

NP Datalog: A logic language for expressing NP search and optimization problems

SERGIO GRECO, CRISTIAN MOLINARO, IRINA TRUBITSYNA
and ESTER ZUMPARO

*Dipartimento di Elettronica Informatica e Sistemistica (DEIS), Università della Calabria,
87036 Rende, Italy*

(e-mail: {greco, cmolinaro, irina, zumpano}@deis.unical.it)

submitted 4 April 2008; revised 2 September 2009; accepted 2 November 2009

Abstract

This paper presents a logic language for expressing NP search and optimization problems. Specifically, first a language obtained by extending (positive) DATALOG with intuitive and efficient constructs (namely, stratified negation, constraints, and exclusive disjunction) is introduced. Next, a further restricted language only using a restricted form of disjunction to define (nondeterministically) subsets (or partitions) of relations is investigated. This language, called NP Datalog, captures the power of DATALOG[∇] in expressing search and optimization problems. A system prototype implementing NP Datalog is presented. The system translates NP Datalog queries into Optimization Programming Language (OPL) programs which are executed by the ILOG OPL Development Studio. Our proposal combines easy formulation of problems, expressed by means of a declarative logic language, with the efficiency of the ILOG System. Several experiments show the effectiveness of this approach.

KEYWORDS: logic languages, stable model semantics, constraint programming, expressivity and complexity of declarative query languages

1 Introduction

It is well known that NP search problems can be formulated by means of DATALOG[∇] (Datalog with unstratified negation) queries under nondeterministic stable model semantics so that each stable model corresponds to a possible solution (Marek and Truszczyński 1991; Saccà 1997). NP optimization problems can be formulated by adding a *max* (or *min*) construct to select the stable model (thus, the solution), which maximizes (resp., minimizes) the result of a polynomial function applied to the answer relation. For instance, consider the *Vertex Cover* problem of the following example.

Example 1

Given an undirected graph $G = \langle N, E \rangle$, a subset V of the vertexes N is a *vertex cover* of G if every edge of G has at least one end in V . The problem can be formulated

in terms of the query $\langle \mathcal{P}_1, v(X) \rangle$, where \mathcal{P}_1 is the following DATALOG[⊃] program:

$$\begin{aligned} v(X) &\leftarrow \text{node}(X), \neg nv(X). \\ nv(X) &\leftarrow \text{node}(X), \neg v(X). \\ c &\leftarrow \text{edge}(X, Y), \neg v(X), \neg v(Y), \neg c. \end{aligned}$$

and the predicates `node` and `edge` define, respectively, the vertexes and the edges of the graph by means of a suitable number of facts. The first two rules define a partition of the relation `node` (v being the vertex cover), whereas the last one enforces every stable model to correspond to some vertex cover as it is satisfied only if the conjunction $\text{edge}(X, Y), \neg v(X), \neg v(Y)$ is false (otherwise the program does not have stable models).

The min vertex cover problem can be expressed by selecting a stable model which minimizes the number of elements in v ; this is expressed by means of the query $\langle \mathcal{P}_1, \min|v(X)| \rangle$. \square

The problem in using DATALOG[⊃] to express search and optimization problems is that the use of unrestricted negation is often neither simple nor intuitive and besides it does not allow the expressive power and complexity of queries to be limited. For instance, in the example above, the use of explicit constraints instead of standard rules would permit the distinction between rules used to infer true atoms and rules used to check properties to be satisfied.

In this paper, in order to enable a simpler and more intuitive formulation for search and optimization problems and an efficient computation of queries, DATALOG-like languages extending positive DATALOG with intuitive and efficient constructs are considered. The first language we present, denoted by DATALOG^{⊃,⊕,⊖}, extends the simple and intuitive structure of DATALOG[⊃] (DATALOG with stratified negation Ullman 1988) with two other types of “controlled” negation: rules with exclusive *disjunctive* heads and *constraint* rules. The same expressive power as DATALOG[⊃] is achieved by such a language. Next, we propose a further restricted language, called *NP Datalog*, where head disjunction is only used to define (nondeterministically) partitions of relations. This language allows us to express, in a simple and intuitive way, both *NP* search and optimization problems. As an example, let us consider again the Vertex Cover problem.

Example 2

The search query of the previous example can be expressed as $\langle \mathcal{P}_2, v(X) \rangle$ with \mathcal{P}_2 defined as follows:

$$\begin{aligned} v(X) \oplus nv(X) &\leftarrow \text{node}(X). \\ \Leftarrow &\text{edge}(X, Y), \neg v(X), \neg v(Y). \end{aligned}$$

where \oplus denotes exclusive disjunction, i.e., if the body of the rule is true, then exactly one atom in the head is true. The rule with empty head defines a constraint, i.e., a rule which is satisfied only if the body is false. The first rule guesses a partition of `node` whereas the second one is a constraint stating that two connected

nodes cannot be both outside the cover, which is defined by the nodes belonging to v . \square

Contribution

The main contribution of this paper is the proposal of a simple and intuitive language where the use of stable model semantics allows us to refrain from uncontrolled forms of unstratified negation¹ and avoid both undefinedness and unnecessary computational complexity.

More precisely, the paper presents the language \mathcal{NP} Datalog, which extends DATALOG ^{\neg} with constraints and head disjunction, where the latter is used only to define (nondeterministically) partitions of “deterministic” relations. This language allows both \mathcal{NP} search and optimization problems to be expressed in a simple and intuitive way.

The simplicity of the \mathcal{NP} Datalog language enables queries to be easily translated into other formalisms such as constraint programming languages, which are well-suited to compute programs defining \mathcal{NP} problems. This paper also shows how \mathcal{NP} Datalog queries can be translated into OPL (*Optimization Programming Language*) (Van Hentenryck 1988; Van Hentenryck *et al.* 1999) programs.

Several examples of queries expressing \mathcal{NP} problems suggest that logic formalisms allow an easy formulation of queries. On the other hand, constraint programming systems permit an efficient execution. Therefore, \mathcal{NP} Datalog can also be used to define a logic interface for constraint programming solvers. We have implemented a system prototype which translates \mathcal{NP} Datalog queries into OPL programs, which are then executed by means of the ILOG OPL Development Studio². The effectiveness of our approach is demonstrated by several experiments comparing \mathcal{NP} Datalog with other systems.

With respect to other logic languages previously proposed (Greco *et al.* 1995; Eiter *et al.* 1997; Cadoli *et al.* 2000; Simons *et al.* 2002; Cadoli and Schaerf 2005), the novelty of the paper is that it considers an answer set programming language able to express the complete set of \mathcal{NP} decision, search and optimization problems, by using a restricted form of unstratified negation.

Organization. The paper is organized as follows. Section 2 introduces syntax and semantics of DATALOG ^{\neg} and its ability to express \mathcal{NP} search and optimization queries under nondeterministic stable model semantics. Section 3 introduces the DATALOG ^{\neg, \oplus, \Leftarrow} language, and shows its ability to express \mathcal{NP} search and optimization problems. Section 4 presents the \mathcal{NP} Datalog language, obtained introducing simple restrictions to DATALOG ^{\neg, \oplus, \Leftarrow} and shows that \mathcal{NP} Datalog has the same expressive power as DATALOG ^{\neg, \oplus, \Leftarrow} and DATALOG ^{\neg} . Section 5 illustrates how \mathcal{NP} Datalog queries can be translated into OPL programs and presents several experiments showing the effectiveness of the proposed approach. Section 6 discusses several related languages and systems recently proposed in the literature. Finally, conclusions are drawn in Section 7.

¹ The constructs here considered, essentially, force the use of a restricted form of unstratified negation.

² ILOG OPL Studio. URL: <http://www.ilog.com/products/oplstudio/>

2 DATALOG[⊖]

It is assumed that the reader is familiar with the basic terminology and notation of relational databases and database queries (Abiteboul *et al.* 1995; Ullman 1988).

Syntax. A DATALOG[⊖] rule r is of the form $A \leftarrow B_1, \dots, B_m, \neg B_{m+1}, \dots, \neg B_n$, where A is an atom (*head* of the rule) and $B_1, \dots, B_m, \neg B_{m+1}, \dots, \neg B_n$ (with $n \geq 0$) is a conjunction of literals (*body* of the rule). A fact is a ground rule with empty body. Generally, predicate symbols are partitioned into two different classes: EDB (or extensional database predicates), i.e., defined by the ground facts of a database, and IDB (or intensional database predicates), i.e., defined by the rules of the program. The definition of a predicate p consists of all the rules (or facts) having p in the head.

A database D consists of all the facts defining EDB predicates, whereas a DATALOG[⊖] program \mathcal{P} consists of the rules defining IDB predicates. It is assumed that programs are *safe* (Ullman 1988), i.e., variables appearing in the head or in negative body literals are range restricted as they appear in some positive body literal, and that possible constants in \mathcal{P} are taken from the database domain. For each rule, variables appearing in the head are said to be universally quantified, whereas the remaining variables are said to be existentially quantified.

The class of all DATALOG[⊖] programs is simply called DATALOG[⊖]; the subclass of all positive (resp. stratified) programs is called DATALOG (resp. DATALOG^{⊖s}) (Abiteboul *et al.* 1995). Observe that $\text{DATALOG} \subseteq \text{DATALOG}^{\ominus s} \subseteq \text{DATALOG}^{\ominus}$ (the class of Datalog queries with possibly unstratified negation).

Semantics. The semantics of a positive program \mathcal{P} is given by the unique minimal model $\mathcal{M}\mathcal{M}(\mathcal{P})$. The semantics of programs with negation \mathcal{P} is given by the set of its stable models $\mathcal{SM}(\mathcal{P})$. An interpretation M is a *stable model* (or *answer set*) of \mathcal{P} if M is the unique minimal model of the positive program \mathcal{P}^M , where \mathcal{P}^M denotes the positive logic program obtained from $\text{ground}(\mathcal{P})$ by removing (i) all rules r such that there is a negative literal $\neg A$ in the body of r and A is in M , and (ii) all the negative literals from the remaining rules (Gelfond and Lifschitz 1988). It is well known that a program may have n stable models with $n \geq 0$. Stratified programs have a unique stable model which coincides with the *perfect model*, obtained by partitioning the program into an ordered number of suitable subprograms (called “strata”) and computing the fixpoints of every stratum in their order (Ullman 1988). Given a set of ground atoms S and an atom $g(t)$, $S[g]$ (resp. $S[g(t)]$) denotes the set of g -tuples (resp. tuples matching $g(t)$) in S .

DATALOG[⊖] Search and Optimization Queries. Search and optimization problems can be expressed using different logic formalisms such as Datalog with unstratified negation.

Definition 1

A DATALOG[⊖] *search query* is a pair $Q = \langle \mathcal{P}, g(t) \rangle$, where \mathcal{P} is a DATALOG[⊖] program and $g(t)$ is an atom s.t. g is an IDB predicate of \mathcal{P} . The answer to Q over a database D is $Q(D) = \{M[g(t)] \mid M \in \mathcal{SM}(\mathcal{P} \cup D)\}$. The answer to the DATALOG[⊖] *optimization query* $\text{opt}(Q) = \langle \mathcal{P}, \text{opt}|g(t)| \rangle$, where *opt* is either *max* or *min*, over a database D ,

consists of the answers in $Q(D)$ with the maximum or minimum (resp., if $opt = max$ or min) cardinality and is denoted by $opt(Q)(D)$. \square

Observe that, for the sake of simplicity, optimization queries computing the maximum or minimum cardinality of the output relation are considered, although any polynomial function might be used. Therefore, the answer here considered is a set of sets of atoms. Possible and certain answers can be obtained by considering the union or the intersection of the sets, respectively. Instead of considering possible and certain reasoning, we introduce nondeterministic answers as follows.

Definition 2

A (nondeterministic) answer to a DATALOG[⊃] search query Q applied to a database D is $Q(D) = S$ where S is a relation selected nondeterministically from $Q(D)$. A (nondeterministic) answer to a DATALOG[⊃] optimization query $opt(Q)$ over a database D is $opt(Q)(D) = S$ where S is a relation selected nondeterministically from $opt(Q)(D)$. \square

It is worth noting that, like for search queries, also for optimization queries the relation with optimal cardinality rather than just the cardinality is returned. Thus, given a search query $Q = \langle \mathcal{P}, g(t) \rangle$ and a database D , the output relation $Q(D)$ consists of all tuples $g(u)$ matching $g(t)$ and belonging to a stable model M of $\mathcal{P} \cup D$, selected nondeterministically. For a given optimization query $OQ = \langle \mathcal{P}, \min|g(t)| \rangle$ (resp. $\langle \mathcal{P}, \max|g(t)| \rangle$), the output relation $OQ(D)$ consists of the set of tuples $g(u)$ matching $g(t)$ and belonging to a stable model M of $\mathcal{P} \cup D$, selected nondeterministically among those which minimize (resp. maximize) the cardinality of the output relation. From now on, we concentrate our attention on nondeterministic queries. An example of a nondeterministic DATALOG[⊃] search query is shown in Example 1; the optimization problem is expressed by rewriting the query goal as $\langle \mathcal{P}_1, \min|v(X)| \rangle$ whose meaning is to further restrict the set of stable models to those for which v has minimum cardinality.

In Saccà (1997) and Greco and Saccà (2002) it has been shown that DATALOG[⊃] search and optimization queries under (nondeterministic) stable model semantics express the class of \mathcal{NP} search and optimization problems (denoted, respectively, by $Q\mathcal{NP}\mathcal{MV}$ and $OPT\ Q\mathcal{NP}\mathcal{MV}$).

3 DATALOG^{⊃,⊕,⊐}

The problem in using DATALOG[⊃] to express search and optimization problems is that the use of unrestricted negation is often neither simple nor intuitive and, besides, it does not allow expressive power (and complexity) to be controlled and in some cases might also lead writing queries having no stable models. In order to avoid these problems, we present a language, called DATALOG^{⊃,⊕,⊐}, where unstratified negation is embedded into built-in constructs, so that the user is forced to write programs using restricted forms of negation without loss of expressive power. Specifically, DATALOG^{⊃,⊕,⊐} extends DATALOG[⊃] with two simple built-in constructs: head (exclusive) disjunction and constraints, denoted by \oplus and \Leftarrow , respectively.

Syntax. In the following rules, $Body(\mathbf{X})$ and $Body(\mathbf{X}, \mathbf{Y}, L)$ are conjunctions of literals, whereas \mathbf{X} and \mathbf{Y} are vectors of range restricted variables.

An (*exclusive*) *disjunctive rule* is of the form:

$$p_1(\mathbf{X}_1) \oplus \cdots \oplus p_k(\mathbf{X}_k) \leftarrow Body(\mathbf{X}), \quad (1)$$

where $\mathbf{X}_i \subseteq \mathbf{X}$ for all $i \in [1..k]$. The intuitive meaning of such a rule is that if $Body(\mathbf{X})$ is true, then exactly one head atom $p_i(\mathbf{X}_i)$ must be true.

A special form of disjunctive rule, called *generalized disjunctive rule*, of the form:

$$\oplus_L p(\mathbf{X}, L) \leftarrow Body(\mathbf{X}, \mathbf{Y}, L) \quad (2)$$

is also allowed. In this rule the number of head disjunctive atoms is not fixed, but depends on the database instance and on the current computation (stable model). The intuitive meaning of this rule is that the relation defined by $\pi_{\mathbf{X}} Body(\mathbf{X}, \mathbf{Y}, L)$ (the projection of the relation $Body(\mathbf{X}, \mathbf{Y}, L)$ on the attributes defined by \mathbf{X}) is partitioned into a number of subsets equal to the cardinality of the relation $\pi_L Body(\mathbf{X}, \mathbf{Y}, L)$ (the number of distinct values for the variable L). Some examples of generalized disjunctive rules will be presented in the next section.

A *constraint (rule)* is of the form:

$$\Leftarrow Body(\mathbf{X}). \quad (3)$$

A ground constraint rule is satisfied w.r.t. an interpretation I if the body of the rule is false in I . We shall often write constraints using rules of the form $A_1 \vee \dots \vee A_k \Leftarrow B_1, \dots, B_m$ (or $B_1, \dots, B_m \Rightarrow A_1 \vee \dots \vee A_k$) to denote a constraint of the form $\Leftarrow B_1, \dots, B_m, \neg A_1, \dots, \neg A_k$ (i.e., negative literals are moved from the body to the head). For instance, the constraint $\Leftarrow \text{edge}(X, Y), \neg v(X), \neg v(Y)$ of Example 2 can be rewritten as $v(X) \vee v(Y) \Leftarrow \text{edge}(X, Y)$ or as $\text{edge}(X, Y) \Rightarrow v(X) \vee v(Y)$. Here the symbol \vee denotes inclusive disjunction and is different from \oplus , as the latter denotes exclusive disjunction. It should be recalled that inclusive disjunction allows more than one atom to be true while exclusive disjunction allows only one atom to be true.

Definition 3

A DATALOG ^{$\neg, s, \oplus, \Leftarrow$} *search query* is a pair $Q = \langle \mathcal{P}, g(t) \rangle$, where \mathcal{P} is a DATALOG ^{$\neg, s, \oplus, \Leftarrow$} program and $g(t)$ is an IDB atom. A DATALOG ^{$\neg, s, \oplus, \Leftarrow$} *optimization query* is a pair $\langle \mathcal{P}, \text{opt}|g(t)| \rangle$, where *opt* is either *max* or *min*. \square

The query $\langle \mathcal{P}_2, v(X) \rangle$ of Example 2 is a DATALOG ^{$\neg, s, \oplus, \Leftarrow$} search query, whereas the query $\langle \mathcal{P}_2, \min|v(X)| \rangle$ is a DATALOG ^{$\neg, s, \oplus, \Leftarrow$} optimization query.

Semantics. The declarative semantics of a DATALOG ^{$\neg, s, \oplus, \Leftarrow$} query is given in terms of an “equivalent” DATALOG ^{\neg} query and stable model semantics. Specifically, given a DATALOG ^{$\neg, s, \oplus, \Leftarrow$} program \mathcal{P} , $st(\mathcal{P})$ denotes the standard DATALOG ^{\neg} program derived from \mathcal{P} as follows:

- (1) Every standard rule in \mathcal{P} belongs to $st(\mathcal{P})$,

- (2) Every disjunctive rule $r \in \mathcal{P}$ of the form (1) is translated into k rules of the form:

$$p_j(\mathbf{X}_j) \leftarrow \text{Body}(\mathbf{X}), \neg p_1(\mathbf{X}_1), \dots, \neg p_{j-1}(\mathbf{X}_{j-1}), \neg p_{j+1}(\mathbf{X}_{j+1}), \dots, \neg p_k(\mathbf{X}_k)$$

with $j \in [1..k]$, plus $((k - 1) \times k)/2$ constraints of the form:

$$\Leftarrow \text{Body}(\mathbf{X}), p_i(\mathbf{X}_i), p_j(\mathbf{X}_j)$$

with $i, j \in [1..k]$ and $i < j$. It is worth noting that the constraints are necessary only if p_j is defined by some other rule.

- (3) Every generalized disjunctive rule of the form (2) is translated into the two rules:

$$\begin{aligned} p(\mathbf{X}, L) &\leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}, L), \neg \text{diff_}p(\mathbf{X}, L), \\ \text{diff_}p(\mathbf{X}, L) &\leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}, L), p(\mathbf{X}, L'), L' \neq L, \end{aligned}$$

where $\text{diff_}p$ is a new predicate symbol and L' is a new variable, plus the constraint:

$$\Leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}, L), p(\mathbf{X}, L_1), p(\mathbf{X}, L_2), L_1 \neq L_2.$$

Here $\text{diff_}p$ is used to avoid inferring two ground atoms $p(x, l_1)$ and $p(x, l_2)$ with $l_1 \neq l_2$. Observe that even in this case the constraint has to be introduced if p is defined by some other rule.

- (4) Every constraint rule of the form (3) is translated into a rule of the form:

$$c \leftarrow \text{Body}(\mathbf{X}), \neg c,$$

where c is a new predicate symbol not appearing elsewhere.

For any $\text{DATALOG}^{\neg, \oplus, \Leftarrow}$ search query $Q = \langle \mathcal{P}, g(t) \rangle$ (resp. optimization query $OQ = \langle \mathcal{P}, \text{opt}[g(t)] \rangle$), $st(Q) = \langle st(\mathcal{P}), g(t) \rangle$ (resp. $st(OQ) = \langle st(\mathcal{P}), \text{opt}[g(t)] \rangle$) denotes the corresponding DATALOG^{\neg} (resp. optimization) query.

Definition 4

Given a $\text{DATALOG}^{\neg, \oplus, \Leftarrow}$ query Q and a database D , the (nondeterministic) answer to the query Q over D is obtained by applying the DATALOG^{\neg} query $st(Q)$ to D , i.e., $Q(D) = st(Q)(D)$. □

It is worth noting that $\text{DATALOG}^{\neg, \oplus, \Leftarrow}$ has the same expressive power of DATALOG^{\neg} , that is both \mathcal{NP} search and optimization problems can be expressed by means of $\text{DATALOG}^{\neg, \oplus, \Leftarrow}$ queries under stable model semantics (Zumpano *et al.* 2004). The further restricted languages $\text{DATALOG}^{\neg, \oplus}$ (Datalog with stratified negation and exclusive disjunction) and $\text{DATALOG}^{\oplus, \Leftarrow}$ (Datalog with exclusive disjunction and constraints) have the same expressive power. A similar result has been presented in (East and Truszczyński 2006), where it has been shown that positive Datalog with constraints and head (inclusive) disjunction, called *PS logic*, has the same expressive power as DATALOG^{\neg} . Clearly, $\text{DATALOG}^{\oplus, \Leftarrow}$ is captured by PS logic, since exclusive disjunction can be emulated by using inclusive disjunction and constraints. It is interesting to observe that analogous results could be obtained for others ASP languages. For instance, (positive) Datalog with cardinality constraints, as proposed in Smodels, captures the expressive power of DATALOG^{\neg} since exclusive disjunction

and denial constraints can be emulated by means of cardinality constraints (Niemela et al. 1999).

4 \mathcal{NP} Datalog

We now present a simplified version of $\text{DATALOG}^{\neg, \oplus, \leftarrow}$ introducing further restrictions on disjunctive rules. The basic idea consists in restricting $\text{DATALOG}^{\neg, \oplus, \leftarrow}$ to obtain, without loss of expressive power, a language which can be executed more efficiently or easily translated in other formalisms.

In the following rules, $\text{Body}(\mathbf{X}, \mathbf{Y})$ denotes a conjunction of literals where \mathbf{X} and \mathbf{Y} are vectors of range restricted variables.

A *partition rule* is a disjunctive rule of the form:

$$p_1(\mathbf{X}) \oplus \cdots \oplus p_k(\mathbf{X}) \leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}) \quad (4)$$

or of the form

$$p_0(\mathbf{X}, c_1) \oplus \cdots \oplus p_0(\mathbf{X}, c_k) \leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}), \quad (5)$$

where p_0, p_1, \dots, p_k are distinct IDB predicates not defined elsewhere in the program and c_1, \dots, c_k are distinct constants. The intuitive meaning of these rules is that the projection of the relation defined by $\text{Body}(\mathbf{X}, \mathbf{Y})$ on \mathbf{X} is partitioned nondeterministically into k relations or k distinct sets of the same relation. Clearly, every rule of form (5) can be rewritten into a rule of form (4) and vice versa.

A *generalized partition rule* is a (generalized) disjunctive rule of the form:

$$\oplus_L p(\mathbf{X}, L) \leftarrow \text{Body}(\mathbf{X}, \mathbf{Y}), d(L), \quad (6)$$

where p is an IDB predicate not defined elsewhere and d is a database domain predicate specifying the domain of the variable L . The intuitive meaning of such a rule is that the projection of the relation defined by $\text{Body}(\mathbf{X}, \mathbf{Y})$ on \mathbf{X} is partitioned into a number of subsets equal to the cardinality of the relation d .

In the following, the existence of *subset rules* is also assumed, i.e., rules of the form

$$s(\mathbf{X}) \subseteq \text{Body}(\mathbf{X}, \mathbf{Y}), \quad (7)$$

where s is an IDB predicate not defined elsewhere in the program. Observe that a subset rule of the form above corresponds to the generalized partition rule with $d = \{0, 1\}$. On the other hand, every generalized partition rule can be rewritten into a subset rule and constraints.

In the previous rules, $\text{Body}(\mathbf{X}, \mathbf{Y})$ is a conjunction of literals not depending on predicates defined by partition or subset rules. We recall that the subset rules used here are based on the proposal in Greco and Saccà (1997). A similar type of subset rules has also been proposed by Gelfond (2002) where the language ASET-Prolog (an extension of A-Prolog with sets) is presented.

Definition 5

An \mathcal{NP} Datalog program consists of three distinct sets of rules:

- (1) *partition* and *subset* rules defining *guess (IDB) predicates*,

- (2) standard stratified datalog rules defining standard (IDB) predicates, and
- (3) constraints rules.

where every guess predicate is defined by a unique subset or partition rule. □

According to the definition above, the set of IDB predicates of an \mathcal{NP} Datalog program can be partitioned into two distinct subsets (namely, guess and standard) depending on the rules used to define them. Clearly, predicates defined by partition or subset rules are not recursive as the body of these rules cannot contain guess predicates or predicates depending on guess predicates.

Example 1

Vertex cover (version 3). The \mathcal{NP} Datalog program:

$$\begin{aligned} v(X) &\subseteq \text{node}(X). \\ \text{edge}(X, Y) &\Rightarrow v(X) \vee v(Y). \end{aligned}$$

is derived from the one presented in Example 2 by replacing the disjunctive rule defining v with a subset rule. □

It is important to note that here a simpler form of $\text{DATALOG}^{\neg, \oplus, \Leftarrow}$ queries is considered. Therefore, $\mathcal{NP} \text{ Datalog} \subseteq \text{DATALOG}^{\neg, \oplus, \Leftarrow}$ and every \mathcal{NP} Datalog query can be rewritten into an equivalent DATALOG^{\neg} query.

Definition 6

An \mathcal{NP} Datalog search query is a pair $Q = \langle \mathcal{P}, g(t) \rangle$, where \mathcal{P} is an \mathcal{NP} Datalog program and $g(t)$ is an IDB atom denoting the output relation. An \mathcal{NP} Datalog optimization query is a pair $\langle \mathcal{P}, \text{opt}|g(t)| \rangle$, where $\text{opt} \in \{max, min\}$. □

Observe that, for the sake of simplicity, our attention is restricted to optimization queries computing the maximum or minimum cardinality of the output relation, although any polynomial function might be used. Moreover, as stated by the following theorem, \mathcal{NP} Datalog captures the complexity classes of \mathcal{NP} search and optimization problems.

Theorem 1

- (1) $\text{search}(\mathcal{NP} \text{ Datalog}) = \mathcal{QNP}$, and
- (2) $\text{opt}(\mathcal{NP} \text{ Datalog}) = \text{OPT } \mathcal{QNP}$.

Proof. Membership is trivial as $\mathcal{NP} \text{ Datalog} \subseteq \text{DATALOG}^{\neg, \oplus, \Leftarrow}$.

To prove hardness the well-known Fagin's result is used (Fagin 1974; see also Johnson 1990; Papadimitriou 1994): it states that every \mathcal{NP} recognizable database collection is defined by an existential second-order formula $\exists R\Phi$, where R is a list of new predicate symbols and Φ is a first-order formula involving predicate symbols in a database schema \mathcal{DS} and in R . As shown in (Kolaitis and Papadimitriou 1991), this formula is equivalent to one of the form (*second-order Skolem normal form*)

$$(\exists \mathcal{S})(\forall \mathbf{X})(\exists \mathbf{Y})(\theta_1(\mathbf{X}, \mathbf{Y}) \vee \dots \vee \theta_k(\mathbf{X}, \mathbf{Y})),$$

where \mathcal{S} is a superlist of R , $\theta_1, \dots, \theta_k$ are conjunctions of literals involving variables in \mathbf{X} and \mathbf{Y} , and predicate symbols in \mathcal{S} and $D\mathcal{B}$. Consider the

program \mathcal{P} :

$$\begin{array}{ll} s_j(W_j) \oplus \hat{s}_j(W_j) & \leftarrow (\forall s_j \in \mathcal{S}), \\ q(X) & \leftarrow \theta_i(X, Y) \quad (1 \leq i \leq k), \\ g & \leftarrow \neg q(X). \end{array}$$

The first group of rules selects a set of constants from the database domain, for each predicate symbol s_j . The second group of rules implements the above second-order formula. The third rule checks if there is some \mathbf{X} for which the formula is not satisfied. Therefore, the formula is satisfied if, and only if, there is a stable model M such that $\neg g \in M$. \square

Thus, \mathcal{NP} Datalog has the same expressive power as both $\text{DATALOG}^{\neg, \oplus, \leftarrow}$ and DATALOG^{\neg} . The idea underlying \mathcal{NP} Datalog is that \mathcal{NP} search and optimization problems can be expressed using partition (or subset) rules to guess partitions or subsets of sets, whereas constraints are used to verify properties to be satisfied by guessed sets or sets computed by means of stratified rules. It is important to observe that the proof of Theorem 1 follows a schema which has been used in other proofs concerning the expressive power of Datalog with negation under stable model semantics (Schlipf 1995; Saccà 1997; Baral 2003). Indeed, in such proofs (see, for instance, the proof of Theorem 6.3.1 in Baral 2003) negation is only used to express exclusive disjunction and constraints.

Although the aim of this work is not the definition of techniques for the efficient computation of queries, we would point out that \mathcal{NP} Datalog programs can be computed following the classical stratified fixpoint algorithm enriched with a guess and check technique.

The advantage of expressing search and optimization problems by using rules with built-in predicates rather than standard DATALOG^{\neg} rules is that the use of built-in atoms preserves simplicity and intuition in expressing problems and allows queries to be easily optimized and translated into other target languages for which efficient executors exist. A further advantage is that the use of built-in predicates in expressing optimization queries permits us to easily identify problems for which “approximate” answers can be found in polynomial time. For instance, maximization problems defined by constraint free \mathcal{NP} Datalog queries where negation is only applied to guess atoms or atoms not depending on guess atoms (called deterministic) are constant approximable (Greco and Saccà 1997). Indeed, these problems belong to the class of constant approximable optimization problems $\text{MAX } \Sigma_1^3$ (Kolaitis and Thakur 1995) and, therefore, \mathcal{NP} Datalog could also be used to define the class of approximable optimization problems, but this is outside the scope of this paper.

\mathcal{NP} Datalog allows us to also use a finite subset of the integer domain and the standard built-in arithmetic operators. More specifically, reasoning and computing over a finite set of integer ranges is possible with the unary predicate **integer**, which

³ This class was firstly introduced in (Papadimitriou and Yannakakis 1991) as $\text{MAX } \mathcal{NP}$.

consists of the facts `integer(x)`, with $\text{MinInt} \leq x \leq \text{MaxInt}$, and the standard arithmetic operators defined over the integer domain.

The following example shows how the arithmetics operators could be used to compute prime numbers

```
composite(X) ← integer(Y), integer(Z), X = Y * Z.
prime(X) ← integer(X), not composite(X).
```

Thus, the language allows arithmetic expressions which involve variables taking integer values to appear as operands of comparison operators (see Example 9).

4.1 Examples

Some examples are now presented showing how classic search and optimization problems can be defined in *NP Datalog*.

Example 4

Max satisfiability. Two unary relations `c` and `a` are given in such a way that a fact `c(x)` denotes that `x` is a clause and a fact `a(v)` asserts that `v` is a variable occurring in some clause. We also have two binary relations `p` and `n` such that the facts `p(x, v)` and `n(x, v)` state that a variable `v` occurs in the clause `x` positively or negatively, respectively. A boolean formula, in conjunctive normal form, can be represented by means of the relations `c`, `a`, `p`, and `n`.

The maximum number of clauses simultaneously satisfiable under some truth assignment can be expressed by the query $\langle \mathcal{P}_{\text{sat}}, \max |f(X)| \rangle$ where \mathcal{P}_{sat} is the following program:

```
s(X) ⊆ a(X).
f(X) ← p(X, V), s(V).
f(X) ← n(X, V), ¬s(V). □
```

Observe that the max satisfiability problem is constant approximable as no constraints are used and negation is applied to guess atoms only.

In the following examples, a database graph $G = \langle N, E \rangle$ defined by means of the unary relation `node` and the binary relation `edge` is assumed.

Example 5

k-Coloring. Consider the well-known problem of `k-colorability` consisting in finding a *k-coloring*, i.e., an assignment of one of `k` possible colors to each node of a graph `G` such that no two adjacent nodes have the same color. The problem can be expressed by means of the *NP Datalog* query $\langle \mathcal{P}_{\text{k-col}}, \text{col}(X, C) \rangle$ where $\mathcal{P}_{\text{k-col}}$ consists of the following rules:

```
⊕C col(X, C) ← node(X), color(C).
← edge(X, Y), col(X, C), col(Y, C).
```

and the base relation `color` contains exactly `k` colors. The first rule guesses an assignment of colors to the nodes of the graph, while the constraint verifies that two joined vertices do not have the same color. □

Example 6

Min Coloring. The query modeling the Min Coloring problem is obtained from the k -coloring example by adding a rule storing the used colors as follows:

$$\begin{aligned} \oplus_C \text{ col}(X, C) &\leftarrow \text{node}(X), \text{color}(C). \\ \Leftarrow \text{ edge}(X, Y), \text{ col}(X, C), \text{ col}(Y, C). \\ \text{used_color}(C) &\leftarrow \text{col}(X, C). \end{aligned}$$

and replacing the query goal with $\text{min}|\text{used_color}(C)|$. \square

Example 7

Min Dominating Set. Given a graph $G = \langle N, E \rangle$, a subset of the vertex set $V \subseteq N$ is a dominating set if for all $u \in N - V$ there is a $v \in V$ such that $(u, v) \in E$. The \mathcal{NP} Datalog query $\langle \mathcal{P}_{\text{ds}}, v(X) \rangle$ expresses the problem of finding a dominating set, where \mathcal{P}_{ds} is the following program:

$$\begin{aligned} v(X) &\subseteq \text{node}(X). \\ \text{connected}(X) &\leftarrow \text{edge}(X, Y), v(Y). \\ \text{node}(X) \wedge \neg v(X) &\Rightarrow \text{connected}(X). \end{aligned}$$

The constraint states that every node not belonging to the dominating set, namely the relation v , must be connected to some node in v . A dominating set is said to be *minimum* if its cardinality is minimum. Therefore, the optimization problem is expressed by replacing the query goal $v(X)$ with $\text{min}|v(X)|$. \square

Note that if an \mathcal{NP} minimization query has an empty answer there is no solution for the associated search problem.

Example 8

Min Edge Dominating Set. Given a graph $G = \langle N, E \rangle$, a subset of the edge set $A \subseteq E$ is an edge dominating set if for all $e_1 \in E - A$ there is an $e_2 \in A$ such that e_1 and e_2 are adjacent. The min edge dominating set problem is defined by the \mathcal{NP} Datalog query $\langle \mathcal{P}_{\text{eds}}, \text{min}|e(X, Y)| \rangle$ where \mathcal{P}_{eds} consists of the following rules:

$$\begin{aligned} e(X, Y) &\subseteq \text{edge}(X, Y). \\ v(X) &\leftarrow e(X, Y). \\ v(Y) &\leftarrow e(X, Y). \\ \text{edge}(X, Y) &\Rightarrow v(X) \vee v(Y). \end{aligned} \quad \square$$

Example 9

N-Queens. This problem consists in placing N queens on an $N \times N$ chessboard in such a way that no two queens are in the same row, column, or diagonal. It can be expressed by the \mathcal{NP} Datalog query $\langle \mathcal{P}_{\text{queen}}, \text{queen}(R, C) \rangle$ where $\mathcal{P}_{\text{queen}}$ consists of the following rules:

$$\begin{aligned} \oplus_C \text{ queen}(R, C) &\leftarrow \text{num}(R), \text{num}(C). \\ \Leftarrow \text{ queen}(R_1, C), \text{ queen}(R_2, C), R_1 \neq R_2. \\ \Leftarrow \text{ queen}(R_1, C_1), \text{ queen}(R_2, C_2), R_1 \neq R_2, R_1 + C_1 = R_2 + C_2. \\ \Leftarrow \text{ queen}(R_1, C_1), \text{ queen}(R_2, C_2), R_1 \neq R_2, R_1 - C_1 = R_2 - C_2. \end{aligned}$$

The database contains facts of the form $\text{num}(1) \dots \text{num}(N)$ for the N -queens problem. The partition rule assigns to each row exactly one queen. The first constraint states that no two different queens are in the same column. The last two constraints state that no two different queens are on the same diagonal. \square

Example 10

Latin Squares. This problem consists in filling an $N \times N$ table with N different symbols in such a way that each symbol occurs exactly once in each row and exactly once in each column. Tables are partially filled. The \mathcal{NP} Datalog query $\langle \mathcal{P}_{1s}, \text{square}(R, C, V) \rangle$ expresses the problem, where \mathcal{P}_{1s} consists of the following rules:

$$\begin{aligned} \oplus_V \text{square}(R, C, V) &\leftarrow \text{num}(R), \text{num}(C), \text{num}(V) \\ &\leftarrow \text{square}(R, C_1, V), \text{square}(R, C_2, V), C_1 \neq C_2 \\ &\leftarrow \text{square}(R_1, C, V), \text{square}(R_2, C, V), R_1 \neq R_2 \\ \text{square}(R, C, V) &\leftarrow \text{preassigned}(R, C, V). \end{aligned}$$

The database contains facts of the form $\text{num}(1) \dots \text{num}(N)$ for an $N \times N$ table and facts of the form $\text{preassigned}(R, C, V)$ whose meaning is that the entry $\langle R, C \rangle$ of the table contains the symbol V (here the symbols used are the numbers from 1 to N). The partition rule assigns exactly one symbol to each entry of the table. The first (resp. second) constraint states that a symbol cannot occur more than once in the same row (resp. column). The last constraint states that preassigned symbols must be respected. \square

5 Translating \mathcal{NP} Datalog queries into OPL programs

Several languages have been designed and implemented for hard search and optimization problems. These include logic languages based on stable models (e.g., *DeRes*, *DLV*⁴, *ASSAT*⁵, *Smodels*⁶, *Cmodels*, *Clasp*⁷) (Cholewinski *et al.* 1996; Simons *et al.* 2002; Lin and Zhao 2004; Lierler 2005a, 2005b; Leone *et al.* 2006; Gebser *et al.* 2007), constraint logic programming systems (e.g., *SICStus Prolog*⁸, *ECLiPSe*, *XSB*⁹, *Mozart*) (Rao *et al.* 1997; Van Roy 1999; Wallace and Schimpf 1999) and constraint programming languages (e.g., ILOG OPL, Lingo) (Van Hentenryck 1988; Finkel *et al.* 2004). The advantage of using logic languages based on stable model semantics with respect to constraint programming is their ability to express complex \mathcal{NP} problems in a declarative way. On the other hand, constraint programming languages are very efficient in solving optimization problems.

As \mathcal{NP} Datalog is a language to express \mathcal{NP} problems, the implementation of the language can be performed by translating queries into target languages specialized in combinatorial optimization problems, such as constraint programming languages. The implementation of \mathcal{NP} Datalog is carried out by means of a

⁴ DLV Web Site. URL: <http://www.dbai.tuwien.ac.at/proj/dlv/>

⁵ ASSAT Web Site. URL: <http://assat.cs.ust.hk/>

⁶ Smodels Web Site. URL: <http://www.tcs.hut.fi/Software/smodels/>

⁷ Clasp Web Site. URL: <http://www.cs.uni-potsdam.de/clasp/>

⁸ SICStus Prolog Web Site. URL: <http://www.sics.se/is1/sicstuswww/site/index.html>

⁹ XSB Web Site. URL: <http://xsb.sourceforge.net/>

system prototype translating \mathcal{NP} Datalog queries into OPL programs. OPL is a constraint programming language well suited for solving both search and optimization problems. OPL programs are computed by means of the ILOG OPL Development Studio. This section shows how \mathcal{NP} Datalog queries are translated into OPL programs.

\mathcal{NP} Datalog programs have an associated database schema specifying the used database domains and for each base predicate the domain associated with each attribute. For instance, the database schema associated with the min coloring query of Example 6 is

```
DOMAINS :    node;color.
PREDICATES : edge(node,node).
```

Starting from the database schema, the compiler also deduces the schema of every derived predicate and introduces new domains, obtained from the database domains. For instance, for the program of Example 6 the schemas associated with the predicates `col` and `used_color` are, respectively, `col(node,color)` and `used_color(color)`. Considering the program of Example 6 and assuming to also have the following rules:

```
p(X) ← node(X).
p(X) ← color(X).
q(X) ← node(X),color(X).
```

the schemas associated with `p` and `q` are $p(D_p)$ and $q(D_q)$, where D_p is the union of the domains `node` and `color`, whereas D_q is the intersection of the domains `node` and `color`. Database domain instances are defined by means of unary ground facts. Integer domains are declared differently. For instance, the database schema associated with the *N-queens* program of Example 9 is as follows:

```
INT-DOMAINS :  num.
```

Moreover, whenever the **integer** predicate is used in a program, the range of considered integers has to be specified in the schema, as shown in the following example:

```
MinInt = 0.
MaxInt = 10.
```

A predicate p is said to be *constrained* if (i) p depends on a guess predicate, and (ii) there is a constraint or an optimized query goal containing p or containing a predicate q which depends on p . Moreover, a constrained predicate is said to be *recursion-dependent* if it is recursive or depends on a constrained recursive predicate.

Every \mathcal{NP} Datalog program \mathcal{P} consists of a set \mathcal{P}_S of standard rules, a set \mathcal{P}_G of rules defining guess predicates and a set \mathcal{P}_C of constraints. A program $\mathcal{P} = \mathcal{P}_S \cup \mathcal{P}_G \cup \mathcal{P}_C$ can be also partitioned into four sets:

- (1) $\mathcal{P}^1 = \mathcal{P}_S^1$ consisting of the set of rules defining standard predicates not depending on guess predicates;
- (2) $\mathcal{P}^2 = \mathcal{P}_G \cup \mathcal{P}_S^2 \cup \mathcal{P}_C^2$ consisting of (i) the set of rules defining guess predicates (\mathcal{P}_G), (ii) the set of standard rules defining constrained predicates which are

- not recursion-dependent (\mathcal{P}_S^2), and (iii) the set of constraints \mathcal{P}_C^2 containing only base predicates and predicates defined in $\mathcal{P}_G \cup \mathcal{P}_S^1 \cup \mathcal{P}_S^2$;
- (3) $\mathcal{P}^3 = \mathcal{P}_S^3 \cup \mathcal{P}_C^3$ consisting of the set of standard rules defining constrained, recursion-dependent predicates (\mathcal{P}_S^3) and the set of constraints \mathcal{P}_C^3 containing predicates defined in \mathcal{P}_S^3 ;
 - (4) $\mathcal{P}^4 = \mathcal{P}_S^4$ consisting of the set of rules defining standard predicates which depend on guess predicates and are not constrained.

The evaluation of an \mathcal{NP} Datalog program \mathcal{P} over a database $D\mathcal{B}$ is carried out by performing the following steps:

- (1) First, the (unique) stable model of $\langle \mathcal{P}_S^1, D\mathcal{B} \rangle$ (say it $D\mathcal{B} \cup M_1$) is computed.
- (2) Next, a stable model of $\langle \mathcal{P}_S^2 \cup \mathcal{P}_G, D\mathcal{B} \cup M_1 \rangle$ satisfying the constraints \mathcal{P}_C^2 (say it $D\mathcal{B} \cup M_1 \cup M_2$) is computed.
- (3) Afterward, if a model $D\mathcal{B} \cup M_1 \cup M_2$ exists, the (unique) stable model of $\langle \mathcal{P}_S^3, D\mathcal{B} \cup M_1 \cup M_2 \rangle$ (say it $D\mathcal{B} \cup M_1 \cup M_2 \cup M_3$) is computed. If this model satisfies the constraints \mathcal{P}_C^3 , then the next step is executed, otherwise the second step is executed again, that is, another stable model of $\langle \mathcal{P}_S^2 \cup \mathcal{P}_G, D\mathcal{B} \cup M_1 \rangle$ satisfying the constraints \mathcal{P}_C^2 is computed.
- (5) Finally, if a model $D\mathcal{B} \cup M_1 \cup M_2 \cup M_3$ satisfying the constraints \mathcal{P}_C^3 exists, the (unique) stable model of $\langle \mathcal{P}_S^4, D\mathcal{B} \cup M_1 \cup M_2 \cup M_3 \rangle$ is evaluated.

It is worth noting that, if there is no constrained recursive predicate, the component \mathcal{P}^3 is empty and then an \mathcal{NP} Datalog program can be evaluated by performing only steps 1, 2, and 4 (i.e., the iteration introduced in step 3 is not needed).

The partition of programs into four components suggests that subprograms \mathcal{P}_S^1 , \mathcal{P}_S^3 , and \mathcal{P}_S^4 can be evaluated by means of the standard fixpoint algorithm. In the following, stratified subprograms, such as \mathcal{P}_S^1 , \mathcal{P}_S^3 , and \mathcal{P}_S^4 , are called *deterministic* as they have a unique stable model, whereas subprograms which may have zero or more stable models are called *nondeterministic*. Thus, given a database $D\mathcal{B}$ and an \mathcal{NP} Datalog query $Q = \langle P, G \rangle$, we have to generate an OPL program equivalent to the application of the query Q to the database $D\mathcal{B}$.

We first show how the database is translated and next consider the translation of queries.

Database translation. An integer domain relation is translated into a set of integers, whereas a noninteger domain relation is translated into a set of strings. The translation of a base relation with arity $n > 0$ consists of two steps: (i) declaring a new tuple type with n fields (whose type is either string or integer, according to the schema), (ii) declaring a set of tuples of this type. For instance, the translation of the database containing the facts `node(a)`, `node(b)`, `node(c)`, `node(d)`, `edge(a, b)`, `edge(a, c)`, `edge(b, c)`, and `edge(c, d)`, consists of the following OPL declarations:

```
{string} node = {a, b, c, d};
tuple edge_type {string a1; string a2;};
{edge_type} edge = {⟨a, b⟩, ⟨a, c⟩, ⟨b, c⟩, ⟨c, d⟩}.
```

The database $\text{num}(1), \text{num}(2), \text{num}(3)$ for the *N-Queens* problem of Example 9 is translated as follows:

$$\{\mathbf{int}\} \text{ num} = \{1, 2, 3\}.$$

When an integer range is specified in the schema, the following set is added to the OPL database:

$$\{\mathbf{int}\} \text{ integer} = \text{asSet}(\text{MinInt} .. \text{MaxInt}),$$

where MinInt and MaxInt are the values specified in the schema.

Query translation. The translation of an *NP Datalog* query $Q = \langle \mathcal{P}, G \rangle$ is carried out by translating the deterministic subprograms into *ILOG OPL Script* programs by means of a function Fixp and the nondeterministic subprograms into OPL programs by means of a function $\mathcal{W}_{\mathcal{P}}$ or a slightly different function $\mathcal{W}_{\mathcal{Q}}$ if the predicate in the query goal is defined in \mathcal{P}^2 . More specifically, $\text{Fixp}(\mathcal{P})$ generates an OPL script program which emulates the fixpoint computation of \mathcal{P} , whereas $\mathcal{W}_{\mathcal{P}}(\mathcal{P})$ (resp. $\mathcal{W}_{\mathcal{Q}}(\mathcal{Q})$) translates the *NP Datalog* program \mathcal{P} (resp. query \mathcal{Q}) into an equivalent OPL program.

It is worth noting that

- (1) If the query goal is not defined over component \mathcal{P}^4 , this component does not need to be evaluated and, therefore, it is not translated into an OPL Script program.
- (2) If the query goal is defined in component \mathcal{P}^1 , we have to check that the components \mathcal{P}^2 and \mathcal{P}^3 admit stable models.
- (3) If the query goal G is defined in component \mathcal{P}^2 , we have to compute the query $\langle \mathcal{P}^2, G \rangle$ over the stable model (which includes the database) obtained from the computation of component \mathcal{P}^1 and check that component \mathcal{P}^3 admits stable models.
- (4) Similarly, if the query goal G is defined in component \mathcal{P}^3 , we have to compute the query $\langle \mathcal{P}^3, G \rangle$ over a stable model of $\mathcal{P}^1 \cup \mathcal{P}^2 \cup D\mathcal{B}$.
- (5) If the query goal G is defined in component \mathcal{P}^4 , first we compute a stable model M for components \mathcal{P}^1 , \mathcal{P}^2 and \mathcal{P}^3 and next compute the fixpoint of component \mathcal{P}^4 over M .

First, we informally present how a deterministic component (\mathcal{P}_S^1 , \mathcal{P}_S^3 , and \mathcal{P}_S^4 in our partition) is translated into an *ILOG OPL Script* program, and next we show how the remaining rules are translated into an OPL program.

Translation of deterministic components. The translation of a stratified program \mathcal{P}_S produces an ILOG OPL Script program which emulates the application of the naive fixpoint algorithm to the rules in \mathcal{P}_S .

The following example shows how a set of stratified rules is translated into an ILOG OPL Script program.

Example 11

Transitive closure. Consider the following \mathcal{NP} Datalog program \mathcal{P}_{tc} computing the transitive closure of a graph:

$$\begin{aligned} tc(X, Y) &\leftarrow \text{edge}(X, Y). \\ tc(X, Y) &\leftarrow \text{edge}(X, Z), tc(Z, Y). \end{aligned}$$

The corresponding OPL Script program $\text{Fixp}(\mathcal{P}_{tc})$ is as follows:

```
// tc declaration
int tc [node][node];

execute{
  // exit rule
  for (var x in edge) {
    tc [x.a1][x.a2] = 1;
  }
  // recursive rule
  var modified = true;
  while (modified) {
    modified = false;
    for (var e in edge)
      for (var y in node)
        if (tc [e.a2][y] == 1 & tc [e.a1][y] == 0) {
          tc [e.a1][y] = 1;
          modified = true;
        }
  }
}
```

□

In the program above we have three sets of statements declaring variables and computing exit and recursive rules. Specifically:

- (1) A two-dimensional integer array `tc` is declared.
- (2) The first **forall** block evaluates the exit rule defining `tc` by inserting each edge into the transitive closure.
- (3) The recursive rule is evaluated by means of the classical *naive fixpoint algorithm* (Ullman 1988). Specifically, the statements inside the **while** block insert a pair $\langle e.a1, y \rangle$ in the transitive closure, if there exist an edge $\langle e.a1, e.a2 \rangle$ and a node y such that the transitive closure contains the pair $\langle e.a2, y \rangle$. The loop ends when no more pairs of nodes can be derived.

If a program contains negated literals, it is possible to apply the stratified fixpoint algorithm, by dividing the rules into strata and computing one stratum at a time, following the order derived from the dependencies among predicate symbols.

Translation of nondeterministic components. The translation of a nondeterministic program \mathcal{P} (denoted by $\mathcal{W}_{\mathcal{P}}(\mathcal{P})$) produces an OPL program. For the sake of simplicity of presentation, it is assumed that \mathcal{P} satisfies the following conditions:

- Guess predicates are defined by either generalized partition rules or subset rules;

- Standard predicates are defined by a unique extended rule of the form:

$$A \leftarrow \text{body}_1 \vee \dots \vee \text{body}_m,$$

where body_i is a conjunction of literals;

- Constraint rules are of the form $A \Leftarrow B$, where A is a disjunction of atoms and B is a conjunction of atoms;
- Rules do not contain two (or more) occurrences of the same variable taking values from different domains;
- Constants appear only in built-in atoms of the form $x\theta y$ where θ is a comparison operator.

It should be noticed that the previous assumptions do not imply any limitation as every program can be rewritten in such a way that it satisfies them. For instance, the two rules defining the predicate f in Example 4 can be rewritten into the rule

$$f(X) \leftarrow (p(X, V), s(V)) \vee (n(X, Z), \neg s(Z)),$$

whereas the rules defining the predicate v in Example 8 can be rewritten in the form

$$v(X) \leftarrow e(X, V) \vee e(U, X).$$

Specifically, the function $\mathcal{W}_{\mathcal{P}}$ receives in input a program \mathcal{P} and gives in output an OPL program consisting of two components $\mathcal{W}_{\mathcal{P}}(\mathcal{P}) = (\mathcal{T}_{\mathcal{Q}}(\mathcal{P}), \mathcal{T}_{\mathcal{P}}(\mathcal{P}))$, where (i) $\mathcal{T}_{\mathcal{Q}}(\mathcal{P})$ consists of the definition of arrays of integers and decision variables, (ii) $\mathcal{T}_{\mathcal{P}}(\mathcal{P})$ translates the *NP Datalog* program into an OPL program. Analogously, the function $\mathcal{W}_{\mathcal{Q}}$ receives in input an *NP Datalog* query $\mathcal{Q} = \langle \mathcal{P}, G \rangle$ and gives in output an OPL program consisting of two components $\mathcal{W}_{\mathcal{Q}}(\langle \mathcal{P}, G \rangle) = (\mathcal{T}_{\mathcal{Q}}(\mathcal{P}), \mathcal{T}_{\mathcal{Q}}(\langle \mathcal{P}, G \rangle))$, where $\mathcal{T}_{\mathcal{Q}}(\langle \mathcal{P}, G \rangle)$ translates the *NP Datalog* query into an OPL program.

The function $\mathcal{T}_{\mathcal{Q}}(\mathcal{P})$ introduces some data structures for each IDB predicate defined in \mathcal{P} . Specifically, for each IDB predicate p defined in \mathcal{P}^2 with arity k , a k -dimensional array of boolean decision variables is introduced as follows:

dvar boolean $p[D_1, \dots, D_k]$;

where D_1, \dots, D_k denote the domains on which the predicate p is defined. For instance, for the binary predicate `col` of Example 6 the declaration

dvar boolean `col[node, color]`;

is introduced. For any other IDB predicate q defined in \mathcal{P} with arity m , a m -dimensional array of integers is introduced as follows:

int $q[D_1, \dots, D_m]$;

where D_1, \dots, D_m denote the domains on which the predicate p is defined.

The function $\mathcal{T}_{\mathcal{Q}}$ and $\mathcal{T}_{\mathcal{P}}$ are defined as follows:

- (1) **Query:** $\mathcal{T}_{\mathcal{Q}}(\langle \mathcal{P}, G \rangle) = \mathcal{T}_{\mathcal{Q}}(G) \mathcal{T}_{\mathcal{P}}(\mathcal{P})$
- (2) **Goal:**
 - (a) $\mathcal{T}_{\mathcal{Q}}(v(X_1, \dots, X_k)) = \emptyset$

- (b) $\mathcal{T}_{\mathcal{P}}(\min|v(X_1, \dots, X_k)|)$
 $= \text{minimize sum}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k)) v[X_1, \dots, X_k];$
- (c) $\mathcal{T}_{\mathcal{P}}(\max|v(X_1, \dots, X_k)|)$
 $= \text{maximize sum}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k)) v[X_1, \dots, X_k];$

(3) **Sequence of rules:** $\mathcal{T}_{\mathcal{P}}(S_1 \dots S_n) = \text{subject to}\{\mathcal{T}_{\mathcal{P}}(S_1) \dots \mathcal{T}_{\mathcal{P}}(S_n)\};$

(4) **Partition rules** of the form

$$\oplus_{\mathbf{1}} s(X_1, \dots, X_k, L) \leftarrow \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n), d(L)$$

are translated into the following OPL statement:

$$\begin{aligned} &\text{forall}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k)) \\ &\quad \mathcal{T}_{\mathcal{P}}(\exists(Y_1, \dots, Y_n) \text{ body}(X_1, \dots, X_k, Y_1, \dots, Y_n)) > 0 \\ &\quad \Rightarrow \text{sum}(L \text{ in } d) s[X_1, \dots, X_k, L] == 1; \\ &\text{forall}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k), L \text{ in } d) \\ &\quad s[X_1, \dots, X_k, L] > 0 \Rightarrow \mathcal{T}_{\mathcal{P}}(\exists(Y_1, \dots, Y_n) \text{ body}(X_1, \dots, X_k, Y_1, \dots, Y_n)) > 0; \end{aligned}$$

(5) **Subset rules** of the form

$s(X_1, \dots, X_k) \subseteq \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n)$ are translated as follows:

$$\begin{aligned} &\text{forall}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k)) \\ &\quad s[X_1, \dots, X_k] > 0 \Rightarrow \mathcal{T}_{\mathcal{P}}(\exists(Y_1, \dots, Y_n) \text{ body}(X_1, \dots, X_k, Y_1, \dots, Y_n)) > 0; \end{aligned}$$

(6) **Standard rules** of the form

$p(X_1, \dots, X_k) \leftarrow \text{Body}_1(X_1, \dots, X_k, Y_1^1, \dots, Y_{n_1}^1) \vee \dots \vee \text{Body}_m(X_1, \dots, X_k, Y_1^m, \dots, Y_{n_m}^m)$
 are translated as follows:

$$\begin{aligned} &\text{forall}(X_1 \text{ in dom}(X_1), \dots, X_k \text{ in dom}(X_k)) \\ &\quad p[X_1, \dots, X_k] > 0 \Leftrightarrow \mathcal{T}_{\mathcal{P}}(\exists(Y_1^1, \dots, Y_{n_1}^1) \text{Body}_1) + \dots + \mathcal{T}_{\mathcal{P}}(\exists(Y_1^m, \dots, Y_{n_m}^m) \\ &\quad \text{Body}_m) > 0; \end{aligned}$$

where $Y_1^1, \dots, Y_{n_1}^1$ is the list of existentially quantified variables in Body_1 .

(7) **Conjunction of literals with existentially quantified variables:** A conjunction of literals with $n > 0$ existentially quantified variables is translated as follows:

$$\mathcal{T}_{\mathcal{P}}(\exists(Y_1, \dots, Y_n) \text{Body}) = (\text{sum}(Y_1 \text{ in } D_1, \dots, Y_n \text{ in } D_n) (\mathcal{T}_{\mathcal{P}}(\text{Body})))$$

where D_j is the domain associated with the variable Y_j .

(8) **Conjunction of literals without existentially quantified variables:**

$$\mathcal{T}_{\mathcal{P}}(A_1, \dots, A_k) = \begin{cases} (\mathcal{T}_{\mathcal{P}}(A_1) * \dots * \mathcal{T}_{\mathcal{P}}(A_k)) & \text{if } k > 0 \\ 1 & \text{if } k = 0 \end{cases}$$

(9) **Literal:**

$$\mathcal{T}_{\mathcal{P}}(q(X_1, \dots, X_k)) = \begin{cases} q[X_1, \dots, X_k], & \text{if } q \text{ is a derived pred.} \\ (\text{sum}(\langle X_1, \dots, X_k \rangle \text{ in } q) 1 > 0), & \text{if } q \text{ is a base pred.} \end{cases}$$

$\mathcal{T}_{\mathcal{P}}(q(X)) = (\text{sum}(X \text{ in } q) 1 > 0)$ if q is a domain predicate,

$\mathcal{T}_{\mathcal{P}}(E_1 \theta E_2) = (E_1 \theta E_2)$, where θ is a comparison operator and E_1, E_2 are either variables or constants or arithmetic expressions,

$$\mathcal{T}_{\mathcal{P}}(\neg A) = (1 - \mathcal{T}_{\mathcal{P}}(A));$$

(10) **Constraints** of the form $A_1 \vee \dots \vee A_m \Leftarrow \text{body}(X_1, \dots, X_k)$ where $\text{body}(X_1, \dots, X_k)$ is a conjunction of atoms are translated as follows:

$$\begin{aligned} &\mathcal{T}_{\mathcal{P}}(A_1 \vee \dots \vee A_m \Leftarrow \text{body}(X_1, \dots, X_k)) \\ &= \text{forall}(X_1 \text{ in } \text{dom}(X_1), \dots, X_k \text{ in } \text{dom}(X_k)) \\ &\quad \mathcal{T}_{\mathcal{P}}(\text{body}(X_1, \dots, X_k)) > 0 \Rightarrow (\mathcal{T}_{\mathcal{P}}(A_1) + \dots + \mathcal{T}_{\mathcal{P}}(A_m)) > 0; \end{aligned}$$

For $m = 0$ the above constraint becomes $\mathcal{T}_{\mathcal{P}}(\text{body}(X_1, \dots, X_k)) > 0 \Rightarrow \text{false}$.

Observe that the OPL code associated with the translation of a rule can be simplified by means of trivial reductions. As an example, an expression of the form: $((c > 0) > 0)$ can be simply replaced by $(c > 0)$, whereas expressions of the form $1 * 1$ are replaced by 1.

The following theorem shows the correctness of our translation. As we partition a program \mathcal{P} into four distinct components $\mathcal{P}^1, \mathcal{P}^2, \mathcal{P}^3$, and \mathcal{P}^4 , where the components $\mathcal{P}^1, \mathcal{P}_S^3$, and \mathcal{P}^4 are computed by means of a fixpoint algorithm, whereas the components \mathcal{P}^2 and P_C^3 are translated into OPL programs, we next show the correctness of the translation of queries $Q = \langle \mathcal{P}, G \rangle$, where $\mathcal{P} = \mathcal{P}^2$, i.e., we assume that components $\mathcal{P}^1, \mathcal{P}^3$, and \mathcal{P}^4 are empty. Thus, such queries consist only of rules defining guess predicates, constraints and standard rules defining constrained predicates which are not recursion-dependent. Programs and queries of this form will be called *R-NP Datalog* (restricted *NP Datalog*).

Theorem 2

For every *R-NP Datalog* query $Q, W_{\mathcal{D}}(Q) \equiv Q$.

Proof. For each *R-NP Datalog* query $Q = \langle \mathcal{P}, G \rangle$, where each standard predicate is defined by a unique extended rule, the query $Q^r = \langle \mathcal{P}^r, G^r \rangle$ is derived as follows:

(1) Every generalized partition rule of the form:

$$\oplus_L s(X_1, \dots, X_k, L) \Leftarrow \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n), d(L)$$

is substituted by the constraints:

$$\begin{aligned} &\text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n) \Rightarrow s(X_1, \dots, X_k, L), \\ &s(X_1, \dots, X_k, L_1), s(X_1, \dots, X_k, L_2) \Rightarrow L_1 = L_2, \\ &s(X_1, \dots, X_k, L) \Rightarrow \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n). \end{aligned}$$

(2) Each subset rule of the form:

$$s(X_1, \dots, X_k) \subseteq \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n)$$

is replaced by the constraint:

$$s(X_1, \dots, X_k) \Rightarrow \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n).$$

(3) Every standard rule of the form:

$$A \Leftarrow \text{Body}_1 \vee \dots \vee \text{Body}_m$$

is replaced by the constraint¹⁰:

$$A \Leftrightarrow \text{Body}_1 \vee \cdots \vee \text{Body}_m.$$

- (4) For each derived predicate p with schema $p(\text{dom}_1, \dots, \text{dom}_k)$ we introduce (i) a new predicate symbol p' with schema $p'(\text{dom}_1, \dots, \text{dom}_k)$, and (ii) a rule of the following form:

$$p(X_1, \dots, X_k) \oplus p'(X_1, \dots, X_k) \leftarrow \text{dom}_1(X_1), \dots, \text{dom}_k(X_k). \quad (8)$$

These rules are introduced to assign, nondeterministically, a truth value to derived atoms.

Clearly, the queries Q and Q' are equivalent as the correct truth value of derived atoms is determined by the constraints. It is worth pointing out that for each (partition, subset, and standard) rule r a constraint of the form $\text{Head}(r) \Rightarrow \text{Body}(r)$ was introduced to guarantee that models contain only “supported atoms”, i.e., atoms derivable from r .

The program $\mathcal{W}_{\mathcal{Q}}(Q)$ is just a translation of Q' into OPL statements where:

- Rules of form (8) do not need to be translated into correspondent OPL statements as each derived predicate p , defined by such a rule, is translated into a boolean k -dimensional array.
- The first two constraints, derived from the rewriting of partition rules, for ensuring (i) the assignment of each element in the body to some class L and (ii) the uniqueness of this assignment, are rewritten into a unique OPL constraint. □

Example 12

Min Coloring. The OPL program corresponding to the (simplified) translation of the min coloring query of Example 6 is as follows:

```

dvar boolean col[node,color];
dvar boolean used_color[color];

minimize
    sum(c in color) used_color[c];
subject to {
    forall (x in node)
        (sum (x in node) 1 > 0) > 0  $\Rightarrow$  sum(c in color) col[x,c] == 1;

    forall (x in node, c in color)
        col[x,c] > 0  $\Rightarrow$  (sum(x in node) 1 > 0) > 0;

    forall (c in color)

```

¹⁰ A shorthand for the two constraints:

$$A \Rightarrow \text{Body}_1 \vee \cdots \vee \text{Body}_m$$

$$A \Leftarrow \text{Body}_1 \vee \cdots \vee \text{Body}_m$$

```

used_color[c] > 0 ⇔ sum(x in node) col[x, c] > 0;
forall (x in node, y in node, c in color)
  (sum(<x, y> in edge) 1 > 0) * col[x, c] * col[y, c] > 0 ⇒ false;
};

```

Code optimization. The number of (ground) constraints can be strongly reduced by applying simple optimizations to the OPL code.

- *Range restriction.* If the OPL code contains constructs of the form:

```

forall(X1 in D1, ..., Xn in Dn)
  (sum(<X1, ..., Xk> in T) 1 > 0) <Statement1> ⇒ <Statement2>

```

the **sum** construct can be deleted so that the constraint can be rewritten as follows:

```

forall(<X1, ..., Xk> in T, Xk+1 in Dk+1, ..., Xn in Dn)
  1 <Statement1> ⇒ <Statement2>

```

If the OPL code contains constructs of the form

```

forall(X1 in D1, ..., Xn in Dn)
  (sum(X1 in D1) 1 > 0) <Statement>

```

the **sum** construct can be deleted so that the constraint can be rewritten as

```

forall(X1 in D1, ..., Xn in Dn)
  1 <Statement>

```

Example 13

By applying the optimizations above, the min coloring problem can be rewritten as follows:

```

dvar boolean col[node, color];
dvar boolean used_color[color];

minimize
  sum(c in color) used_color[c];

subject to {
  forall (x in node)
    1 > 0 ⇒ sum(c in color) col[x, c] == 1;

  forall (x in node, c in color)
    col[x, c] > 0 ⇒ 1 > 0;

  forall (c in color)
    used_color[c] > 0 ⇔ sum(x in node) col[x, c] > 0;

  forall (<x, y> in edge, c in color)
    col[x, c] * col[y, c] > 0 ⇒ false;
};

```

□

- *Constraint optimization.* A very simple optimization consists in deleting the OPL constraints whose head is always true (e.g., the head consists of the constant 1) as they are always satisfied. For instance, in the above example the second OPL constraint can be deleted as its head consists of the constant 1. An additional simple optimization can be performed by “pushing down” conditions defined inside the OPL constraints. For instance, the following code:

$$Q (X_1 \text{ in } D_1, \dots, X_k \text{ in } D_k) \langle \text{statement}_1 \rangle X_i \theta X_j \Rightarrow \langle \text{statement}_2 \rangle,$$

where Q is either **forall** or **sum**, X_i and X_j are either variables or constants or arithmetic expressions, θ is a comparison operator, can be rewritten as

$$Q (X_1 \text{ in } D_1, \dots, X_k \text{ in } D_k : X_i \theta X_j) \langle \text{statement}_1 \rangle \Rightarrow \langle \text{statement}_2 \rangle$$

- *Arrays reduction.* A further optimization can be performed by reducing the dimension of the arrays (of decision variables) corresponding to some guess predicates. Specifically, given a guess predicate s defined by generalized partition rules of the form:

$$\oplus_L s(X_1, \dots, X_k, L) \leftarrow \text{body}(X_1, \dots, X_k, Y_1, \dots, Y_n), \text{dom}(L)$$

instead of declaring a $(k+1)$ -dimensional array of boolean decision variables, it is possible to introduce a k -dimensional array s of integer decision variables ranging in $\{0, \dots, |\text{dom}|\}$ and map each value in dom to $\{1, \dots, |\text{dom}|\}$ by means of a one-to-one function. The meaning of $s[X_1, \dots, X_k] = c$ is that if $c \neq 0$ then the atom $s(X_1, \dots, X_k, c')$ is *true*, where c' is the value in dom corresponding to the integer c ; if $c = 0$ then the atom $s(X_1, \dots, X_k, c')$ is *false* for any value c' in dom . Clearly, to make consistent the OPL program, every instance of $s[X_1, \dots, X_k, C]$ must be substituted with $(s[X_1, \dots, X_k] == C)$ and in each **forall** or **sum** statement containing variables ranging in dom the condition $C \neq 0$ must be verified.

Example 14

The application of the previous optimizations to the program of Example 13 gives the following OPL program:

```

int cardcolor = card(color);
range intcolor = 0..cardcolor;
dvar int col[node] in intcolor;
dvar boolean used_color[intcolor];

minimize
    sum(c in intcolor : c ≠ 0) used_color[c];
subject to {
    forall (x in node)
        1 > 0 ⇒ sum(c in intcolor : c ≠ 0) (col[x]== c) > 0;

    forall (c in intcolor : c ≠ 0)
        used_color[c] > 0 ⇔ sum(x in node) (col[x]== c) > 0;

    forall ((x,y) in edge, c in intcolor : c ≠ 0)
        (col[x]==c) * (col[y]==c) > 0 ⇒ false;
};
    
```

□

Variable deletion. A further optimization regards the deletion of unnecessary variables and the reduction of domains. For instance, in the last constraint in the OPL program of the previous example, the variable c can be deleted as it is just used to define the matching between $\text{col}[X]$ and $\text{col}[Y]$. Thus, this constraint can be rewritten as:

```
forall (( $x, y$ ) in edge)
  ( $\text{col}[X] == \text{col}[Y]$ ) > 0  $\Rightarrow$  false;
```

Observe that, if the body of the partition rule only contains database domains, the integer decision variables of the guess predicate can range in the set of integers $\{1, \dots, |\text{dom}|\}$ as the head atom is true for all possible values of its variables X_1, \dots, X_k . This means that under such circumstances, it is not necessary to introduce the additional condition stating that the value of the variable cannot be 0. Under this rewriting, the first constraint can be deleted as its head is always satisfied.

The following example shows the final version of the min coloring program, obtained by applying the optimizations above.

Example 15

Min Coloring (*optimized version*).

```
int cardcolor = card(color);
range intcolor = 1..cardcolor;
dvar int col[node] in intcolor;
dvar boolean used_color[intcolor];

minimize
  sum( $c$  in intcolor) used_color[ $c$ ];
subject to {
  forall ( $c$  in intcolor)
    used_color[ $c$ ] > 0  $\Leftrightarrow$  sum( $x$  in node) ( $\text{col}[x] == c$ ) > 0;
  forall (( $x, y$ ) in edge)
    ( $\text{col}[x] == \text{col}[y]$ ) > 0  $\Rightarrow$  false;
};
```

The OPL program corresponding to the k -coloring problem consists of only one constraint, namely the second OPL constraint in the previous example.

Aggregates. The current version of the paper does not include aggregates, although the language could be easily extended with stratified aggregates which can be effortlessly translated into OPL programs.

Consider, for instance, a digraph stored by means of the two relations *node* and *edge* and the following logic rule with aggregates¹¹:

$$\text{out}(X, C) \leftarrow \text{edge}(X, Y), \text{count}((X), C)$$

¹¹ The syntax used refers to the proposal presented in (Greco 1999).

computing for each node X the number C of outgoing arcs. Such a rule could be easily translated into the following OPL script code:

```

{string} node = ...;
tuple edge{string a; string b;};
{edge} edges = ...;
int out[node];
execute{
    for (var e in edges)
        out[e.a] = out[e.a] + 1;
};

```

In this paper, we have not considered aggregates since we would like to define more efficient translations which allow us to express and efficiently compute greedy and dynamic programming algorithms. In the literature, there have been several proposals to extend Datalog with aggregates. For instance, the proposal of (Greco 1999) allows us to write rules with stratified aggregates and evaluate programs so that the behavior of dynamic programming is captured (see also Greco and Zaniolo 2001 for greedy algorithms).

Consider the query $\langle SP, stc(X, Y, C) \rangle$ computing the shortest paths of a weighted digraph, where SP consists of the following rules:

$$stc(X, Y, C) \leftarrow tc(X, Y, C), \min((X, Y), C).$$

$$tc(X, Y, C) \leftarrow edge(X, Y, C).$$

$$tc(X, Y, C) \leftarrow edge(X, Z, C1), tc(Z, Y, C2), C = C1 + C2.$$

and weights associated with arcs are positive integers. A standard translation and execution has two main problems: (i) the computation is not efficient since for each pair of nodes all paths with different weights are considered, and (ii) if the graph is cyclic the computation never terminates (or terminates with an error). Since shortest paths can be obtained by considering other shortest paths, an OPL Script computing them could be as follows:

```

// declarations
tuple edge{string a; string b; int c;};
{edge} edges = ...;
{string} node = ...;
int tc[x in node][y in node] = maxint;
int stc[x in node][y in node] = maxint;

execute {
    // exit rule
    for (var e1 in edge) {
        tc[e1.a][e1.b] = e1.c;
        stc[e1.a][e1.b] = e1.c;
    }
}

```

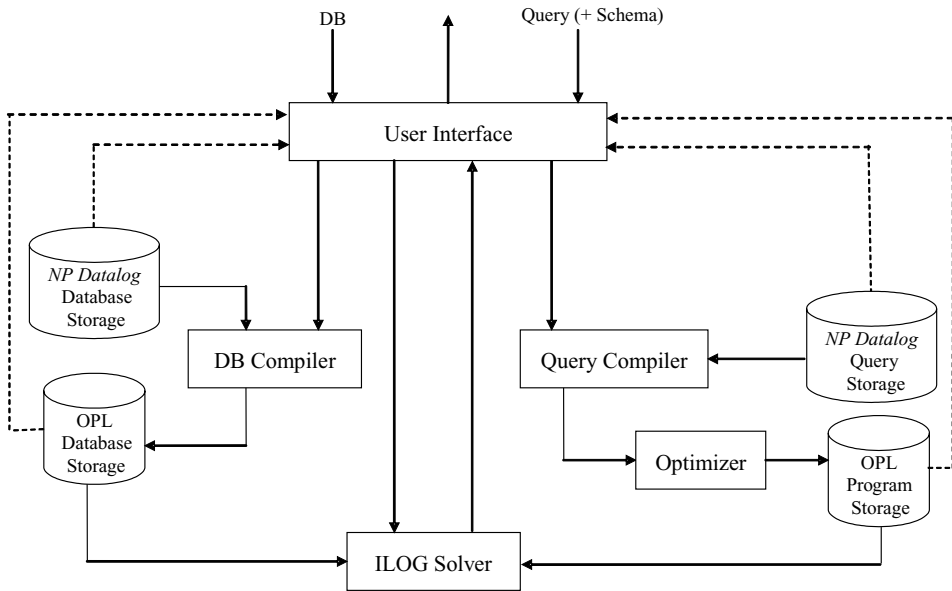


Fig. 1. System architecture.

```

// recursive rule
var modified = true;
while (modified) {
  modified = false;
  for (var e in edges)
    for (var y in node) {
      tc[e.a][y] = e.c + stc[e.b][y];
      if(tc[e.a][y] < stc[e.a][y]) {
        modified = true;
        stc[e.a][y] = tc[e.a][y];
      }
    }
}
}

```

5.1 Implementation and experiments

A system prototype translating *NP Datalog* queries into OPL programs and executing the target code using the ILOG OPL Development Studio has been implemented. The system architecture, depicted in Figure 1, consists of five main modules whose functionalities are next briefly discussed.

- *User interface* – This module receives in input a pair of strings identifying the file containing the source database and the file containing the query. If both the database and the query have already been translated, then the User interface (UI) asks the module *ILOG Solver* to execute the query. If the database (resp. query) has not been translated, then the UI sends the name of the file containing the source database (resp. query) to the module *Database*

Compiler (resp. *Query Compiler*) to be translated. Moreover, this module is in charge of visualizing the answer to the input query.

- *Database compiler* – This module translates the source database into an OPL database.
- *Query compiler* – This module receives in input an *NP Datalog* query and gives in output the corresponding OPL code. In order to check the correctness of the query and generate the target code, the module uses information on the schema of predicates.
- *Optimizer* – This module rewrites the OPL code received from the module *Query Compiler* and gives in output the target (optimized) OPL code.
- *Query executor* – This module consists of the ILOG OPL Development Studio which executes the query stored by the module *Optimizer* into the *OPL program storage*, over a database stored into the *OPL database storage*. The module *Query executor* interacts with the module *User interface* by providing it the obtained result.

Therefore, *NP Datalog* can be also used to define a logic interface for constraint programming solvers such as ILOG. The experiments presented in this subsection show that the combination of the two components is effective so that constraint solvers (as well as SAT solvers) can be used as an efficient tool for computing logic queries whose semantics is based on stable models.

In order to assess the efficiency of our approach, we have performed several experiments comparing the performance obtained by implementing *NP Datalog* over the ILOG OPL Development Studio against Answer Set Programming systems. Specifically, *NP Datalog/OPL* has been compared with *DLV*, *Smodels*, *ASSAT*, *Clasp*, and *XSB*. The following version of the aforementioned systems have been used:

- ILOG OPL Development Studio 6.1,
- DLV release 2007-10-11,
- Smodels 2.33 (and *lparse* 1.1.1),
- ASSAT 2.02 (*lparse* 1.1.1 and *zChaff* 2007.3.12¹²),
- Clasp 1.2.1 (and *lparse* 1.1.1),
- XSB version 3.2 March 15, 2009.

The performances of the systems have been evaluated by measuring the time necessary to find one solution of the following problems: 3-Coloring, Hamiltonian Cycle, Transitive Closure, Min Coloring, N-Queens, and Latin Squares.

For each system, we have used efficient encodings of the problems which exploit efficient built-in constructs provided by the systems. Every encoding and database used in the experiments can be downloaded from the *NP Datalog* web site¹³.

All the experiments were carried out on a PC with a processor Intel Core Duo 1.66 GHz and 1 GB of RAM under the Linux operating system. In the sequel of this section the experimental results are presented.

¹² zChaff. URL: <http://www.princeton.edu/~chaff/zchaff.html>

¹³ NP datalog web site. URL: <http://www.info.deis.unical.it/npdatalog/>




Fig. 2. Structured graphs.

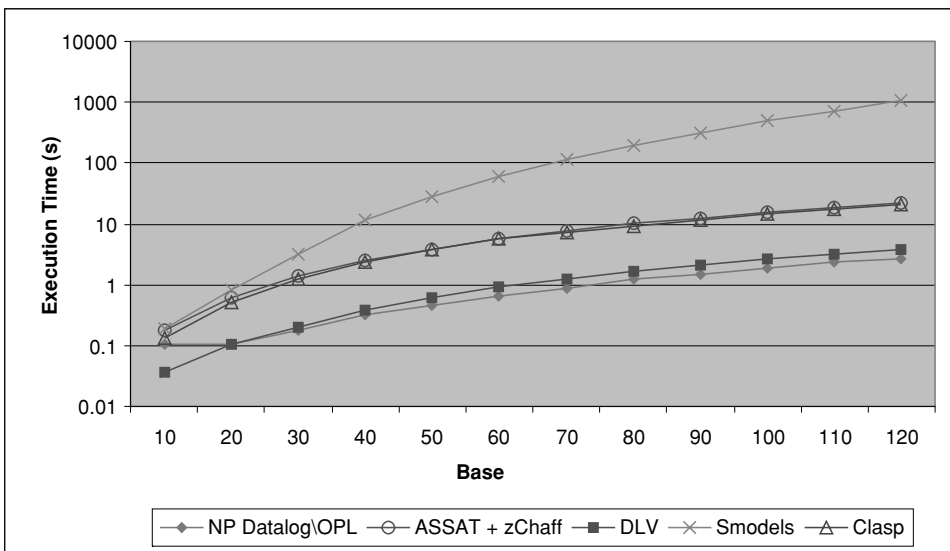


Fig. 3. Execution time for the 3-coloring problem on structured graphs.

3-Coloring. The 3-Coloring query has been evaluated on structured graphs of the form reported in Figure 2(i) and random graphs. Specifically, structured graphs with $base = height$ have been used (here $base$ denotes the number of nodes in the same row, $height$ the number of nodes in the same column; the total number of nodes in the graph is $base * height$). The random graphs have been generated by means of Culberson’s graph generator¹⁴. Specifically, the following parameters have been used: K -coloring scheme equal to *Equi-partitioned*, *Partition number* equal to 3, *Graph type* is *IID* (*independent random edge assignment*). Both structured and random graphs are all 3-colorable; the results, showing the execution times (in seconds) as the size of the graph increases, are reported in Figures 3 and 4, respectively.

¹⁴ K-Colorable graph generator.
 URL: <http://web.cs.ualberta.ca/~joe/Coloring/Generators/generate.html>

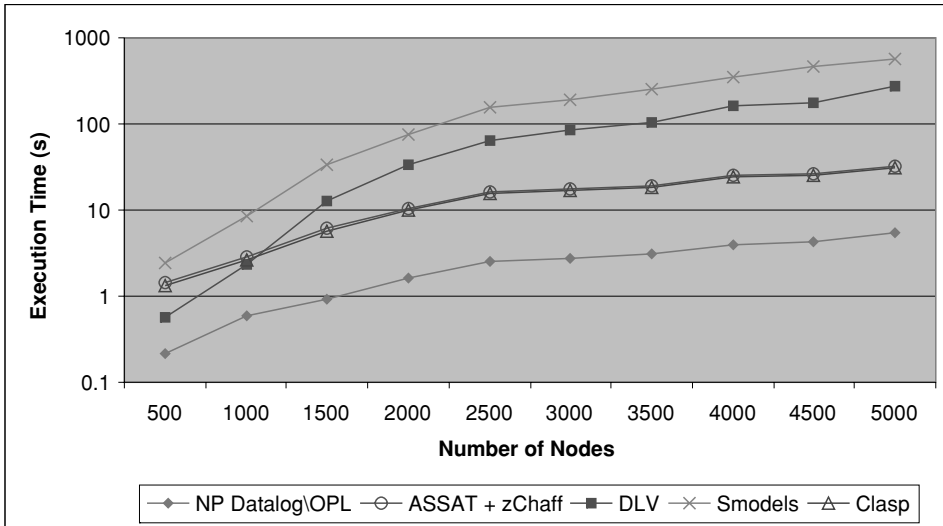


Fig. 4. Execution time for the 3-coloring problem on random graphs.

As for structured graphs, the x -axis reports the number of nodes in the same layer (i.e., the value of *base*). *NP Datalog* and DLV are faster than the other systems; ASSAT and Clasp have almost the same execution times (observe that the scale of the y -axis is logarithmic).

Regarding random graphs, it is worth noting that we have considered, for each number of nodes, five different graphs. Thus, the execution times reported in Figure 4 have been obtained by evaluating the query five times (over different graphs with the same number of nodes) and computing the mean value. *NP Datalog/OPL* is faster than the other systems; again, ASSAT and Clasp have almost the same execution times.

Hamiltonian Cycle. The Hamiltonian Cycle problem has been evaluated over benchmark graphs used to test other systems¹⁵ and random graphs generated by means of Culberson's graph generator¹⁶. All the graphs have a Hamiltonian cycle. The *NP Datalog* encoding (as well as the encodings for the other systems) can be found on *NP Datalog* web site. The results are reported in Figures 5 and 6. The x -axis reports the used graphs: a label $nwXaY$ refers to a graph with X nodes and Y arcs. Observe that, in Figure 5, a missing value means that the system has not answered in 30 min. Clasp is the fastest system for both types of graphs. DLV and Smodels are on average faster than the remaining systems. For large "dense" graphs Smodels outperforms DLV, but on some benchmark instances it runs out of time.

Transitive Closure. The Transitive Closure problem has been evaluated over directed structured graphs such as those reported in Figure 7. Specifically, instances with

¹⁵ Hamiltonian Cycle Instances. URL: <http://assat.cs.ust.hk/Assat-2.0/hc-2.0.html>

¹⁶ Hamiltonian Cycle Program Archive.
URL: <http://web.cs.ualberta.ca/~joe/Theses/Hcarchive/main.html>

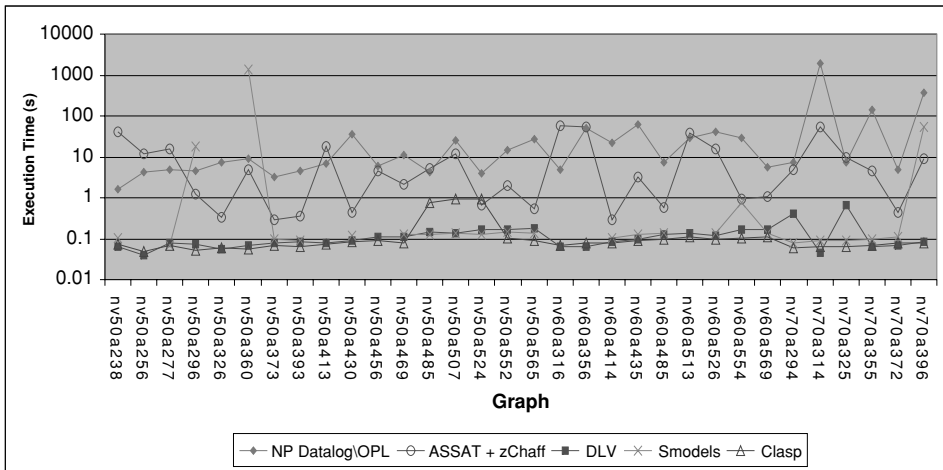


Fig. 5. Execution time for the Hamiltonian Cycle problem on benchmark graphs.

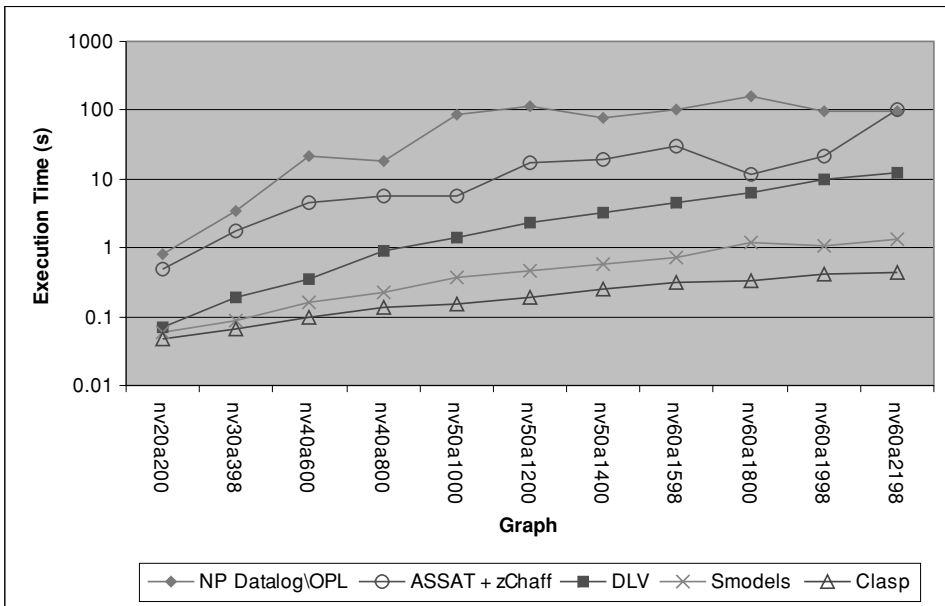


Fig. 6. Execution time for the Hamiltonian Cycle problem on random graphs.

base = *height* have been used (*base* denotes the number of nodes in the same row, *height* the number of nodes in the same column).

The results, which are reported in Figure 8, show that DLV and XSB are faster than the other systems; ASSAT, Clasp, and Smodels almost have the same execution times.

Min Coloring. As for the Min Coloring optimization problem, we have used structured graphs such as those of Figure 2. Instances having the structure reported in Figure 2(i) need at least three colors to be colored, whereas instances having the structure reported in Figure 2(ii) need at least four colors to be colored. The number




Fig. 7. Directed structured graphs.




Fig. 8. Execution time for the Transitive Closure problem on directed structured graphs.

of colors available in the database has been fixed for the two structures, respectively, to four and five (one more than the number of colors necessary to color the graph). The results are reported in Figure 9 and show that *NP Datalog* outperforms DLV. A missing time means that DLV runs out of time (also in this case we had a 30 min time limit).

N-Queens. We have considered empty chessboards to be filled with *N* Queens for increasing values of *N*. The results are reported in Figure 10. Clasp is faster than *NP Datalog* which is in turn faster than ASSAT; for a high number of queens, DLV and Smodels become slower than the other systems.

Latin Squares. We have considered partially filled tables which have been generated randomly. In every table, 60% of the squares are empty. We have considered, for each table size, five different instances. Thus, the execution times reported in Figure 11 have been obtained by evaluating the query five times (over different tables of the

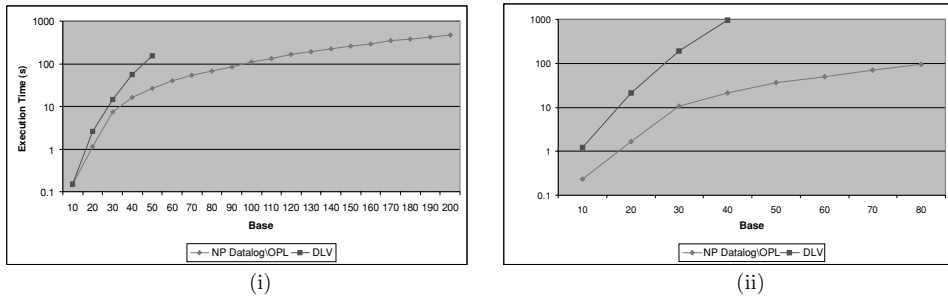


Fig. 9. Execution time for the Min Coloring problem on structured graphs.

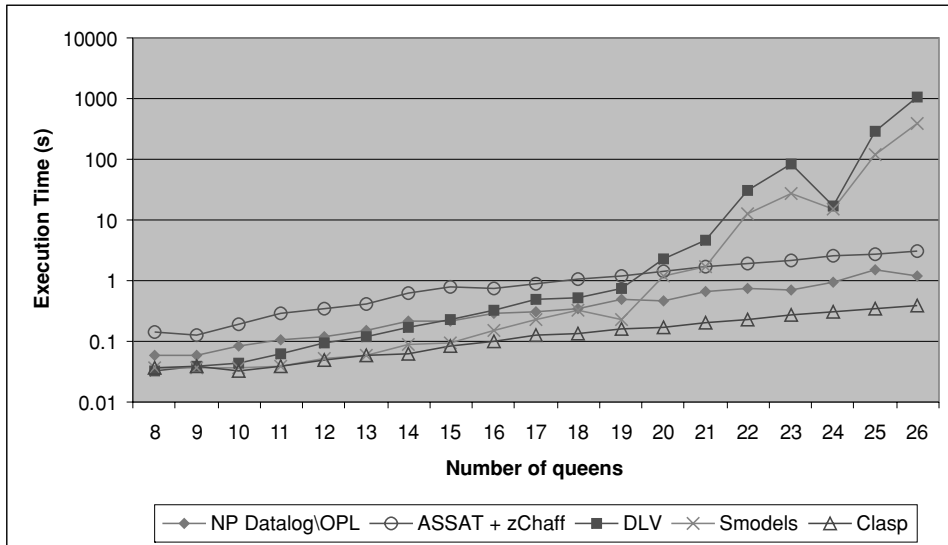


Fig. 10. Execution time for the N-Queens problem.

same size) and computing the mean value. The results show that *NP Datalog* and *Clasp* are faster than the other systems.

The experimental results reported above show that our system only seems to suffer with programs where the evaluation of the deterministic components is predominant. The reason is that deterministic components (often consisting of recursive rules) are translated into OPL scripts, which correspond to the evaluation of such components by means of the naive fixpoint algorithm, whereas problems which can be expressed without recursion (or in which the nondeterministic components are predominant) are executed efficiently. The implementation of our system prototype could be enhanced by making more efficient the translation of stratified (sub)programs or by using a different evaluator for these components. For instance, they could be evaluated by means of ASP systems, thus combining their efficiency in the computation of deterministic components with the efficiency of OPL in the computation of nondeterministic components.




Fig. 11. Execution time for the Latin Squares problem on random squares.

6 Related languages and systems

Several languages have been proposed for solving \mathcal{NP} problems. Here we have analyzed three different classes of languages: specification languages, constraint and logic programming languages, and answer set logic languages.

6.1 Specification languages

Specification languages are highly declarative and allow the user to specify problems in terms of guess and check techniques.

NP-SPEC (Cadoli *et al.* 2000; Cadoli and Schaerf 2005) is a logic-based specification language allowing the built-in second-order predicates *Subset*, *Partition*, *Permutation*, and *IntFunc*. The semantics of an NP-SPEC program is based on the notion of model minimality and the language upon which this semantics relies on is $DATALOG^{CIRC}$, i.e., an extension of DATALOG in which only some predicates are minimized and the interpretation of the other is left open. An NP-SPEC program consists of two sections: the DATABASE section, specifying the instance and the SPECIFICATION section specifying the question. To make NP-SPEC executable, specifications are translated into SAT instances and then executed using a SAT solver.

KIDS (Kestrel Interactive Development System) (Smith 1990) is a semi-automatic program development system that, starting from an initial specification of the problem, produces an executable code through a set of consistency-preserving transformations. The problem is written in a logic based language augmented with set-theoretic data types and functional constraints on the input/output behavior. To make the language executable, specifications are firstly translated into *CommonLisp* and then into machine code. Before the compilation task,

the user may select an optimization technique, such as simplification or partial evaluation, to obtain a more efficient target code.

SPILL-2 (SPecifications In a Logic Language) is the second version of an executable typed logic language that is an extension of the Prolog-like language Goedel (Kluzniak and Milkowska 1997). A specification in *SPILL-2* consists of a set of type declarations, a set of function declarations, a set of predicate declarations and a number of logical expressions (queries) that are used to test the specification. A specification in *SPILL* is required to be “executable” in the sense that it is possible to “test” whether a provided solution is feasible w.r.t. a given specification. The execution of a program consists in evaluating each query in the context of the specification and reporting the result (*true* if the query succeeds and *false* otherwise).

6.2 Constraint and logic programming languages

The basic idea of constraint programming (CP) is to model and solve a problem by exploring the set of constraints that fully characterize the problem. Almost all computationally hard problems, such as planning, scheduling and graph theoretic problems, fall into this category. A large number of systems (more than 40) for solving CP problems have been developed in computer science and artificial intelligence:

- *Constraint Logic Programming*, an extension of logic programming able to manage constraints, started about 20 years ago by Jaffar *et al.* (Jaffar *et al.* 1992). Several constraint logic languages allowing the formulation of constraints over different domains exist. Basically, all these languages embed efficient constraint solvers in logic-based programming languages, such as Prolog. Here we cite, among the others, CLP (Marriott and Stuckey 1998), SICStus Prolog, BProlog (Zhou 2002), ECLiPSe (Wallace and Schimpf 1999), and Mozart (Van Roy 1999).
- *ILOG OPL Development Studio*, an integrated development environment for mathematical programming and combinatorial optimization applications. The syntax of OPL is well suited to express optimization problems defined in the mathematical programming style (Van Hentenryck 1988; Van Hentenryck *et al.* 1999).
- *Constraint LINGO* (Finkel *et al.* 2004), a high-level logic-programming language for expressing tabular constraint-satisfaction problems such as those found in logic puzzles and combinatorial problems such as graph coloring.

Several languages extending Prolog have been proposed as well. Most of these languages have been designed to provide powerful capabilities to represent and solve general problems and not to solve \mathcal{NP} problems. Here we mention:

*BinProlog*¹⁷, a fast and compact Prolog compiler, based on the transformation of Prolog to binary clauses. BinProlog is based on the BinWAM abstract machine, a specialization of the WAM for the efficient execution of binary logic programs.

¹⁷ BINPROLOG. URL: <http://www.binnetcorp.com/BinProlog/>

XSB (Rao *et al.* 1997), an extension of Prolog supporting the well-founded semantics (Van Gelder *et al.* 1991) and including implementations of OLDT (tabling) and HiLog terms. OLDT resolution is extremely useful for recursive query computation, allowing programs to terminate correctly in many cases where Prolog does not. HiLog supports a type of higher order programming in which predicate symbols can be variable or structured.

An extension of classical first-order logic, called *ID-Logic*, has been proposed in (Denecker 2000). Basically, in an ID-Logic theory, we can distinguish four different components describing (i) data, (ii) open predicates, (iii) definitions, and (iv) assertions (or constraints). The relationships between ID-Logic and ASP has been studied in (Marien *et al.* 2004), where it has been also presented how ID-Logic theories can be translated into DATALOG⁻ programs under ASP semantics.

6.3 Answer set programming languages and systems

Several deductive systems based on stable model semantics have been developed too. Here we discuss some of the more interesting answer-set-based systems and languages:

DLV (Vienna University of Technology and University of Calabria) (Leone *et al.* 2006; Eiter *et al.* 1997) is a deductive database system, based on disjunctive logic programming. DLV extends Datalog with general negation, inclusive head disjunction and two different forms of constraints: strong constraints, which must be satisfied, and weak constraints, which are satisfied if possible (preferred models are those which minimize the number of ground weak constraints which are not satisfied). For instance, the program of Example 1 is a DLV program, whereas by replacing exclusive disjunction with inclusive disjunction in the program of Example 2 we get a DLV program (the minimality of the models guarantees that every node cannot belong to both relations v and nv). The optimization query of examples 1 and 2 can be defined by adding the weak constraint $:\sim v(X)$ which minimizes the number of ground false weak constraints (i.e., v -tuples).

Smodels (Helsinki University of Technology) (Simons *et al.* 2002) is a system for answer set programming consisting of Smodels, an efficient implementation of the stable model semantics for normal logic programs and *lparse*, a front-end that transforms user programs so that they can be understood by Smodels.

Besides standard rules *lparse* also supports a number of extended rules: *choice*, *constraint* and *weight rules*. The formal semantics of all three types of rules can be defined through the use of *weight constraints* and *weight constraint rules*. In *lparse* the weight constraints are implemented as special literal types. Basically, a weight constraint is of the form: $L \leq l_1 = w_1, \dots, l_n = w_n \leq U$ where l_1, \dots, l_n , are literals, L and U are the integral lower and upper bounds, and w_1, \dots, w_n are weights of the literals. The intuitive semantics of a weight constraint is that it is satisfied exactly when the sum of weights of satisfied literals l_1, \dots, l_n is between L and U , inclusive. A weight constraint rule is of the form $C_0 \leftarrow C_1, \dots, C_n$ where C_0, \dots, C_n are weight constraints. Besides the use of literals, *lparse* also enhances

the use of *conditional literals* having the form: $p(X) : q(X)$ where $p(X)$ is any basic literal and $q(X)$ is a domain predicate.

Datalog Constraint. (East and Truszczyński 2000) proposed a new nonmonotonic logic, called *Datalog with constraints* or *DC*. A *DC* theory consists of constraints and Horn rules (*Datalog program*). The language is determined by a set of atoms $At = At_C \cup At_H$ where At_C and At_H are disjoint. Formally, a *DC* theory is a triple $T = (T_C, T_H, T_{PC})$ where T_C is a set of constraints over At_C , T_H is a set of Horn rules whose head atoms belong to At_H and T_{PC} is a set of constraints over At (*post constraints*). The problem of the existence of an answer set, for a finite propositional *DC* theory T , is \mathcal{NP} -complete (East and Truszczyński 2000).

ASSAT. (Lin and Zhao 2004) proposed a translation from normal logic programs with constraints under the answer set semantics to propositional logic. The peculiarity of this technique consists in the fact that for each loop in the program, a corresponding loop formula to the program's completion is added. The result is a one-to-one correspondence between the answer sets of the program and the models of the resulting propositional theory. As in the worst case the number of loops in a logic program can be exponential, the technique proposes to add a few loop formulas at a time, selectively. Based on these results, a system called *ASSAT(X)*, depending on the SAT solver X used, has been implemented for computing answer sets of a normal logic program with constraints.

Cmodels (Lierler 2005a, 2005b) is an answer set programming system that uses the frontend *lparse* and whose main computational characteristic is that it computes answer sets using a SAT solver for search. *Cmodels* deals with programs that may contain disjunctive, choice, cardinality and weight constraint rules. The basic execution steps of the system can be outlined as follows: (1) the program's completion is produced; (2) a model of the completion is computed using a SAT solver; (3) if the model is indeed an answer set, then the model is returned, otherwise the system goes back to Step 2. The idea is thus to use a SAT solver for generating model candidates and then check if they are indeed the answer sets of a program. The way Step 3 is implemented depends on the class of a logic program.

Clasp (Gebser et al. 2007) is an answer set solver for (extended) normal logic programs. It combines the high-level modeling capacities of answer set programming (ASP) with state-of-the-art techniques from the area of Boolean constraint solving. In fact, the primary *Clasp* algorithm relies on conflict-driven learning, a technique that proved successful for satisfiability checking (SAT). Unlike other ASP solvers that use conflict-driven learning, *Clasp* does not rely on legacy software, such as a SAT solver or any other existing ASP solver.

A-Prolog (Gelfond 2002) is a logic language whose semantics is based on stable models, designed to represent defaults (i.e., statements of the form "Elements of a class C normally satisfy property P "), exceptions and causal effects of actions ("statement F becomes true as a result of performing an action A "). In the same work, an extension of the language, called *ASET-Prolog*, is presented. Such an extension enriches the language with two new types of atoms: *s-atoms*, which allows us to define subsets of relations, and *f-atoms*, which allows us to

express constraints on the cardinality of sets. An interesting application showing how declarative programming in A-Prolog can be used to describe the dynamic behavior of digital circuits is presented in (Balduccini *et al.* 2000).

6.4 Comparison with the other approaches proposed in the literature

NP Datalog is related (i) to specification languages, for the style of defining problems, (ii) to answer set languages, for the syntax and declarative semantics, and (iii) to constraint programming.

6.4.1 Specification languages

The problem with specification languages is the tradeoff between the expressiveness of the formal notation and its execution. In general, specifications can be executed only by blind search through the space of all proofs. A possible solution consists in adding (to specifications) refinements which improve the execution, but the result could be a longer specification, containing details and, consequently, hard to understand.

NP-SPEC programs have a structure similar to the one of *NP Datalog* programs, although from the syntax point of view, the use of meta-predicates, in some cases, does not make programs shorter and more intuitive (see, for instance, the N-Queen problem reported in (Cadoli and Schaerf 2005)). NP-SPEC uses in addition to standard Datalog rules, also meta-predicates and set operators, whereas *NP Datalog* uses only standard Datalog rules with shortcuts for limited forms of (unstratified) negation. Moreover, although there is no difference in expressivity, guesses in NP-SPEC are defined over base relations, whereas in *NP Datalog* they are defined over general ‘deterministic’ relations defined by stratified Datalog programs. As a further difference, the partition mechanism is more general and flexible in *NP Datalog* w.r.t. NP-SPEC as in the latter the number of partitions is fixed. Concerning the semantics aspects, the declarative semantics of NP-SPEC programs is based on the notion of model minimality, whereas those of *NP Datalog* is based on stable models.

KIDS results are “sensitive” to the implementation issue. Indeed, the *KIDS* system is semiautomatic: the user is asked to interact with the system in order to transform high level declarative specification into an efficient, correct and executable program. Moreover, the complexity of the final implementation in *KIDS* can result in dramatic improvements if specialized techniques are used. On the other hand, *NP Datalog* is a fully declarative language whose execution process is automatically optimized by the ILOG OPL Development Studio.

SPILL-2 is not meant to use the specification of a problem in order to compute a solution, but to test the specification against some specific case, i.e., to verify whether a given specification implies certain intended properties or, in other words, if a specified property is consistent with the specification. As for differences, specification and queries in *SPILL* are compiled to Prolog, whereas our approach introduces specifications using *NP Datalog* and then performs the translation

of queries into OPL programs. Moreover, *NP Datalog* is based on stable model semantics, whereas SPILL uses a pure first-order semantics, i.e., it does not include any form of model minimization operations. As for a further difference, it is worth noting that SPILL does not provide a characterization of its expressive power and its complexity.

6.4.2 Constraint and logic programming languages

Constraint Logic Languages, such as SICStus Prolog, ECLiPSe and BProlog, are extensions of Prolog and, therefore, they are not fully declarative. Their semantics is based on top-down evaluation of queries (SLDNF resolution), whereas answer-set programming is based on bottom-up evaluation. *XSB* is an extension of Prolog with a declarative semantics (namely the well-founded semantics) based on top-down evaluation of queries (OLDT resolution) with tabling. Moreover, while answer-set languages permit *NP* problems to be easily expressed (it suffices to translate their logic definition into logic programming rules), constraint logic programming languages are procedural and the efficient implementation of *NP* problems is hard and time-consuming.

The relationship between ID Logic and *NP Datalog* is strong since in our language we can also distinguish components describing data, guess predicates, standard rules and constraints, which correspond, respectively, to the ID logic components describing data, open predicates, definitions and constraints. Moreover, the aim of *NP Datalog* is also the easy translation into different formalisms (other than ASP), including constraint programming languages.

6.4.3 Answer set programming languages

The main difference of *NP Datalog* with respect to DLV and Smodels is that only restricted forms of (unstratified) negations, embedded into built-in constructs, are allowed. As a consequence, *NP Datalog* is less expressive than DLV since the latter also uses (inclusive) disjunction and permits expression of problems in the second level of the polynomial hierarchy. The use of simpler languages such as *NP Datalog* allows us to avoid writing nonintuitive queries which are difficult to optimize or translate in other formalisms for which efficient executors exist. It is important to observe that cardinality constraints and conditional literals of Smodels allows us to express both subset and (generalized) partition rules as defined in *NP Datalog*. This means that in Smodels it is also possible to avoid using unstratified negation without losing expressiveness. We also note that *s-atoms* and *f-atoms* of ASET-Prolog enable us to express subset and (generalized) partition rules.

NP Datalog is also strongly connected to Datalog Constraint (DC), which is also based on stable model semantics. The main difference between *NP Datalog* and DC consists in the fact that *NP Datalog* forces users to write queries in a more disciplined form. In particular, DC guesses are expressed by means of constraints (a guess is any set of atoms in At_C satisfying the constraint in T_C) and there is no clear separation between T_C (constraints used to *guess*) and T_{PC} (constraints used to *check*). Moreover, DC only uses positive rules to infer true atoms, whereas

NP Datalog uses stratified rules. The expressive power of both languages captures the first level of the polynomial hierarchy. The experiments reported in (East and Truszczyński 2000) show that the guess and check style of expressing hard problems can be further optimized.

The philosophy of ASSAT is similar to that of *NP Datalog*: while in the ASSAT approach programs are translated in propositional logic and then executed by means of a SAT solver, *NP Datalog* programs are translated into OPL programs and then executed by using the ILOG OPL Development Studio. An approach similar to the one of ASSAT is adopted by Cmodels.

A-Prolog is a general logic language (i.e., it allows function symbols, classical negation, head disjunction, and subset rules) whose semantics is based on stable models. The aim of A-Prolog is the design of a general language for knowledge representation and causal reasoning, whereas *NP Datalog* is a simpler language (similar to the restricted FA-Prolog) which can be easily efficiently executed and translated in other formalisms.

7 Conclusion

NP search and optimization problems can be formulated as DATALOG^\neg queries under nondeterministic stable model semantics. In order to enable a simpler and more intuitive formulation of these problems, the *NP Datalog* language has been proposed. It is obtained by extending stratified Datalog with constraints and two constructs for expressing partitions of relations, so that search and optimization queries can be expressed using only simple forms of unstratified negation. It has also been shown that *NP Datalog* captures the class of *NP* search and optimization problems and that *NP Datalog* queries can be easily translated into OPL programs. An algorithm for the translation of *NP Datalog* programs into OPL statements has been provided and its correctness has been proved. The proposed algorithm has been implemented by a system prototype which takes in input an *NP Datalog* query and gives in output an equivalent OPL program which is then executed using the ILOG OPL Development Studio. Consequently, *NP Datalog* can also be used to define a logic interface for constraint programming solvers. Several experiments comparing the computation of queries by different systems have shown the validity of our approach.

References

- ABITEBOUL, S., HULL, R. AND VIANU, V. 1995. *Foundations of Databases*. Addison-Wesley.
- BALDUCCINI, M., GELFOND, M. AND NOGUEIRA, M. 2000. A-Prolog as a tool for declarative programming. In *Proc. of the International Conference on Software Engineering and Knowledge Engineering* 63–72.
- BARAL, C. 2003. *Knowledge Representation Reasoning and Declarative Problem Solving*. Cambridge University Press.
- CADOLI, M., IANNI, G., PALOPOLI, L., SCHAEFER, A. AND VASILE, D. 2000. NP-SPEC: An executable specification language for solving all problems in NP. *Computer Languages* 26(2–4), 165–195.

- CADOLI, M. AND SCHAERF, A. 2005. Compiling problem specifications into SAT. *Artificial Intelligence* 162(1–2), 89–120.
- CHOLEWINSKI, P., MAREK, V. W. AND TRUSZCZYNSKI, M. 1996. Default reasoning system DeReS. In *Proc. of the International Conference on Principles of Knowledge Representation and Reasoning*, L. C. Aiello et al., Eds. Morgan Kaufmann, 518–528.
- DENECKER, M. 2000. Extending classical logic with inductive definitions. In *Proc. of the International Conference on Computational Logic*, J. W. Lloyd et al., Eds. LNCS, vol. 1861. Springer, 703–717.
- EAST, D. AND TRUSZCZYNSKI, M. 2000. DATALOG with constraints – An answer-set programming system. In *Proc. of the National Conference on Artificial Intelligence*, AAAI Press/The MIT Press, 163–168.
- EAST, D. AND TRUSZCZYNSKI, M. 2006. Predicate-calculus based logics for modeling and solving search problems. *ACM Transaction on Computational Logic* 7(1), 38–83.
- EITER, T., LEONE, N., MATEIS, C., PFEIFER, G. AND SCARCELLO, F. 1997. A deductive system for non-monotonic reasoning. In *Proc. of the International Conference on Logic Programming and Nonmonotonic Reasoning*, J. Dix et al., Eds. LNCS, vol. 1265. Springer, 363–374.
- FAGIN, R. 1974. Generalized first-order spectra and polynomial-time recognizable sets. In *Complexity of Computation*, R. Karp, Ed. SIAM-AMS Proc., vol. 7, 43–73.
- FINKEL, R. A., MAREK, V. W. AND TRUSZCZYNSKI, M. 2004. Constraint Lingo: Towards high-level constraint programming. *Software Practice and Experience* 34(15), 1481–1504.
- GEBSER, M., KAUFMANN, B., NEUMANN, A. AND SCHAUB, T. 2007. Conflict-driven answer set solving. In *Proc. of the International Joint Conference on Artificial Intelligence*, M. M. Veloso, Ed., 386–392.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *Proc. of the International Conference on Logic Programming*, R. A. Kowalski and K. A. Bowen, Eds. The MIT Press, 1070–1080.
- GELFOND, M. 2002. Representing knowledge in a-prolog. *Computational Logic: Logic Programming and Beyond*, A. C. Kakas and F. Sadri, Eds. LNCS, vol. 2408. Springer, 413–451.
- GRECO, S. 1999. Dynamic programming in datalog with aggregates. *IEEE Transactions on Knowledge and Data Engineering* 11(2), 265–283.
- GRECO, S. AND SACCÀ, D. 2002. Search and optimization problems in datalog. *Computational Logic: Logic Programming and Beyond*, A. C. Kakas and F. Sadri, Eds. LNCS, vol. 2408. Springer, 61–82.
- GRECO, S. AND SACCÀ, D. 1997. NP optimization problems in datalog. In *Symposium on Logic Programming*, J. Maluszynski, Ed. The MIT Press, 181–195.
- GRECO, S., SACCÀ, D. AND ZANIOLO C. 1995. Datalog with stratified negation and choice: from P to D^P . In *Proc. of the International Conference on Database Theory*, G. Gottlob and M. Y. Vardi, Eds. LNCS, vol. 893. Springer, 574–589.
- GRECO, S. AND ZANIOLO, C. 2001. Greedy algorithms in datalog. *Theory and Practice of Logic Programming* 1(4), 381–407.
- JAFFAR, J., MICHAYLOV, S., STUCKEY, P.J. AND YAP, R. 1992. The CLP(R) language and system. *ACM Transaction on Programming Languages and Systems* 14(3), 339–395.
- JOHNSON, D. S. 1990. A Catalog of Complexity Classes. In *Handbook of Theoretical Computer Science*, Vol. 1, J. van Leewen, Ed. North-Holland, 67–161.
- KLUZNIAK, F. AND MILKOWSKA, M. 1997. Spill – A logic language for writing testable requirements specifications. *Science of Computer Programming* 28(2–3), 193–223.
- KOLAITIS, P. AND PAPADIMITRIOU, C. 1991. Why not negation by fixpoint? *Journal of Computer and System Sciences* 43(1), 125–144.

- KOLAITIS, P. G. AND THAKUR, M. N. 1995. Approximation properties of NP minimization classes. *Journal of Computer and System Sciences* 50 (3), 391–411.
- LEONE, N., PFEIFER, G., FABER, W., EITER, T., GOTTLÖB, G., PERRI, S. AND SCARCELLO, F. 2006. The DLV System for knowledge representation and reasoning. *ACM Transactions on Computational Logics* 7(3), 499–562.
- LIERLER, Y. 2005. Disjunctive answer set programming via satisfiability. In *Proc. of the International Workshop on Answer Set Programming*, M. De Vos and A. Proveti, Eds. CEUR Workshop Proceedings 142, 85–97.
- LIERLER, Y. 2005. Cmodels – SAT-based disjunctive answer set solver. In *Proc. of the International Conference on Logic Programming and Nonmonotonic Reasoning*, C. Baral et al., Eds. LNCS, vol. 3662. Springer, 447–451.
- LIN, F. AND ZHAO, Y. 2004. ASSAT: Computing answer sets of a logic program by SAT solvers. *Artificial Intelligence* 157(1–2), 115–137.
- MAREK, W. AND TRUSZCZYNSKI, M. 1991. Autoepistemic logic. *Journal of the ACM* 38(3), 588–619.
- MARIEN, M., GILIS, D. AND DENECKER, M. 2004. On the relation between ID-logic and answer set programming. *Logics in Artificial Intelligence*, J. J. Alferes and J. A. Leite, Eds. LNCS, vol. 3229. Springer, 108–120.
- MARRIOTT, K. AND STUCKEY, P. J. 1998. *Programming with Constraints: An Introduction*. MIT Press.
- NIEMELA, I., SIMONS, P. AND SOININEN, T. 1999. Stable model semantics of weight constraint rules. In *Proc. of the International Conference on Logic Programming and Nonmonotonic Reasoning*, M. Gelfond et al., Eds. LNCS, vol. 1730. Springer, 317–331.
- PAPADIMITRIOU, C. H. AND STEIGLITZ, K. 1982. *Combinatorial Optimization – Algorithms and Complexity*. Prentice-Hall.
- PAPADIMITRIOU, C. H. AND YANNAKAKIS, M. 1991. Optimization, approximation, and complexity classes. *Journal Computer and System Sciences* 43(3), 425–440.
- PAPADIMITRIOU, C. H. 1994. *Computational Complexity*. Addison-Wesley.
- RAO, P., SAGONAS, K. F., SWIFT, T., WARREN, D. S. AND FREIRE, J. 1997. XSB: A system for efficiently computing WFS. In *Proc. of the International Conference on Logic Programming and Nonmonotonic Reasoning*, J. Dix et al., Eds. LNCS, vol. 1265. Springer, 431–441.
- SACCÀ, D. 1997. The expressive powers of stable models for bound and unbound queries. *Journal of Computer and System Sciences* 54(3), 441–464.
- SCHLIPF, J. S. 1995. The expressive powers of the logic programming semantics. *Journal of Computer and System Sciences* 51(1), 64–86.
- SMITH, D. R. 1990. KIDS: A semiautomatic program development system. *IEEE Transaction on Software Engineering* 16(9), 1024–1043.
- SIMONS, P., NIEMELA, I. AND SOININEN, T. 2002. Extending and implementing the stable model semantics. *Artificial Intelligence* 138(1–2), 181–234.
- ULLMAN, J. K. 1988. *Principles of Data and Knowledge-Base Systems*, Vol. 1. Computer Science Press, New York.
- VAN GELDER, A., ROSS, K. A. AND SCHLIPF, J. S. 1991. The well-founded semantics for general logic programs. *Journal of ACM* 38 (3), 620–650.
- VAN HENTENRYCK, P. 1988. *The OPL Optimization Programming Language*. MIT Press.
- VAN HENTENRYCK, P., MICHEL, L., PERRON, L. AND REGIN, J. C. 1999. Constraint programming in OPL. In *Proc. of the International Conference on Principles and Practice of Declarative Programming*, G. Nadathur, Ed. LNCS, vol. 1702. Springer, 98–116.
- VAN ROY, P. 1999. Logic programming in Oz with Mozart. In *Proc. of the International Conference on Logic Programming*, D. De Schreye, Ed. The MIT Press, 38–51.

- WALLACE, M. AND SCHIMPF, J. 1999. ECLiPSe: declarative specification and scaleable implementation. In *Proc. of the International Workshop on Practical Aspects of Declarative Languages*, G. Gupta, Ed. LNCS, vol. 1551. Springer, 365–366.
- ZHOU, N. 2002. Implementing constraint solvers in B-prolog. *Intelligent Information Processing*, M. A. Musen et al., Eds. IFIP Conference Proceedings, 249–260.
- ZUMPARO, E., GRECO, S., TRUBITSYNA, I. AND VELTRI, P. 2004. On the semantics and expressive power of Datalog-like languages for NP search and optimization problems. In *Proc. of the ACM Symposium on Applied Computing*, H. Haddad et al., Eds. ACM, 692–697.