Magic Sets with Full Sharing

Paulo J Azevedo
Departamento de Informatica
Universidade do Minho
Campus de Gualtar,
4710 Braga, Portugal
pja@di.uminho.pt

Abstract

This paper discusses the relationship between tabulation and goal-oriented bottom-up evaluation of logic programs. Some differences emerge when one tries to identify features of one evaluation method in the other. We show that to obtain the same effect as tabulation in top-down, one has to perform a careful *adornment* in programs to be evaluated bottom-up. Furthermore we propose an efficient algorithm to perform subsumption checking over adorned *magic facts*. Soundness and completeness of the subsumption algorithm are proved. With the aim of substantiating the claimed improvements yield by this proposal several program evaluations are presented.

1 Introduction

Much has been said about the relations between goal oriented bottom-up and tabulated top-down evaluation e.g. [Warren92, Tamaki&Sato86, Ullman89, Ramak91]. One example of these relations is the equivalence between magic facts of bottom-up and subgoals in top-down. Another is the equivalence between the facts that can be derived by a specific magic fact and the stored solutions in the tables for a specific subgoal. The order in which magic facts are derived is also commonly referred to in the literature as the "order of subgoals evaluation". One important issue is the relations identified by Seki in [Seki89]. The semi-naive procedure, [Balbin&Ramam87], includes subsumption checking to prevent the derivation of duplicate facts. Seki observed that the subsumption checking in semi-naive has a counterpart in tabulated top-down evaluation in two ways. First, when subsumption is applied in magic facts it corresponds to the subsumption test of tabulation (admissibility test of SLD-AL). Secondly, subsumption applied in facts (non magic) derived during bottom-up evaluation corresponds to the duplicate elimination performed in tabulation when a new solution is inserted in the tables. We investigate the use of subsumption, as described in the first case, to eliminate redundancy in the derivation of magic facts.

This paper is concerned with the recomputation that arises in the bottom-up evaluation of magic rewritten programs [Beeri&Ramak91]. The starting point is to observe what features of tabulation appear in goal-oriented bottom-up evaluation. We will show that these features do not appear entirely. Namely, we observed that a magic atom representative of a subgoal that subsume another subgoal does not necessarily subsume the magic fact representative of the subsumed subgoal. Since the adornment process yields syntactically different variants of the same predicate, the traditional implementation of subsumption cannot cope with adorned magic atoms. Consequently, the derivation of facts

triggered by the magic fact $mag_p^{bb}(a,a)$ is repeated by the derivation of facts triggered by $mag_p^{bf}(a)$. Thus, goal-oriented bottom-up evaluation does not exhibit the full benefits of tabulation. We propose forms of overcoming this fault in Magic Sets by introducing two new techniques. First we suggest a different way of dealing with adornments. Secondly, we propose a new subsumption checking algorithm for detecting redundancy among adorned magic facts.

Throughout the paper, since we will deal with Magic Sets, the language will be Datalog and the described rules will be range restricted. Also, evaluation will be always performed by the semi-naive method since it is considered the standard bottom-up strategy [Naughton&Ramak91]. In this paper we will interchangeably refer to the term subgoal meaning magic fact and to order of evaluation meaning order in which magic facts are derived.

The paper is organized as follows: Section 2 describes the idea of Magic Sets and the role of adornments, section 3 shows the relations between tabulation and bottom-up evaluation and section 4 arises the problem of full sharing of answers between adorned magic facts. In section 5 the algorithm for subsumption checking between adorned atoms is introduced. Some examples comparing performance between using our proposals for obtaining sharing of answers and the standard magic sets are presented in section 6. Finally, section 8 concludes the paper.

2 Adornments in Magic Sets

One advantage of the bottom-up methods over top-down methods is that computation is guaranteed to *terminate* for finite model programs i.e. programs with finite minimal Herbrand models. However, one disadvantage is that bottom-up methods compute too many irrelevant facts in relation to the query to be answered. Bottom-up strategies do not consider the query during evaluation i.e. computation is not driven by the query, as happens in top-down methods.

Several approaches exist to include goal orientation in bottom-up evaluation of which the best known is the Magic Sets method [Bancilhon et al. 86]. The basic idea behind these methods is to simulate the propagation of bindings that would be generated in the top-down evaluation of the query. This propagation is achieved by rewriting the original program. In each original rule a new condition is added to restrict the application of the rule. These are known as *filter relations*. A new set of rules is created to simulate the propagation of bindings. In general, given a query Q and a database $DD =< IDB; EDB >$ we transform $Q \cup IDB$ into IDB^* to reflect the propagation of bindings.

Let us take the following example to illustrate how Magic Sets work. Consider the query ?p(a) and the rule p(X) $\leftarrow$ q(X). To apply magic sets rewriting one adds to the original program a seed that represents the query. In our case we add the seed mag_p(a). This new relation (mag_p) works as a filter relation. Thus, and as described before, we add to each original rule a filter relation to restrict the derivation of tuples according to the requirements expressed in the query. In our case we get the new rule p(X) $\leftarrow$ q(X) & mag_p(X). To simulate the propagation of bindings one generates new rules (magic rules), according to the literals in the body of the original rules. For instance, suppose that q is an IDB defined predicate. Then, from our rule one would generate the following magic rule for q: mag_q(X) $\leftarrow$ mag_p(X). The generated rule derives the subgoal ?q(a)

that would be generated if the original rule with the query was processed top-down. Note that the magic rule propagates the bindings in the query to the 'subgoal' ?q. Suppose the rule defining q is $\mathsf{q(X)} \leftarrow \mathsf{s(X)}$. The transformed rule is $\mathsf{q(X)} \leftarrow \mathsf{s(X)}$ & $\mathsf{mag_q(X)}$. In this sense, the q tuples are generated according to the bindings in the original query which are further propagated to the rule defining q. The magic rules are said to *materialize* the corresponding subgoals that would be generated if the original program were processed top-down. The magic relations generated during the bottom-up evaluation are known as *Magic Sets*. The general recipe for the magic rewriting can be summarized as follows. Given a query $?q(\vec{c})$ and a database $P =< IDB; EDB >$ generate the magic version $Mag(P)$:

- For each rule in IDB of P, add the modified version to $Mag(P)$. If the rule has head $p(\vec{t})$ then the modified version is obtained by adding the literal $mag_p(\vec{t})$ to the body.
- For each rule in IDB of P with head $p(\vec{t})$, and for each literal $q_i(\vec{t_i})$ in its body, where q_i is an IDB-defined predicate, add a magic rule to $Mag(P)$. The rule's head is $mag_q_i(\vec{t_i})$ and the body contains the literal $mag_p(\vec{t})$ and all the literals that precede q_i in the original rule.
- create a seed fact $mag_q(\vec{c})$ to represent the query.

For the following example, where s and r are EDB defined predicates:

$\mathsf{p(X,Y)} \leftarrow \mathsf{s(X)}$ & $\mathsf{q(X,Y)}$
$\mathsf{q(X,Y)} \leftarrow \mathsf{r(X,Y)}$

and the query $?\mathsf{p(a,Y)}$, the magic rewriting generates the program:

$\mathsf{mag_p(a,Y)}$
$\mathsf{p(X,Y)} \leftarrow \mathsf{mag_p(X,Y)}$ & $\mathsf{s(X)}$ & $\mathsf{q(X,Y)}$
$\mathsf{q(X,Y)} \leftarrow \mathsf{mag_q(X,Y)}$ & $\mathsf{r(X,Y)}$
$\mathsf{mag_q(X,Y)} \leftarrow \mathsf{mag_p(X,Y)}$ & $\mathsf{s(X)}$

One can see that in this magic program some redundant passing of bindings occurs. For instance, the Y variable in the rule defining mag_q is redundant since this variable is not bound by any constant. Thus, this parameter is redundant and could be eliminated from the magic predicate. The same happens with the seed representing the original query. In general, one disadvantage of this rewriting is that range restriction is not preserved i.e. programs that were originally range restricted lose this property after the rewriting and consequently tuples with variables will be generated during the evaluation of the program. Another point to be made about this program is that a fixed strategy for passing bindings was chosen. If Magic Sets are concerned with implementing the passing of bindings during bottom-up as happens in top-down evaluation then one should be able to specify different strategies of passing the bindings. This corresponds to a computation rule in top-down evaluation. However, as the rewriting process occurs at compile time, our strategy must be static contrasting with a computation rule that can be dynamic. Note that this strategy will influence the way in which the magic rules are constructed, since these rules implement the passing of bindings.

In the Magic Sets method, it is assumed that all rules are range restricted. Moreover, each rule has assigned a sideways information passing (SIP) strategy. This strategy represents a decision about the order in which the conditions of the rule will be evaluated and how values for variables are passed from conditions to other conditions during evaluation. In other words, a SIP describes how we evaluate a rule when a given set of head arguments is bound. There are two techniques for the implementation of these strategies. One is the generation of magic rules. The other is the *adornment process*, where through a set of strings a representation of the expected pattern is attached to each literal. Adornments also ensure that range restriction is preserved.

Let us consider the definition of adornments following [Beeri&Ramak91]. As is known, in the Magic Sets method each rule has a sideways information passing strategy (SIP) assigned to it. This strategy represents a decision about the order in which the conditions of the rule will be evaluated and how values for variables are passed from one condition to another during evaluation. The specification of the SIP corresponds to specifying the computation rule in top-down evaluation. An *adornment* is a string from the alphabet {b,f} that represents the expected pattern of bound (b) and free (f) variables in the arguments of a predicate. Intuitively, an adorned occurrence of a predicate corresponds to a computation of that predicate with some arguments bound to constants and other arguments free. For instance, $\mathsf{p}^{\mathsf{bff}}$ corresponds to computing p with the first argument bound and the other two free. Notice that each SIP implicitly determines a pattern of bound/unbound arguments for each predicate to be evaluated. The task of the adornment process is to make this implicit pattern explicit.

Consider the following example where the SIP strategy coincides with left-to-right Prolog strategy for passing bindings. Consider a rule: $\mathsf{p}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{s}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{t}(\mathsf{Z},\mathsf{Y})$. The call is $?\mathsf{p}(\mathsf{a},\mathsf{X})$, corresponding to the adornment p^{bf}. If p has adornment {bf} then s has the adornment {bf} because it receives the variable X bound from the head. Also t has adornment {bf} since it receives Z from predicate s. The adorned version of the rule is then: $\mathsf{p}^{\mathsf{bf}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{s}^{\mathsf{bf}}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{t}^{\mathsf{bf}}(\mathsf{Z},\mathsf{Y})$. Now, if there is no rule for s with adornment {bf} then the adornment process must also be applied to the rules defining s. Of course, different SIPs would determine different adornments for the same rule. The definition of magic rewriting with adornments as in [Beeri&Ramak91] is the following: Let P^{ad} be the adorned version of program (database) P following a given SIP-strategy (we still assume the same left-to-right strategy) and a query $q(\vec{x})$.

1. create a new predicate $mag_p(\vec{t^b})$ for each $p(\vec{t})$ in P^{ad}; $\vec{t^b}$ means the bound arguments of $\vec{t}$.

2. for each rule in P^{ad} add the modified rule to P^{magic} which is the original rule with the body extended with the literal $mag_p(\vec{t^b})$ if the head is $p(\vec{t})$ (i.e. only the bound (b) arguments are in the magic literal).

3. For each rule $p(\vec{t}) \leftarrow q_1(\vec{t_1})\ \&\ \ldots\ \&\ q_n(\vec{t_n})$ in P^{ad} generate several magic rules: $mag_q_i(\vec{t_i^b}) \leftarrow mag_p(\vec{t^b})\ \&\ q_1(\vec{t_1})\ \&\ \ldots\ \&\ q_{i-1}(\vec{t}_{i-1})$ is added to P^{magic} for each $1 \leq i \leq n$ and the order of i respects the order on the SIP,

4. add the seed fact $mag_q(\vec{x^b})$ representing the query $q(\vec{x})$.

As previously identified in the literature, e.g. [Seki89, Warren92, Naughton&Ramak91, Ramak91], one consequence of this rewriting is that calls in top-down are represented in

bottom-up by magic facts. Considering what adornments represent and following step 3 of the magic rewriting, the head of a magic rule has the same adornment as the literal that gives rise to the magic rule. Consequently, the adornments in a magic fact represent the pattern of bound and free arguments of a call. The example described before is now:

$\mathsf{mag_p}^{\mathsf{bf}}(\mathsf{a})$
$\mathsf{p}^{\mathsf{bf}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{mag_p}^{\mathsf{bf}}(\mathsf{X})\ \&\ \mathsf{s}(\mathsf{X})\ \&\ \mathsf{q}^{\mathsf{bf}}(\mathsf{X},\mathsf{Y})$
$\mathsf{q}^{\mathsf{bf}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{mag_q}^{\mathsf{bf}}(\mathsf{X})\ \&\ \mathsf{r}(\mathsf{X},\mathsf{Y})$
$\mathsf{mag_q}^{\mathsf{bf}}(\mathsf{X}) \leftarrow \mathsf{mag_p}^{\mathsf{bf}}(\mathsf{X})\ \&\ \mathsf{s}(\mathsf{X})$

3 Tabulation & Bottom-Up Evaluation

Elimination of redundant computation is a subject that has been studied by several authors. Redundancy arises due to the repetition of the same call to a procedure in a program. Recursive programs are a common example where this phenomenon appears. In the special case of logic programming, several proposals exist in the literature which address this issue, e.g. [Tamaki&Sato86, Vieille89]. The basic idea common to all these proposals, known as *tabulation techniques*, is to eliminate subquery redundancy by storing intermediate results that are subsequently used by identical queries. One can see this technique as a form of bottom-up programming [Warren92, Tamaki&Sato86], since some results are preserved along the evaluation.

The way tabulation is approached in logic programming can be summarized in the following way. Certain predicates are specified to be tabulated i.e. both their answers and calls are to be stored during computation. Whenever a call to one of these predicates arises, one of two processes occurs. If the call is an instance of a previously computed one then the latter is answered by retrieving answers from the tables. Otherwise, a new table is constructed for this call to be answered and, as answers are obtained from the computation in the program they are stored in the constructed table. Future calls that are instances of the stored one reuse the answers in the table. We determine that a call is an instance of another through subsumption.

Definition 1 *Let Q_1 and Q_2 be atoms. We say that $Q_1 \sqsupseteq Q_2$ (reads Q_1* **subsumes** *Q_2) if there exists a substitution θ such that $Q_1\theta = Q_2$. The subsuming atom (Q_1) is the more general one and the subsumed atom (Q_2) the more specific.*

In the case of tabulation, the subsumption test checks whether a literal is an instance of another literal. Subsumption checking (also referred to as admissibility checking [Vieille87]) between calls determines whether a call is answered from the table (is redundant in relation to a stored one) or from the program.

The matter of tabulation is also important for ensuring termination. In fact, the use of tabulation yields that the classes of programs where tabulated top-down proof procedures (with subgoal generalization [Vieille89]) and bottom-up procedures terminate coincide [Tamaki&Sato86].

Several authors observe that these described features of tabulated top-down proof procedures also appear in goal-oriented bottom-up evaluaters. In [Seki89], more detailed relations between these two forms of computation are put forward. The author identifies relations between SLD-AL [Vieille87], a tabulated proof procedure, and the Alexander

Templates rewriting which is a variant of Supplementary Magic Sets without adornments. Magic predicates are *call* predicates in Alexander templates and derived facts are *sol* facts. Supplementary Magic Sets [Beeri&Ramak91], avoid redundant joins by deriving supplementary relations. Seki establishes the relation between the admissibility test and the subsumption checking in a bottom-up evaluation. The latter subscribes the need in bottom-up evaluation (for instance in the semi-naive strategy) to check whether a newly derived fact is subsumed by a previously derived one. This subsumption checking can be reduced to simple duplicate elimination if only ground facts are derived. If a goal $\leftarrow \mathsf{q}$ is admissible (is not subsumed by a call stored in the tables) then correspondingly the subsumption checking in bottom-up determines that the fact $\mathsf{mag_q}$ (or $\mathsf{call_q}$ in Alexander templates) is a newly derived one. On the other hand a newly derived lemma L in SLD-AL corresponds to a newly derived fact in Alexander Templates sol_L or simply L in Magic Sets.

The conclusion that one should draw from these two remarks is that subsumption checking in bottom-up has a counterpart in top-down in two ways. First, in the admissibility test on calls and second in the duplicate elimination performed on lemmas. However, one should notice that the introduction of adornments can corrupt these relations. Consider the case where the fact $\mathsf{mag_q}$ was previously derived and it is 'compared' with the fact $\mathsf{mag_q^b}$. Since syntactically they are unrelated, one cannot establish any subsumption relation between the two magic facts.

It is interesting to notice that tabulation systems like XSB [Sagonas et al. 94] do not incorporate a subsumption checking mechanism but rather perform variant checking[1]. For reasons related to the way the answers to a call (and the stored call) are indexed, XSB uses a much simpler method to eliminate redundancy. The price to pay is that not all recomputation is eliminated. For instance if the call $?\mathsf{p}(\mathsf{a},\mathsf{Y})$ is stored then only variants of this call e.g. $?\mathsf{p}(\mathsf{a},\mathsf{Z})$ are considered as having their answers in the tables. Thus if a call that is an instance of our stored one is derived e.g. $?\mathsf{p}(\mathsf{a},\mathsf{b})$ it is not identified as having answers in the table and consequently is recomputed in the program. Notice however that this phenomenon is equivalent to the one that arises within magic sets and semi-naive evaluation. The calls $?\mathsf{p}(\mathsf{a},\mathsf{Y})$ and $?\mathsf{p}(\mathsf{a},\mathsf{Z})$ are equivalent to the magic fact $\mathsf{mag_p^{bf}}(\mathsf{a})$. Thus, when the latter call (magic fact) is derived it is identified as already answered. Consequently it is eliminated by the subsumption checking performed on derived facts by the semi-naive procedure. However, the second described case corresponds in magic sets to derive first the magic fact $\mathsf{mag_p^{bf}}(\mathsf{a})$ and then $\mathsf{mag_p^{bb}}(\mathsf{a},\mathsf{b})$. Applying subsumption checking between these two facts returns failure because they are syntactically unrelated. Therefore, the magic fact $\mathsf{mag_p^{bb}}(\mathsf{a},\mathsf{b})$ is derived and the computation associated with it is redone. The aim of the following section is to explore the details involved with adornments and the desirable feature of sharing answers among computations of related magic facts.

4 Adornments and 'Sharing'

An interesting phenomenon occurs with adornments in the context of tabulation. Consider in top-down evaluation the following order of calls: first $?\mathsf{p}(\mathsf{X},\mathsf{Y})$ and then $?\mathsf{p}(\mathsf{a},\mathsf{Z})$. In tabulated top-down the admissibility (subsumption) test would force the second call to reuse the answers that were computed and stored by the first. In the bottom-up evaluation

[1] One atom is a *variant* of another if they are the same up to variable renaming.

of the corresponding magic rewritten program, the fact $\mathsf{mag_p^{ff}}$ would be generated first which would lead to the computation of the complete extension of predicate p i.e. all its solutions. Then, $\mathsf{mag_p^{bf}(a)}$ would be generated leading to the computation of p facts that have a as first parameter. Notice that, according to the described subgoals, the magic fact $\mathsf{mag_p^{bf}(a)}$ is redundant in relation to $\mathsf{mag_p^{ff}}$: the facts 'computed' by $\mathsf{mag_p^{bf}(a)}$ are included in the facts 'computed' by $\mathsf{mag_p^{ff}}$.

Subsumption in the semi-naive procedure is used to prevent the derivation of duplicate facts. Now, if goal orientation is included, as is the case of programs transformed by the magic rewriting, subsumption must also prevent redundant computation by eliminating the derivation of redundant magic facts. As mentioned earlier, subsumption in the semi-naive strategy applied to magic programs corresponds to the subsumption test between subgoals performed in tabulated top-down evaluation. Eliminating subgoal redundancy avoids redundant computation in tabulated top down evaluation. Consequently, subsumption in semi-naive should avoid redundant computation by eliminating redundant magic facts. However, subsumption does not work on adorned facts because syntactically the magic facts are unrelated. Furthermore, following the earlier example, the computed answers for the predicate $\mathsf{p^{ff}}$ cannot be shared with predicate $\mathsf{p^{bf}}$ since both are now different predicates. Thus due to adornments, in semi-naive evaluation of magic programs, the subsumption test cannot check that answers derived with the first magic fact should be used to answer the requirements of the second fact. In this way, adornments remove from bottom-up one of the most desirable features of tabulation: the sharing of answers between similar calls.

Our aim is to have a bottom-up evaluation that preserves the sharing of solutions among common calls, as happens in top-down. To achieve this, a program will be adorned in a different way. Adornments are used in several query optimization techniques, helping to cut down the relevant search space, e.g. [Morris88, Kemp et al. 90, Naughton et al 89]. But for the magic rewriting it is only necessary to consider adornments in the magic literals. In this way all adorned versions of a predicate will generate answers that potentially can be used by all the different adorned literals present in the body of rules. An implicit adornment is considered instead of an explicit 'renaming' of literals in rules. Consider the following rule.

$$\mathsf{p(X,Y)} \leftarrow \mathsf{a(X,Z)} \ \& \ \mathsf{b(Y,Z)}$$

Assuming the query $?\mathsf{p(a,Y)}$ and a left-to-right SIP strategy, standard adornment together with magic rewriting produces:

$$\mathsf{p^{bf}(X,Y)} \leftarrow \mathsf{a^{bf}(X,Z)} \ \& \ \mathsf{b^{fb}(Y,Z)} \ \& \ \mathsf{mag_p^{bf}(X)}$$

However, according to our idea of implicit adornments, it is sufficient to adorn the magic predicate only, which leads to:

$$\mathsf{p(X,Y)} \leftarrow \mathsf{a(X,Z)} \ \& \ \mathsf{b(Y,Z)} \ \& \ \mathsf{mag_p^{bf}(X)}$$

On generating magic rules the same idea applies. For instance, the magic rules for a and for b are:

$$\mathsf{mag_a^{bf}(X)} \leftarrow \mathsf{mag_p^{bf}(X)}$$

$$\mathsf{mag_b^{fb}(Z)} \leftarrow \mathsf{mag_p^{bf}(X)} \ \& \ \mathsf{a(X,Z)}$$

In this way we gain a generation of facts of the same predicate that enables the sharing of answers between literals of the same predicate in the body of rules. A single rule can still generate several different adorned versions as happens in the standard adornment, since the information provided by the SIP strategy is still followed. Furthermore the main aim of adornments is still considered i.e. to implement the SIP strategy. Observe that by omitting adornments from literals in the body of rules we do not lose the benefits provided by the adornment process. The adornment is implicit in the way that rules are processed, so the well known optimizations that make use of adornments can still be applied e.g. [Kemp et al. 90, Naughton et al 89, Ramak. et al. 88]. This is due to the fact that at compile time, when these optimizations are applied, it is known (through the SIP-strategy) which adornment a literal would have if the standard adornment is applied.

Combined with this new rewriting we need a subsumption test on the generated magic facts capable of identifying redundant magic facts. In the next section, an efficient algorithm for performing this task will be described.

5 Subsumption Checking over Adorned Atoms

Consider the following example, which is the magic program of the right linear version of the recursive definition of ancestor.

$\mathsf{anc}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Y}).$
$\mathsf{anc}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{anc}(\mathsf{Z},\mathsf{Y}).$

The predicate par corresponds to the following chain: $a \longrightarrow b \longrightarrow c \longrightarrow d$, which is the EDB:

$\mathsf{par}(\mathsf{a},\mathsf{b}).$
$\mathsf{par}(\mathsf{b},\mathsf{c}).$
$\mathsf{par}(\mathsf{c},\mathsf{d}).$

The transformed program according to the standard adornment process and the query $\mathsf{anc}^{\mathsf{fb}}$ is:

$\mathsf{anc}^{\mathsf{fb}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Y})\ \&\ \mathsf{mag_anc}^{\mathsf{fb}}(\mathsf{Y}).$
$\mathsf{anc}^{\mathsf{fb}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{anc}^{\mathsf{bb}}(\mathsf{Z},\mathsf{Y})\ \&\ \mathsf{mag_anc}^{\mathsf{fb}}(\mathsf{Y}).$
$\mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{Z},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{mag_anc}^{\mathsf{fb}}(\mathsf{Y}).$

$\mathsf{anc}^{\mathsf{bb}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Y})\ \&\ \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{X},\mathsf{Y}).$
$\mathsf{anc}^{\mathsf{bb}}(\mathsf{X},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{anc}^{\mathsf{bb}}(\mathsf{Z},\mathsf{Y})\ \&\ \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{X},\mathsf{Y}).$
$\mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{Z},\mathsf{Y}) \leftarrow \mathsf{par}(\mathsf{X},\mathsf{Z})\ \&\ \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{X},\mathsf{Y}).$

If the query is ?anc(X, d) then we add the magic fact $\mathsf{mag_anc}^{\mathsf{fb}}(\mathsf{d})$. The semi-naive evaluation is:

$\mathsf{T}^1 = \mathsf{EDB} \cup \{\mathsf{mag_anc}^{\mathsf{fb}}(\mathsf{d})\}$

$\mathsf{T}^2 = \mathsf{T}^1 \cup \{\mathsf{anc}^{\mathsf{fb}}(\mathsf{c},\mathsf{d}), \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{b},\mathsf{d}), \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{c},\mathsf{d}), \mathsf{mag_anc}^{\mathsf{bb}}(\mathsf{d},\mathsf{d})\}$

$\mathsf{T}^3 = \mathsf{T}^2 \cup \{\mathsf{anc}^{\mathsf{bb}}(\mathsf{c},\mathsf{d})\}$

$\mathsf{T}^4 = \mathsf{T}^3 \cup \{\mathsf{anc}^{\mathsf{bb}}(\mathsf{b},\mathsf{d}), \mathsf{anc}^{\mathsf{fb}}(\mathsf{b},\mathsf{d})\}$

$\mathsf{T}^5 = \mathsf{T}^4 \cup \{\mathsf{anc}^{\mathsf{fb}}(\mathsf{a},\mathsf{d})\}$

Notice that all the derived magic facts in step 2 (T^2) are redundant in relation to the magic fact representing the initial query, because they are subsumed by the latter. Furthermore, the $\mathsf{anc}^{\mathsf{fb}}$ fact derived at step 2 cannot be used by the $\mathsf{anc}^{\mathsf{bb}}$ literal in the body of the second rule defining $\mathsf{anc}^{\mathsf{fb}}$. Although syntactically $\mathsf{anc}^{\mathsf{bb}}(\mathsf{c},\mathsf{d})$ is different from $\mathsf{anc}^{\mathsf{fb}}(\mathsf{c},\mathsf{d})$, both represent that c is an ancestor of d.

A new algorithm is required to identify subsumption relations between the adorned magic facts. This is analogous to the *admissibility test* for top-down tabulation referred to in [Vieille87]. As shown in the example, although syntactically unrelated, semantically (based on the information contained in the adornments) one adorned magic fact can subsume another. For instance, the magic fact $\mathsf{mag_p}^{\mathsf{bff}}(\mathsf{a})$ subsumes the fact $\mathsf{mag_p}^{\mathsf{bbf}}(\mathsf{a},\mathsf{b})$, since the former corresponds to a goal $?\mathsf{p}(\mathsf{a},\mathsf{Y},\mathsf{Z})$ and the latter to $?\mathsf{p}(\mathsf{a},\mathsf{b},\mathsf{X})$. Without such a subsumption test, the full benefits associated with tabulation cannot be obtained in bottom-up evaluation.

5.1 A new definition of subsumption

First, we define subsumption in adorned magic facts. We rely on a translation from magic facts into the corresponding subgoals in top down evaluation. Since we are dealing with magic sets, we assume that no *aliasing* of variables [Ullman89] occurs (i.e. all magic facts represent atoms with distinct variables) and no derived fact contains function symbols in its arguments, which simplifies subsumption checking.

Definition 2 *The translation of a magic fact $mag_S^{\alpha}(\vec{c})$, where α is the adornment sequence of 'b's and 'f's, is the term $S(\vec{x})$ where $\vec{x}$ is composed of the constants that appear in $\vec{c}$ for the parameters that are 'b' in α and a distinct variable for each parameter that are 'f' in α.*

For instance, the magic fact $mag_p^{bfbbf}(a,b,c)$ is translated into the term $?p(a,X,b,c,Y)$. Subsumption between two magic facts is reduced to the subsumption between the corresponding subgoals resulting from the translation described above.

Definition 3 *A magic fact M_1 subsumes a magic fact M_2 if the corresponding term S_1 of M_1 subsumes the term S_2 of M_2 i.e. $S_1 \sqsupseteq S_2$.*

The idea is that instead of translating adorned magic facts into corresponding atoms and checking subsumption between these atoms, one can make use of the information in the adornments to directly determine whether an adorned magic fact subsumes another. Since adornments in magic facts represent the pattern of bound/free variables in their arguments, subsumption checking can be reduced to operations over adornment sequences ('bf' sequences).

Recall that G subsumes S (denoted $G \sqsupseteq S$) if there is a substitution θ for the variables in G such that $G\theta = S$. Thus, the subsumption test should check if such a substitution θ exists, succeeding if it does, failing otherwise. An alternative way to define subsumption is the following [Pereira&Shieber87]:

Definition 4 *$G \sqsupseteq S$ if $\exists\theta = m.g.u(G,S)$ and $S\theta = S$.*

This is equivalent to say that G subsumes S if the most general unifier (m.g.u) of S and G does not bind any variable in S. We can show that the latter definition is equivalent to the former by assuming that both atoms unify. If S and G unify then $\exists\theta \;\; S\theta = G\theta$. If θ does not bind any variable in S (i.e. $S = S\theta$) then it is the same as saying $\exists\theta \;\; S = G\theta$. But this last statement is precisely the first definition of subsumption.

Considering the latter definition of subsumption we can think of subsumption checking as reduced to operations with arguments of the atoms to be checked. Notice that since we are dealing with magic sets, all the described programs are Datalog and no *aliasing* of variables occurs. To optimize the operation with adornments we translate the 'bf' adornment sequences.

Definition 5 *The translation of an adornment sequences is a binary number obtained through the following substitution:*

- *each position 'b' in the original adornment is substituted by the digit '1'.*
- *and each 'f' by the digit '0'.*

Thus, an initial adornment has now a translation into a sequence of bits (binary number), e.g. the sequence bfbf is translated into the sequence of bits '1010'. The advantage of such a translation is that one can reduce the operations over arguments that occur in subsumption checking into logical operations on bits i.e. logical operations with binary numbers. For convenience and since we are operating with the adornment sequences, each adornment is an extra argument of the corresponding magic fact. For instance, the original magic fact $mag_p^{bfbf}(c,a)$ is now the term $mag_p(1010,c,a)$ where the adornment sequence is the first argument of the magic fact. The full new rewriting can now be presented:

Definition 6 *Let* P^{ad} *be the adorned version of program (database)* P *following a given SIP-strategy (we still assume the same left-to-right strategy) and a query* $\mathsf{q}(\vec{\mathsf{x}})$.

1. *create a new predicate* $\mathsf{mag_p}(\mathsf{bit},\vec{\mathsf{t}}^{\mathsf{b}})$ *for each* $\mathsf{p}^{\mathsf{ad}}(\vec{\mathsf{t}})$ *in* P^{ad}*;* $\vec{\mathsf{t}}^{\mathsf{b}}$ *means the bound arguments of* $\vec{\mathsf{t}}$ *and* bit *is the translation of the adornment* ad *according to definition 5.*

2. *for each rule in* P^{ad} *add the modified rule to* $\mathsf{P}^{\mathsf{magic}}$ *which is the original rule with the body extended with the literal* $\mathsf{mag_p}(\mathsf{bit},\vec{\mathsf{t}}^{\mathsf{b}})$ *if the headis* $\mathsf{p}^{\mathsf{ad}}(\vec{\mathsf{t}})$ *(i.e. only the bound (b) arguments are in the magic literal).*

3. *For each rule* $\mathsf{p}^{\mathsf{ad}_0}(\vec{\mathsf{t}}) \leftarrow \mathsf{q}_1^{\mathsf{ad}_1}(\vec{\mathsf{t}}_1)\ \&\ \ldots\ \&\ \mathsf{q}_{\mathsf{n}}^{\mathsf{ad}_{\mathsf{n}}}(\vec{\mathsf{t}}_{\mathsf{n}})$ *in* P^{ad} *generate several magic rules:* $\mathsf{mag_q_i}(\mathsf{bit_i},\vec{\mathsf{t}}_{\mathsf{i}}^{\mathsf{b}}) \leftarrow \mathsf{mag_p}(\mathsf{bit}_0,\vec{\mathsf{t}}^{\mathsf{b}})\ \&\ \mathsf{q}_1(\vec{\mathsf{t}}_1)\ \&\ \ldots\ \&\ \mathsf{q}_{\mathsf{i}-1}(\vec{\mathsf{t}}_{\mathsf{i}-1})$ *is added to* $\mathsf{P}^{\mathsf{magic}}$ *for each* $1 \leq \mathsf{i} \leq \mathsf{n}$ *and the order of* i *respects the order on the SIP. Again* $\mathsf{bit_i}$ *is the translation of the adornment* $\mathsf{ad_i}$*,*

4. *add the seed fact* $\mathsf{mag_q}(\mathsf{bits},\vec{\mathsf{x}}^{\mathsf{b}})$ *representing the query* $\mathsf{q}(\vec{\mathsf{x}})$*,where* bits *is the translation of the adornment associated with* $\vec{\mathsf{x}}$*.*

Theorem 1 *(Preservation of answers) Let* $< p^a, P^{ad} >$ *be a query and an adorned program transformed by standard magic sets rewriting [Beeri&Ramak91]. Let* $< p^b, P^{bits} >$ *be the same query and program transformed following definition 6.* $< p^a, P^{ad} >$ *and* $< p^b, P^{bits} >$ *are equivalent i.e. the two programs produce the same answer for the resulting queries on* p*.*

Proof *This can be proved by considering the correspondence between the binary numbers and the adornment sequences.*

Theorem 2 *(Efficiency) Let P be a program and q a query. Let P^{mg} be P and q with the original magic rewriting applied [Beeri&Ramak91]. Let P^{bits} be P and q with the rewriting of definition 6 applied. Let $\mathcal{S}n(P)$ be a function that determines the number of facts derived during Standard Semi-Naive evaluation [Balbin&Ramam87] of program P. $\mathcal{S}n(P^{bits}) \leq \mathcal{S}n(P^{mg})$.*

Proof *Straightforward by considering that now duplicate elimination of facts can be truly obtained in P^{bits}.*

Consider again the procedures that are performed during subsumption checking. To check whether $G \sqsupseteq S$ these procedures must check whether the variables of S are bound by any of the ground parameters of G. Consider adornments as binary numbers. Consider also that adornments represent the parameters of atoms, indicating the bound/free positions. The required procedure can be implemented by a simple logical operation on sequences of bits. Bearing in mind that '0' represents a position of free variable, performing an binary **or** operation over two adornment sequences yields another adornment sequence that represents the bound/free position in the parameters of both atoms after being unified. If a position is bound in one atom then after unification the same position is bound on both atoms. Recall that subsumption can be reduced to checking whether the m.g.u between G and S does not bind any variable in S. Thus, if the resultant adornment sequence of the logical **or** operation matches the sequence representative of S then $G \sqsupseteq S$.

Consider the following example; the atoms $\mathsf{p(a,Y,Z)}$ and $\mathsf{p(a,b,c)}$ have as m.g.u the substitution $\{\mathsf{Y/b,Z/c}\}$. The atoms after unification are both $\mathsf{p(a,b,c)}$ which corresponds to the adornment '111'. The first atom is represented by the adornment '100' and the second by '111'. Performing 100 **or** 111 yields 111. The first fact subsumes the second, because the resulting adornment sequence obtained from the **or** operation is equal to the sequence of the second fact. The fact that these two sequences are equal means that the m.g.u. does not perform any substitution on the variables of the second atom. Recall that since Magic Sets are used, no aliasing of variables occurs (sharing of variables in one term) [Ullman89]. This justifies why binary numbers can be used to check subsumption.

We can summarize the subsumption algorithm in the following way. Consider that we want to check whether $G^{\alpha} \sqsupseteq S^{\beta}$, where α and β are the adornments. Then *subsumes* is defined as:

$$subsumes(G^{\alpha}, S^{\beta}) \leftrightarrow \mathbf{or}(\alpha, \beta, \beta) \;\&\; match(G, S).$$

When two magic facts succeed in the logical **or** test, one has to confirm whether the bound positions of both facts that coincide represent parameters that match. In other words, one has to perform pattern matching between the bound parameters of both facts. Since we are dealing with adorned magic facts, the parameters in these facts are all ground, corresponding to the bindings to be passed. In the definition described above this corresponds to the predicate *match*. However, this procedure *match* must be adjusted because we need to know the adornments to determine again which arguments in G correspond to which in S. Consider two magic facts $mag^{sp}(\vec{s})$ and $mag^{ge}(\vec{g})$, of which the first is more specific and the second more general. To perform pattern matching one compares the adornments ge and sp. From this comparison one matches only the

positions on $\vec{s}$ and $\vec{g}$ that have '1' on both sp and ge, assuming that we already work with the translated sequences. As an example, consider the magic facts mag_p(0010, a) and mag_p(1011, j, a, c). Comparing the adornment sequences tells us that it is only necessary to match the third position in both facts. This is equivalent to comparing the first (and only) argument from the former (which is the constant a) with the second argument of the latter fact (constant a also).

We can also determine the positions to be compared through binary operations with the translated adornment sequences into binary numbers. Checking the bits in both sequences that are **on** i.e. assigned with 1, can be performed by successive operations of shifting and binary conjunctions. First we assign a variable with a binary number that has the same number of bits as the adornment sequences and all the bits turned off (i.e. 0) except the left-most one. Thus, if the numbers representing the adornment sequences have five bits (i.e. the number of parameters in the translated atom is five) then the variable is assigned with '10000'. We assume that there is a pointer for each magic fact. These pointers point to the list of parameters. Then, we perform two binary conjunctions between the two adornment sequences and the variable. Matching between the pointed parameters is only performed if both conjunctions yield non-zero results. Now, for each conjunction that gives non zero result we increment the respective pointer. Finally we perform a one bit shifting operation to the right on the used variable. This process is repeated while the pointer of the most general fact does not point to nil i.e. the list of parameters is not totally visited. This ensures that the number of comparisons between arguments of the two magic facts coincides with the number of parameters of the most general magic fact.

Let us consider an example with the magic facts mag_p(001, c) and mag_p(101, a, c). The auxiliary variable is assigned with '100'. Initially the pointer of the first magic fact points to the parameter c and the second to the parameter a. The conjunction '100 & 001 = 000' and '100 & 101 = 100' do not respect the first requirement. Thus no matching is performed and only the pointer for the second magic fact is incremented, pointing now to the constant c. Shifting the variable gives the binary number '010'. Both conjunctions yield zero as result. Therefore no matching is performed and no pointer is incremented. After the shifting, the variable has the value '001'. The operations are repeated and this time both conjunctions yield non zero results i.e. '001 & 001 = 001' and '001 & 101 = 001'. Thus, matching between pointed parameters is performed, which corresponds to apply matching between the constant c from the first magic fact with the constant c from the second magic fact.

5.2 The Algorithm

Finally, we are in position to present the complete *subsumes* algorithm. We use **or** and & to denote the binary operations of disjunction and conjunction, respectively. Two adorned magic facts, $mag_p^{sp}(\vec{s})$ and $mag_p^{ge}(\vec{g})$, participate in the algorithm. The adornments sp' and ge' are now the translations into sequences of bits of the original adornment sequences, following definition 5. The original magic facts are translated into respectively $mag_p(sp', \vec{s})$ and $mag_p(ge', \vec{g})$. P_g is the pointer to the list of parameters in $\vec{g}$ and P_s is the pointer to the list of parameters in $\vec{s}$. Initially both point to the first argument of each magic atom. The algorithm checks whether $mag_p^{ge}(\vec{g}) \sqsupseteq mag_p^{sp}(\vec{s})$.

Algorithm *Subsumes*

1. if $sp' \neq sp'$ **or** ge' then fail and exit.

2. else check pattern matching between $\vec{s}$ and $\vec{g}$.

   {The algorithm goes through ge' and sp', from left to right, to determine the positions to be matched.}

   $Aux := 1 << (n-1)$ where n is the number of bits in the sequences sp' and ge'. {shift to the left $n-1$ times the number 1 in binary format}

   Do while $P_g \neq nil$ {does not point to nil}

   $G := ge'$ & Aux;

   $S := sp'$ & Aux;

   if $G \neq 0$ and $S \neq 0$ then

   if not match(P_g, P_s) then fail and exit;

   if $G \neq 0$ then make P_g point to next position;

   if $S \neq 0$ then make P_s point to next position;

   $Aux >> 1$ {shift once to the right};

   Endwhile

3. succeed.

The first step of the algorithm works as a preliminary test. The second step performs pattern matching. Note that the algorithm stops when all the arguments of the more general atom are visited ($\vec{g}$).

Let us consider some examples in the application of subsumption to the elimination of redundant magic facts derived during semi-naive evaluation: The calls ?$\mathsf{p}(\mathsf{a},\mathsf{Y},\mathsf{Z})$ and ?$\mathsf{p}(\mathsf{a},\mathsf{b},\mathsf{Z})$ correspond to the magic facts $\mathsf{mag_p}^{\mathsf{bff}}(\mathsf{a})$ and $\mathsf{mag_p}^{\mathsf{bbf}}(\mathsf{a},\mathsf{b})$, respectively. Suppose the former is a previously derived fact and the latter is a new fact. We want to check whether $\mathsf{mag_p}^{\mathsf{bff}}(\mathsf{a}) \sqsupseteq \mathsf{mag_p}^{\mathsf{bbf}}(\mathsf{a},\mathsf{b})$. Performing '100 or 110' results in '110' which is equal to the sequence in the new fact. Next, both sequences of arguments match since the first binding of the first fact (a) matches the first binding of the second (a). Therefore $\mathsf{mag_p}^{\mathsf{bff}}(\mathsf{a}) \sqsupseteq \mathsf{mag_p}^{\mathsf{bbf}}(\mathsf{a},\mathsf{b})$. In a semi-naive evaluation the new fact would be eliminated, meaning that redundant computation associated with this fact would be avoided.

Consider now the case where both facts are identical, say both are the fact $\mathsf{mag_p}^{\mathsf{bff}}(\mathsf{a})$. '100 or 100' is equal to '100' and pattern matching succeeds as well, so the newly derived fact would be eliminated which shows that the algorithm performs equality checking as required. Consider the case where neither of the atoms subsumes the other. For instance the queries ?$\mathsf{p}(\mathsf{a},\mathsf{Y})$ and ?$\mathsf{p}(\mathsf{X},\mathsf{a})$, are represented by the magic facts $\mathsf{mag_p}^{\mathsf{bf}}(\mathsf{a})$ and $\mathsf{mag_p}^{\mathsf{fb}}(\mathsf{a})$. The operation '10 or 01' gives '11' as result. Thus, the algorithm returns failure. Consider finally an example with different bindings. Assume the magic facts $\mathsf{mag_p}^{\mathsf{fbf}}(\mathsf{c})$ and $\mathsf{mag_p}^{\mathsf{fbb}}(\mathsf{a},\mathsf{c})$. The adornments checking succeeds since 010 or 011 = 011. However comparing the bindings gives failure because $\mathsf{c} \neq \mathsf{a}$.

The algorithm complexity is characterized by a $\mathcal{O}(m)$ behaviour where m is the number of arguments of the more general magic atom i.e. $m = length(\vec{g})$. Here, m also represents the number of comparisons performed during pattern matching i.e. the second step of the

subsumption algorithm. The logical operations over adornments are negligible because they can be implemented at a machine register level. Notice that when subsumption checking returns failure the number of comparisons performed by the algorithm is always less than or equal to m. For instance, performing the second step of the algorithm between the magic atoms $\mathsf{mag_p}^{\mathsf{bbffff}}(\mathsf{f},\mathsf{a})$ and $\mathsf{mag_p}^{\mathsf{bbfffb}}(\mathsf{a},\mathsf{a},\mathsf{c})$ (which returns failure) requires only one comparison, but $m = 2$.

The complexity of the standard implementation in Prolog of subsumption is $\mathcal{O}(2n)$, where n is the arity of the intervening atoms (which would be the translation of the intervening magic atoms, according to definition 2) e.g. [Pereira&Shieber87].

5.3 Soundness and Completeness

In this section we prove soundness and completeness of the proposed subsumption algorithm.

Lemma 1 *Assume that the translations of the magic facts M_1^{ge}, M_2^{sp}, according to definition 2, are respectively S_1 and S_2. Assume also that no aliasing of variables occurs in S_1, S_2. If $\exists\theta$ such that $\theta = m.g.u(S_1, S_2)$ and $S_2\theta = S_2$ then $sp' = sp'$ **or** ge', (where sp' and ge' are the binary translation of the adornments sp and ge, according to Definition 5).*

Proof

*If ge and sp are the pattern of bound/free arguments of M_1, M_2 then they are also the patterns for S_1, S_2. Suppose t is the pattern associated with the atom $S_2\theta$. Assume that t' is the translation, according to definition 5, of t. Since $\theta = m.g.u.(S_1, S_2)$ a position in t' is '1' if the same position is '1' in sp' or in ge'. Thus $t' = sp'$ **or** ge'. Therefore, since $S_2 = S_2\theta$ yields $sp' = t'$, we have $sp' = sp'$ **or** ge'.* □

Theorem 3 *(Soundness)*

Assume that the translations of the magic facts M_1^{ge}, M_2^{sp}, according to definition 2, are respectively S_1 and S_2. If $subsumes(M_1^{ge}, M_2^{sp}) = true$ then $S_1 \sqsupseteq S_2$

Proof

If $subsumes(M_1^{ge}, M_2^{sp}) = true$ then two facts hold:

*1. $sp' = sp'$ **or** ge' (where sp' and ge' are sp and ge following definition 5) and*

2. $match(M_1, M_2) = true$

*Since $match(M_1, M_2)$ succeeds then all bound arguments of M_1 and M_2 coincide and consequently the same happens with S_1 and S_2. We assume no aliasing of variables. This means that S_1, S_2, have all distinct variables. The fact that $sp' = sp'$ **or** ge' holds, means that no free position of M_2 coincides with a bound position of M_1. The same happens with S_2 and S_1. Since no aliasing of variables among S_1, S_2 occurs and $match(M_1, M_2) = true$ ensures that S_1 unifies with S_2. Thus, $\exists\theta$ such that $\theta = m.g.u(S_1, S_2)$. But no variable argument of S_2 coincides with a bound argument of S_1.*

Hence $S_2\theta = S_2$, and therefore $S_1 \sqsupseteq S_2$. □

Theorem 4 *(Completeness)*

Assume that the translations of the magic facts M_1^{ge}, M_2^{sp}, according to definition 2, are respectively S_1 and S_2. If $S_1 \sqsupseteq S_2$ then $subsumes(M_1^{ge}, M_2^{sp}) = true$.

Proof

Following definition 4, if $S_1 \sqsupseteq S_2$ *then* $\exists\theta$ *such that* $S_2\theta = S_2$ *and* $\theta = m.g.u(S_1, S_2)$. *We assume that* S_1 *and* S_2 *have no* aliasing *of variables.*
Since $S_1\theta = S_2\theta$ *then* $match(M_1, M_2) = true$ *due to the fact that, the translations of* M_1^{ge}, M_2^{sp}, *according to Definition 2, are* S_1, S_2. *Following Lemma 1, having* $S_2\theta = S_2$ *implies that* $sp' = sp'$ **or** ge' *(where* sp' *and* ge' *are the translations of the adornments* sp *and* ge, *according to definition 5).*
We know that $subsumes(M_1^{ge}, M_2^{sp}) \leftrightarrow$ **or**$(ge', sp,' sp') \;\&\; match(M_1, M_2)$.
Therefore $subsumes(M_1^{ge}, M_2^{sp}) = true$ □

In practical terms the problem that one has to address is how to efficiently perform subsumption between one newly derived adorned magic fact and a *set* of previously derived adorned magic facts. Thus, we have to extend the proposed algorithm to include an proper mechanism for the indexing of derived adorned magic facts. In [Rao et al. 94] a *trie*-like structure was proposed to index calls and their computed answers in a tabulated top-down procedure (XSB Prolog). Given a fixed order of term traversal, tries can be used to index terms (in our case magic facts). The major advantages of these structures is that it gives a collapsed check/insert operation. In our case, performing subsumption requires one traversal for each binary sequence in the trie that satisfies step 1 of our algorithm. Insertion is collapsed with one of these traversals performed during subsumption checking. When traversing the trie, the described bits operation of our algorithm are executed: The first step of the subsumption algorithm is performed according to the first parameter of each term (which is the adornment sequence); The remaining path is traversed according to the bits operations described in step 2.

6 Examples

We take the previous example of right-recursive definition of ancestor of section 5 for demonstrating the benefits of the proposed adornment process and the new subsumption checking algorithm. The example will be executed by semi-naive evaluation incorporating the new subsumption checking to determine whether newly generated magic facts should be eliminated. These two proposals overcome the redundancy in the evaluation observed in section 5. Applying the rewriting of definition 6 to this example yields:

$\mathsf{anc(X,Y)} \leftarrow \mathsf{par(X,Y)} \;\&\; \mathsf{mag_anc(01,Y)}.$
$\mathsf{anc(X,Y)} \leftarrow \mathsf{par(X,Z)} \;\&\; \mathsf{anc(Z,Y)} \;\&\; \mathsf{mag_anc(01,Y)}.$
$\mathsf{mag_anc(11,Z,Y)} \leftarrow \mathsf{par(X,Z)} \;\&\; \mathsf{mag_anc(01,Y)}.$

$\mathsf{anc(X,Y)} \leftarrow \mathsf{par(X,Y)} \;\&\; \mathsf{mag_anc(11,X,Y)}.$
$\mathsf{anc(X,Y)} \leftarrow \mathsf{par(X,Z)} \;\&\; \mathsf{anc(Z,Y)} \;\&\; \mathsf{mag_anc(11,X,Y)}.$
$\mathsf{mag_anc(11,Z,Y)} \leftarrow \mathsf{par(X,Z)} \;\&\; \mathsf{mag_anc(11,X,Y)}.$

Semi-naive evaluation, which includes our subsumption checking algorithm, for the same query is:

$\mathsf{T}^1 = \mathsf{EDB} \cup \{\mathsf{mag_anc(01,d)}\}$

$\mathsf{T}^2_{\mathsf{before}} = \mathsf{T}^1 \cup \{\mathsf{anc(c,d)}, \mathsf{mag_anc(11,b,d)}, \mathsf{mag_anc(11,c,d)}, \mathsf{mag_anc(11,d,d)}\}$

$\mathsf{T}^2_{\mathsf{after}} = \mathsf{T}^1 \cup \{\mathsf{anc(c,d)}\}$

$T^3 = T^2 \cup \{anc(b, d)\}$

$T^4 = T^3 \cup \{anc(a, d)\}$

We split the relevant steps of the semi-naive evaluation into T_{before} and T_{after}, meaning respectively the facts derived **before** subsumption is applied and the facts that remain **after** subsumption checking. Evaluation for the original magic rewriting of the same program with the same query was:

$T^1 = EDB \cup \{mag_anc^{fb}(d)\}$

$T^2 = T^1 \cup \{anc^{fb}(c, d), mag_anc^{bb}(b, d), mag_anc^{bb}(c, d), mag_anc^{bb}(d, d)\}$

$T^3 = T^2 \cup \{anc^{bb}(c, d)\}$

$T^4 = T^3 \cup \{anc^{bb}(b, d), anc^{fb}(b, d)\}$

$T^5 = T^4 \cup \{anc^{fb}(a, d)\}$

In the evaluation of the second version of the program only the first two rules are fired. Redundant computation of step 2 in the evaluation of the first version of this example is eliminated because the redundant magic facts are subsumed by the initial query. One consequence of this checking and of the way adornments are performed is that derivation of duplicate facts for **anc** is eliminated. This derivation of duplicates appears in the semi-naive evaluation of the first version of this program, on steps 3 and 4.
Consider the same example but now for a cyclic graph which is represented by the EDB:

$par(a, b).$
$par(b, c).$
$par(c, d).$
$par(d, e).$
$par(e, a).$

Evaluation of the original magic rewriting of the same program for the query $?anc(X, e)$ is:

$T^1 = EDB \cup \{mag_anc^{fb}(e)\}$

$T^2 = T^1 \cup \{anc^{fb}(d, e), mag_anc^{bb}(b, e), mag_anc^{bb}(c, e), mag_anc^{bb}(d, e),$
$mag_anc^{bb}(e, e), mag_anc^{bb}(a, e)\}$

$T^3 = T^2 \cup \{anc^{bb}(d, e)\}$

$T^4 = T^3 \cup \{anc^{bb}(c, e), anc^{fb}(c, e)\}$

$T^5 = T^4 \cup \{anc^{bb}(b, e), anc^{fb}(b, e)\}$

$T^6 = T^5 \cup \{anc^{bb}(a, e), anc^{fb}(a, e)\}$

$T^7 = T^6 \cup \{anc^{bb}(e, e), anc^{fb}(e, e)\}$

Semi-naive evaluation with our subsumption checking algorithm for the same query is:

$T^1 = EDB \cup \{mag_anc(01, e)\}$

$T^2_{before} = T^1 \cup \{anc(d, e), mag_anc(11, b, e), mag_anc(11, c, e), mag_anc(11, d, e),$
$mag_anc(11, e, e), mag_anc(11, a, e)\}$

$T^2_{after} = T^1 \cup \{anc(d, e)\}$

$T^3 = T^2 \cup \{anc(c, e)\}$

$T^4 = T^3 \cup \{anc(b, e)\}$

$T^5 = T^4 \cup \{anc(a, e)\}$

$T^6 = T^5 \cup \{anc(e, e)\}$

Again, in step 2, subsumption checking prevents the use of redundant magic facts, namely the ones with the mag_anc^{bb} adornment. Consequently and due to the way we apply adornments, the repeated answers for anc with the different adornments are not derived. Let us now consider the same example for a graph with alternative paths. The EDB is now:

$par(a, b)$.
$par(b, c)$.
$par(b, d)$.
$par(c, e)$.
$par(c, d)$.
$par(d, e)$.

Evaluation for the original magic rewriting and the same query is:

$T^1 = EDB \cup \{mag_anc^{fb}(e)\}$

$T^2 = T^1 \cup \{anc^{fb}(d, e), anc^{fb}(c, e), mag_anc^{bb}(b, e), mag_anc^{bb}(c, e),$
$mag_anc^{bb}(d, e), mag_anc^{bb}(e, e)\}$

$T^3 = T^2 \cup \{anc^{bb}(d, e), anc^{bb}(c, e)\}$

$T^4 = T^3 \cup \{anc^{fb}(b, e), anc^{bb}(b, e)\}$

$T^5 = T^4 \cup \{anc^{fb}(a, e)\}$

Semi-naive evaluation for our version of the program is:

$T^1 = EDB \cup \{mag_anc(01, e)\}$

$T^2 = T^1 \cup \{anc(d, e), anc(c, e)\}$

$T^3 = T^2 \cup \{anc(b, e)\}$

$T^4 = T^3 \cup \{anc(a, e)\}$

7 Discussion

In [Sagiv90], it is shown that within Magic Sets the idea that more bound parameters in a query is always better than fewer is not correct. In other words, computing $?p(a,b)$ is not always better than computing the query $?p(a,Y)$ and checking whether b is in the answer. Sagiv shows that in some examples having the first query with the adornment p^{bb} leads to the appearance of the adornment p^{bf} in the body of the rules defining p. However, this implies the derivation of magic rules to the adornment p^{bf} and also to the adornment p^{bb}. Thus, recomputation will arise. Furthermore the same answers will be generated for the adornment p^{bb} and p^{bf}. This seems to be an evidence that our work and [Sagiv90] address a similar problem. Sagiv proposes a new program transformation to factorize predicates into new ones that correspond to the bound and free arguments described in the adornments. We address the same problem by simply introducing a new subsumption checking algorithm with an adornment process that is only applied to magic literals.

It is generally accepted by the Deductive Databases community e.g. [Sagiv90], that the number of derived facts in a computation is a good indication of the relative efficiency of the evaluation method. With the examples of the last section, we have shown that the efficiency of the bottom-up evaluation is improved. It can be proved that our proposal reduces evaluation from $\mathcal{O}(n^2)$ complexity to $\mathcal{O}(n)$, where n is the number of EDB facts (which is actually what happens in the presented examples), for non subsumption-free[2] [Maher&Ramak89] magic programs. Obviously, with subsumption-free programs our techniques perform poorly and worst than the standard combination of semi-naive evaluation and magic sets rewriting due to the burden of the new "semantic subsumption" of magic facts. Another important overhead is related by the removal of adornments from the literals in the bodies of rules. Without adornments no indexing of answers can be applied and consequently irrelevant facts can be tried in the bodies of rules.

Other techniques exist to improve standard magic sets as for instance *factoring* [Naughton et al 89] and the proposal in [Kemp et al. 90]. In general, factoring a program is an undecidable problem and the application of the proposal in [Kemp et al. 90] is restricted to left- right- and multi-linear programs. Actually, factoring could not be applied to the ancestor example of section 6 with a fb query. However, it remains to be investigated what is the inter-relation between these proposals and ours.

With our proposal an efficient tabulation technique is obtained in bottom-up evaluation, since now the total reuse of previous computation occurs. Our bottom-up mechanism can be related to the OLDT proof procedure [Tamaki&Sato86] but where no indexing of answers occurs. The proposed subsumption checking algorithm is equivalent to the *instance* checking included in the OLDT procedure.

8 Conclusions

In this paper we considered the relations present in the literature between the bottom-up and the top-down evaluation strategies. It was identified that the characteristic features of tabulation were not present in bottom-up with goal orientation. Namely, we observed that subsumption checking between subgoals (magic facts) was not implemented and sharing of derived facts between literals of the same predicate was not obtained. The desirable fea-

[2]Subsumption-free programs are defined in [Maher&Ramak89]. Here we assume magic rewritten programs according to definition 6 where subsumption is defined through the algorithm of section 5.2

tures of tabulation were rectified by proposing a new adornment process and an algorithm for checking subsumption over adorned magic facts. An efficient tabulation technique is obtained in bottom-up evaluation, since now the total reuse of previous computation occurs.

With these two proposals (the new subsumption algorithm and the new way of adorning rules) we introduce mechanisms to identify magic facts that were already answered (thus they are redundants) and to enable a sharing of answers among literals of the same predicate. Clearly, performing subsumption checking carries additional costs. However, as previously shown, first in the literature for the case of subsumption checking in tabulated top down evaluation e.g. [Vieille89, Tamaki&Sato86], and here with the examples, these overheads are negligible when compared with redundant computation that can (possibly) be avoided. Furthermore, the proposed algorithm was shown to have a reasonable complexity which indicates that it is efficient enough to overcome the burden associated with subsumption checking.

Adornments are a vital process to preserve range restriction among magic rewritten programs. Having non range restricted rules leads to the derivation of non ground facts. In a relational database this requires some sort of *Domain Closure Axiom* to obtain finite answers.

The proposed algorithm should be implemented in a way that enables the switching on/off of the subsumption checking, before an evaluation is performed. This implementation policy follows, for instance, the way other optimization techniques appear in the deductive database system CORAL [Ramak. et al. 92]. In this way, one could switch on in situations where different instances of the same magic fact are derived and switch off when only variants are derived.

Acknowledgements: Thanks to Mario Florido, Alipio Jorge, Jorge Sousa Pinto and Ian Mackie for the comments given in earlier drafts of this paper. Paulo Azevedo is supported by Junta Nacional de Investigação Científica e Tecnológica - Programa PRAXIS XXI grant number BD/2832/94.

References

[Balbin&Ramam87] Balbin I., Ramamohanarao K.,
A Generalization of the Differential Approach to Recursive Query Evaluation
in Journal of Logic Programming 1987, pp 259-262.

[Bancilhon et al. 86] Bancilhon F., Maier D., Sagiv Y., Ullman J.
Magic Sets and other strange ways to Implement Logic Programs
in Proceedings of the 5th Symposium on Principles of Databases Systems, (PODS) 1986.

[Beeri&Ramak91] Beeri C., Ramakrishnan R.
On the Power of Magic
in Journal of Logic Programming vol 10, pp 255-299, 1991.

[Kemp et al. 90] Kemp D., Ramamohanarao K., Somogyi Z.,
Right-, left- and multi-linear rule transformation that maintain context information

in Proceedings of the 16th International Conference on
Very Large Databases (VLDB), Australia 1990.

[Maher&Ramak89] Maher M., Ramakrishnan R.
Deja Vu in Fixpoints of Logic Programs
in Proceedings of the North American Conference on Logic
Programming, Cleveland, Ohio 1989.

[Morris88] Morris K.,
An Algorithm for Ordering Subgoals in NAIL!
in Proceedings of the 7th Symposium on Principles of Databases
Systems (PODS), pp 82-88, Austin 1988.

[Naughton&Ramak91] Naughton J., Ramakrishnan R.,
Bottom-Up Evaluation of Logic Programs
in Computational Logic: Essays in Honor of Alan Robinson
1991.

[Naughton et al 89] Naughton J., Ramakrishnan R., Sagiv Y., Ullman
Argument Reduction Through Factoring
in Proceedings of the 15th International Conference
on Very Large Databases (VLDB), Amsterdam 1989.

[Pereira&Shieber87] Pereira F., Shieber S.,
Prolog and Natural Language Analysis
Language & Information no. 10, Stanford, CA: Center for Study
of Language and Information, 1987.

[Ramak. et al. 88] Ramakrishnan R., Beeri C., Krishnamurthy R.,
Optimizing Existential Datalog Queries
in Proceedings of the 7th Symposium on Principles of Databases
Systems (PODS), Austin 1988.

[Ramak91] Ramakrishnan R.,
Magic Templates: A Spellbinding Approach to Logic Programs
in Journal of Logic Programming pp 189-216, vol 11, 1991.

[Ramak. et al. 92] Ramakrishnan R., Srivastava D., Sudarshan S.,
CORAL - Control, Relations and Logic
in Proceedings of the 18th Very Large DataBases Conference,
Vancouver, Canada 1992.

[Rao et al. 94] P. Rao, I. Ramakrishnan, T. Swift, D.S. Warren.
Dynamic Argument Reduction for the In-Memory Data Queries
in Proc. Second ICLP-Workshop on Deductive Databases
and Logic Programming of the International Conference on
Logic Programming, pag 117-129, Genova, 1994.

[Sagiv90] Sagiv Y.,
Is There Anything Better than Magic?
in Proceedings of the North American Conference on
Logic programming, 1990.

[Sagonas et al. 94] Sagonas K., Swift T., Warren D.S.,
XSB as an Efficient Deductive Database Engine
in Proceedings of SIGMOD 1994 Conference ACM.

[Seki89] Seki H.
On The Power of Alexander Templates
in Proceedings of the 8th Symposium on Principles of Databases Systems, (PODS) 1989.

[Tamaki&Sato86] Tamaki H., Sato T.,
OLD Resolution with Tabulation
in Proceedings of the 3rd International Conference on Logic Programming, London U.K. 1986, pp 84-98.

[Warren92] Warren D. S.,
Memoing for Logic Programs
in Communications of the ACM Vol 35, No 3, March 1992, pp 93-111.

[Ullman89] Ullman J.,
Bottom-up beats Top-down for Datalog
in Proceedings of the 8th Symposium on Principles of Databases Systems, (PODS) 1989.

[Vieille87] Vieille L.,
Database-Complete proof procedures based on SLD Resolution.
in Proceedings of the 4th International Conference on Logic Programming, Melbourne, Australia 1987, pp 74-103.

[Vieille89] Vieille L.,
Recursive Query Processing: The Power of Logic
in Theoretical Computer Science, vol 69,
Elsevier Science Publishing 1989, pp 1-53.