

Improving Query Optimization for Disjunctive Datalog^{*}

Chiara Cumbo¹, Wolfgang Faber², and Gianluigi Greco³

¹ Dipartimento di Matematica, Università della Calabria, 87030 Rende, Italy
`cumbo@mat.unical.it`

² Institut für Informationssysteme, TU Wien, 1040 Wien, Austria
`faber@kr.tuwien.ac.at`

³ DEIS, Università della Calabria, 87030 Rende, Italy
`ggreco@si.deis.unical.it`

Abstract. In this paper we present a technique for the optimization of (partially) bound queries over disjunctive deductive databases. In particular, we extend the magic-set optimization technique (originally defined for non-disjunctive deductive databases) to the disjunctive case. The method presented in this paper improves a similar approach presented in [7] by reducing the number of additionally introduced predicates and rules.

One drawback, which is intrinsic to both techniques, is that redundant rules may be created frequently. In order to overcome this, we present a method for identifying such superfluous rules, which may subsequently be deleted without changing the semantics, thereby reducing the size of the rule-base, an important performance factor.

1 Introduction

There is a growing body of work on answering queries over disjunctive deductive databases [5, 8] that have been proposed in the literature. In particular, the efficient evaluation of queries in disjunctive deductive databases and the definition of efficient evaluation algorithms for assigning semantics to disjunctive databases have been the subject of several proposals. Most of these proposals are based on the definition of efficient fixpoint algorithms and optimization is achieved by using heuristics.

The first attempt to use a different optimization technique for the evaluation of bound disjunctive Datalog queries was introduced in [6]. This approach stems from the observation that bottom-up query answering methods tend to explore a much larger search space than what is strictly necessary, whereas top-down processing uses the information provided by the query to prune the search space, thus performing a more focused search, by extending binding propagation methods defined for Datalog queries to disjunctive Datalog queries.

^{*} This work was supported by the EC under projects IST-2002-33570 INFOMIX and IST-2001-37004 WASP.

Its main idea is to adapt the well-known magic-set method for the case of disjunctive programs. The magic-set method is a strategy for simulating the top-down evaluation of a query by modifying the original program by means of certain rules, acting as filters for the relevant information needed for the query. It essentially consists of three separate steps: (1) An *Adornment step* in which the relationship between a bound argument in the rule head and the bindings in the rule body is made explicit. (2) A *Generation step* in which the adorned program is used to generate the *magic rules* which simulate the top-down evaluation scheme. (3) A *Modification step* in which the adorned rules are modified by the magic rules generated in step (2); these rules are called *modified rules*.

An elaborated extension of the magic-set technique to the disjunctive case has been published in [7]. In this paper we present an alternative extension, which needs fewer additional predicates and rules, by exploiting a direct coupling with the DLV system [2, 3]. Due to space limits, we will not describe the architecture of the system in detail and refer to [4].

A crucial point of both approaches is that during the rewriting redundant rules are created frequently. Since the number of rules in a program is a critical performance factor, these redundancies can deteriorate runtime behavior. In extreme cases, this overhead alone can outweigh the benefits of the optimization.

Unfortunately, it is infeasible to identify all cases of redundancies because of complexity results for the subsumption problem. However, we formulate a technique which can identify some of these cases in polynomial time. In principle, it consists of a greedy approach to detecting subsumption. Informal performance tests have shown a positive impact on a number of problems. In particular, the proposed approach yields a speedup on a number of problems and reduces overhead for other problems, while its introduced overhead is only mild.

2 Preliminaries on Disjunctive Datalog

A *disjunctive rule* (*rule*, for short) r is a formula

$$a_1 \vee \dots \vee a_n \text{ :- } b_1, \dots, b_k, \text{ not } b_{k+1}, \dots, \text{ not } b_m. \quad (1)$$

where $a_1, \dots, a_n, b_1, \dots, b_m$ are classical literals and $n \geq 0, m \geq k \geq 0$. The disjunction $a_1 \vee \dots \vee a_n$ is the *head* of r , while the conjunction $b_1, \dots, b_k, \text{ not } b_{k+1}, \dots, \text{ not } b_m$ is the *body* of r . Moreover, $H(r) = \{a_1, \dots, a_n\}$ is the set of the literals in the head and $B(r) = B^+(r) \cup B^-(r)$ is the set of the body literals, where $B^+(r)$ (the *positive body*) is $\{b_1, \dots, b_k\}$ and $B^-(r)$ (the *negative body*) is $\{b_{k+1}, \dots, b_m\}$. Finally, if the body is empty (i.e. $k = m = 0$), r is called a fact, and we usually omit the “:-”

sign. A predicate defined only by facts is called *EDB predicate*, while a predicate defined by the rules of the program is called *IDB predicate*.

A *disjunctive datalog program* (alternatively, *disjunctive logic program*, *disjunctive deductive database*) \mathcal{P} is a finite set of rules. A Datalog^V program \mathcal{P} (i.e., such that $\forall r \in \mathcal{P} : B^-(r) = \emptyset$) is called *positive*. A Datalog^V program such that each rule has exactly one head predicate is referred to as Datalog program. Note that in this paper only programs are considered, for which $\forall r \in \mathcal{P} : H(r) \neq \emptyset$ holds.

The model-theoretic semantics assigns to any program \mathcal{P} the set $\mathcal{SM}(\mathcal{P})$ of its *stable models*. Let \mathcal{P} be a disjunctive datalog program and \mathcal{F} be a set of facts. Then, a *query* Q for \mathcal{P} over the set of facts \mathcal{F} is a conjunction of atoms, that defines a mapping from the facts of $\mathcal{P} \cup \mathcal{F}$ to a finite (possibly empty) set of finite (possibly empty) tuples of atoms for Q . Given a query Q and an interpretation M of \mathcal{P} , $\vartheta(Q, M)$ denotes the set of substitutions for the variables in Q such that Q is true in M . The answer to a query Q over the fact \mathcal{F} , under the *brave* semantics, denoted by $Ans_b(Q, \mathcal{F})$, is the set $\cup_M \vartheta(Q, M)$, such that $M \in \mathcal{SM}(\mathcal{P} \cup \mathcal{F})$. The answer to a query Q over the fact \mathcal{F} , under the *cautious* semantics, denoted by $Ans_c(Q, \mathcal{F})$, is the set $\cap_M \vartheta(Q, M)$, such that $M \in \mathcal{SM}(\mathcal{P} \cup \mathcal{F})$.

3 Magic-Set for Datalog Queries

In this section we briefly review the main steps of the magic-set method, for the case of Datalog queries. For details, we refer to [1].

Adornment Step. An *adorned program* is a program whose predicate symbols have associated a string α , defined on the alphabet $\{b, f\}$, of length equal to the arity of the predicate. A character b (resp. f) in the i -th position of the adornment associated with a predicate p means that the i -th argument of p is bound (resp. free), i.e., there (resp. not) is a way for passing some binding information from the query to it, by simulating a top-down evaluation.

The adornment consists of generating an adorned program, that can be used for deriving how the binding of the query propagates. The predicates occurring in the query are adorned by marking each constant argument with b , all others with f . This is accomplished by a function **Adorn-Query** which outputs a tuple $\langle adornedQuery, S \rangle$, where *adornedQuery* is the adorned query and S is a stack of the adorned predicates.

Example 1. Consider the query $p(1), q(2, X), r(X, Y)?$, where p and q are IDB predicates, while r is an EDB predicate. Then, the function **AdornQuery** outputs the query $p^b(1), q^{bf}(2, X), r(X, Y)$, and pushes on the stack S the adorned predicates p^b and q^{bf} . \square

Once the stack of adorned predicates S has been created, we use it for adorning all rules of the program by tracing the propagation of the binding of the arguments of the query into the rules. In particular, given an adorned predicate \mathbf{p}^α , each rule r having the literal $\mathbf{p}(\mathbf{t})$ in the head is adorned by exploiting a particular strategy that is known as *SIPS* (Sideways Information Passing Strategy), which simulates the data flow occurring in the top-down evaluation of the query.

In the following, we will refer to the algorithm implementing the SIPS as the **AdornRule** algorithm. Roughly, it takes as input a rule r to be adorned w.r.t. \mathbf{p}^α , the stack S of adorned predicates to be used for propagating the binding (constructed by **AdornQuery**), and an additional argument *adornedPredicates*, containing all the adorned predicates generated so far. Initially, after the call to **AdornQuery**, we have that *adornedPredicates* and S are equal.

Then, iteratively, a *preferred* literal among the ones that are *admissible* in the rule, that is among the ones that can be adorned respecting the SIPS, is chosen. The iterations continue until all literals have been considered or no more *admissible* literal is left.

Thus, a natural way (the one usually described in the literature) for generating the adorned program consists of the following steps: (1) extract an adorned predicate \mathbf{p}^α from S , (2) use \mathbf{p}^α for propagating the binding into the program, and collect in S all the other adorned predicates generated, (3) if S is not empty return to Step (1).

Example 2. Continuing from Ex. 1, consider the rules $\mathbf{p}(X) :- \mathbf{q}(X, X)$. and $\mathbf{q}(X, Y) :- \mathbf{r}(X, Y), \mathbf{q}(X, Y)$. First, \mathbf{p}^b is popped from S , generating the rule $\mathbf{p}^b(X) :- \mathbf{q}^{bb}(X, X)$. and pushing \mathbf{q}^{bb} onto the stack. Then, \mathbf{q}^{bb} is popped, generating the rule $\mathbf{q}^{bb}(X, Y) :- \mathbf{r}(X, Y), \mathbf{q}^{bb}(X, Y)$., and finally \mathbf{q}^{bf} is popped, giving rise to $\mathbf{q}^{bf}(X, Y) :- \mathbf{r}(X, Y), \mathbf{q}^{bb}(X, Y)$. \square

Generation and Modification Steps. We use the adorned program for the generation of the magic rules. First, we define a transformation over the adorned literals. For each adorned literal \mathbf{p}^α , its *magic version* consists of the literal $\mathbf{magic}\text{-}\mathbf{p}^\alpha$ in which we eliminate the variables which are free w.r.t. α . In the following, given an adorned predicate \mathbf{p}^α , we often denote its magic version by $\mathbf{magic}(\mathbf{p}^\alpha)$; moreover, given a conjunction of n adorned literals $\mathbf{p}_1^{\alpha_1}, \dots, \mathbf{p}_n^{\alpha_n}$, we denote by $\mathbf{magic}(\mathbf{p}_1^{\alpha_1}, \dots, \mathbf{p}_n^{\alpha_n})$ the conjunction $\mathbf{magic}(\mathbf{p}_1^{\alpha_1}), \dots, \mathbf{magic}(\mathbf{p}_n^{\alpha_n})$. Then, for each adorned predicate \mathbf{p}^α in the body of any adorned rule r_a we generate a magic rule r_m such that (i) the head of r_m consists of its magic version, and (ii) the body of r_m consists of the magic version of the head literal of r_a , and those predicates of r_a which are relevant for propagating the binding on \mathbf{p}^α .

In the following we denote by $\mathbf{Generate}(r_a)$ the set of magic rules, generated by the procedure outlined above.

Finally, the modification step consists in the modification of each adorned rule; specifically, for each adorned rule whose head is $p^\alpha(X)$, we extend the rule body by inserting $\mathbf{magic}(p^\alpha)$. The final program will contain only the rules which are needed to answer the query.

Example 3. Continuing from Ex. 2, the corresponding magic rules are $\mathbf{magic_q}^{bb}(X, X) :- \mathbf{magic_p}^b(X)$., $\mathbf{magic_q}^{bb}(X, Y) :- \mathbf{magic_q}^{bb}(X, Y)$., $\mathbf{magic_q}^{bb}(X, Y) :- \mathbf{magic_q}^{bf}(X), r(X, Y)$.

The associated modified rules will be $p^b(X) :- \mathbf{magic_p}^b(X), q^{bb}(X, X)$., $q^{bb}(X, Y) :- \mathbf{magic_q}^{bb}(X, Y), r(X, Y), q^{bb}(X, Y)$., and $q^{bf}(X, Y) :- \mathbf{magic_q}^{bf}(X), r(X, Y), q^{bb}(X, Y)$. \square

Query Rules. In the steps described so far, the query is used only for deriving the original set of adorned predicates. In fact, for each adorned predicate of the query $g_i^{\alpha_i}$, (i) the fact $\mathbf{magic}(g_i^{\alpha_i})$ has to be asserted, (ii) a rule of the form $g_i(t_i) :- g_i^{\alpha_i}(t_i)$ has to be introduced.

Note that Step (1) is necessary in order to provide the starting point for the top-down evaluation, while Step (2) is necessary in order to provide the answer to the original query. We assume that this is done inside a function called *BuildQueryRules*.

Example 4. Continuing from Ex. 3, the generated query rules are $\mathbf{magic_p}^b(1)$., $\mathbf{magic_q}^{bf}(2)$., $p(1) :- p^b(1)$., and $q(2, X) :- q^{bf}(2, X)$. \square

4 Magic-Set Method for Datalog[∨] Programs

In this section we present a technique for the optimization of disjunctive Datalog queries, which has been implemented and integrated into the DLV system [2, 3]. One proposal for extending the magic-set technique to disjunctive Datalog queries has already been proposed in [7], and, in fact, the spirit of the present paper is close to this approach.

Yet, there are important differences; in particular, we do not use any additional predicates (called *collecting predicates* in [7]) or any additional rules for adorning disjunctive rules. Indeed, such predicates and rules are intrinsically needed in [7], as that technique consists of a rewriting into an equivalent program to be evaluated by DLV; conversely, our technique has been designed for being integrated into the core of the system, and, hence, it may use many optimization strategies resulting in a significant speed-up of the computation.

As a result of the application of these new ideas, of the great attention in the implementation issues we obtain an efficient implementation

(especially if compared with the prototypes proposed in the literature), whose designing has been left as a material for further research in [7].

The main algorithm, called **DisjunctiveMagicSet**, is reported in Figure 1. We point out that despite all the traditional implementations of the magic-set technique, our proposal is not conceptually divided into the traditional *adornment*, *generation*, and *modification* step; indeed, it exploits a stack S of predicates, for storing all the adorned predicates to be used for propagating the binding of the query: at each step, an element is removed from S , and it is used for processing each rule r of a given program \mathcal{P} , by contextually generating the associated modified and magic rules. Moreover, we point out that the algorithm can be used for propagating the binding into non-disjunctive rules as well.

In more detail, the algorithm takes in input a positive disjunctive datalog program \mathcal{P} (without integrity constraints) and a query $\mathbf{g}_1(\mathbf{t}_1), \dots, \mathbf{g}_n(\mathbf{t}_n)$. If the query contains some IDB predicates that can be used for propagating the binding, it outputs a program \mathcal{P}' consisting of a set of *modified* rules, of *magic* rules, and of *query* rules. The adorned query and the corresponding magic rules are built in the steps 2 and 3 by means of the functions **AdornQuery** and **BuildQueryRules**, respectively.

Example 5. Consider the disjunctive datalog program

$$\begin{aligned} \mathbf{p}(\mathbf{X}) \vee \mathbf{p}(\mathbf{Y}) &:- \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}(\mathbf{X}). \\ \mathbf{r}(\mathbf{Y}) &:- \mathbf{p}(\mathbf{Y}). \end{aligned}$$

and the query $\mathbf{p}(1)$? over the set of facts $\{\mathbf{a}(1, 2), \mathbf{a}(2, 3)\}$. In the algorithm of Figure 1, the function **AdornQuery** outputs the query $\mathbf{p}^b(1)$, and pushes on the stack S the adorned predicate \mathbf{p}^α . Moreover, the function **BuildQueryRules** outputs the rule $\mathbf{p}(\mathbf{X}) :- \mathbf{p}^b(\mathbf{X})$ and the fact `magic_pb(1)`. \square

The central part of the algorithm consists of the steps 4-13, that are repeated until the stack S becomes empty, i.e., until there is no further adorned predicate to be propagated. Next, we briefly explain the main ideas that are used for propagating the binding into disjunctive rules, and that are exploited in the implementation.

Let \mathbf{p}^α be an adorned predicate that has been removed from the stack S in step 5, and consider a disjunctive rule r in \mathcal{P} of the form

$$r : \mathbf{p}(\mathbf{t}) \vee \mathbf{p}_1(\mathbf{t}_1) \vee \dots \vee \mathbf{p}_n(\mathbf{t}_n) :- \mathbf{q}_1(\mathbf{s}_1), \dots, \mathbf{q}_m(\mathbf{s}_m).$$

```

Input: A DatalogV program  $\mathcal{P}$ , and a query  $g_1(t_1), \dots, g_n(t_n)$ .
Output: The optimized program  $\mathcal{P}'$ .
var  $S$ : stack of adorned predicates; adornedQuery: set of literals;
      modifiedRules, magicRules: set of rules;
begin
1. if  $g_1(t_1), \dots, g_n(t_n)$  has some IDB predicate then
2.    $\langle \text{adornedQuery}, S \rangle := \text{AdornQuery}(g_1(t_1), \dots, g_n(t_n), \mathcal{P})$ ;
3.    $\langle \text{modifiedRules}, \text{magicRules} \rangle := \text{BuildQueryRules}(\text{adornedQuery})$ ;
4.   while  $S \neq \emptyset$  do
5.      $p^\alpha := S.\text{pop}()$ ;
6.     for each rule  $r \in \mathcal{P}$ :  $p(t) \vee p_1(t_1) \vee \dots \vee p_n(t_n) :- q_1(s_1), \dots, q_m(s_m)$  do
7.       let  $r_s$  be the rule of the form
            $p(t) :- \text{not } p_1(t_1), \dots, \text{not } p_n(t_n), q_1(s_1), \dots, q_m(s_m)$ ;
8.       let  $r_a := \text{Adorn}(r_s, p^\alpha, S)$  of the form
            $p^\alpha(t) :- \text{not } p_1^{\alpha_1}(t_1), \dots, \text{not } p_n^{\alpha_n}(t_n), q_1^{\beta_1}(s_1), \dots, q_m^{\beta_m}(s_m)$ ;
9.        $\text{magicRules} := \text{magicRules} \cup \text{Generate}(r_a)$ ;
10.      let  $r_m$  be the modified rule of the form
            $p^\alpha(t) \vee p_1^{\alpha_1}(t_1) \vee \dots \vee p_n^{\alpha_n}(t_n) :-$ 
                $\text{magic}(p^\alpha(t), p_1^{\alpha_1}(t_1), \dots, p_n^{\alpha_n}(t_n), q_1(s_1)^{\beta_1}, \dots, q_m(s_m)^{\beta_m})$ ;
11.       $\text{modifiedRules} := \text{modifiedRules} \cup \{r_m\}$ ;
12.    end for
13.  end while
14.   $\mathcal{P}' := \text{magicRules} \cup \text{modifiedRules}$ ;
15.  return  $\mathcal{P}'$ ;
16. end if
end.

```

Fig. 1. Algorithm *DisjunctiveMagicSet*

in which the binding of p^α must be propagated due to the presence of a predicate $p(t)$ in its head. The difference w.r.t. the case of normal programs is that the binding have to be propagated also in the predicates $p_1(t_1), \dots, p_n(t_n)$ occurring in the head. A simple idea for reusing the standard procedures consists of replacing r with a rule r_s without disjunction of the form

$$r_s : p(t) :- \text{not } p_1(t_1), \dots, \text{not } p_n(t_n), q_1(s_1), \dots, q_m(s_m).$$

that can be adorned in the usual manner, since it is not disjunctive. The reason for pushing the predicates $p_1(t_1), \dots, p_n(t_n)$ in the body by negating them, lies in the fact that these predicates cannot be used for binding variables, as they actually occur only in the head of the original rule.

The rule r_s is then adorned by means of the function *AdornRule* informally described in the previous section. Here we recall that it takes in input a rule and an adorned predicate and outputs the adorned rule, while pushing on the stack S all the adorned predicates that have been generated during its execution. Hence, in step 8 we assume that the rule

r_a obtained by means of this adornment is of the form

$$r_a : \mathbf{p}^\alpha(\mathbf{t}) :- \text{not } \mathbf{p}_1^{\alpha_1}(\mathbf{t}_1), \dots, \text{not } \mathbf{p}_n^{\alpha_n}(\mathbf{t}_n), \mathbf{q}_1^{\beta_1}(\mathbf{s}_1), \dots, \mathbf{q}_m^{\beta_m}(\mathbf{s}_m).$$

where $\alpha_1, \dots, \alpha_n, \beta_1, \dots, \beta_m$ are generic adornment strings.

Example 6. Consider again the program and the query of Example 5, and let \mathbf{p}^b be the adorned predicate to be processed. The only rule in which the binding can propagate is the disjunctive one $\mathbf{p}(\mathbf{X}) \vee \mathbf{p}(\mathbf{Y}) :- \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}(\mathbf{X})$. Hence, when propagating on the predicate $\mathbf{p}(\mathbf{X})$ we preliminarily construct the standard rule $\mathbf{p}(\mathbf{X}) :- \text{not } \mathbf{p}(\mathbf{Y}), \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}(\mathbf{X})$. whose adornment is $\mathbf{p}^b(\mathbf{X}) :- \text{not } \mathbf{p}^b(\mathbf{Y}), \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}^b(\mathbf{X})$. Moreover, it is important to note that the binding of \mathbf{p}^b will also propagate on $\mathbf{p}(\mathbf{Y})$ by producing the adorned rule $\mathbf{p}^b(\mathbf{Y}) :- \text{not } \mathbf{p}^b(\mathbf{X}), \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}^b(\mathbf{X})$. Note that in both cases the predicate \mathbf{r}^b is pushed onto the stack S , and this predicate will generate the adorned rule $\mathbf{r}^b(\mathbf{Y}) :- \mathbf{p}^b(\mathbf{Y})$. \square

The algorithm subsequently uses the adorned rule r_a for generating the magic rules in the step 9. As r_a is not a disjunctive rule, this is essentially the same for the technique in the case of Datalog queries, and, hence, is carried out by means of the function *Generate* sketched in the previous section.

Example 7. In the program of Example 6 the propagation of the predicate \mathbf{p}^b generates the following magic rules

$$\begin{aligned} \text{magic_r}^b(\mathbf{X}) & :- \text{magic_p}^b(\mathbf{X}), \mathbf{p}^b(\mathbf{Y}), \mathbf{a}(\mathbf{X}, \mathbf{Y}). \\ \text{magic_p}^b(\mathbf{Y}) & :- \text{magic_p}^b(\mathbf{X}), \mathbf{r}^b(\mathbf{X}), \mathbf{a}(\mathbf{X}, \mathbf{Y}). \end{aligned}$$

$$\begin{aligned} \text{magic_r}^b(\mathbf{X}) & :- \text{magic_p}^b(\mathbf{Y}), \mathbf{p}^b(\mathbf{X}), \mathbf{a}(\mathbf{X}, \mathbf{Y}). \\ \text{magic_p}^b(\mathbf{X}) & :- \text{magic_p}^b(\mathbf{Y}), \mathbf{r}^b(\mathbf{X}), \mathbf{a}(\mathbf{X}, \mathbf{Y}). \end{aligned}$$

where the first group is generated from $\mathbf{p}^b(\mathbf{X}) :- \text{not } \mathbf{p}^b(\mathbf{Y}), \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}^b(\mathbf{X})$, while the second one from $\mathbf{p}^b(\mathbf{Y}) :- \text{not } \mathbf{p}^b(\mathbf{X}), \mathbf{a}(\mathbf{X}, \mathbf{Y}), \mathbf{r}^b(\mathbf{X})$. Moreover, by considering also the rule $\mathbf{r}^b(\mathbf{Y}) :- \mathbf{p}^b(\mathbf{Y})$ we derive $\text{magic_p}^b(\mathbf{Y}) :- \text{magic_r}^b(\mathbf{Y})$. \square

Finally, step 10 is devoted to the generation of the modified rule. The problem is that we need to reconstruct a disjunctive rule by pushing back each adorned predicate \mathbf{p}_i into the head, that has been translated into the body in the *adornment* step. Hence, we construct the rule r_m of the form

$$\mathbf{p}^\alpha(\mathbf{t}) \vee \mathbf{p}_1^{\alpha_1}(\mathbf{t}_1) \vee \dots \vee \mathbf{p}_n^{\alpha_n}(\mathbf{t}_n) :- \text{magic}(\mathbf{p}^\alpha(\mathbf{t}), \mathbf{p}_1^{\alpha_1}(\mathbf{t}_1), \dots, \mathbf{p}_n^{\alpha_n}(\mathbf{t}_n)), \mathbf{q}_1^{\beta_1}(\mathbf{s}_1), \dots, \mathbf{q}_m^{\beta_m}(\mathbf{s}_m).$$

and outputs a conjunction of magic predicates associated to the ones in input; for instance given $p^{bfb}(X, Y, Z)$ it outputs $magic_p^{bfb}(X, Z)$

Example 8. By means of the application of the generation step to our running example, we derive the following set of modified rules

$$\begin{aligned} p^b(X) \vee p^b(Y) &:- magic_p^b(X), magic_p^b(Y), a(X, Y), r^b(X). \\ p^b(Y) \vee p^b(X) &:- magic_p^b(Y), magic_p^b(X), a(X, Y), r^b(X). \end{aligned}$$

Finally, for the sake of completeness we point out that the rewritten program \mathcal{P}' is constituted by the above rules, plus the magic rules of Example 7 and the following rule $r^b(X) :- magic_r^b(X), p^b(X)$. obtained by considering the propagation of r^b . \square

As pointed out before, the spirit of this approach is the same of the one proposed in [7] (even though further optimized), and, hence, the soundness and completeness of the above algorithm follow easily.

Proposition 1. *Let \mathcal{P} be a Datalog[∨] program, let Q be a query, and let \mathcal{P}' be the result of the Algorithm **DisjunctiveMagicSet**. Then, $\mathcal{P} \equiv_Q \mathcal{P}'$ under both the cautious and brave semantics.*

5 Identifying Redundant Rules

In this section we address one drawback of the magic-set method when extended to disjunctive Datalog programs, consisting of the generation of a number of redundant, useless rules, which may reduce the optimization effect. For a better understanding of this point, let us return to Example 8; as the careful reader may have noticed, the first two modified rules coincide, even though they are generated by adorning different head atoms of the same rule. Such a redundancy is, indeed, intrinsic in any possible implementation; nonetheless, we point out that compared to [7], our rewriting reduces drastically the presence of redundant rules, since it does not require any additional predicates or any additional rules, which are an additional source of redundancy.

In our implementation, great attention has been devoted to this issue, since we have experimentally observed that it is crucial for the whole optimization process for several problem domains. In particular, we propose an efficient (polynomial time) heuristic for the individuation of such rules, whose relevance is not restricted to the magic-set method in itself.

In fact, this heuristic has been integrated into the core of the DLV system, in order to identify redundancy in any type of program. However,

redundant rules are less frequently found in user-specified programs. We point out that the complexity results obtained for subsumption problems implies that such an heuristic cannot be complete. Still, we are able to show its soundness, i.e., all the rules identified by the heuristic are indeed subsumed by some other rule.

Before presenting the algorithm we introduce some concepts and notations. Let C be a set of literals, and X_1, \dots, X_n the variables occurring in C . Then, a *substitution* for C is a finite set of pairs $\vartheta = \{X_1/t_1, \dots, X_n/t_n\}$ where t_1, \dots, t_n are terms (either variables or constants). Moreover, we denote by $\vartheta(C)$ the set of literals obtained from C by simultaneously replacing each occurrence of X_i by t_i ($1 \leq i \leq n$).

Definition 1. Let r_1 and r_2 be two rules of \mathcal{P} . Then, r_1 is *subsumed* by r_2 (denoted by $r_1 \sqsubseteq r_2$) if there exists a substitution ϑ for $H(r_1) \cup B(r_1)$, such that $\vartheta(H(r_1)) \subseteq H(r_2)$ and $\vartheta(B(r_1)) \subseteq B(r_2)$. A rule r_1 is *redundant* if there exists a rule r_2 such that $r_1 \sqsubseteq r_2$. \square

Definition 2. Let L_1 and L_2 be two sets of literals. Then, a *matcher* for L_1 w.r.t. L_2 is a substitution ϑ for L_1 such that $\vartheta(L_1) = L_2$. \square

Given two rules r_1 and r_2 of a logic program, the concept of subsumption can be restated in terms of existence of matchers.

Definition 3. Let r_1 and r_2 be two disjunctive rules, l be a literal of r_1 . Then, l is a *candidate for the subsumption verification* for r_1 w.r.t. r_2 if there exists a substitution ϑ such that (i) if $l \in B(r_1)$, then there exists $l' \in B(r_2)$ such that $\vartheta(l) = \{l'\}$, and (ii) if $l \in H(r_1)$, then there exists $l' \in H(r_2)$ such that $\vartheta(l) = \{l'\}$. Moreover, ϑ is called *matcher for the candidate* l . \square

Note that in the above definition, each matcher ϑ satisfying conditions 1 and 2 gives rise to at least one literal of r_2 that can be matched with l ; indeed, we denote by $matchable(l, r_1, r_2)$ the set of all such literals of r_2 .

We point out that given a literal l , three different representative situations may occur: i) there exists exactly one candidate literal, i.e., $|matchable(l, r_1, r_2)| = 1$, (in such a case the literal is said *deterministic*), ii) $|matchable(l, r_1, r_2)| = 0$, or iii) $|matchable(l, r_1, r_2)| > 1$, i.e., the selection is non-deterministic. Obviously, the source of the complexity of the problem lies in literals of the latter type. Note that in the point ii) above l cannot be a candidate.

In the following, given a rule r , and a set of literals L of r , we denote by $r - L$ the rule obtained by removing from r each literal in L . Moreover,

given a substitution ϑ we denote by $\vartheta(r)$ the rule obtained by the application of ϑ to both the head and the body of r . Our heuristic is based on the following observation.

Proposition 2. *Let r_1 and r_2 be two disjunctive rules. Then, r_1 is subsumed by r_2 if and only if there exists an ordering of all the literals in $H(r_1) \cup B(r_1)$ of the form l_1, \dots, l_m and a sequence of substitutions $\vartheta_1, \dots, \vartheta_m$, such that*

1. *each l_i is a candidate for the subsumption verification for $\vartheta(r_1 - \{l_1, \dots, l_{i-1}\})$ w.r.t. r_2 , where $\vartheta = \vartheta_1 \cup \dots \cup \vartheta_{i-1}$, and*
2. *ϑ_i is a matcher for the candidate l_i .*

Essentially, our heuristic tries to construct the sequence l_1, \dots, l_m of the above proposition and the associated substitutions $\vartheta_1, \dots, \vartheta_n$ in an incremental way. At the first step we choose l'_1 in r_1 and l''_1 in r_2 such that there exists a matcher ϑ_1 for l'_1 and l''_1 . Then, at each successive step, say i , we choose a literal l'_i of r_1 that is a candidate for the subsumption verification. Moreover, in case there is more than one candidate, we select the one which is *preferred* according to the following definition.

Definition 4. Let r_1 and r_2 be two disjunctive rules, let l_1, \dots, l_j be a sequence of distinct literals of r_1 and $\vartheta_1, \dots, \vartheta_j$ be a sequence of matchers. Then, a candidate literal l of $r_1 - \{l_1, \dots, l_j\}$ is *preferable* to a candidate literal l' of $r_1 - \{l_1, \dots, l_j\}$ w.r.t. l_1, \dots, l_j if i) $|\text{matchable}(l, r_1, r_2)| < |\text{matchable}(l', r_1, r_2)|$ or ii) $|\text{matchable}(l, r_1, r_2)| = |\text{matchable}(l', r_1, r_2)|$ and the number of distinct variables not yet matched w.r.t. $\vartheta_1, \dots, \vartheta_j$ of l is greater than that of l' . Moreover, a literal l in r_1 is *preferred* if it does not exist a literal preferable to it. \square

Note that a deterministic literal is always preferable to a non-deterministic one. Moreover, we point out that, in general, deciding whether l is a preferred literal can be done in polynomial time.

The heuristic is based on the idea of selecting the preferred literal as described in the above definition. Moreover, each time a literal l is selected, it is deleted from the rule, while other than storing the sequence $\vartheta_1, \dots, \vartheta_m$, we simply update a unique substitution ϑ . The soundness of this approach can be easily proved.

It is worth noting that the algorithm runs in polynomial time, while the problem is NP-hard, Hence completeness cannot hold; nevertheless, several experimental results confirmed its effectiveness.

6 Conclusions

We have presented a new technique for the optimization of (partially) bound queries, that extends the magic-set method to the case of disjunctive databases. As one intrinsic drawback of this method is the generation of many redundant rules, we have also introduced a strategy for their efficient identification. We point out that the proposed heuristic can be applied also for identifying redundant rules in programs in which no optimization is performed, albeit this situation is likely to occur less frequently. We are currently working on the definition of more focused prevention strategies that can be directly integrated into the optimization module of the DLV system.

Finally, our long-term objective is to extend the application of the magic set method to the case of disjunctive programs with constraints, and, more ambitiously, to the case of general unstratified programs.

References

1. Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
2. Thomas Eiter, Nicola Leone, Cristinel Mateis, Gerald Pfeifer, and Francesco Scarcello. A Deductive System for Nonmonotonic Reasoning. In Jürgen Dix and Ulrich Furbach and Anil Nerode, editor, *Proceedings of the 4th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR'97)*, number 1265 in Lecture Notes in AI (LNAI), pages 363–374, Berlin, 1997. Springer.
3. Thomas Eiter, Nicola Leone, Cristinel Mateis, Gerald Pfeifer, and Francesco Scarcello. The KR System `dlv`: Progress Report, Comparisons and Benchmarks. In Anthony G. Cohn, Lenhart Schubert, and Stuart C. Shapiro, editors, *Proceedings Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 406–417. Morgan Kaufmann Publishers, 1998.
4. Wolfgang Faber and Gerald Pfeifer. DLV homepage, since 1996. <http://www.dlvsystem.com/>.
5. M. Gelfond and V. Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9:365–385, 1991.
6. Sergio Greco. Optimization of Disjunction Queries. In Danny De Schreye, editor, *Proceedings of the 16th International Conference on Logic Programming (ICLP'99)*, pages 441–455, Las Cruces, New Mexico, USA, November 1999. The MIT Press.
7. Sergio Greco. Binding Propagation Techniques for the Optimization of Bound Disjunctive Queries. *IEEE Transactions on Knowledge and Data Engineering*, 15(2):368–385, March/April 2003. Extended Abstract appeared as [6].
8. Teodor C. Przymusiński. Stable Semantics for Disjunctive Programs. *New Generation Computing*, 9:401–424, 1991.