p(b,X)$, we can view this as query $p(\{a,b\},X)$. We call such an object a *generalized query*. The state memorized by the algorithm is a couple $\langle Q,R \rangle$, where Q is a set of generalized queries, and R is a set of derived relations, together with their current values.

The iterative interpreter is as follows:

```
Initial state is  $\langle \{query(X)\}, \{\} \rangle$ 
while the state changes do
  for all generalized queries in  $Q$  do
    for all rules whose head matches the
      generalized query do
      begin
        unify rule with the generalized query,
        (i.e. propagate the constants)
        this generates new generalized queries
        for each derived predicate in the body
        by looking up the base relations,
        generate new tuples, (by replacing
        each base predicate on the right by
        its value and every derived predicates
        by its current temporary value)
        add these new tuples to  $R$ ,
        add these new generalized queries to  $Q$ 
      end
```

Let us now run this interpreter against our example logic database.

The initial state is $\langle \{query(X)\}, \{\} \rangle$

Step 1

We try to solve $query(X)$. Only rule $r3$ applies. The unification produces the generalized query $ancestor(\{aa\},X)$. This generates temporary relations for $query$ and $ancestor$ with empty set values. Attempts at generating tuples for this generalized query fail.

The new state vector is

$\langle \{query(X), ancestor(aa,X)\}, \{ancestor=\{\}, query=\{\}\} \rangle$

Step 2

A new generalized query has been generated, so we go on. We try to evaluate each of the generalized

queries $query(X)$ does not give anything new, so we try $ancestor(\{aa\},X)$
 Using rule r2, and unifying, we get $parent(aa,X)$
 This is a base relation, so we can produce a set of tuples
 Thus we generate a value for ancestor which contains all the tuples of $parent(aa,X)$ and the new state vector is

$\langle \{query(X), ancestor(\{aa,X\}),$
 $\{ancestor=\{(aa,aaa),(aa,aab)\}, query=\{\}\} \rangle$

We now solve $ancestor(aa,X)$ using r1 Unification produces the expression

$parent(aa,Z), ancestor(Z,Y)$

We try to generate new tuples from this expansion and the current ancestor value but get no tuples
 We also generate new generalized queries by looking up parent and instantiating Z This produces the new expression

$parent(aa,\{aaa,aab\}), ancestor(\{aaa,aab\},Z)$

This creates two new queries which are added to the generalized query and the new state is

$\langle \{query(X), ancestor(\{aa,aaa,aab\},X),$
 $\{ancestor=\{(aa,aaa),(aa,aab)\}, query=\{\}\} \rangle$

Step 3

New generalized queries and new tuples have been generated so we continue We first solve $query(X)$ using r3 and get the value $\{(aa,aaa), (aa,aab)\}$ for query The resulting new state is

$\langle \{query(X), ancestor(\{aa,aaa,aab\},X),$
 $\{ancestor=\{(aa,aaa),(aa,aab)\},$
 $query=\{(aa,aaa),(aa,aab)\}\} \rangle$

We now try to solve $ancestor(\{aa,aaa,aab\},X)$
 Using r2, we get $parent(\{aa,aaa,aab\},X)$ which is a base relation and generates the following tuples in ancestor $\{(aa,aaa),(aa,aab),(aaa,aaaa)\}$ This produces the new state

$\langle \{query(X), ancestor(\{aa,aaa,aab\},X),$
 $\{ancestor=\{(aa,aaa),(aa,aab),(aaa,aaaa)\},$
 $query=\{(aa,aaa),(aa,aab)\}\} \rangle$

We now solve $ancestor(\{aa,aaa,aab\},X)$ using r1 and we get $p(\{aa,aaa,aab\},Z), ancestor(Z,Y)$ We bind Z by going to the parent relation, and we get

$p(\{aa,aaa,aab\},\{aaa,aab,aaaa\}),$
 $ancestor(\{aaa,aab,aaaa\},Y)$

This generates the new generalized query $ancestor(\{aaa,aab,aaaa\},Y)$ and the new state

$\langle \{query(X), ancestor(\{aa,aaa,aab,aaaa\},X),$
 $\{ancestor=\{(aa,aaa),(aa,aab),(aaa,aaaa)\},$
 $query=\{(aa,aaa),(aa,aab)\}\} \rangle$

Step 4

A new generalized query has been generated, so we continue Solving the ancestor queries using r2 will not produce any new tuples, and solving it with r3 will not produce any new generalized query nor any tuples The algorithm terminates

Concerning the performance of the method, one can note that (i) the set of potentially relevant facts is better than for naive (in this example it is optimal), and (ii) QSQI has the same duplication problem as naive evaluation each step entirely duplicates the previous strategy

Iterative Query/Subquery is presented in [Vielle 85 and 86] To our knowledge it has not been implemented

3.2.4. Recursive Query/Subquery or Extension Tables

Recursive Query/Subquery (QSQR) is a top-down interpreted recursive strategy

The application domain is the set of range restricted rules without evaluable predicates

It is of course a recursive version of the previous strategy As before, we maintain temporary values of derived relations and a set of generalized queries The state memorized by the algorithm is still a couple $\langle Q,R \rangle$, where Q is a set of generalized queries and R is a set of derived relations together with their current values However, besides this explicit state, the recursion mechanism stores at each level in the stack the tuples returned by the evaluation of the query, but this seems to have been solved reasonably in the existing implementation The algorithm uses a selection function which, given a rule, can choose the first and the next derived predicate in the body to be "solved"

The recursive interpreter is as follows

procedure evaluate(q)

(* q is a generalized query *)

begin

while "new tuples are generated" **do**
 for all rules whose head matches the generalized query **do**

begin

unify the rule with the generalized query, (i.e., propagate the constants)

until there are no more derived predicate on the right **do**

begin

choose the first/next derived predicate according to the selection function, generate the corresponding generalized query, (This is done by replacing in the

rule each base predicate by its value
 and each previously solved derived
 predicate by its current value)
 eliminate from that generalized query
 the queries that are already in Q,
 this produces a new generalized query q',
 add q' to Q,
 evaluate(q')
end,

replace each evaluated predicate by its
 value and evaluate the generalized query q,
 (This can be done in some order without
 waiting for all predicates to be evaluated)
 add the results in R,
 return the results
end

end

Initial state is $\langle \{query(X)\}, \{\} \rangle$
 evaluate(query(X))

It is important to note that this version of QSQ is very similar to Prolog. It solves goals in a top-down fashion using recursion, and it considers the literals ordered in the rule (the order is defined by the selection function). The important differences with Prolog are (i) the method is set-at-a-time instead of tuple-at-a-time, through the generalized query concept, and (ii) as pointed out in [Dietrich and Warren 86], the method uses a dynamic programming approach of storing the intermediate results and re-using them when needed. This dynamic programming feature also solves the problem of cycles in the facts while Prolog will run in an infinite loop in the presence of such cycles, QSQR will detect them and stop the computation when no new tuple is generated. Thus, QSQR is complete over its application domain whereas Prolog is not.

Here is the ancestor example

```

evaluate(query(X))
  use rule r3
  query(X) - ancestor(aa,X)
  this generates the query ancestor({aa},X)
  new state is
   $\langle \{ancestor(\{aa\},X), query(X)\}, \{\} \rangle$ 
  evaluate(ancestor({aa},X))
    Step 1 of the iteration
    use rule r1
    ancestor({aa},Y) - parent({aa},Z),
                        ancestor(Z,Y)
    by looking up parent we get the
    bindings {aaa,aab} for Z
    this generates the query
    ancestor({aaa,aab},X)
    new state is
  
```

```

   $\langle \{ancestor(\{aa,aaa,aab\},X), query(X)\}, \{\} \rangle$ 
  evaluate (ancestor({aaa,aab},X))
  (this is a recursive call)
    Step 1 1
    use r1
    ancestor({aaa,aab},Y) -
      parent({aaa,aab},Z),ancestor(Z,Y)
    by looking up parent we get the
    binding {aaaa} for Z
    new state is
     $\langle \{ancestor(\{aa,aaa,aab,aaaa\},X), query(X)\}, \{\} \rangle$ 
    evaluate(ancestor({aaaa},X))
    (this is a recursive call)
      Step 1 1 1
      use r1
      ancestor({aaaa},Y) -
        parent({aaaa},Z),ancestor(Z,Y)
      by looking up parent we get no
      binding for Z
      use r2
      ancestor({aaaa},Y) -
        parent({aaaa},Y)
      this fails to return any tuple
      end of evaluate(
        ancestor({aaaa},X))
      Step 1 1 2
      nothing new is produced
      end of evaluate(
        ancestor({aaaa},Y))
    use r2
    ancestor({aaa,aab},Y) -
      parent({aaa,aab},Y)
    this returns the tuple
    ancestor(aaa,aaaa)
    new state is
     $\langle \{ancestor(\{aa,aaa,aab,aaaa\},X), query(X)\}, \{ancestor=\{(aaa,aaaa)\}\} \rangle$ 
    Step 1 2
    same as Step 1, nothing new produced
    end of evaluate (
      ancestor({aaa,aab},X))
  use rule r2
  ancestor({aa},X) - parent({aa},Y)
  returns the tuples
  ancestor(aa,aaa) and ancestor(aa,aab)
  new state is
   $\langle \{ancestor(\{aa,aaa,aab,aaaa\},X), query(X)\}, \{ancestor=\{(aaa,aaaa), (aa,aaa),(aa,aab)\}\} \rangle$ 
  Step 2
  nothing new produced
  end of evaluate({aa},X)
  generate tuples from r3
  new state is
   $\langle \{ancestor(\{aa,aaa,aab,aaaa\},X),$ 

```

```

query(X)}, {ancestor={{(aaa,aaaa),
(aa,aaa),(aa,aab)}, query=(aaa,aaaa),
(aa,aaa),(aa,aab)}} >
end of evaluate(query(X))

```

Recursive Query/Subquery is described in [Vielle 85 and 86] A compiled version has been implemented on top of the INGRES relational system [Vielle 86] In [Dietrich and Warren 85], along with a good survey of some of these strategies, a method called "extension tables" is presented It is, up to a few details, the same method

3.2.5. Henschen-Naqvi

Henschen-Naqvi is a top-down, compiled and iterative method

The application domain is that of linear range restricted rules

The method has a compilation phase which generates an iterative program That iterative program is then run against the data base The general strategy is fairly complex to understand, and we shall restrict ourselves to describing it in the "typical case" which is

```

p(X,Y) - up(X,XU),p(XU,YU),down(YU,Y)
p(X,Y) - flat(X,Y)
query(X) - p(a,X)

```

Note that the relation names *up* and *down* are not to be confused with the notions "top-down" or "bottom-up", which are characteristics of evaluation strategies Let us introduce some simple notation, which will make reading the algorithm much simpler Since we are only dealing with binary relations, we can view these as set-to-set mappings Thus, the relation *r* associates with each set *A* a set *B*, consisting of all the elements related to *A* by *r* We denote *A r* the image of *A* by *r*, and we have

$$A r = \{ y \mid r(x,y) \text{ and } x \in A \}$$

If we view relations as mappings, we can compose them, and we shall denote *r s* the composition of *r* and *s* Therefore

$$A (r s) = (A r) s$$

This approach is similar to the formalism described in [Gardarin and Mandreville 86] We shall denote the composition of relation *r* *n* times with itself *rⁿ* Finally we shall denote set union by '+' Once this notation is introduced, it is easy to see that the answer to the query is

```

{a} flat + {a} up flat down
+ {a} up up flat down down +
+ {a} upn flat downn +

```

The state memorized by the algorithm is a couple $\langle V, E \rangle$, where *V* is the value of a unary relation and *E* is an expression At each step, using *V* and *E*, we compute some new tuples and compute the new values of *V* and *E*

The iterative program is as follows

```

V = {a},
E = λ, /* the empty string */
while "new tuples are generated in V" do
  begin
    /* produce some answer tuples */
    answer = answer + V flat E,
    /* compute the new value */
    V = V up ,
    /* compute the new expression */
    E = E | down,
  end

```

Note that *E* is an *expression*, and is augmented each time around the loop by concatenating "down" to it through the "cons" operator As can be seen from this program, at step 1, the value *V* represents {a} up' and the expression *E* represents down' Therefore the produced tuples are

{a} up' flat down'

This is not meant to be a complete description of the method, but a description of its behavior in the typical case

The Henschen-Naqvi method is described in [Henschen and Naqvi 84] The method has been implemented in the case described here This implementation can be found in [Laskowski] An equivalent strategy is described using a different formalism in [Gardarin and Mandreville 86] The performance of the strategy is compared to Semi-Naive evaluation and another method (not described here) in [Han and Lu 86]

3.2.6. Prolog

Prolog is a top-down, interpreted and recursive method

The application domain of Prolog is difficult to state precisely (1) it is data dependent in the sense that the facts have to be acyclic for the interpreter to terminate, and (2) there is no simple syntactic characterization of a terminating Prolog program The job of characterizing the "good" rules is left to the programmer

use parent(aa,aab)
 the new rule is
 "parent(aa,aab),ancestor(aab,Y)"
 solve(ancestor(aab,Y),answer2)
 (this call is not described)
 this returns nothing
 process ancestor(Z,Y)
 we instantiate parent and ancestor
 with the useful facts
 this produces ancestor(aa,aaaa)
 we add it to the relevant facts
 relevant =
 {parent(aa,aaa),parent(aa,aab),
 ancestor(aa,aaaa)},

rule r2 "ancestor(X,Y) - parent(X,Y)"
 using the relevant facts we produce
 {ancestor(aa,aaa),ancestor(aa,aab)}
 we add these to relevant
 relevant =
 {parent(aa,aaa),parent(aa,aab),
 ancestor(aa,aaa), ancestor(aa,aab),
 ancestor(aa,aaaa)},
 this rule does not produce any subquery

Step 2

will not produce anything new,
 and so the algorithm stops

The APEX method is described in [Lozinski 85 and 85a] The method has been implemented

3.3. Optimization Strategies

We now turn to the description of the second class of strategies the optimization strategies

The main drawbacks of the naive evaluation method are

- 1 The potential set of relevant facts is too big (In other words, it does not make good use of the query bindings), and
- 2 It generates a lot of duplicate computation

A number of optimization strategies have recently been proposed to overcome those two difficulties

3.3.1. Aho-Ullman

Aho and Ullman [Aho and Ullman 79] present an algorithm for optimizing recursive queries by computing selections with the least fixpoint operator (LFP) The input is an expression

$$\sigma_F(\text{LFP}(r=f(r)))$$

where $f(r)$ is a monotonic relational algebra expression (under the ordering of set inclusion) and con-

tains at most one occurrence of r The output is an equivalent expression where the selection has been pushed through as far as possible

We introduce their notation and ideas through an example Consider

$$\begin{aligned}
 a(X,Y) &- a(X,Z), p(Z,Y) \\
 a(X,Y) &- p(X,Y) \\
 q(X) &- a(\text{john},X)
 \end{aligned}$$

Aho-Ullman write this as

$$\sigma_{a_1=\text{john}}(\text{LFP}(a = a p \cup p))$$

In this definition, a is a relation which is defined by a *fixpoint* equation in relational algebra, and p is a base relation If we start with a empty and repeatedly compute a using the rule $a = a p \cup p$, at some iteration, there is no change (since the relation p is finite) Because the function used in the fixpoint equation is monotonic, this is the *least fixpoint* of the fixpoint equation [Tarski 55] It is the smallest relation a which satisfies the equation, i.e. contains every tuple which can be generated by using the fixpoint rule, and no tuple which cannot The query is simply the selection $a_1=\text{john}$ applied to this relation Thus, the query is a selection applied to the transitive closure of p

We now describe how the Aho-Ullman algorithm optimizes this query We use $'$ to denote composition, which is a join followed by projecting out the join attributes We begin with the expression

$$\sigma_{a_1=\text{john}}(a)$$

and by replacing a by $f(a)$ we generate

$$\sigma_{a_1=\text{john}}(a p \cup p)$$

By distributing the selection across the join, we get

$$\sigma_{a_1=\text{john}}(a p) \cup \sigma_{a_1=\text{john}}(p)$$

Since the selection in the first subexpression only involves the first attribute of a , we can rewrite it as

$$\sigma_{a_1=\text{john}}(a) p$$

We observe that this contains the subexpression

$$\sigma_{a_1=\text{john}}(a)$$

which was the first expression in the series If we denote this by E , the desired optimized expression is then

$$\text{LFP}(E = E p \cup \sigma_{a_1=\text{john}}(p))$$

This is equivalent to the Horn Clause query

$$a(\text{john},Y) - a(\text{john},Z), p(Z,Y)$$

$a(\text{john}, Y) - p(\text{john}, Y)$
 $q(X) - a(\text{john}, X)$

The essence of the strategy is to construct a series of equivalent expressions starting with the expression $\sigma_F(r)$ and repeatedly replacing the single occurrence of r by the expression $f(r)$. Note that each of these expressions contains just one occurrence of R . In each of these expressions, we push the selection as far inside as possible. Selection distributes across union, commutes with another selection and can be pushed ahead of a projection. However, it distributes across a Cartesian product $Y \times Z$ only if the selection applies to components from just one of the two arguments Y and Z . The algorithm fails to commute the selection with the LFP operator if the (single) occurrence of r is in one of the arguments of a Cartesian product across which we cannot distribute the selection. We stop when this happens or when we find an expression of the form $h(g(\sigma_F(r)))$ and one of the previous expressions in the series is of the form $h(\sigma_F(r))$. In the latter case, the equivalent expression that we are looking for is $h(\text{LFP}(s=g(s)))$, and we have succeeded in pushing the selection ahead of the LFP operator.

We note in conclusion that the expression $f(r)$ must contain no more than one occurrence of r . For instance, the algorithm does not apply in this case

$$\sigma_{a_1=\text{john}}(\text{LFP}(a = a \cup p))$$

Aho and Ullman also present a similar strategy for commuting projections with the LFP operator, but we do not discuss it here.

3.3.2. Kifer-Lozinskii

The Kifer-Lozinskii algorithm is an extension of the Aho-Ullman algorithm described above. However, rules are represented as rule/goal graphs rather than as relational algebra expressions, and the strategy is described in terms of *filters* which are applied to the arcs of the graph. It is convenient to think of the data as flowing through the graph along the arcs. A *filter* on an arc is a selection which can be applied to the tuples flowing through that arc, and is used to reduce the number of tuples that are generated. Transforming a given rule/goal graph into an equivalent graph with (additional) filters on some arcs is equivalent to rewriting the corresponding set of rules.

The execution of a query starts with the nodes corresponding to the base relations sending all their tuples through all arcs that leave them. Each axiom node that receives tuples generates tuples for its head predicate and passes them on through

all its outgoing arcs. A relation node saves all new tuples that it receives and passes them on through its outgoing arcs. Computation stops (with the answer being the set of tuples in the query node) when there is no more change in the tuples stored at the various nodes at some iteration. We note that this is simply Semi-Naive evaluation.

Given filters on all the arcs leaving a node, we can 'push' them through the node as follows. If the node is a relation node, we simply place the disjunction of the filters on each incoming arc. If the node is an axiom node, we place on each incoming arc the strongest consequence of the disjunction that can be expressed purely in terms of the variables of the literal corresponding to this arc.


The objective of the optimization algorithm is to place the "strongest" possible filters on each arc. Starting with the filter which represents the constant in the query, it repeatedly pushes filters through the nodes at which the corresponding arcs are incident. Since the number of possible filters is finite, this algorithm terminates. It stops when further pushing of filters does not change the graph, and the graph at this point is equivalent to the original graph (although the graph at intermediate steps may not). Note that since the disjunction of 'true' with any predicate is 'true', if any arc in a loop is assigned the filter 'true', all arcs in the loop are subsequently assigned the filter 'true'.

Consider the transitive closure example that we optimized using the Aho-Ullman algorithm. We would represent it by the following axioms


r1 $a(X, Y) - a(X, Z), p(Z, Y)$
r2 $a(X, Y) - p(X, Y)$
r3 $q(X) - a(\text{john}, X)$

Given below is the corresponding system graph, before and after optimization. (We have omitted the variables in the axioms for clarity.)

Before



After



We begin the optimization by pushing the selection through the relation node a . Thus the arcs from $r1$ to a and from $r2$ to a both get the filter ' $1=John$ ' (We have simplified the conventions for keeping track of variables - ' 1 ' refers to the first attribute of the corresponding head predicate) We then push these filters through the corresponding axiom nodes, $r1$ and $r2$. Pushing ' $1=John$ ' through node $r2$ puts the filter ' $p_1=John$ ' on the arc from p to $r2$. Pushing ' $1=John$ ' through node $r1$ puts the filter ' $a_1=John$ ' on the arc from a to $r1$. Note that it does not put anything on the arc from p to $r1$ (empty filters are equivalent to 'true'). There are no arcs entering p , and the filter on the arc from a to $r1$ does not change the disjunction of the filters on arcs leaving a (which is still ' $a_1=John$ ') So the algorithm terminates here.

The analogy with the Aho-Ullman algorithm is easily seen when we recognize that a filter is a

selection, pushing through a relation node is distribution across a \cup and pushing through an axiom node is distribution across a Cartesian product. In general, the optimizations achieved by the two algorithms are identical. However, the Kifer-Lozinski algorithm is more general in that it successfully optimizes some expressions containing more than one occurrence of the defined predicate. An example is the expression

$$\sigma_{a_1=John}(LFP(a = (a \ p \cup a \ q \cup p)))$$

The Aho-Ullman algorithm does not apply in this case because there are two occurrences of R in $f(R)$. The Kifer-Lozinski algorithm optimizes this to

$$LFP((\sigma_{a_1=John}(a \ p) \cup (\sigma_{a_1=John}(a \ q) \cup (\sigma_{a_1=John}(p))))$$

Essentially, it improves upon the Aho-Ullman algorithm in that it is able to distribute selection across some unions where both arguments contain r .

Further, the algorithm can work directly upon certain mutually recursive rules, for example

$$\begin{aligned} r1 & \ r(X,Y) - b(X), s(X,Y) \\ r2 & \ s(X,Y) - c(X), r(X,Y) \\ r3 & \ q(X) - r(X,John) \end{aligned}$$

Before applying the Aho-Ullman algorithm, these rules must be rewritten as follows

$$\begin{aligned} r1 & \ r(X,Y) - b(X), c(X), r(X,Y) \\ r2 & \ q(X) - r(X,John) \end{aligned}$$

Note that the Kifer-Lozinski algorithm fails to optimize both

$$\begin{aligned} \sigma_{a_1=John}(LFP(a = a \ a \cup p)), \text{ and} \\ \sigma_{a_1=John}(LFP(a = a \ p \cup p \ a \cup p)) \end{aligned}$$

3.3.3. Magic Sets

The idea of the Magic Sets optimization is to simulate the sideways passing of bindings a la Prolog by the introduction of new rules. This cuts down on the number of potentially relevant facts.

The application domain is the set of bottom-up evaluable rules.

We shall describe the strategy in detail, using as an example a modified version of the same-generation rule set

$$\begin{aligned} sg(X,Y) - p(X,XP),p(Y,YP),sg(YP,XP) \\ sg(X,X) \\ query(X) - sg(a,X) \end{aligned}$$

Note that in this version the two variables XP and

YP have been permuted Note also that the second rule is not range restricted The first step of the magic set transformation is the introduction of adornments and the generation of adorned rules

Given a system of rules, the *adorned rule system* [Ullman 85] is obtained as follows

For each rule r and for each adornment a of the predicate on the left, generate an adorned rule Define recursively an argument of a predicate in the rule r to be *distinguished* [Henschen and Naqvi 84] if either it is bound in the adornment a , or it is a constant, or it appears in a base predicate occurrence that has a distinguished variable Thus, the sources of bindings are (i) the constants and (ii) the bindings in the head of the rule These bindings are propagated through the base predicates If we consider each distinguished argument to be bound, this defines an adornment for each derived literal on the right The adorned rule is obtained by replacing each derived literal by its adorned version

If we consider the rule

$$sg(X, Y) - p(X, XP), p(Y, YP), sg(YP, XP)$$

with adornment bf on the head predicate, then X is distinguished because bound in $sg(X, Y)$, XP is distinguished because X is distinguished and $p(X, XP)$ is a base predicate and these are the only distinguished variables Thus the new adorned rule is

$$sg^{bf}(X, Y) - p(X, XP), p(Y, YP), sg^{fb}(YP, XP)$$

If we consider a set of rules, this process generates a set of adorned rules The set of adorned rules has size $K R$ where R is the size of the original set of rules and K is a factor exponential in the number of attributes per derived predicate So, for instance, if every predicate has three attributes, then the adorned system is eight times larger than the original system However, we do not need the entire adorned system and we only keep the adorned rules which derive the query In our example the reachable adorned system is

$$\begin{aligned} &sg^{bf}(X, Y) - p(X, XP), p(Y, YP), sg^{fb}(YP, XP) \\ &sg^{fb}(X, Y) - p(X, XP), p(Y, YP), sg^{bf}(YP, XP) \\ &sg^{bf}(X, X) \\ &sg^{fb}(X, X) \\ &query^f(X) - sg^{bf}(a, X) \end{aligned}$$

Clearly, this new set of rules is equivalent to the original set in the sense that it will generate the same answer to the query

The magic set optimization consists in generating from the given set of rules a new set of rules, which are equivalent to the original set with respect to the query, and such that their bottom-

up evaluation is more efficient This transformation is done as follows (i) for each occurrence of a derived predicate on the right of an adorned rule, we generate a magic rule (ii) For each adorned rule we generate a modified rule

Here is how we generate the magic rule (i) choose an adorned literal predicate p on the right of the adorned rule r , (ii) erase all the other derived literals on the right, (iii) in the derived predicate occurrence replace the name of the predicate by $magic\ p\ 1^a$ where a is the literal adornment, and erase the non distinguished variables, (iv) erase all the non distinguished base predicates, (v) in the left hand side, erase all the non distinguished variables and replace the name of the predicate by $magic\ p\ 1^a$, where p_1 is the predicate on the left, and a' is the adornment of the predicate p_1 , and finally (vi) exchange the two magic predicates

For instance the adorned rule

$$sg^{bf}(X, Y) - p(X, XP), p(Y, YP), sg^{fb}(YP, XP)$$

generates the magic rule

$$magic^{fb}(XP) - p(X, XP), magic^{bf}(X)$$

Note that the magic rules simulate the passing of bound arguments through backward chaining (We have dropped the suffix "sg" in naming the magic predicates since it is clear from the context)

Here is how we generate the modified rule For each rule whose head is $p\ a(X)$ where X is the list of distinguished variables in that occurrence of p For instance the adorned rule

$$sg^{bf}(X, Y) - p(X, XP), p(Y, YP), sg^{fb}(YP, XP)$$

generates the modified rule

$$sg^{bf}(X, Y) - p(X, XP), p(Y, YP), magic^{bf}(X), sg^{fb}(YP, XP)$$

Finally the complete modified set of rules for our example is

$$\begin{aligned} &magic^{fb}(XP) - p(X, XP), magic^{bf}(X) \\ &magic^{bf}(YP) - p(Y, YP), magic^{fb}(Y) \\ &magic^{bf}(a) \\ &sg^{bf}(X, Y) - \\ &\quad p(X, XP), p(Y, YP), magic^{bf}(X), sg^{fb}(YP, XP) \\ &sg^{fb}(X, Y) - \\ &\quad p(X, XP), p(Y, YP), magic^{fb}(Y), sg^{bf}(YP, XP) \\ &sg^{bf}(X, X) - magic^{bf}(X) \\ &sg^{fb}(X, X) - magic^{fb}(X) \\ &query^f(X) - sg^{bf}(a, X) \end{aligned}$$

The idea of the magic set strategy was presented in [Bancilhon et al 86] and the precise algorithm is described in [Bancilhon et al 86a] The "Alexan-

dre" strategy described in [Rohmer and Lescoeur 85] appears to be based on the same idea. To our knowledge, the strategy is not implemented.

3.3.4. Counting and Reverse Counting.

Counting and Reverse Counting are derived from the magic set optimization strategy.

They apply under two conditions: (i) the data is acyclic and (ii) there is at most one recursive rule for each predicate, and it is linear.

We first describe counting using the "typical" single linear rule system.

```
p(X,Y) - flat(X,Y)
p(X,Y) - up(X,XU),p(XU,YU),down(YU,Y)
query(Y) - p(a,Y)
```

The idea consists in introducing magic sets (called *counting* sets) in which elements are numbered by their distance to the element *a*. Remember that the magic set essentially marks all the *up* ancestors of *a* and then applies the rules in a bottom-up fashion to only the marked ancestors. In the counting strategy, at the same time we mark the ancestors of *john*, we number them by their distance from *a*. Then we can "augment" the *p* predicate by numbering its tuples and generate them by levels as follows:

```
counting(a,0)
counting(X,I) -counting(Y,J),up(Y,X),I=J+1
p'(X,Y,I) - counting(X,I),flat(X,Y)
p'(X,Y,I) - counting(X,I),up(X,XU),
p'(XU,YU,J),down(YU,Y),I=J-1
query(X) - p'(a,X,0)
```

Thus at each step, instead of using the entire magic set, we only use the tuples of the correct level, thus minimizing the set of relevant tuples. But in fact, it is useless to compute the first attribute of the *p* predicate. Thus the system can be further optimized into:

```
counting(a,0)
counting(X,I) -counting(Y,J),up(Y,X),I=J+1
p''(Y,I) - counting(X,I),flat(X,Y)
p''(Y,I) - p''(YU,J),down(YU,Y),I=J-1,J>0
query(X) - p''(Y,0)
```

It is interesting to notice that this new set of rules is in fact simulating a stack.

Reverse counting is another variation around the same idea. It works as follows: (i) first compute the magic set, then (ii) for each element *b* in the magic set number all its *down* descendants and its *up* descendants and add to the answer all the *down* descendants having same number as *a* (because *a* is in the *up* descendants). This gives the following

equivalent system

```
magic(a)
magic(Y) - magic(X),up(X,Y)
des up(X,X,0) - magic(X)
des down(X',Y,0) - magic(X'),flat(X',Y)
des up(X',X,I) - des up(X',Y,J),
up(X,Y),I=J+1
des down(X',X,I) - des down(X',Y,J),
down(Y,X),I=J+1
query(Y) - des up(X',a,Y),des down(X',Y,I)
```

This can be slightly optimized by limiting ourselves to the *b*'s which will join with *flat* and restricting the *down* des's to be in the magic set. This generates the following system:

```
magic(a)
magic(Y) - magic(X),up(X,Y)
des up(X,X,0) - magic(X),flat(X,Y)
des down(X',Y,0) - magic(X'),flat(X',Y)
des up(X',X,I) - magic(X),des up(X',Y,J),
up(X,Y),I=J+1
des down(X',X,I) - des down(X',Y,J),
down(Y,X),I=J+1
sg(a,Y) - des up(X',a,Y),des down(X',Y,I)
```

Note that we still have the problem of a "late termination" on *down* because we number *all* the descendants in *down*, even those of a lower generation than *a*.

The idea of counting was presented in [Bancilhon et al 86] and a formal description of counting and of an extension called "magic counting" was presented in the single rule case in [Sacca and Zaniolo 86]. An extension to the fully general case of Horn Clauses with function symbols is described in [Sacca and Zaniolo 86a]. We did not cover this extension here. Reverse counting is described in [Bancilhon et al 86]. They have not been implemented.

3.4. Summary of Strategy Characteristics.

A summary of the characteristics of each strategy is presented in Table 1.

Table 1: Summary of Strategy Characteristics

<i>Method</i>	<i>Application Range</i>	<i>Top down vs Bottom Up</i>	<i>Compiled vs Interpreted</i>	<i>Iterative vs Recursive</i>
Naive Evaluation	Bottom-up Evaluable	Bottom Up	Compiled	Iterative
Semi-Naive Evaluation	Bottom-up Evaluable	Bottom Up	Compiled	Iterative
Query/Subquery	Range Restricted No Arithmetic	Top Down	Interpreted	Iterative
Query/Subquery	Range Restricted No Arithmetic	Top Down	Interpreted	Recursive
APEX	Range Restricted No Arithmetic Constant Free	Mixed	Mixed	Recursive
Prolog	User responsible	Top Down	Interpreted	Recursive
Henschen-Naqvi	Linear	Top Down	Compiled	Iterative
Aho-Ullman	Strongly Linear	Bottom Up	Compiled	Iterative
Kifer-Lozinski	Range Restricted No Arithmetic	Bottom Up	Compiled	Iterative
Counting	Strongly Linear	Bottom Up	Compiled	Iterative
Magic Sets	Bottom-up evaluable	Bottom Up	Compiled	Iterative

4. Performance Comparisons

In this section, we present the results of a comparative performance evaluation of the various strategies. To perform such a comparison we must

- (1) Choose a set of rules and queries which will represent our benchmark.
- (2) Choose some test data which will represent our extensional database.
- (3) Choose a cost function to measure the performance of each strategy.
- (4) Evaluate the performance of each query against the extensional databases.

We first describe the four queries used as “typical” intensional databases. Then, we present our characterization of the data. Each relation is characterized by four parameters and it is argued that a number of familiar data structures, e.g. trees, can be described in this framework. We describe our cost metric, which is the size of the intermediate results before duplicate elimination. We present analytical cost functions for each query evaluation strategy on each query. The cost functions are plotted for three sets of data - tree, inverted tree

and cylinder. We discuss these results informally.

The performance issue was addressed informally through the discussion of a set of examples in [Bancilhon et al 86a]. Han and Lu [Han and Lu 86] have reported a study of the performance of a set of four evaluation strategies (including Naive and Henschen-Naqvi and two others not considered here) on the same generation example, using randomly generated data. Their model is based on the selectivity of the join and select operations and the sizes of the data relations. They consider both CPU and IO cost. We have chosen to concentrate on one aspect of the problem, which is the number of successful firings (measured using the sizes of the intermediate relations) and have studied a wider range of strategies, queries and data.

4.1. Workload: Sample Intensional Databases and Queries

Instead of generating a general mix, we have chosen four queries which have the properties of exercising various important features of the strategies. We are fully aware of the fact that this set

is insufficient to provide a complete benchmark, but we view this work as a first step towards a better understanding of the performance behavior of the various strategies

The queries are three different versions of the ancestor query and a version of the same-generation query. The first one is just a classical ancestor rule and query with the first attribute bound

Query 1 $a(X,Y) - p(X,Y)$
 $a(X,Y) - p(X,Z),a(Z,Y)$
 $query(X) - a(john,X)$

Because most strategies are representation dependent, we have studied the same example with the second attribute bound instead of the first. This will allow us to determine which strategies can solve both cases

Query 2 $a(X,Y) - p(X,Y)$
 $a(X,Y) - p(X,Z),a(Z,Y)$
 $query(X) - a(X,john)$

The third version of the ancestor example specifies ancestor using recursive doubling. This enables us to see how the strategies react to the non linear case. This example being fully symmetric, it is sufficient to test it with its first attribute bound

Query 3 $a(X,Y) - p(X,Y)$
 $a(X,Y) - a(X,Z),a(Z,Y)$
 $query(X) - a(john,X)$

Finally to study something more complex than transitive closure, we have chosen a generalized version of the same generation example, bound on its first attribute

Query 4 $p(X,Y) - flat(X,Y)$
 $p(X,Y) - up(X,XU),p(XU,YU),$
 $down(YU,Y)$
 $query(X) - p(john,X)$

4.2. Characterizing Data: Sample Extensional Databases

Because we decided on an analytical approach, we had to obtain tractable formulae for the cost of each strategy against each query. Therefore, each relation must be characterized by a *small* set of parameters. Fortunately, because of the choice of our workload, we can restrict our attention to binary relations

We represent every binary relation by a directed graph and view tuples as arcs and domain elements as nodes. Nodes are arranged in layers and each arc goes from a node in one layer to a node in the next. Note that in these graphs each node has at

least one in-arc or one out-arc. Nodes in the first layer have no incoming arcs and nodes in the last layer have no outgoing arcs

Let R be a binary relation and A be a set. Recall that we denote by $A.R$ the set

$$A.R = \{y \mid x \in A \text{ and } R(x,y)\}$$

We characterize a binary relation R by

- (1) F_R the *fan-out factor*,
- (2) D_R the *duplication factor*,
- (3) h_R the *height*, and
- (4) b_R the *base*


F_R and D_R are defined as follows given a "random" set A of n nodes from R , the size of $A.R$ is $n.F_R$ (We use $'$ here to denote multiplication. It should be clear from the context whether $'$ denotes multiplication or composition) before duplicate elimination. D_R is the duplication factor in $A.R$, i.e. the ratio of the size of $A.R$ before and after duplicate elimination. Thus the size of $A.R$ after duplicate elimination is $n.F_R/D_R$

We call $E_R = F_R/D_R$ the *expansion factor* of R

The base b_R is the number of nodes which do not have any antecedents. The height h_R is the length of the longest chain in R

When no confusion is possible, we shall simply use F, D, h and b instead of F_R, D_R, h_R and b_R

The typical structure consists of a number of layers. There are $(h_R + 1)$ layers of nodes in the structure, numbered from top to bottom (as 0 to h). There are b_R nodes in level 0



This "parametrized structure" is fairly general and can represent a number of typical configurations. A binary balanced tree of height 1 is defined by

$$F=2, D=1, h=1, b=1$$

The same binary tree upside down is defined by

$$F=1/2, D=2, h=1, b=2^1$$

A list of length 1 is defined by

$$F=1, D=1, h=1, b=1$$

A set of n lists of length l is defined by

$$F=1, D=1, h=1, b=n$$

A parent relation, where each person has two children and each child has two parents, is defined by

$$F=2, D=2, h=\text{number of generations}, \\ b=\text{number of people of unknown parentage}$$

However, this formalism does not represent cycles. Nor does it represent short cuts, were a short cut is the existence of two paths of different length going from one point to another. Clearly, they would violate our assumption that nodes were arranged in layers with arcs going from nodes in one layer to the next. We also emphasize that we assume the data to be *random*, with a uniform distribution. Thus, the values F and D are average values.

Our assumption that the duplication factor is independent of the size is a very crude approximation. For instance it implies that if you start from one node you still generate some duplicates. Obviously the duplication factor increases with the size of the start set. Therefore, our approximation overestimates the number of duplicates. However, it becomes reasonable as the size of the start set becomes large. It is also dependent upon our assumption that the data is *random* (with a uniform distribution) and not *regular*.

Let us now turn to the problem of characterizing inter-relation relationships. Let A and B be two sets. The *transfer ratio* of A with respect to B , denoted $T_{A,B}$ is the number such that given a random set of n nodes in A , the size of $A \cap B$ after duplicate elimination is $n T_{A,B}$. Note that $0 \leq T \leq 1$.

This definition can be extended to binary relations by considering only the columns of the relations. We shall denote the i -th column of R by R_i . Thus, given two binary relations R and S , the number of tuples in the (ternary) result of the join of R and S is $n T_{R_2, S_1}$, where n is the number of tuples in R .

4.3. The Cost Metrics

We have chosen for our cost measure the number of successful inferences performed by the system.

The simplest way to obtain this cost function is to measure the size of the intermediate results *before* duplicate elimination.

Note that in this model the measure of complexity of the join, the cartesian product, intersection and selection is the size of the result, the measure of complexity of union is the sum of the sizes of the arguments (each tuple present in both argument is going to fire twice), and the measure of complexity of projection is the size of the argument. Readers



familiar with performance evaluation of relational queries might be surprised by these measures. However, it is argued in [Bancilhon 85] that they are meaningful. In essence, our cost is a measure of one important factor in the performance of a query evaluation system, the number of successful firings, rather than a measure of the actual run-time performance.

4.4. Cost Evaluation

For each strategy and for each query, we have analytically evaluated the cost of computing the given query using the given strategy. The cost is expressed as a function of the data parameters F , D , h and b . The formulae are listed in Appendix 1, and their derivations are contained in [Bancilhon and Ramakrishnan 86]. To compare these fairly complex formulae, we have plotted a number of curves, some of which are included in Appendix 2.

4.5. Graphical Comparison of the Costs

The curves shown in Appendix 2 show the relative performance of the various strategies on each of the sample queries for three sets of data. They are relations in which the tuples are arranged in a tree structure, an inverted tree structure, and a "cylinder". A cylinder is a structure in which each layer has b nodes and each node has on the average two incoming and two outgoing arcs. We present below a sample relation of each type.



The choice of these structures was made in order to study the effects of uneven distribution of the data and the effects of duplication. We have fixed the sizes of all relations at 100,000 tuples. For the tree structure, we vary the shape by changing the fan-out F while keeping the number of arcs (which is the number of tuples) constant. Clearly, decreasing the fan-out increases the depth of the structure and vice-versa. Similarly, the shape of the inverted tree is varied by varying the duplication factor. The shape of the cylinder is varied by varying the ratio of breadth b to height h , again keeping the number of arcs constant.

For each query and data structure, we plot the cost of each strategy against the shape of the data (measured in terms of the parameter used to vary it). Thus, for each query, we plot cost vs F for the tree, cost vs D for the inverted tree, and cost vs b/h for the cylinder. We do this for each strategy. The cost is computed using the cost functions listed in the appendix. We have often displayed a subset of the curves (for the same query and data structure) over a different range, to allow a better comparison.

For the ancestor queries, we plot the cost of each strategy for the cases when the parent relation has 100,000 tuples and the data in it has the shape of a tree, an inverted tree and a cylinder.

For the same generation example, we have assumed that the relations *up* and *down* are identical and that the fan-out and duplication for the relation *flat* are both equal to 1. We have also assumed that the transfer ratio from up to flat is equal to the transfer ratio from flat to down. We have assumed that all three relations (*up*, *flat* and *down*) have 100,000 tuples. We plot the cost of each strategy as the shape of *up* and *down* varies for a total of six cases: the cases when the structure is a tree, an inverted tree and a cylinder, with the transfer ratio equal to 1 and 0.01 (100% and 1% respectively).

4.6. Summary of the Curves

There are several important points to be seen in the curves. For a given query, there is a clear ordering of the various strategies which usually holds over the entire range of data. The difference in performance between strategies is by orders of magnitude, which emphasizes the importance of choosing the right strategy. The cost of the optimal strategy is less than 10,000 in each of the queries we have considered, over the entire range of data. The size of the data is 100,000 tuples. This indicates that recursive queries can be implemented efficiently.

We present a summary of the ordering of the strategies, as seen in the corresponding curves. We use \ll to denote an order of magnitude or greater difference in performance, and for a given query, we list in parentheses those strategies that perform identically for all data. We refer to the various strategies using the following acronyms for brevity: HN (Henschen-Naqvı), C (Counting), MS (Magic Sets), QSQR, QSQI, APEX, P (Prolog), SN (Semi-Naive), N (naive) and KL (Kifer-Lozinskiı).

Query 1 (Ancestor bf)

Tree
 $(HN, C) \ll (MS, QSQR, APEX) = P \ll QSQI \ll (SN, KL) \ll N$

Inverted tree
 $(HN, C) \ll (MS, QSQR, APEX) \ll P \ll QSQI \ll (SN, KL) \ll N$

Cylinder
 $(HN, C) \ll (MS, QSQR, APEX) \ll QSQI \ll (SN, KL) \ll N \ll P$

Query 2 (Ancestor fb)

All data
 $(HN, C) \ll (MS, QSQR, KL) \ll QSQI \ll APEX \ll SN \ll N \approx P$

Query 3 (Ancestor bf, non-linear)

All data
 $QSQR \ll QSQI \ll APEX \ll (SN, MS, KL) \ll N$
 (HN, Counting and Prolog do not apply)

Query 4 (Same Generation bf)

Tree
 $C \ll HN \approx (MS, QSQR, APEX) = P \ll QSQI \ll (SN, KL) \ll N$

Inverted tree
 $C \ll HN \approx (MS, QSQR, APEX) \ll P \ll QSQI \ll (SN, KL) \ll N$

Cylinder
 $C \ll HN \approx (MS, QSQR, APEX) \ll QSQI \ll (SN, KL) \ll N \ll P$

To summarize the ancestor results, the following order is seen to hold for the ancestor queries:

$(HN, C) \ll (MS, QSQR) \ll QSQI \ll APEX \ll SN \ll N$

There are some exceptions and additions to the above ordering. In the non-linear case, Henschen-Naqvı and Counting do not apply, and Magic Sets reduces to Semi-Naive. Kifer-Lozinskiı performs

like Semi-Naive, except in the case where the second argument is bound, and in this case it performs like QSQR. APEX performs like QSQR in the case where the first argument is bound. Prolog performs poorly when it cannot propagate the constant in the query (the case where the second argument is bound), as expected. When it can propagate the constant, its performance degrades sharply with duplication, especially as the depth of the data structure increases. This is readily seen from the curves for the cylinder.

To summarize the same generation results, we have

$$C \ll HN \approx (MS, QSQR, APEX) \ll QSQR \ll (SN, KL) \ll (P, N)$$

Prolog behaves like QSQR when there is no duplication (tree). With duplication, its performance degrades so sharply with an increase in the depth of the data structure that we have classified it with Naive, although it performs better than Semi-Naive over a wide range.

4.7. Interpreting the Results

These results indicate that the following three factors greatly influence the performance

- 1 The amount of *duplication* of work,
- 2 The size of the set of *relevant facts*, and
- 3 The use of *unary* vs *binary* intermediate relations

By duplication of work, we refer to the repeated firing of a rule on the same data. This can occur due to duplication in the data (e.g. Prolog), or due to an iterative control strategy that does not remember previous firings (e.g. QSQR and Naive). Relevant facts have been defined earlier, and their significance in reducing the number of useless firings has been explained. The third factor is hard to define precisely. Strategies which only look at sets of nodes rather than sets of arcs perform better than those that look at sets of arcs, by an order of magnitude or more. They are less generally applicable since this often involves a loss of information. This usually leads to non-termination unless the database has certain properties, such as linearity of rules and acyclicity of the extensional database. The following discussion is intended to clarify these concepts, as well as to explain the performance of the various strategies in terms of these three factors.

4.7.1. The Ancestor Queries

We begin by looking at the ancestor queries. The effect of duplication is seen by considering Prolog and QSQR, both of which do duplicate work, for different reasons. When the first argument is bound, Prolog performs like QSQR on a tree data structure, where exactly one arc enters each node (equivalently, there is exactly one way of deriving a given answer). With duplication (i.e. on the average more than one arc enters a given node) performance degrades dramatically. Prolog's performance for the same query on a cylinder is comparable to Naive evaluation, a difference of several orders of magnitude! We note that the set of relevant facts is comparable in the two cases, being the set of nodes reachable from the node denoting the constant in the query (which will henceforth be referred to as the query node). However, in the case of the cylinder, these nodes can be reached along several paths and Prolog infers them afresh along each path. QSQR performs duplicate computation for a different reason, which is that its iterative control strategy does not remember previous firings. Essentially, there are as many steps (executions of the control loop) as the longest path from the query node, and all nodes reached by a path of length less than or equal to i are recomputed at all steps after the i th. This can be seen by comparing QSQR and QSQR and noting that QSQR is orders of magnitude worse in all cases. QSQR uses the same set of relevant facts (the reachable nodes) and differs only in that it has a recursive control strategy that avoids precisely this duplication. Naive evaluation also does a lot of duplicate work, for the same reason as QSQR, i.e. it does not remember previous firings. Semi-Naive differs from Naive only in that it remembers all previous firings and does not repeat them. Thus, the effect of duplication can also be seen in the difference between Naive and Semi-Naive.

The effect of a smaller set of relevant facts can be seen in the vast difference between Magic Sets and Semi-Naive. Magic Sets is simply Semi-Naive applied to the set of relevant facts, which is determined to be the set of reachable nodes except in the doubly recursive case. In this case, the first phase of the Magic Sets strategy, which computes the set of relevant facts, fails and the Magic Sets strategy degenerates to Semi-Naive. This effect can also be seen in the behavior of Prolog on a tree data structure (which means we eliminate the effect of duplication) when the first argument is free. Prolog's depth first strategy is unable to propagate the constant in the second argument of the query. In other words, it must consider all facts in the

database, and its performance degrades by several orders of magnitude. Similarly, the Kifer-Lozinski strategy degenerates to Semi-Naive when the optimization algorithm fails to push down the constant in the query. We note that pushing the constant (equivalently, the selection that it represents) is equivalent to cutting down on the number of relevant facts.

QSQR succeeds in restricting the set of relevant facts to the set of nodes reachable from the query node even in the non-linear version of ancestor. QSQR also succeeds in doing this, but performs a lot of duplicate computation. The Magic Sets algorithm uses the entire parent relation for the set of relevant facts and so degenerates to Semi-Naive. APEX, for reasons explained below, also uses a much larger set of relevant facts. So, although it improves upon Semi-Naive computation in this case, it is much worse than QSQR. Henschen-Naqv1 and Counting do not apply and Prolog does not terminate. Thus QSQR is the only strategy that succeeds in both restricting the set of relevant facts and avoiding duplicate work. It does this at the cost of implementing the recursive control, which is a cost that we do not understand at this stage.

The behavior of APEX illustrates the interesting distinction between the set of relevant facts and the set of *useful* facts. The first step in the APEX strategy is to find what APEX calls the set of relevant facts (which is actually a subset of the set of relevant facts as we have defined it, since it does not include all facts that could derive an answer). In the ancestor examples, these are facts from the relation *parent*, and the firing of the first rule adds them to the ancestor relation. Subsequently, these facts are substituted (in turn) into both the parent and ancestor predicates in the body of the second rule. Except in the first case, this leads to subqueries whose answers are not relevant. For example, in the case where the second argument is bound to *john*, the set of relevant (a la APEX) facts is the set of facts $p(X, \text{john})$. By substituting these into the parent predicate in the second rule, we generate the query $a(\text{john}, ?)$. This computes the ancestors of *john*, whereas the given query $a(?, \text{john})$ asks for the descendants of *john*. This is because APEX does not make the distinction that facts of the form $p(X, \text{john})$ are relevant to the query $a(?, \text{john})$ only when substituted into the ancestor predicate in the second rule. This is a distinction that the Magic Sets strategy makes, and it thereby reduces the number of useless firings.

We now consider the third factor, the arity of the intermediate relations. The two strategies which

use unary intermediate relations are the Henschen-Naqv1 and Counting strategies. In essence, at step 1 they compute the set of relevant facts which is at a distance 1 from the query node. Let us denote this set by S_1 . At the next step, they compute the set of those nodes in *parent* to which there is an arc from a node in S_1 . Thus, they compute all nodes reachable from *john*, and further they compute each node at most D times where D is the duplication factor. However, the unary relations strategy fails to terminate if the query node is in a cycle. Also, neither the Henschen-Naqv1 nor the Counting strategy applies when there are non-linear rules.

Magic Sets computes exactly the same set of relevant facts and does no duplicate work. However, in the second phase at step 1 it computes all arcs in the transitive closure of *parent* (restricted to the set of relevant facts) of length 1. In particular, this includes all arcs of length 1 rooted at *john*. This is the answer, and this is essentially all that the more specialized methods, Henschen-Naqv1 and Counting, compute. Everything else that the Magic Sets strategy does is useless computation. Thus, the cost of the Magic Sets strategy is the number of arcs in the transitive closure of the subtree rooted at *john* (i.e. the subtree of nodes reachable from *john*).

The recursive control of QSQR generates subqueries using precisely the nodes in set S_1 at step 1, and the answer to each of these subqueries is the set of all nodes in the subtree rooted at that node. By induction, it is easy to see that the total cost involved in computing a query is the number of arcs in the transitive closure of the subgraph rooted at that query node. The intermediate relations here are the (binary) sets of answers to each subquery. This seems to indicate the power of a recursive control strategy since it succeeds in reducing both the set of relevant facts and the amount of duplicate work.

4.7.2. The Same Generation Query

We conclude this discussion by explaining the performance of the various strategies in the same generation query in terms of these three factors. Counting has the best performance since it uses the smallest set of relevant facts (the nodes of *up* which are reachable from the query node), does not do duplicate computation, and further, uses unary intermediate relations. It executes the query in two phases. In the first phase, at step 1, it computes the set of all nodes in *up* that are reachable from the query node via a path of length 1. In the second phase, it first computes the nodes of *down* that are reachable from this set via an arc of flat,

still retaining the distance of each set from the query node. In subsequent iterations, it steps through *down* once each time, such that each node in a set that is i steps away from the query node in *up* is the root of paths of length i in *down*.

Henschen-Naqvı uses the same set of relevant facts, and is a unary strategy, but it does a lot of duplicate work. It is a single phase algorithm, which does the same amount of work as the first phase of Counting in computing sets of *up* nodes along with their distances from the query node. However, it steps through *down* i times for each set at a distance i from the query node in *up*. Since it does not keep track of the work it does in step i at step $i+1$, it repeats a lot of the work in stepping through *down*.

The set of relevant facts for Magic Set, QSQR and APEX is again the set of *up* nodes reachable from the query node. They do not perform duplicate computation. However, they work with binary relations, in effect computing all paths with equal lengths in *up* and *down* linked by a single arc in *flat*. Thus, their performance is inferior to that of Counting. Our graphs show their performance to be identical to that of Henschen-Naqvı. It is to be expected that they perform similarly since the duplicate work done by Henschen-Naqvı is offset by the fact that they work with binary relations. However, their performance is not really identical. It appears to be so in our curves for two reasons. The first is our approximation of the number of arcs of length l to $n(l) \text{ gsum}(E, h-l)$. The second is the fact that we plot the curves for cases where *up* and *down* are identical. Under these conditions, the expressions for the performance of these methods become identical.

QSQI is similar to QSQR except that at each step, it duplicates the work of the previous steps, and so it is inferior to Magic Set, QSQR and APEX. Semi-Naive uses binary relations, and although it does not do duplicate work, this is outweighed by the fact that the set of relevant facts is all the nodes in *up*. So it performs worse than QSQI. Kifer-Lozinski degenerates to Semi-Naive since the optimization strategy fails to make any improvements to the system graph. Prolog is similar to QSQR when there is no duplication in the data, but its cost increases exponentially with the depth of the data structure when there is duplication. Naive evaluation uses the entire set of nodes in *up* as relevant facts, does duplicate work since it does not remember firings, and uses binary intermediate relations. With the exception of Prolog over a certain range, it is clearly the worst strategy.

Finally, we note that when the transfer ratio T is 0.01 (1%), the cost of computing the answer by Naive or Semi-Naive evaluation is essentially that of computing all arcs in the relation *flat*, and so the two methods perform almost identically.

4.8. Summary and Caveats

Our conclusions may be summarized as follows:

- 1 For a given query, there is a clear ordering of the strategies.
- 2 The more specialized strategies perform significantly better.
- 3 Recursion is a powerful control structure which reduces the number of relevant facts and eliminates duplicate work.
- 4 The choice of the right strategy is critical since the differences in performance are by orders of magnitude.
- 5 Three factors which greatly influence performance are (i) duplication of work, (ii) the set of relevant facts, and (iii) the arity of the intermediate relations.

The results seem robust in that the performance of the various strategies usually differ by orders of magnitude, which allows a wide latitude for the approximations in the model and cost evaluation. Also, the curves rarely intersect, which means that the relative ordering of the strategies is maintained in most cases over the entire range of data.

However, it must be emphasized that our cost function makes some crude approximations. The cost of join is linear in the size of the result, a consequence of our using the size of intermediate relations as the cost measure. We also ignore the cost of disk accesses, and the cost of implementing a recursive control strategy. Our model suffers from the approximation that duplication is independent of the size of the start set.

Finally, our sample data and queries are limited, and the results must be extrapolated to other data and queries with caution, especially since the results show some variance in the relative performance of the strategies for different sets of data and queries. In particular, our benchmark is limited to the type of data and query where there is a *large* amount of data and the size of the answer to the query is *small*. This clearly favors the "smart" strategies and obscures, for instance, the fact that Semi-Naive performs as well as any other strategy when computing the entire transitive closure of a relation [Bancilhon 85]. Further, our data contains no cycles or shortcuts. This is an important limitation since it favours some of the specialized strategies. For instance, there are cases where Counting performs worse than Magic Sets [Bancilhon et

al 86] This is not shown by our results since these cases involve shortcuts in the data

We have also assumed in this paper that methods should strive for generality, i.e. we have not addressed the problem of finding a set of specialized operators which would solve the "real life" cases of recursion. Other authors have addressed this problem, mainly by concentrating on the transitive closure operator [Valduriez and Boral 85] or extensions of it [Dayal et al 85, Rosenthal et al 85]

5. Conclusions

In this paper, we have given a description and comparative evaluation of the major strategies for processing logic queries without function symbols

We have tried to identify the exact application domain for each method. We have also tried to describe the strategies in a uniform manner. Unfortunately, we have only been partially successful at that. We have identified a set of major characteristics of the strategies: method vs optimization strategy, top-down vs bottom-up, recursive vs iterative and compiled vs interpreted. But some of these characteristics are somewhat arbitrary for the same strategy: it is sometimes possible to have a compiled or interpreted version. For instance, we have presented a compiled version of naive evaluation, while SNIP is an interpreted version of it. It seems also reasonable to design a compiled version of iterative QSQ. We also argued that the distinction between optimization strategy and method was mainly of a pedagogical interest. However, the top-down vs bottom-up and recursive vs iterative distinction seems to capture intrinsic properties of the strategies. But we consider that the problem of finding a good taxonomy of strategies is still wide open.

We have presented a performance comparison of ten methods. Even though the "benchmark" we have used is incomplete, the cost measure too elementary and the approximations crude, we found the results to be valuable. The robustness of the results (at least on our workload), both in terms of the order of magnitude differences between the costs of the strategies and in terms of invariance of the results to the parameters which we varied, was a surprise. We have also been able to explain most of our results through three factors: duplication, relevant facts and unary vs binary. While the first two factors were well known, the third one came also as a surprise, even though it was probably already understood in [Sacca and Zaniolo 86]

Acknowledgements

We wish to thank Bill Alexander and Patrick Valduriez for careful proofreading of parts of the manuscript. We are extremely grateful to Paris Kanellakis, Eliezer Lozinski, Jeff Ullman, Laurent Vieille and Carlo Zaniolo who provided enlightening comments and suggested many corrections and improvements to the paper.

References

- [Afrati et al 86]
"Convergence of Sideways Query Evaluation," F. Afrati, C. Papadimitriou, G. Papageorgiou, A. Roussou, Y. Sagiv and J. Ullman, *Proc 5th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1986*
- [Aho and Ullman 79]
"Universality of Data Retrieval Languages," A. Aho and J. Ullman, *Proc 6th ACM Symposium on Principles of Programming Languages, 1979*
- [Apt and Van Emden 82]
"Contributions to the Theory of Logic Programming," *JACM, 1982*
- [Bancilhon 85]
"Naive Evaluation of Recursively Defined Relations," F. Bancilhon, in *On Knowledge Base Management Systems - Integrating Database and AI Systems*, Brodie and Mylopoulos, Eds, Springer-Verlag
- [Bancilhon 85a]
"A Note on the Performance of Rule Based Systems," F. Bancilhon, *MCC Technical Report DB-022-85, 1985*
- [Bancilhon et al 86]
"Magic Sets and Other Strange Ways to Implement Logic Programs," F. Bancilhon, D. Maier, Y. Sagiv and J. Ullman, *Proc 5th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1986*
- [Bancilhon et al 86a]
"Magic Sets: Algorithms and Examples," F. Bancilhon, D. Maier, Y. Sagiv and J. Ullman, *Unpublished Manuscript, 1986*
- [Bancilhon and Ramakrishnan 86]
"Performance Evaluation of Data Intensive Logic Programs," F. Bancilhon and R. Ramakrishnan, *Unpublished Manuscript, March 1986*
- [Bayer 85]
"Query Evaluation and Recursion in Deductive Database Systems," R. Bayer, *Unpub-*

- ished Manuscript, 1985
- [Chang 81]
 "On the Evaluation of Queries Containing Derived Relations in Relational Databases," C Chang, In *Advances in Data Base Theory, Vol 1*, H Gallaire, J Minker and JM Nicolas, Plenum Press, New York, 1981, pp 235-260
- [Dayal et al 85]
 "PROBE- a Research Project in Knowledge-Oriented Database Systems Preliminary Analysis," U Dayal, A Buchmann, D Goldhirsch, S Heiler, F Manola, J Orenstein and A Rosenthal, *Technical Report, CCA-85-03, July 1985*
- [Delobel 86]
 "Bases de Donnees et Bases de Connaissances Une Approche Systemique a l'Aide d'une Algebre Matricielle des Relations," C Delobel *Journées Francophones, Grenoble, January 1986*
- [Dietrich and Warren 85]
 "Dynamic Programming Strategies for the Evaluation of Recursive Queries," S W Dietrich and D S Warren, *Unpublished Report, 1985*
- [Gallaire et al 84]
 "Logic and Data Bases A Deductive Approach," H Gallaire, J Minker and J-M Nicolas, *Computing Surveys, Vol 16, No 2, June 1984*
- [Gardarin and Maindreville 85]
 "Evaluation of Database Recursive Logic Programs as Recurrent Function Series," G Gardarin and Ch de Maindreville, *Proc SIGMOD 86, Washington, D C, May 1986*
- [Han and Lu 86]
 "Some Performance Results on Recursive Query Processing in Relational Database Systems," J Han and H Lu *Proc Data Engineering Conference, Los Angeles, February 1986*
- [Henschen and Naqvi 84]
 "On Compiling Queries in Recursive First-Order Data Bases," L Henschen and S Naqvi, *JACM, Vol 31, January 1984, pp 47-85*
- [Kifer and Lozinski 85]
 "Query Optimization in Logic Databases," M Kifer and E Lozinski, *Technical Report, SUNY at Stonybrook, June 1985*
- [Laskowski 84]
 "Compiling Recursive Axioms in First Order Databases," K Laskowski, *Masters Thesis, Northwestern University, 1984*
- [Lozinski 83]
 "A Problem-Oriented Inferential Database System," E Lozinski, *Tech Report 83-17, The Hebrew University of Jerusalem, May 1983*
- [Lozinski 85]
 "Evaluating Queries in Deductive Databases by Generating," E Lozinski, *Proc 11th International Joint Conference on Artificial Intelligence, 1985*
- [Lozinski 85a]
 "Inference by Generating and Structuring of Deductive Databases," E Lozinski, *Unpublished Manuscript, 1985*
- [Marque-Pucheu 83]
 "Algebraic Structure of Answers in a Recursive Logic Database," G Marque-Pucheu, *To appear in Acta Informatica*
- [Marque-Pucheu et al 84]
 "Interfacing Prolog and Relational Database Management Systems," G Marque-Pucheu, J Martin-Gallaussaux and G Jomier, in *New Applications of Databases, Gardarin and Gelenbe Eds, Academic Press, London, 1984*
- [McKay and Shapiro 81]
 "Using Active Connection Graphs for Reasoning with Recursive Rules," D McKay and S Shapiro, *Proc 7th International Joint Conference on Artificial Intelligence, 1981*
- [Morris et al 86]
 "Design Overview of the NAIL! System," K Morris, J Ullman and A Van Gelder, *Proceedings of the 3rd International conference on Logic Programming, London, July 1986*
- [Naqvi and Fishman 81]
 "An Improved Compiling Technique for First Order Databases," Shamim A Naqvi and Daniel H Fishman, *Proc Formal Bases for Databases, Toulouse, October 1981*
- [Naqvi and Henschen 83]
 "Synthesizing Least Fixed Point Queries into Non-recursive Iterative Programs," S Naqvi and L Henschen, *Proc 9th International Joint Conference on Artificial Intelligence, Karlsruhe, 1983*
- [Reiter 78]
 "Deductive Question Answering on Relational Data Base," R Reiter, In *Logic and Data Bases, H Gallaire and J Minker Ple-*

- num Press, New York, 1978, pp 149-177
- [Rohmer and Lescoeur 85]
 "La Methode Alexandre une solution pour traiter les axiomes recursifs dans les bases de donnees deductives ," *Colloque Reconnaissance de Formes et Intelligence Artificielle, Grenoble, November 1985*
- [Rosenthal et al 85]
 "Traversal Recursion A Practical Approach to Supporting Recursive Applications," A Rosenthal, S Heiler, U Dayal, F Manola, *Unpublished Report, CCA, December 1985*
- [Sacca and Zaniolo 86a]
 "On the Implementation of a Simple Class of Logic Queries for Databases," D Sacca and C Zaniolo, *Proc 5th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1986*
- [Sacca and Zaniolo 86b]
 "Implementing Recursive Logic Queries with Function Symbols," *Unpublished Manuscript, April 1986*
- [Sagiv and Ullman 84]
 "Complexity of a Top-Down Capture Rule," Y Sagiv and J Ullman, *Technical Report, Stanford University, STAN-CS-84-1009, 1984*
- [Shapiro and McKay 80]
 "Inference with Recursive Rules," S Shapiro and D McKay, *Proc 1st Annual National Conference on Artificial Intelligence, August, 1980*
- [Shapiro et al 82]
 "Bi-Directional Inference," S Shapiro, J Martins and D McKay, *Proc 4th Annual Conference of the Cognitive Science Society, Ann Arbor, Michigan, 1982*
- [Tarski 55]
 "A Lattice Theoretical Fixpoint Theorem and its Applications" A Tarski, *Pacific Journal of Mathematics 5, 1955, pp 285-309*
- [Ullman 85]
 "Implementation of Logical Query Languages for Databases," J Ullman, *TOLD, Vol 10, No 3, pp 289-321, 1985*
- [Ullman and Van Gelder 85]
 "Testing Applicability of Top-Down Capture Rules," J Ullman and A Van Gelder, *Technical Report, Stanford University, STAN-CS-85-1046, 1985*
- [Van Emden and Kowalski 76]
 "The Semantics of Predicate Logic as a Programming Language," M Van Emden and R Kowalski, *JACM, Vol 23, No 4, October 1976, pp 733-742*
- [Valduriez and Boral 86]
 "Evaluation of Recursive Queries Using Join Indices," P Valduriez and H Boral, *Proc First Intl Conference on Expert Database Systems, Charleston, 1986*
- [Vieille 85]
 "On Handling Recursively Defined Virtual Relations in Deductive Databases," L Vieille, *Unpublished Report, ECRC, Munich*
- [Vieille 86]
 "Recursive axioms in Deductive Databases The Query/Subquery Approach," L Vieille, *Proc First Intl Conference on Expert Database Systems, Charleston, 1986*
- [Zaniolo 85]
 "The Representation and Deductive Retrieval of Complex Objects," C Zaniolo, *Proc 11th Int Conference on Very Large Data Bases, Stockholm, September 1985*
- [Zaniolo 86]
 "Safety and Compilation of Non-Recursive Horn Clauses," C Zaniolo *Proc First Intl Conference on Expert Database Systems, Charleston, 1986*

Appendix 1: The Cost Functions

We first explain the notation and terminology used in analytically deriving the cost functions. We have derived expressions for the cost of each strategy on each of the four queries we have defined. We refer the reader to [Bancilhon and Ramakrishnan 86] for the derivation of these expressions.

We denote the number of nodes at level i in relation R by $n_R(i)$, and the total number of arcs in R (which is the number of tuples in R) by A_R . Where no confusion is possible, we drop the subscript.

We denote the sum of the $(h+1)$ st elements of the geometric series of ration E by $gsum(E,h)$, thus

$$gsum(E,h) = (1 + E + E^2 + E^3 + \dots + E^h)$$

We define the length of an arc in the transitive closure of R (which we denote by R^*) to be the length of the path of R that generates it. (Note that this is well defined because there are no short-cuts.)

We denote by $a_{R^*}(l)$ the number of arcs of length exactly l in R^* . Where the context is clear, we write $a(l)$.

$$a(l) = n(l) + n(l+1) + \dots + n(h) = n(l) \text{ gsum}(E, h-l)$$

We denote by h' the *average* level

$$h' = h - \left\lfloor \frac{\left(\sum_{i=1}^h (i n(i)) \right)}{N} \right\rfloor$$

It denotes the mean level at which we pick a node, assuming nodes are uniformly distributed. We have actually defined h' as the distance of the mean level from the highest level h for notational convenience, since this is a quantity we use extensively.

Query 1 (Ancestor.bf)

1.1 *Naive evaluation* $D \sum_{i=1}^h (h-i+1) a(i) + E \text{ gsum}(E, h'-1)$

1.2 *Semi-Naive Evaluation* $D \sum_{i=1}^h a(i) + E \text{ gsum}(E, h'-1)$

1.3 *QSQ, Iterative* $E \text{ gsum}(E, h'-1) + F \sum_{i=1}^{h'} (h'-i+1) E^{i-1}$

1.4 *QSQ, Recursive* $(F+E) \text{ gsum}(E, h'-1) + D \sum_{i=1}^{h'} E^i \text{ gsum}(E, h'-i)$

1.5 *Henschen-Naqvi* $(F+E) \text{ gsum}(E, h'-1)$

1.6 *Prolog* $\text{gsum}(F, h') + E \text{ gsum}(E, h'-1) + \sum_{i=1}^{h'} (F^i) \text{ gsum}(F, h'-i)$

1.7 *APEX* $(F+E) \text{ gsum}(E, h'-1) + D \sum_{i=1}^{h'} E^i \text{ gsum}(E, h'-i)$

1.8 *Kifer-Lozinski* $D \sum_{i=1}^h a(i) + E \text{ gsum}(E, h'-1)$

1.9 *Magic Sets* $(F+E) \text{ gsum}(E, h'-1) + D \sum_{i=1}^{h'} E^i \text{ gsum}(E, h'-i)$

1.10 *Counting* $(F+E) \text{ gsum}(E, h'-1)$

Query 2 (Ancestor.fb)

2.1 *Naive evaluation* $D \sum_{i=1}^h (h-i+1) a(i) + (1/E) \text{ gsum}(1/E, h-h'-1)$

2.2 *Semi-Naive Evaluation* $D \sum_{i=1}^h a(i) + (1/E) \text{ gsum}(1/E, h-h'-1)$

2 3 <i>QSQ, Iterative</i>	$(1/E) \text{gsum}(1/E, h-h'-1) + D \sum_{i=1}^{h-h'} (h-h'-i+1) i (1/E)^{i-1}$
2 4 <i>QSQ, Recursive</i>	$1 + (1/E) \text{gsum}(1/E, h-h'-1) + F \sum_{i=1}^{h-h'} (1/E)^i \text{gsum}(1/E, h-h'-i)$
2 5 <i>Henschen-Nagvi</i>	$(D+1/E) \text{gsum}(1/E, h-h'-1)$
2 6 <i>Prolog</i>	$(1/E) \text{gsum}(1/E, h-h'-1) + \sum_{i=1}^h n(i) \text{gsum}(F, h-i)$
2 7 <i>APEX</i>	$(1/E)^{(h-h')} (E \text{gsum}(E, h-1) + D \sum_{i=1}^h E^i \text{gsum}(E, h-i))$
2 8 <i>Kifer-Lozinski</i>	$(D+1/E) \text{gsum}(1/E, h-h'-1)$
2 9 <i>Magic Sets</i>	$1 + (1/E) \text{gsum}(1/E, h-h'-1) + F \sum_{i=1}^{h-h'} (1/E)^i \text{gsum}(1/E, h-h'-i)$
2 10 <i>Counting</i>	$(D+1/E) \text{gsum}(1/E, h-h'-1)$

Query 3 (Ancestor.bf, Non-Linear Version)

3 1 <i>Naive evaluation</i>	$E \text{gsum}(E, h'-1) + D \sum_{i=1}^h (\log(h/i)+1) (i-1) a(i)$
3 2 <i>Semi-Naive Evaluation</i>	$E \text{gsum}(E, h'-1) + D \sum_{i=1}^h (i-1) a(i)$
3 3 <i>QSQ, Iterative</i>	$E \text{gsum}(E, h'-1) + F \sum_{i=1}^{h'} (h'-i+1) i E^{i-1}$
3 4 <i>QSQ, Recursive</i>	$F+E \text{gsum}(E, h'-1) + D \sum_{i=2}^{h'} (i-1) E^i$
3 5 <i>Henschen-Nagvi</i>	Does not apply
3 6 <i>Prolog</i>	Does not terminate
3 7 <i>APEX</i>	$E \text{gsum}(E, h'-1) + (1/E)^{h-h'} (D \sum_{i=1}^h (i-1) E^i \text{gsum}(E, h-i))$ $+ E^{h'} (F \sum_{i=1}^h (i-1) (1/E)^i \text{gsum}(1/E, h-i))$
3 8 <i>Kifer-Lozinski</i>	$E \text{gsum}(E, h'-1) + D \sum_{i=1}^h (i-1) a(i)$
3 9 <i>Magic Sets</i>	$E \text{gsum}(E, h'-1) + D \sum_{i=1}^h (i-1) a(i)$
3 10 <i>Counting</i>	Does not apply

Query 4 (Same Generation.bf)

In the following expressions, $h'_{up\ down} = \min(h'_{up}, h'_{down})$, and $h_{up\ down} = \min(h_{up}, h_{down})$

4.1 Naive evaluation

$$A_{flat} + T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} D_{down} \sum_{i=1}^{h_{up\ down}} (h_{up\ down} - i + 1) a_{up}(i) E_{down}^i +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4.2 Semi-Naive Evaluation

$$A_{flat} + T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} D_{down} \sum_{i=1}^{h_{up\ down}} a_{up}(i) E_{down}^i +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4.3 QSQ, Iterative

$$(h'_{up\ down} + 1) F_{flat} +$$

$$T_{up\ 2\ flat\ 1} F_{flat} \sum_{i=1}^{h'_{up\ down}} (h'_{up\ down} - i + 1) E_{up}^i +$$

$$E_{flat} T_{up\ 2\ flat\ 1} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (h'_{up\ down} - i + 1) E_{up}^i \text{gsum}(E_{down}, i-1) +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4.4 QSQ, Recursive

$$F_{up} \text{gsum}(E_{up}, h'_{up}-1) + E_{up} \text{gsum}(E_{up}, h'_{up}-1) T_{up\ 2\ flat\ 1} F_{flat} +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} D_{down} \sum_{i=1}^{h'_{up\ down}} E_{up}^i \text{gsum}(E_{up}, h'_{up}-1) E_{down}^i +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4.5 Henschen-Naqui

$$F_{up} \text{gsum}(E_{up}, h'_{up}-1) +$$

$$\sum_{i=1}^{h'_{up\ down}} (E_{up}^i T_{up\ 2\ flat\ 1} F_{flat} + T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} E_{up}^i \text{gsum}(E_{down}, i-1)) +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4.6 Prolog

$$\text{gsum}(F_{up}, h'_{up}-1) + F_{up} \text{gsum}(F_{up}, h'_{up}-1) T_{up\ 2\ flat\ 1} F_{flat} +$$

$$T_{up\ 2\ flat\ 1} F_{flat} T_{flat\ 2\ down\ 1} \sum_{i=1}^{h'_{up\ down}} F_{up}^i \text{gsum}(F_{up}, h'_{up}-1) F_{down}^i +$$

$$T_{up\ 2\ flat\ 1} E_{flat} T_{flat\ 2\ down\ 1} F_{down} \sum_{i=1}^{h'_{up\ down}} (E_{up} E_{down})^i$$

4 7 APEX

$$\begin{aligned}
 & F_{up} \text{gsum}(E_{up}, h'_{up}-1) + E_{up} \text{gsum}(E_{up}, h'_{up}-1) T_{up 2.flat 1} F_{flat} + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} D_{down} \sum_{i=1}^{h'_{up down}} E_{up}^i \text{gsum}(E_{up}, h'_{up}-1) E_{down}^i + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} F_{down} \sum_{i=1}^{h'_{up down}} (E_{up} E_{down})^i
 \end{aligned}$$

4 8 Kifer-Lozinski

$$\begin{aligned}
 & A_{flat} + T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} D_{down} \left(\sum_{i=1}^{h_{up down}} (a_{up}(i) E_{down}^i) \right) + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} F_{down} \sum_{i=1}^{h'_{up down}} (E_{up} E_{down})^i
 \end{aligned}$$

4 9 Magic Sets


$$\begin{aligned}
 & F_{up} \text{gsum}(E_{up}, h'_{up}-1) + E_{up} \text{gsum}(E_{up}, h'_{up}-1) T_{up 2.flat 1} F_{flat} + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} D_{down} \sum_{i=1}^{h'_{up down}} E_{up}^i \text{gsum}(E_{up}, h'_{up}-1) E_{down}^i + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} F_{down} \sum_{i=1}^{h'_{up down}} (E_{up} E_{down})^i
 \end{aligned}$$

4 10 Counting


$$\begin{aligned}
 & F_{up} \text{gsum}(E_{up}, h'_{up}-1) + \\
 & T_{up 2.flat 1} F_{flat} (1 + E_{up} \text{gsum}(E_{up}, h'_{up}-1)) + \\
 & \sum_{i=1}^{h'_{up down}} T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} D_{down} (E_{up} E_{down})^i + \\
 & T_{up 2.flat 1} E_{flat} T_{flat 2.down 1} F_{down} \sum_{i=1}^{h'_{up down}} (E_{up} E_{down})^i
 \end{aligned}$$

Appendix 2: The Curves


A.b.f. Tree




A.b.f. Inverted tree




A.b.f. Cylinder




A.fb. Tree




Sg.bf. Tree. T = 1



Sg.bf. Cylinder. T = 1



A.dr. Cylinder



Sg.bf. Cylinder. T = 0.01

