

A Framework for Testing Safety and Effective Computability of Extended Datalog (Extended Abstract)

Ravi Krishnamurthy,¹
Raghu Ramakrishnan,²
Oded Shmueli^(1,3)

ABSTRACT

This paper presents a methodology for testing a general logic program containing function symbols and built-in predicates for *safety* and *effective computability*. Safety is the property that the set of answers for a given query is finite. A related issue is whether the evaluation strategy can effectively compute all answers and terminate. We consider these problems under the assumption that queries are evaluated using a bottom-up fixpoint computation. We also approximate the use of function symbols by considering Datalog programs with infinite base relations over which *finiteness constraints* and *monotonicity constraints* are considered. One of the main results of this paper is a recursive algorithm, *check_clique*, to test the safety and effective computability of predicates in arbitrarily complex cliques. This algorithm takes certain procedures as parameters, and its applicability can be strengthened by making these procedures more sophisticated. We specify the properties required of these procedures precisely, and present a formal proof of correctness for algorithm *check_clique*. This work provides a framework for testing safety and effective computability of recursive programs, and is based on a clique by clique analysis. The results reported here form the basis of the safety testing for the LDL language, being implemented at MCC.

1. Introduction

There has been much work recently on efficient bottom-up evaluation of logic programs. (See [Bancilhon and Ramakrishnan 86] for a survey.) Since these bottom-up strategies promise to be an efficient approach to evaluating recursive queries, it is important to investigate their properties. In the context of databases, the set of answers is always finite in the absence of recursion and negation. It is widely expected that while the addition of recursion to the query language will make it more expressive, the typical query will still require a finite set of answers. Thus, an important question is whether the set of answers for a given query is indeed finite. If not, the program is probably incorrect. This is the *safety* problem, and has been shown to be undecidable for Horn Clause programs with function symbols [Shmueli 87, Gaifman 86]. A related issue is whether the evaluation strategy (in particular, bottom-up evaluation, possibly after rewriting the program) computes all answers and terminates.

(1) MCC, Austin. (2) University of Wisconsin-Madison. Work partly done while visiting MCC. (3) Technion, Israel.

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.

While the questions of safety and termination are important questions for general logic programs, they are particularly of interest for programs expressing database queries. We address these questions in this paper. We consider programs which do not contain terms constructed using function symbols (i.e. all arguments are constants or variables), also known as *Datalog*. However, we allow infinite base relations, and these can be used to model terms constructed with function symbols and evaluable functions such as arithmetic operations. We use the name *Extended Datalog* to denote Datalog with this extension.

One of the main results of this paper is a recursive algorithm, *check_clique*, to test the safety and effective computability of predicates in arbitrarily complex cliques. This algorithm takes certain procedures as parameters, and its applicability can be strengthened by making these procedures more sophisticated. We specify the properties required of these procedures precisely, and present a formal proof of correctness for algorithm *check_clique*. While the complexity of this algorithm is exponential in the size of the clique, we believe that most cliques in practice are likely to be sufficiently small that this is acceptable. (The procedures that are input to algorithm *check_clique* include procedures for inferring whether certain *monotonicity constraints* hold over the arguments of derived predicates and whether the set of values that can appear in a given argument position is *bounded*. We discuss this issue in Section 5.)

Our framework for testing safety and effective computability is based on a simple bottom-up model of execution. We assume that the given program has already been rewritten, if necessary, according to the desired sideways information passing strategies. Thus, in the safety testing phase, we only check the bottom-up fixpoint computation of the given program, assuming that it will not be further transformed or otherwise optimized. Thus, if a query is declared to be safe and effectively computable, bottom-up execution according to our model is an effective procedure for computing it. An important aspect of our approach to testing is that it is based on a clique by clique analysis. We develop algorithms for testing the safety of predicates in a clique. We test a given program by topologically sorting its cliques, and then testing them in topological order starting from the leaf cliques. This makes the analysis of large programs tractable since the complexity of our algorithms is dominated by the size of recursive cliques, and we expect these to be small independent of the size of the program.

The *capture rules* framework proposed in [Ullman 85] also addressed the issues of safety and effective computability: if a query can be "captured" using some known capture rules, then there is an algorithm which computes all answers. The results described in this paper allow us to detect the safety and effective computability of all queries that can be "captured" using the capture rules described in the literature [Ullman 85, Sagiv and Ullman 85, Afrati et al. 86, Ullman and Van Gelder 85]. On the other hand, several of the examples presented here cannot be handled using the previously known capture rules.

The rest of this paper is organized as follows. We present some definitions in Section 2. In Section 3, we present an overview of our approach to safety analysis. In Section 4, we provide an introduction to the problems involved in testing recursive cliques for safety. In Section 5, we develop algorithm *check_clique* and prove its correctness. We discuss related work in Section 6, and present our conclusions in Section 7.

2. Definitions

In this section, we review a number of basic concepts from the literature. We assume the standard definitions of *Horn Clause rules*, *literals*, *(non-) ground terms*, *queries*, and *answers*. We assume that queries only contain variables. †

We only consider programs in which all arguments are variables or constants. Thus, terms constructed using function symbols are not allowed. We approximate such terms by allowing infinite base relations. We use the name *Extended Datalog* to denote this

† Queries containing constants are dealt with by first rewriting the program. See Section 6 for a discussion.

class of programs. The following example illustrates how a program with function symbols can be approximated as an *Extended Datalog* program.

Example 2.1 The following is a program that concatenates two lists.

```
concat([X|Y],Z,[X|U]) :- concat(Y,Z,U).
concat([],Z,Z) :-.
```

We approximate it as follows in *Extended Datalog*:

```
concat(V,Z,W) :- concat(Y,Z,U), h(X,Y,V), h(X,U,W).
concat([],Z,Z) :-.
```

The predicate *h* is an infinite base predicate corresponding to the function symbol *l*. Conceptually, it contains all triples (element, list1, list2) such that list2 is the term [element|list1]. □

We partition the given *Extended Datalog* program into a set of *facts* which are stored in the *base predicates* (the *Extensional Database* or *EDB*), and a set of rules defining *derived predicates* (the *Intensional Database* or *IDB*). A derived predicate is one which appears in the head of a rule. Without loss of generality, we assume that the *EDB* and *IDB* are disjoint sets of predicates. In translating a *Horn Clause* program with function symbols into an *Extended Datalog* program, there is some loss of information. In order to retain more information, and thus improve the approximation, we allow a set of *integrity constraints* (*IC*) to be specified over the *EDB* predicates. The set of facts in the *EDB* predicates must satisfy these constraints. Thus a database is a tuple (*EDB*, *IDB*, *IC*).

Predicates which are empty for all *EDB* instances unnecessarily complicate the safety analysis, and w.l.g., we assume that the program does not contain such predicates. †

A given query is *safe* if it has a finite set of answers for all instances of the *EDB* which satisfy all integrity constraints. Thus, a query is *unsafe* if there is some instance of the *EDB* which satisfies the integrity constraints and is such that the query has an infinite set of answers. The integrity constraints we consider in this paper are *finiteness constraints* and *monotonicity constraints*.

A *finiteness constraint* (*FC*) over a base predicate *r* is of the form $X \rightarrow Y$, where *X* and *Y* are sets of arguments. An instance of *r* satisfies this constraint if and only if the following property holds: For each tuple *t* in *r*, the set of tuples $\{s[Y] \mid s[X] = t[X]\}$, is finite. Note that this definition is strictly weaker than the traditional definition of a functional dependency; it holds trivially for all finite predicates (and in particular, for all finite *EDB*

† A polynomial time algorithm for identifying such predicates is presented in [Ramakrishnan et al. 87].

predicates).

Example 2.3 Consider the predicate $f(X,Y)$. If we define $f(X,Y)$ by $Y = 2*X$, then we have the finiteness constraints $f_1 \rightarrow f_2$ and $f_2 \rightarrow f_1$. If we define $g(X,Y)$ by $X < 0$ and $Y > 0$, there is no finiteness constraint between the arguments of 'g'. If we define $h(X,Y)$ by $X > 0$ and $Y = 0$ or 5 , we have the finiteness constraint $h_1 \rightarrow h_2$. \square

Let r_i and r_j be arguments of predicate r . A *monotonicity constraint* is a couple $r_i > r_j$. The constraint $r_i > r_j$ holds in an instance of r if and only if the value in the i th argument is strictly greater than the value in the j th argument in every tuple. We may also specify a monotonicity constraint $r_i > c$ (resp. $r_i < c$) where 'c' is a constant. This constraint holds in an instance of r if and only if the value in the i th argument is strictly greater (lesser than) than the constant 'c'. Of course, this assumes that the values are drawn from a domain with a partial order. Although it is possible to consider many different orders, we assume that there is a single order. (This is only for ease of exposition, and is discussed further after Example 4.2.)

Example 2.4 Consider the predicate $f(X,Y)$. If we define $f(X,Y)$ by $Y = 2*X$, then we have the finiteness constraints $f_1 \rightarrow f_2$ and $f_2 \rightarrow f_1$. If we define $f(X,Y)$ by $Y = 2*X$, $X > 0$, $Y > 0$, we have the finiteness constraints $f_1 \rightarrow f_2$ and $f_2 \rightarrow f_1$, and the monotonicity constraint $f_2 > f_1$. If we define $f(X,Y)$ by $X < 0$ and $Y > 0$, there is no finiteness constraint between the arguments of 'f', but the monotonicity constraint $f_2 > f_1$ holds. If we define $f(X,Y)$ by $X > 0$ and $Y = 0$ or 5 , we have the finiteness constraint $f_1 \rightarrow f_2$, but there is no monotonicity constraint relating X and Y . \square

A predicate p is defined using predicate q if there is a rule containing p in the head and q in the body. We denote this as $p \leftarrow q$. We say that p depends on q if p is defined using q or p depends on r and r depends on q . We denote this as $p * \leftarrow q$. p is a recursive predicate if $p * \leftarrow p$. Two predicates p and q are mutually recursive if $p * \leftarrow q$ and $q * \leftarrow p$. A clique is a maximal set of mutually recursive predicates. We assume that each non-recursive predicate forms a singleton clique.

We begin by introducing several concepts and definitions. For each fact that belongs to a derived predicate, there exists a finite *derivation tree*, which describes how it is derived from base facts using rules of the program. Let $p(c)$ be a fact in the derived predicate p . Then the tree has $p(c)$ at its root, the leaves are base facts, and each internal node is labeled by a fact, and by a rule which generates this fact from the facts labeling its children. A base fact may be viewed as a derivation tree of height one.

In order to reason with monotonicity constraints, we need the notion of an *argument mapping*. Our definition is similar to the definitions presented in [Afrati et al. 86, Ramakrishnan et al. 87]. Let p be the head and q be a derived literal occurrence in the body of a rule r . An argument mapping (p, q, r) is a graph with the set of nodes being the argument places of p and q .[†] For each rule, we obtain an argument mapping between the head literal and each derived literal occurrence in the body. We draw an undirected *edge* between two nodes if the same variable occurs in the corresponding argument positions. We draw an *arc* from one node, say n_1 , to another, say n_2 if variables X and Y appear in the corresponding argument positions, and we can infer that $X > Y$ from the monotonicity constraints in this rule.

Argument mappings (p, q, r_1) and (q, m, r_2) can be composed to yield a composite mapping $(p, q, r_1). (q, m, r_2)$ by merging the corresponding nodes in q . In particular, the composite mapping $(p, q_1, r). (q_1, q_2, r_1) \dots (q_n, p, r_n)$ represents a *cyclic composite argument mapping*, and we can complete the cycle by joining the corresponding nodes in the two instances of p with edges. Further, if $q_i \triangleleft p, i = 1$ to n , then this is a *simple cycle*.

Given a database $DB = (EDB, IDB, IC)$, consider the set of all derivation trees for facts in DB . (Henceforth, when we refer to derivation trees, it is understood that we refer to a tree in this set.) We observe that every path in a derivation tree induces a rule sequence. In fact, every path in a derivation tree induces a composite argument mapping. (In addition to the rule sequence, we must consider, at each node in the path, which literal in the rule body was expanded.) We thus often speak of the composite mapping *corresponding* to a path in a derivation tree. Also, since many paths in derivation trees may induce the same composite mapping, we include all such paths in the set of paths *corresponding* to the composite mapping.

3. Overview of the Safety Analysis

In addition to safety, we are also interested in whether the set of answers to the query is *effectively computable*, that is, whether the evaluation strategy computes all answers and terminates, and every intermediate relation is of finite size. Effective computability implies safety, but the converse is not true. In order to discuss effective computability, we define a stronger notion of safety.

[†] Without loss of generality, we assume that each literal occurrence has been given a unique name. A simple way to do this would be to use the rule number and the position of the literal in the rule, together, as the literal name. We also assume that all variable names are unique. We note that the third component of an argument mapping is redundant, but we include it for ease of exposition.

A predicate p is *strongly safe* if every derived predicate q such that $p \leftarrow q$ is safe. A query is *strongly safe* if the query predicate is strongly safe. A program is strongly safe if every predicate in it is strongly safe. A clique is strongly safe if every predicate in it is strongly safe. A rule is strongly safe if every predicate which appears in the body is strongly safe.

We now present an overview of our approach to testing a program. Clique C1 is a *child* (resp. *descendant*), of clique C2 if $p \leftarrow q$ (resp. $* \leftarrow$) for some p in C2 and q in C1. We denote this as $C2 \leq C1$ (resp. $C2 * \leq C1$). We observe that $* \leq$ is a partial order. A clique C is a *leaf* clique if there is no clique C1 such that $C \leq C1$. We test program P for strong safety according to this ordering of cliques. We begin with cliques that are leaves. All predicates in such a clique are either base predicates or predicates that belong to it.

Let C1 be a leaf clique, and let q be a predicate in C1 which is used to define some predicate in another clique C2. After analyzing clique C1, if it is strongly safe, we replace all occurrences of q by a base predicate. (If we cannot show this clique to be strongly safe, then we cannot show the program - which contains predicates from this clique - to be strongly safe.) After doing this for all predicates in C1 which are used to define predicates outside C1, we remove C1 from the set of cliques. We then consider another leaf clique, and proceed in this way until all cliques have been considered.

In order to reason about effective computability, we must first specify the model of execution. Our *model of execution* is the following: We assume that a given program P is first transformed to P^{ms} , which is then evaluated bottom-up. Thus, the rules are repeatedly applied (in any order) until no new tuples are produced by the application of any rule. A rule is *applied* by taking the join of (the current instances of) the relations in the body of the rule and projecting out tuples corresponding to the head. We assume a left-to-right order in computing the join of the body relations.

We make the following assumptions about programs:

1. All head variables in a rule also appear in the body.
2. If an infinite base relation f appears in the body of a rule, Let $\{f_i, \dots, f_j\}$ be the set of argument positions such that the variables in them appear to the left of this occurrence of f in the body of the rule. Then, these arguments must determine all arguments of f through the FCs given to hold over f .
3. Consider an infinite base relation f with an FC $f_i, \dots, f_j \rightarrow f_k$. Given an assignment of constants to argument positions f_i, \dots, f_j , we can compute the (finite) set of values appearing in the argument position

f_k in a finite amount of time.

We conclude this section by formally presenting some important properties of this model of execution. In particular, theorem 3.2 characterizes the effective computability of strongly safe programs.

We define an application of a rule to be effectively computable if it is possible to compute the join of the body predicates in left-to-right order without constructing infinite intermediate results, and the set of tuples thus produced for the head is also finite.

Lemma 3.1: Each application of a rule produces only a finite number of tuples for the head predicate and is effectively computable, given the three assumptions about programs. []

We now consider the notion of effective computability in the context of the left-to-right bottom-up model of execution presented above. A query was defined to be effectively computable if its evaluation terminated after computing all answers, and only constructed finite intermediate relations. For the bottom-up model of execution, the intermediate relations are of two kinds - the derived predicates in the program, and the temporary relations created in the application of a rule (to evaluate the joins of body predicates). From Lemma 3.1, given our three assumptions about programs, we know that the temporary relations created in any rule application are finite. Thus, to show a query to be effectively computable in our model of execution, we only need to show that every derived predicate is safe, and that the computation terminates after a finite number of rule applications.

We have the following theorem characterizing strongly safe queries.

Theorem 3.2: If a query is strongly safe, then it is effectively computable according to the above model of execution, given the three assumptions about programs. []

To summarize, our approach is to test for strong safety on a clique by clique basis. Theorem 3.2 then ensures effective computability according to our model of execution given our three assumptions about programs.

4. An Introduction to Testing Recursive Cliques for Strong Safety

We now introduce some technical definitions used in the safety analysis, and through illustrative examples on recursive cliques containing a single rule, we outline the safety check for such simple recursion.

Consider a property *prop*, defined over sequences, and a predicate p . Consider a path in a derivation tree. We say that this path satisfies property *prop* over predicate p if the sequence of facts of the form $p()$ on this path satisfy property *prop*. We also say

that the path satisfies $(p, prop)$, for brevity. For example, $prop$ could be defined over a sequence of k -ary facts to mean that for any pair of facts, if fact 1 precedes fact 2 in the sequence, the i th argument of fact 1 is $<$ the i th argument of fact 2, for some i . If the i th argument in any fact is an integer, $<$ could simply be the usual ordering defined over integers. As another example, $prop$ could be the property that the i th argument in any fact in a sequence of k -ary facts should be a positive integer less than 10. For a third example, $prop$ could be defined as the conjunction of the previous two properties.

A property is *well-founded* if the following holds. Let S be a sequence which satisfies the property, and let D be the domain of elements in S . Then: (1) We can define an ordering \ll such that $ei \ll ei+1$, for all adjacent elements $ei, ei+1$ in the sequence S . (2) For any $c1, c2$ in D , if $c1 \ll c2$, then there are only a finite number of values c in D such that $c1 \ll c \ll c2$. (3) There is an element c in D such that $c \ll c1$ for all $c1$ in D .

A large class of well-founded properties can be specified by simply specifying an argument position, an ordering, and a bound. In fact, all examples in this paper (with the exception of Example 4.2) use only a subset of this class, with a single ordering of values. The consideration of more sophisticated orderings, possibly over combinations of arguments, or the use of several different orderings in showing safety of a single program, present no difficulties.

Given a rule sequence, we say that the rule sequence satisfies $(p, prop)$ if every path in the set of paths corresponding to this rule sequence satisfies $(p, prop)$.

We say that a variable X is *bounded* if it can only take on values from a finite domain. An argument position in a literal occurrence is *bounded* if it contains a constant or the variable in it is bounded. Similarly, we say that X is *bounded above* (resp. *below*) if it can only take on values from a domain that is bounded above (resp. below). A cycle in an argument mapping is bounded above (below) if it contains at least one node that is bounded above (below).

Consider a cyclic (composite) argument mapping $(p, q1, r).(q1, q2, r1) \dots (qn, p, rn)$. The corresponding rule sequence is $r, r1, \dots, rn$. A cycle in an argument mapping is a cycle in the corresponding graph when directions of arcs are ignored. An *increasing* (resp. *decreasing*) cycle in the argument mapping is a cycle such that applying the sequence of rules $rn, \dots, r1, r$ assigns *increasing* (resp. *decreasing*) values to the nodes on the cycle in the argument mapping. (Clearly, such a cycle must contain at least one arc, and it cannot contain two oppositely directed arcs. Our choice of directions for *increasing* and *decreasing* cycles reflects our bottom-up application of rules.)

Note that every path corresponding to an increasing (resp. decreasing) cycle satisfies a simple increasing (resp. decreasing) property. Further, if the cycle is bounded, then the property is well-founded. Consequently, any single rule clique is safe if a well-founded property is associated with it. These observations are made precise in Section 5, but we illustrate the ideas in the following example.

Example 4.1 Consider the following program:

1. $p(X, Y) :- p(U, V), X=U-1, Y=V+1, b(X)$
2. $q(X, Y) :- q(U, V), X=U-1, Y=V+1, b(X)$
3. $p(X, Y) :- c(X, Y)$
4. $q(X, Y) :- d(X, Y)$

There are two recursive cliques, say $C1$ and $C2$, containing q and p respectively. Consider the argument mapping corresponding to the rule cycle $(1,1)$ in $C2$: $p(X, Y), p(U, V)$. There is an arc from U to X , and one from V to Y . Thus, we have an increasing cycle and a decreasing cycle in this argument mapping. Further, X is bounded since it appears in a finite base relation b , and so the first cycle of rules can be applied only a finite number of times. Each rule application can only produce a finite number of values. The clique $C2$ is therefore strongly safe. A similar analysis shows $C1$ to be strongly safe as well. Now, let us add the following rule:

5. $p(X, Y) :- q(Y, X)$

This changes the clique structure by making $C2$ depend on $C1$. However, $C1$ and $C2$ are still the only cliques, and the previous analysis holds. Indeed, q is strongly safe, and the tuples it contributes to p may be thought of as being in a finite base relation. We now add a sixth rule to the program:

6. $q(U, V) :- p(U, V)$

This changes the clique structure drastically. The two cliques are merged into a single clique containing both p and q . In addition to the two rule cycles considered before, there is the rule cycle $(5,6)$, whose argument mapping is $(p(X, Y), q(Y, X), p(Y, X))$. The exchange of arguments enables repeated applications of this rule cycle, in conjunction with rule cycles $(1,1)$ and $(2,2)$, to produce an infinite number of tuples. \square

The purpose of this section has been to give the reader some intuition into the problem of testing recursive cliques for safety and effective computability. In the next section, we develop these ideas rigorously.

5. Testing Arbitrary Cliques for Strong Safety

In the previous section we considered how a single recursive rule cycle could be tested for strong safety. If there are two or more such cycles involving the predicates in a clique, the prob-

lem becomes more complex because these rule cycles could interact in such a way that the tuples produced by one of them invalidate the assumptions made in declaring one of the others strongly safe. (We encountered this problem when we added rule 6 to Example 4.1.)

Consider the argument mappings corresponding to two rule cycles:

1. The monotonic (increasing or decreasing) cycles considered in the two mappings may involve different argument positions. Thus, the application of one of these rule cycles may introduce arbitrary values into argument positions not on the monotonic cycle in its own argument mapping (but possibly on the monotonic cycle used to check one of the other rule cycles, thus invalidating that check).
2. The monotonic cycles considered in the two mappings may involve different orderings. Thus, even if the same argument positions are involved, a rule cycle may introduce arbitrary values into these positions with respect to any ordering other than that used in checking it.

Example 5.1 The following example illustrates the second point above:

1. $p(X,Y) :- p(U,V), X = U-1, Y = V+1, b(X)$
2. $p(X,Y) :- p(U,V), X = U+1, Y = V+1, b(X)$
3. $p(X,Y) :- c(X,Y)$
4. $c(1,11)$ 5. $c(2,12)$ 6. $b(1)$ 7. $b(2)$

The two simple rule cycles are (1,1) and (2,2). The first argument is part of a decreasing cycle in rule cycle (1,1), and part of an increasing cycle in rule cycle (2,2). In both, the first argument is bounded because X occurs in the finite base relation b . So each of these cycles is safe if it is the only rule cycle. However, taken together, they are unsafe. It is possible to alternate between the two cycles so that the first argument alternates between the values 1 and 2 (which are the only values in b), and on each application the second argument increases by 1. Thus, the second argument can be any integer greater than 10 (since the least value in c is 11). []

In order to deal with arbitrary cliques, we must analyze the interactions between the simple rule cycles generated by the rules in the clique. The following theorem underlies our approach.

Theorem 5.1: Consider a database $DB = (EDB, IDB, IC)$. If the height of all derivation trees is bounded for all EDBs, then every predicate is strongly safe. []

We present an algorithm to test safety, and prove its correctness by establishing that if the algorithm certifies a program, then all

derivation trees are of bounded height (Theorem 5.1). The details of the algorithm are as follows.

First, we identify all the simple rule cycles in the clique and generate the corresponding (cyclic) argument mappings. Then we test each cyclic mapping $(p, _ , _), \dots, (_ , p, _)$ as if it were the only cyclic mapping in the clique, using the results of the previous section. In doing this, we check that there is some property defined over the arguments of p , say $prop_p$, which is well-founded with respect to this argument mapping. Having done this for all mappings, we verify that the interaction of rule cycles does not invalidate the assumptions made in checking them individually. For example, if $(p, _ , _), \dots, (_ , p, _)$ is considered well-behaved because the first argument decreases on each application of this cycle of rules and is bounded below, we must check that applications of every other rule cycle either leaves the value of the first argument unchanged or decreases it.

We make use of two procedures called *check* and *tag_prop*. We now specify them formally.

check($C, prop$): Let cycle $C = (p, q_1, r_1), (q_1, q_2, r_2), \dots, (q_n, p, r_n)$. The procedure *check*($C, prop$) returns *true* only if the following holds: Every path $P = (p(), q_1(), q_2(), \dots, q_n(), p())$ in a derivation tree satisfies $(p, prop)$.

For example, we can show that the values in a given argument position (in a path of a derivation tree) are monotonically decreasing by showing that this argument appears in a decreasing cycle in the corresponding argument mapping. Often, the constraints given to hold on the base predicates suffice to show this. However, in general it is necessary to be able to infer that certain monotonicity constraints hold (between arguments of a derived predicate) from the constraints given to hold on the base predicates. This is a significant problem in its own right, and we do not address it in this paper.

tag_prop($C, prop, q_j, prop_j$): Let cycle $C = (p, q_1, r_1), (q_1, q_2, r_2), \dots, (q_n, p, r_n)$. Then, $prop_1, prop_2, \dots, prop_n$ are some n properties such that the following holds:

Let $Q_i = (q_i(), \dots, q_i())$, $i = 1$ to n , be any n paths such that Q_i satisfies $(q_i, prop_i)$, $i = 1$ to n . Let $P = (p(), q_1(), q_2(), \dots, q_n(), p())$ be any path in a derivation tree which satisfies $(p, prop)$. Then the path $(p(), Q_1(), Q_2(), \dots, Q_n(), p())$ also satisfies $(p, prop)$.

For example, if $prop$ is the property that some argument of p is increasing, then there is some increasing cycle in the argument mapping C which contains occurrences of this argument. If some argument of q_j appeared in this cycle, then $prop_j$ is the property that this argument of q_j should increase or remain the same. (Note that this property is *not* well-founded.)

We present a recursive procedure *check_clique* which examines the interactions between rule cycles. We prove that if this procedure returns *true*, then all paths in a derivation tree built using rules from the given clique are bounded in length, where the bound is a function of the DB. Thus, the number of derivation trees for facts in DB is bounded. (This implies that every predicate is safe, and it immediately follows that every predicate is strongly safe as well. Effective computability then follows from Theorem 3.2.)

We note that a given rule cycle may be used to expand any predicate which occurs on it, and so if the cycle contains two predicates p and q , the corresponding argument mapping can be denoted as either $(p, _ , _), \dots , (_ , p, _)$ or $(q, _ , _), \dots , (_ , q, _)$. While the fact that they refer to the same rule cycle should be recognized and utilized in testing whether these mappings satisfy some property, for the purposes of the subsequent analysis it is convenient to assume that both representations are generated. Thus, by referring to all simple cycles of the form (p, \dots , p) we are assured that indeed all simple rule cycles containing p are considered.

The following procedure checks an arbitrary clique for strong safety.

proc check_clique: boolean

1. For each predicate pi in the given clique, associate a well-founded property $prop_i$.
2. FLAG := true; $i = 1$;
while FLAG and $i \leq n$ **do** /* $n = \#$ of predicates in C */
 FLAG := FLAG & verify(pi , { pi }, $prop_i$);
 $i := i + 1$
od.
3. return FLAG.

end. /* *check_clique* */

proc verify(p , IGNORE, prop): boolean

/* Does the detailed work. Checks if $prop$ is a well-founded property for each simple composite argument mapping $(p, _ , _), \dots , (_ , p, _)$. To do this, it must check for interactions with other simple argument mappings. IGNORE is a set of predicates. In testing for interactions, we ignore all argument mappings in which some predicate in IGNORE appears at some position other than the endpoints. This is required for the algorithm to terminate, and is discussed later. */

/* Assumes that programs *check* and *tag_prop* are given. */

1. Let $\{C_1, \dots, C_m\}$ be the set of simple cyclic argument mappings of the form $(p, _ , _), \dots , (_ , p, _)$. We first verify that $prop$ is a well-founded property for each of these cycles considered separately.

FLAG := true;

for all i **in** $\{1, \dots, m\}$ **do**

 /* C_i' is the set of predicates on $C_i =$

$(p, _ , _), \dots , (_ , p, _)$, except for p . */

if $C_i' \cap \text{IGNORE} = \{\}$ **then**

 /* *check* verifies that $(p, prop)$ is preserved by the

 predicates on C_i if they are not expanded further. */

 FLAG := FLAG & check(C_i , prop);

 Let q_1, \dots, q_k be the recursive predicates in C_i' .

for all j **in** $\{1, \dots, k\}$ **do**

 tag_prop(C_i , prop, q_j , prop $_j$);

od

od

2. Now, for each cycle C_i , we ensure that the other cycles preserve the well-foundedness of $prop$. (That is, when these cycles are used to expand predicates on cycles for p .)

for all i **in** $\{1, \dots, m\}$ **do**

if $C_i' \cap \text{IGNORE} = \{\}$ **then**

 Let q_1, \dots, q_k be the recursive predicates in C_i' .

for all j **in** $\{1, \dots, k\}$ **do**

 IGNOREQ := IGNORE $\cup \{q_j\}$;

 /* $prop_j$ is the property with which q_j is "tagged"

 by tag_prop(C_i , prop, q_j , prop $_j$) in Step 1. */

 FLAG := FLAG & verify(q_j , IGNOREQ, prop $_j$)

od

od

3. return (FLAG).

end /* *verify* */

We now present our main theorem.

Theorem 5.2: If *check_clique* returns *true*, then there is a constant $L_{(C, DB)}$ such that the longest path in any derivation tree is less than or equal to $L_{(C, DB)}$. []

We prove that all paths to a leaf in a derivation tree are bounded in length if *check_clique* returns *true*. We argue by induction on the number of recursive predicates which appear on such a path. The following lemma is central to the proof.

Lemma 5.3: If *check_clique* returns *true*, then every segment between two occurrences of p in a path of a derivation tree preserves the well-founded property associated with p . []

Lemmas 5.4 and 5.5 are technical lemmas used to establish Lemma 5.3, and are omitted from this abstract. The following corollary of Theorem 5.2 summarizes our approach.

Corollary 5.6: Consider a clique C in a database $DB = (EDB, IDB, IC)$. If *check_clique* returns *true*, then every predicate in the clique is strongly safe. []

We now consider some limitations of Theorem 5.2 and show how it can be strengthened to overcome them.

Example 5.2: Theorem 5.2 sometimes fails to determine that certain well-founded properties hold, and we present two illustrative programs:

1. $p(X, Y) :- q1(X, Y), q2(X, Y)$
2. $q1(X, Y) :- b(X, Y)$
3. $q2(X, Y) :- f(X, U), p(U, Y)$

b is a finite base predicate and f is an infinite base predicate. The first argument of p is bounded since it can only take values in b . However, *check_clique* considers whether this property is preserved by the rule cycle (1,3) corresponding to the argument mapping $(p, q2, p)$. Since this cycle considered alone does not preserve boundedness of the first argument of p , using Theorem 5.2 we cannot show that the first argument of p is bounded. \square

The previous example brought out a limitation of Theorem 5.2. The intuition behind this can be explained as follows. The theorem can establish that no derivation tree constructed using the rules in the clique contains an infinite path, which is sufficient to ensure that no rule can be applied an unbounded number of times. This in turn ensures that the predicate labeling the root is safe. However, a weaker condition is sufficient to establish safety: in every derivation tree, there is a set of leaves such that they collectively contain all values which appear in the fact labeling the root, and such that the paths from these leaves to the root are finite. (Recall that we establish strong safety by establishing the safety of every predicate. That is, we consider all derivation trees.) We now use this intuition to present a stronger form of Theorem 5.2. This stronger version can easily be used to refine algorithm *check_clique*, as we later explain.

Let us define the head of a rule to be *covered* by the body if every variable in the head also appears in the body.

Theorem 5.6: Let P be a program containing a rule r , which is of the form: $p() :- q1(), q2(), \dots, qn()$. Therefore, there is an argument mapping (p, qi, r) generated from each literal occurrence qi in the body. Let *prop* be a well-founded property associated with p .

IF there is some subset of body literals which covers the head, and further is such that for every literal $qi()$ in this subset, every simple cycle $(p, qi, r), \dots, (_, p, _)$ satisfies property *prop*,

THEN we may ignore all simple cycles $(p, qj, r), \dots, (_, p, _)$, qj not in the subset, in verifying whether *prop* is preserved. \square

Using the above result, procedure *check_clique* can be strengthened to deal with the problem identified in Example 5.2. In fact, the improved version is also considerably more efficient. For example, consider the case when the body literals covering the head are non-recursive. The safety test then becomes trivial,

since, given our three assumptions about programs, non-recursive programs are always safe and effectively computable.[†]

We would like to emphasize the extensible nature of this approach. As we noted earlier, two programs, *check* and *tag_prop* are assumed to be given as inputs to *check_clique*, and using more sophisticated versions of these programs that can deal with a richer class of well-founded orderings clearly increases the scope of *check_clique*. Also, we have not addressed the issue of how to infer monotonicity constraints. Some work in this direction has been presented in [Ullman and Van Gelder 85], and we can take advantage of those (and other, similar) results to infer constraints. A more subtle point is that we need to identify bounded arguments to check whether orderings are well-founded. There are some obvious ways of doing this. For example, the literal $X < 10$ in a rule tells us that X is bounded below, and $b(X)$, where b is a finite base relation, tells us that X is bounded (equivalently, safe). Recall that the methods presented in this paper focus on *strong safety*. Other methods for inferring safety, for example, those presented in [Ramakrishnan et al. 87] for inferring safety using only FCs, can be used in our analysis to identify arguments that are bounded.

6. Related Work

Some results relating to these problems have been presented in the context of testing the applicability *top-down capture rules* [Sagiv and Ullman 84, Ullman and Van Gelder 85, Afrati et al. 86]. These problems were also addressed in [Ramakrishnan et al. 87]. In this paper, we make use of, and extend, this earlier work.

This work can be seen as providing a framework for testing safety and effective computability of recursive programs, in some ways analogous to the *capture rules* framework of Ullman [Ullman 85]. The capture rules framework considers the program represented as a *rule/goal graph*, and allows the design of *capture rules* which describe how nodes in the graph can be "captured". Intuitively, a node can be captured if an effective procedure can be constructed to compute the predicate (or rule) denoted by that node. A capture rule states that a given node may be captured if certain other nodes in the graph have already been captured, and is *substantiated* by a procedure which computes the node given procedures for capturing these other nodes. Thus, if the query node in the rule/goal graph can be captured, we know that it is safe and effectively computable.

Thus, our framework is similar to the capture rules framework in the following sense: if a query is declared to be safe and effectively computable (analogous to "captured"), there is a known

[†] The refinement is not discussed in this abstract.

way to effectively compute it. The analogy goes further. Both frameworks must be supplemented by specific procedures in practice. In the capture rules framework, these procedures are the list of known capture rules, and auxiliary procedures for inferring monotonicity constraints etc. [Ullman and Van Gelder 85]. In our framework, these procedures include algorithm *check_clique*, the two auxiliary procedures called *check_clique*, and other procedures for inferring monotonicity constraints and *boundedness* of predicates.

We now examine the relative merits of the two approaches briefly. Different capture rules can be expressed in terms of a set of sideways information passing strategies (sips) which are implemented by rewriting the program [Beeri and Ramakrishnan 87]. The applicability of a capture rule must be shown by presenting a *substantiation*, that is, a program which computes the (sub)query according to the capture rule. The bottom-up evaluation of the rewritten program provides a potential *substantiation*. The correctness of the rewriting transformations implies that the set of answers is correctly computed by the bottom-up evaluation, but we must still establish safety of all intermediate results and termination. This is the focus of the present paper. Thus, in our framework, the range of sideways information passing strategies (sips) describes the set of possible capture rules, and there is a single procedure (based on *check_clique*) which tests for the applicability of all capture rules by testing the rewritten program for safety and effective computability. Further, the substantiation is always a program which is evaluated bottom-up. In contrast, in the capture rules framework, each new capture rule must be explicitly represented along with its substantiation. Testing applicability is specific to a given capture rule, and the substantiation need not always be a program which is evaluated bottom-up.

As we remarked earlier, the results here can be used to show the “applicability” of capture rules in the literature, and can also be used to show the safety and effective computability of some programs which cannot be “captured” using known capture rules.

We now consider the testing of programs produced by rewriting algorithms like the Magic Sets algorithm. The Magic Sets strategy [Bancilhon et al. 86] restricts computation by rewriting the program so that it only computes facts relevant to the query. The original version achieved this restriction only for certain kinds of rules (essentially, linear recursive rules). An extended version, Generalized Magic Sets, which is based on the notion of *sideways information passing graphs*, achieves this restriction for all programs in which variables that appear in the head of a rule also appear in the body [Beeri and Ramakrishnan 87].

Conceptually, we first rewrite the program and then test it for safety and effective computability. A potential problem is that

the rewritten program typically has much larger cliques than the original program since the rewriting introduces rules which often combine cliques in the original program. Thus, even if the original program has relatively small cliques (as we expect), it is likely that the rewritten program has large cliques. Given the complexity of algorithm *check_clique*, the cost of directly analyzing the rewritten program could be prohibitive. However, we can take advantage of the structure of the transformation to efficiently test rewritten programs for safety and effective computability. The *adorned program* is an intermediate program produced by the Magic sets algorithm, and has a clique structure which is similar to that of the given program. We show how the rewritten program can be analyzed by considering the cliques in the adorned program, and without directly considering the cliques in the rewritten program [Krishnamurthy et al. 87]. We have not presented these results here due to space constraints, but they suggest that the framework here is robust and can provide the basis for a practical system.

In this paper, we do not consider the problem of how to establish that a node is bounded. That is, our algorithms test for safety in conjunction with effective computability, but they cannot be used to determine if a query is simply safe (but, possibly, not effectively computable). Our results clearly provide sufficient conditions for safety, but it is possible to devise weaker conditions if only safety is of interest. (We note that some of the results in the literature for detecting safety, for example [Ramakrishnan et al. 87, Zaniolo 85, Van Gelder and Topor 87], could be used for this purpose.)

Finally, we refer to some interesting work by Kifer and Lozinskii [Kifer and Lozinskii 88]. They present sufficient conditions for safety in terms of a formalism called *term circulation graphs*. We conjecture that their result can be incorporated in our framework as a basis for writing the procedures *check* and *tag_prop*. We intend to examine the relationship between the two approaches more carefully.

7. Conclusions

The important contribution of this paper is an algorithm (*check_clique*) for checking safety and effective computability in a bottom-up model of execution of an arbitrary recursive clique using monotonicity and finiteness constraints. This algorithm provides the central component of a framework for testing programs. The framework makes the testing of large programs tractable by organizing programs into cliques and permitting a clique by clique analysis. The *check_clique* algorithm is parametrized by procedures which are assumed to be available, such as *check*, *tag_prop*, and auxiliary procedures for inferring boundedness and monotonicity constraints. The procedures

check and *tag_prop* represent abstractions of operations that are implicit in any safety analysis involving monotonicity constraints, and in this paper they are isolated and specified rigorously. By providing more sophisticated versions of these procedures, the results presented here can be used to show safety and effective computability of a larger class of programs.

Further work is needed in inferring monotonicity constraints, possibly interacting with the programmer. There are two aspects to this problem. In general, if we cannot show a clique to be strongly safe, it is not clear whether additional constraints will help us to do so, and if so, exactly what these additional constraints are. Also, once we identify a potentially useful constraint, we need to determine whether this constraint holds. A first step in these directions is taken in [Ullman and Van Gelder 85]. We also need to explore well-founded orderings that can be used in testing the safety of cliques, in order to find orderings that are easy to test, and widely applicable. Finally, we emphasize that the framework presented here can be used to test safety and effective computability under various sideways information passing strategies by suitably rewriting the program and then analyzing the rewritten program.

8. Acknowledgements

We thank Michael Kifer and the referees for useful comments.

9. References

[Afrati et al. 86] "Convergence of Sideways Query Evaluation", F. Afrati, C. Papadimitriou, G. Papageorgiou, A. Roussou, Y. Sagiv and J.D. Ullman, *Proc. 5th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1986*, pp 24-30.

[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*, pp 1-16.

[Bancilhon and Ramakrishnan 86] "An Amateur's Introduction to Recursive Query Processing Strategies", F. Bancilhon and R. Ramakrishnan, *Proc. SIGMOD 86*, pp 16-52.

[Bancilhon and Ramakrishnan 87] "Performance Evaluation of Data Intensive Logic Programs", F. Bancilhon and R. Ramakrishnan, *To appear in Foundations of Deductive Databases and Logic Programming, Ed. J. Minker, Morgan Kaufman*.

[Beerl and Ramakrishnan 87] "On the Power of Magic", C. Beerl and R. Ramakrishnan, *Proc. 6th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1987*, pp 269-283.

[Gaifman 86] *Personal communication, reported by Y. Sagiv*.

[Kifer and Lozinskii 88] "SYGRAF: Implementing Logic Programs in a Database Style", M. Kifer and E.L. Lozinskii, *To appear, IEEE Transactions on Software Engineering, 1988*.

[Krishnamurthy et al. 87] "Testing Safety and Effective Computability of Magic Programs", R. Krishnamurthy, R. Ramakrishnan and O.

Shmueli, *Manuscript*.

[Lloyd 84] "Foundations of Logic Programming," J. W. Lloyd, *Springer-Verlag, 1984*.

[Ramakrishnan et al. 87] "Safety of Horn Clauses with Infinite Relations", R. Ramakrishnan, F. Bancilhon and A. Silberschatz, *Proc. 6th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1987*, pp 328-339.

[Sagiv and Ullman 84] "Complexity of a Top-Down Capture Rule" Y. Sagiv and J.D. Ullman, *STAN-CS-84-1009, Department of Computer Science, Stanford University, 1984*.

[Shmueli 87] "Decidability and Expressiveness Aspects of Logic Queries", O. Shmueli, *Proc. 6th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1987*, pp 237-249.

[Ullman 85] "Implementation of Logical Query Languages for Databases," J. Ullman, *TODS, Vol. 10, No. 3, pp. 289-321, 1985*.

[Ullman and Van Gelder 85] "Testing Applicability of Top-Down Capture Rules", J.D. Ullman and A. Van Gelder, *Technical Report, Stanford University, STAN-CS-85-1046, 1985*.

[Van Gelder and Topor 87] "Safety and Correct Translation of Relational Calculus Formulas", A. Van Gelder and R.W. Topor, *Proc. 6th ACM SIGMOD-SIGACT Symposium on Principles of Database Systems, 1987*, pp 313-327.

[Zaniolo 86] "Safety and Compilation of Non-Recursive Horn Clauses", C. Zaniolo, *Proc. First Intl. Conference on Expert Database Systems, Charleston, 1986*.