Enhancing the Expressive Power of the U-Datalog Language

ELISA BERTINO
University of Milano
Via Comelico 39, 20135 Milano, Italy
e-mail: bertino@dsi.unimi.it

BARBARA CATANIA
University of Genova
Via Dodecaneso 35, 16146 Genova, Italy
e-mail: catania@disi.unige.it

ROBERTA GORI
University of Pisa
Corso Italia, 40
56125 Pisa, Italy
e-mail: gori@di.unipi.it

Abstract

U-Datalog has been developed with the aim of providing a set-oriented logical update language, guaranteeing update parallelism in the context of a Datalog-like language. In U-Datalog, updates are expressed by introducing constraints ($+p(X)$, to denote insertion, and $-p(X)$, to denote deletion) inside Datalog rules. A U-Datalog program can be interpreted as a CLP program. In this framework, a set of updates (constraints) is satisfiable if it does not represent an inconsistent theory, that is, it does not require the insertion and the deletion of the same fact. This approach resembles a very simple form of negation. However, on the other hand, U-Datalog does not provide any mechanism to explicitly deal with negative information, resulting in a language with limited expressive power. In this paper, we provide a semantics, based on stratification, handling the use of negated atoms in U-Datalog programs and we show which problems arise in defining a compositional semantics.

1 Introduction

Deductive database technology represents an important step towards the goal of developing highly-declarative database programming languages. Several approaches for the inclusion of update capabilities in deductive languages have been proposed. In general, all those proposals are based on including in rules, in addition to usual atoms, special atoms denoting updates. In most of those proposals, an update execution consists of a query component, identifying the data to be modified, and an update component, performing the actual modification on the selected data. A way to classify deductive update languages is with respect to the approach adopted for handling possible interferences between the query and update component of the same update execution. In particular, updates can be performed as soon as they are generated, as side-effect of the query evaluation, thus, by applying an *immediate* semantics. Languages based on an immediate semantics include $\mathcal{LDL}$ (Naqvi & Tsur, 1989), TL (Abiteboul & Vianu, 1991), DL (Abiteboul & Vianu, 1991), DLP (Manchanda & Warren, 1988), Statelog (Lausen *et al.*, 1998). The immediate semantics is in contrast with the *deferred update semantics*, by which updates are not applied as soon as they are generated during the query evaluation; rather, they are executed only when the query evaluation is completed. Languages based on a deferred semantics include U-Datalog (Bertino *et al.*, 1998b), Update Calculus (Chen, 1995; Chen, 1997), and ULTRA (Wichert & Freitag, 1997; Wichert *et al.*, 1998). Other languages, such as Transaction Logic (Bonner & M.Kifer, 1994), provide both policies.

In this paper, we consider U-Datalog, a language based on a deferred semantics. Even if more expressive and flexible frameworks exist (see for example (Bonner & M.Kifer, 1994; Wichert *et al.*, 1998)), the choice of U-Datalog is motivated by the fact that it represents an immediate extension of Datalog to deal with updates. This aspect makes this language quite suitable for analyzing properties related to logical update languages (Bertino & Catania, 1996). In U-Datalog, updates are expressed by introducing constraints inside Datalog rules. For example, $+p(a)$ states that in the new state $p(a)$ must be true where $-p(a)$ states that in the new state $p(a)$ must be false. Thus, U-Datalog programs are formally modeled as Constraint Logic Programming (CLP) programs (Jaffar *et al.*, 1998).

In CLP, any answer to a given goal (called a *query*, in the database context) contains a set of constraints, constraining the resulting solution. In U-Datalog, each solution contains a substitution for the query variables and a set of updates. The execution of a goal is based on a deferred semantics. In particular, given a query, all the solutions are generated in the so-called *marking phase*, using a CLP answering mechanism. All the updates, contained in the various solutions, are then executed in the *update phase*, by using an operational semantics. The set of all updates generated during the marking phase forms a constraint theory which can be inconsistent. From a logical point of view, this means that the update set contains constraints of the form $+p(a)$, $-p(a)$, requiring the insertion and the deletion of the same fact. The U-Datalog computational model rejects any form of conflict, both locally, i.e., inside a single solution, and globally, from different solutions. Thus, the set of updates to be executed is always consistent.

Besides the marking and update semantics phases, it is often useful to devise an additional semantics, known as *compositional semantics* (Bertino *et al.*, 1998b). This semantics, which is orthogonal with respect to the one defined above, characterizes the semantics of the intensional database independently from the semantics of the extensional one and is based on the notion of open programs (Bossi *et al.*, 1994). The compositional semantics is quite important in the context of deductive databases since it provides a theoretical framework for analyzing the properties of intensional databases. Indeed, it is always recursion free, even if it is not always finite. Therefore, when it is finite, it also represents a useful pre-compilation technique for intensional databases. However, since this semantics is usually expensive to compute, it is mainly used for analysis purposes.

Even if U-Datalog allows us to easily specify updates and transactions, its expressive power is limited since no negation mechanism is provided, even if, due to update inconsistency, some limited form of negation on the extensional database is provided. This kind of negation is obviously not sufficient to support a large variety of user requests. In this paper we provide an operational mechanism handling negated atoms in U-Datalog programs, providing a marking phase and a compositional semantics. The proposed extension is based on the notion of stratification, first proposed for logic programming and deductive databases (Abiteboul & Vianu, 1991; Lausen *et al.*, 1998; Manchanda & Warren, 1988; Naqvi & Tsur, 1989). This extension is not, however, a straightforward extension of previously defined stratification-based semantics for two main reasons. First of all, U-Datalog rules are not range restricted (Ceri *et al.*, 1990) but are required to be *safe through query invocation*,[1] resulting in a non-ground semantics. Note that, even if this is a typical assumption in a real context, most of the other deductive update languages require range restricted rules or interpret free variables as *generation of new values* (Abiteboul & Vianu, 1991). A second difference is that an atom may fail not only because an answer substitution cannot be found but also because it generates an inconsistent set of updates.

In the following, we first introduce U-Datalog in Section 2 and we extend it to deal with negation in Section 3. Finally, in Section 4 we present some conclusions and outline future work. Due to space limitations, we assume the reader to be aware of the basic notions of (constraint) logic programming (Jaffar *et al.*, 1998; Lloyd, 1987) and deductive databases (Ceri *et al.*, 1990). For additional details on U-Datalog, see (Bertino *et al.*, 1998b).

2 U-Datalog

2.1 Syntax

A U-Datalog database consists of: (i) an extensional database (or simply database) EDB, that is, a set of ground atoms (*extensional atoms*); (ii) an intensional database IDB (or simply program), that is, a set of rules of the form:[2]

$$H \leftarrow b_1, \ldots, b_k, u_1, \ldots, u_s, B_1, \ldots, B_t$$

where $H, B_1, \ldots, B_t$ are atoms, $b_1, \ldots, b_k$ are equality constraints, i.e. constraints of type $X = t$ (denoted by $\tilde{b}$), where X is a variable and t is a term, and $u_1, \ldots, u_s$ are *update constraints* (denoted by $\tilde{u}$), also called *update atoms*. An update constraint is an extensional atom preceded by the symbol $+$, to denote an insertion, or by the symbol $-$, to denote a deletion.

In the following, the set of extensional predicates is denoted by Π_{EDB}, the set of intensional predicates is denoted by Π_{IDB}, and the Herbrand Universe is denoted

[1] See Section 2.1 for the formal definition of this property.
[2] In the following, we assume that constants and multiple occurrences of the same variable inside each atom are expressed by equality constraints between terms.

by $\mathcal{H}$ (Lloyd, 1987). Moreover, a conjunction of equality and update constraints is simply called *constraint*.[3] As usual in deductive databases, Π_{EDB} and Π_{IDB} are disjoint. Note that a U-Datalog program can be seen as a CLP program where constraints are represented by equalities and update atoms.

A U-Datalog transaction is a goal. In order to guarantee a finite answer to each goal and the generation of a set of ground updates, we assume that rules are *safe through query invocation*. This means that, given a U-Datalog database IDB and a goal G, each variable appearing in the head or in the update constraint of a rule, used in the evaluation of the goal, either appears in an atom contained in the body of the same rule, or is bound by a constant present in the goal. In this case, G is *admissible* for IDB.

Example 1
The following program is a U-Datalog intensional database:

$r1 : rem\_man(X, Y) \leftarrow -dep\_A(Y), emp\_man(X, Y)$
$r2 : rem\_man(X, Y) \leftarrow -dep\_A(Y), emp\_man(X, Z), rem\_man(Z, Y)$
$r3 : ins\_man(X) \leftarrow +dep\_A(X), rem\_man(X, Y)$

An atom $emp\_man(a, b)$ is true if 'b' is a manager of 'a'. An atom $rem\_man(a, b)$ is true if 'b' is a (possibly indirect) manager of 'a'. As a side effect, it requires the removal of 'b' from department A. An atom $ins\_man(a)$ is true if 'a' has at least one manager and, as side effect, requires the insertion of 'a' in department A. At the same time, it requires the deletions of all the (possibly indirect) managers of 'a' from department A. ◇

2.2 Semantics

U-Datalog constraints are interpreted over the Herbrand Universe $\mathcal{H}$. In this domain, equalities have the usual meaning: $+p(\tilde{X})$ is interpreted as the atom $p(\tilde{X})$ and $-p(\tilde{X})$ is interpreted as the negated atom $\neg p(\tilde{X})$. If a constraint $\tilde{b} \wedge \tilde{u}$ is $\mathcal{H}$-solvable, i.e. if $\mathcal{H} \models \tilde{b}, \tilde{u}$, there exists at least one substitution that makes the constraint true. Thus, for no atom both an insertion and a deletion are simultaneously required. When this is not true, updates are said to be *inconsistent*. The execution of ground inconsistent update atoms (e.g., $+p(a), -p(a)$) may lead to different extensional databases, with respect to the chosen execution order.

The generation of inconsistent updates is avoided as follows: (i) *locally*: a solution containing an inconsistent set of updates (i.e., an unsolvable set of constraints) is not included in the resulting set of solutions for the given goal; (ii) *globally*: if an inconsistency is generated due to two consistent solutions, the goal aborts, no update is executed, and the database is left in the state it had before the goal evaluation.

The semantics of U-Datalog programs is given in two main steps. In the first

[3] In the following, conjunction between constraints is represented by using ',' inside body rules, and by using '$\wedge$' in other contexts.

step, all solutions for a given goal are determined by applying a CLP evaluation method (*marking phase*, see Section 2.2.2). Each solution contains a set of bindings for the query variables and a set of consistent update atoms. In the second step (*update phase*, see Section 2.2.3), the updates gathered in the various solutions are executed only if they are consistent.

Besides the marking and the update semantics, an additional semantics is sometimes introduced in the database context, which is called *compositional semantics* (see Section 2.2.1). Such semantics characterizes the intensional database independently from the semantics of the extensional one. The compositional semantics is quite important in the context of deductive databases since it provides a theoretical framework for analyzing the properties of intensional databases.

2.2.1 Compositional semantics

Since the extensional database is the only time-variant component of a U-Datalog database, for analysis purposes, it is useful to define the semantics of a U-Datalog intensional database independently from the current extensional database. Such semantics is called *compositional semantics* and is always represented by a recursion free set of rules. Therefore, when it is finite, or when an equivalent finite set of rules can be detected, it also represents a useful pre-compilation technique for intensional databases.

The compositional semantics can be defined assuming the intensional database to be an *open program* (Bossi *et al.*, 1994), i.e., a program where the knowledge regarding some predicates is assumed to be incomplete. Under this meaning, a U-Datalog intensional database can be seen as a program that is open with respect to the extensional predicates. The semantics of an open program is a set of rules, whose bodies contain just open predicates. In order to define the compositional semantics of a U-Datalog intensional database, we introduce the following set:

$$ID_{EDB} = \{p(\tilde{X}) \leftarrow p(\tilde{X}) \mid p \in \Pi_{EDB}\}.$$

In the previous expression, $\tilde{X}$ denotes a list of distinct variables. Similarly to (Bertino *et al.*, 1998b), we now introduce an *unfolding operator*. Such operator, given programs P and Q, replaces an atom $p(\tilde{X})$, appearing in the body of a rule in P, with the body of a rule defining p in Q.

Definition 1
Let P and Q be U-Datalog programs. Then[4]

$$\begin{aligned} Unf_P(Q) = \{p(\tilde{X}) \leftarrow \tilde{b}', \tilde{u}', \tilde{H}_1, \ldots, \tilde{H}_n \mid & \exists \text{ a renamed rule} \\ & p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, p_1(\tilde{X}_1), \ldots, p_n(\tilde{X}_n) \in P \\ & \exists p_i(\tilde{Y}_i) \leftarrow \tilde{b}_i, \tilde{u}_i, \tilde{H}_i \in Q \ (i = 1, ..., n), \text{ which share no variables,} \\ & \tilde{b}' \equiv \textstyle\bigwedge_i (\tilde{b}_i \wedge (\tilde{X}_i = \tilde{Y}_i)) \wedge \tilde{b} \\ & \tilde{u}' \equiv \textstyle\bigwedge_i \tilde{u}_i \wedge \tilde{u} \\ & \tilde{b}' \wedge \tilde{u}' \text{ is } \mathcal{H}\text{-solvable}\} \end{aligned}$$

□

[4] In the following, we use the notation $\bigwedge_i c_i$ to represent the conjunction of constraints $c_1 \wedge ... \wedge c_n$, where n is clear from the context. The symbol $\equiv$ denotes syntactic equality.

The compositional semantics of a U-Datalog intensional database IDB is obtained by repeatedly applying the unfolding operator until no new rules are generated.

Definition 2
The compositional semantics $\mathcal{U}_{IDB}$ of IDB with respect to Π_{EDB} is defined as the least fixpoint of $T^c_{IDB}(I) = Unf_{IDB}(I \cup ID_{EDB})$. □

The previous definition is based on the following result, taken from (Bertino *et al.*, 1998b).

Theorem 1
T^c_{IDB} is continuous. □

Theorem 2
For any extensional database EDB, for any admissible goal G, the evaluation of G in $IDB \cup EDB$ generates the same answer constraints than the evaluation of G in $\mathcal{U}_{IDB} \cup EDB$. □

Note that $\mathcal{U}_{IDB}$ is always recursion free. If IDB is a recursive program then $\mathcal{U}_{IDB}$ in general is not finite. However, under specific assumptions, it is equivalent to a finite set of rules (see Section 3.2).

2.2.2 Marking phase

The answers to a U-Datalog query can be computed in a top-down or in an equivalent bottom-up style (Jaffar *et al.*, 1998). Here, we introduce only the bottom-up semantics. The Constrained Herbrand Base $\mathcal{B}$ for a U-Datalog program is defined as the set of *constrained* atoms of the form $p(\tilde{X}) \leftarrow b_1, ..., b_k, u_1, ..., u_n$ where $u_1, ..., u_n$ are update atoms, $b_1, ...b_k$ are equality constraints, $p \in \Pi_{EDB} \cup \Pi_{IDB}$, and $\tilde{X}$ is a tuple of distinct variables. An interpretation is any subset of the Constrained Herbrand Base. Given a U-Datalog database $DB = IDB \cup EDB$,[5] operator $T_{DB} : 2^{\mathcal{B}} \rightarrow 2^{\mathcal{B}}$ is defined as follows:[6]

$$
\begin{aligned}
T_{DB}(I) = \{p(\tilde{X}) \leftarrow \tilde{b}', \tilde{u}' \mid \exists \text{ a renamed rule} \\
& p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, p_1(\tilde{Y}_1), \ldots, p_n(\tilde{Y}_n) \in DB \\
& \exists p_i(\tilde{X}_i) \leftarrow \tilde{b}_i, \tilde{u}_i \in I \ (i = 1, ..., n), \text{ which share no variables} \\
& \tilde{b}' \equiv \textstyle\bigwedge_i (\tilde{b}_i \wedge (\tilde{X}_i = \tilde{Y}_i)) \wedge \tilde{b} \\
& \tilde{u}' \equiv \textstyle\bigwedge_i \tilde{u}_i \wedge \tilde{u} \\
& \tilde{b}' \wedge \tilde{u}' \text{ is } \mathcal{H}\text{-solvable } \}.^7
\end{aligned}
$$

Theorem 3

[5] Even if ground atoms contained in the extensional database should be represented as constrained atoms, we still write them as ground atoms to simplify the notation.
[6] $2^{\mathcal{B}}$ is the set of all the subsets of the Constrained Herbrand Base $\mathcal{B}$.
[7] We assume that all constraints generated by a fixpoint computation are projected onto the set of the head and update atom variables. Moreover, we assume that a constrained atom is inserted in the set being constructed only if it is not redundant.

Let DB be a U-Datalog database. T_{DB} is continuous and admits a unique least fixpoint $\mathcal{FIX}_{DB}$ and $\mathcal{FIX}_{DB} = T_{DB}\uparrow \omega$. Such fixpoint represents the bottom-up semantics of DB.[8] □

Given $\mathcal{FIX}_{DB}$ and a goal $G \equiv\leftarrow \tilde{b}, \tilde{u}, p_1(\tilde{X}_1), ..., p_n(\tilde{X}_n)$, the *solutions* or *answer constraints* for G are all constraints $\tilde{X}_1 = \tilde{Y}_1, ..., \tilde{X}_n = \tilde{Y}_n, \tilde{b}', \tilde{u}'$ such that $p_i(\tilde{Y}_i) \leftarrow \tilde{b}_i, \tilde{u}_i \in \mathcal{FIX}_{DB}$ $(i = 1, ..., n)$, $\tilde{u}' \equiv \tilde{u} \wedge \tilde{u}_1 \wedge ... \wedge \tilde{u}_n$, $\tilde{b}' \equiv \tilde{b} \wedge \tilde{b}_1 \wedge ... \wedge \tilde{b}_n$, and $\tilde{u}' \wedge \tilde{b}' \wedge \tilde{X}_1 = \tilde{Y}_1 \wedge ... \wedge \tilde{X}_n = \tilde{Y}_n$ is $\mathcal{H}$-solvable. Let $\tilde{b}'' \equiv \tilde{X}_1 = \tilde{Y}_1 \wedge ... \wedge \tilde{X}_n = \tilde{Y}_n \wedge \tilde{b}'$. In this case, we write $G, DB \overset{*}{\longmapsto} \langle \tilde{b}'', \tilde{u}'\tilde{b}'' \rangle$, where $\tilde{u}'\tilde{b}''$ denotes the result of the application of the equalities specified in $\tilde{b}''$ to $\tilde{u}'$. We assume that $\tilde{b}''$ is restricted to the variables of G. Note that $\tilde{u}'$ has to be consistent.

Example 2
Consider $EDB_i = \{emp\_man(b, b),\ emp\_man(b, c),\ dep\_A(b),\ dep\_A(c),\ dep\_B(b)\}$ and the intensional database of Example 1. Transaction $T_1 \equiv\leftarrow ins\_man(X)$ evaluated in $EDB_i \cup IDB$ computes the consistent solution $X = b, Y = c, -dep\_A(Y), +dep\_A(X)$. The additional solution $X = b, Y = b, -dep\_A(Y), +dep\_A(X)$ is not consistent and therefore is discarded by the marking phase. ◇

2.2.3 Update phase

The *update phase* atomically executes the updates collected by the marking phase. Updates gathered by the different solutions for a given predicate are executed only if no inconsistency arises. This guarantees that only order independent executions are performed. Formally, let $u = \bigcup\{\tilde{u}_j \mid G, DB \overset{*}{\longmapsto} \langle \tilde{b}_j, \tilde{u}_j \rangle\}$. Let EDB_i be the current database state. If u is a consistent and ground set of updates, the new database EDB_{i+1} is computed as follows: $EDB_{i+1} = (EDB_i \setminus \{p(\tilde{t}) \mid -p(\tilde{t}) \in u\}) \cup \{p(\tilde{t}') \mid +p(\tilde{t}') \in u\}$. In this case, we say that G *commits*, returning the tuple $\langle \{\tilde{b}_j \mid G, DB \overset{*}{\longmapsto} \langle \tilde{b}_j, \tilde{u}_j \rangle\}, EDB_{i+1}, Commit \rangle$. If u is inconsistent or contains at least one non-ground update atom,[9] we let $EDB_{i+1} = EDB_i$ and say that G *aborts*. In this case, the evaluation returns the tuple $\langle \{\}, EDB_i, Abort \rangle$.

Example 3
Consider EDB_i as in Example 2 and the intensional database of Example 1. The execution of transaction $T_1 \equiv\leftarrow ins\_man(X)$ generates the new extensional database $EDB_{i+1} = \{emp\_man(b, b),\ emp\_man(b, c),\ dep\_A(b),\ dep\_B(b)\}$. ◇

3 Introducing negation in U-Datalog

Since solutions containing inconsistent updates are not returned by the marking phase, the U-Datalog semantics models some kind of negation. This form of negation is however very weak with respect to the ability to model arbitrary negation. Indeed, it has been proved that, with respect to the returned substitutions, U-Datalog is

[8] We recall that $T_P \uparrow 0 = \emptyset$, $T_P \uparrow i = T_P(T_P \uparrow i-1)$, $T_P\uparrow \omega = \bigcup_{i<\omega} T_P \uparrow i$.
[9] Note that an unground set of updates can only be generated by a non-admissible goal.

equivalent to Datalog extended with negation on extensional predicates and open with respect to a subset of extensional predicates (Bertino & Catania, 2000). In the following, in order to increase the expressive power of U-Datalog, we introduce *negated atoms* in the bodies of U-Datalog rules. The resulting language is called U-Datalog$^\neg$. Then, we assign a semantics to such language when the considered programs are stratified. A stratified U-Datalog$^\neg$ program is defined as follows.

Definition 3

A U-Datalog$^\neg$ program IDB is stratified if it is possible to find a sequence $P_1, \ldots, P_n$, $P_i \subseteq IDB$ $(i = 1, ..., n)$, (also called stratification) such that the following conditions hold (in the following, we denote with $Pred_i$ the set of predicates defined in P_i):

1. $P_1, \ldots, P_n$ is a partition of the rules of IDB. Each P_i is called "stratum".
2. For each predicate $q \in Pred_j$, all the rules defining q in IDB are in P_j.
3. If $q(u) \leftarrow \ldots, q'(v), \ldots \in IDB$, $q' \in Pred_j$, then $q \in Pred_k$ with $j \leq k$.
4. If $q(u) \leftarrow \ldots, \neg q'(v), \ldots \in IDB$, $q' \in Pred_j$, then $q \in Pred_k$ with $j < k$. □

The previous definition can be extended to deal with a U-Datalog database $DB = IDB \cup EDB$. In this case, all extensional facts belong to the first level.

In order to assign a semantics to stratified U-Datalog$^\neg$ programs, we assume that each rule in the program is safe through query invocation. Due to the introduction of negation, the notion of safety is extended by requiring that each variable appearing in a rule head, in a negated literal contained in a rule body, or in an update atom also appears in a positive literal in the rule body or is bound by a constant present in the goal.

The main differences between the bottom-up semantics we are going to present and the bottom-up semantics defined for Stratified Datalog$^\neg$ programs (Ceri *et al.*, 1990; Chandra & Harel, 1985) are the following. Due to the condition of safety through query invocation, the semantics of a U-Datalog$^\neg$ program may contain non-ground constrained atoms that, however, will be made ground by the goal. Thus, negated atoms cannot be used, as usually done, as conditions to be satisfied by a solution. Indeed, some variables inside the generated solutions may be made ground by the goal. A solution to this problem is to explicitly represent, during the bottom-up computation, the solutions for which a negated atom $\neg B$ is true. In this way, we maintain all the conditions that the solutions have to satisfy but the check will be executed only when a match with a query goal is performed. To represent such solutions, the underlying constraint theory must be extended to deal with inequality constraints of type $X \neq a$, where X is a variable and a is a constant. For example, if $X = a$ is the only solution for $p(X)$, then $X \neq a$ is the solution for $\neg p(X)$.[10]

[10] Note that, due to the Closed Word Assumption (Ceri *et al.*, 1990), this is only a difference at the presentation level that allows us to treat in an homogeneous way positive and negative literals during the bottom-up computation.

A second aspect is related to the semantics of $\neg B$ with respect to the updates collected by B. B, in fact, can also fail due to the generation of inconsistent updates. Thus, all solutions containing inconsistent updates represent solutions for $\neg B$. Solutions for $\neg B$ in a database $IDB \cup EDB$ are therefore obtained by evaluating B in $IDB \cup EDB$, and complementing not only the computed constraints but also the constraints which ensure the consistency of the updates generated by evaluating B.

Finally, we assume that the derivation of $\neg B$ does not generate any update. This assumption is motivated by the fact that the evaluation of $\neg B$ should be considered as a test with respect to the bindings generated by positive atoms.

In the following, we present the marking phase and the compositional semantics for Stratified U-Datalog¬ programs. Note that no modification to the update phase is required. Proofs of the presented results can be found in (Bertino *et al.*, 1999).

3.1 The marking phase semantics

As a natural extension of the constraints domain presented in (Bertino *et al.*, 1998b), the Constrained Herbrand Base for U-Datalog$^\neg$ (denoted by $\mathcal{B}^\neg$) consists of constrained literals of the form $L \leftarrow \tilde{b}, \tilde{u}$, where $\tilde{b}$ is a conjunction of equality and inequality constraints, $\tilde{u}$ is a conjunction of update atoms, and L is a literal. If L is a negated atom, $\tilde{u}$ is empty. In the following, the set of all conjunctions of equalities and inequalities constraints, constructed on the Herbrand Universe $\mathcal{H}$, is denoted by $\mathcal{C}$.

Definition 4
Let $DB = IDB \cup EDB$ be a U-Datalog$^\neg$ database. The bottom-up operator $T^\neg_{DB} : 2^{\mathcal{B}^\neg} \to 2^{\mathcal{B}^\neg}$ is defined as follows:

$$
\begin{aligned}
T^\neg_{DB}(I) = \{p(\tilde{X}) \leftarrow \tilde{b}', \tilde{u}' \mid \exists \text{ a renamed rule} \\
& p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, L_1(\tilde{Y}_1), \ldots, L_n(\tilde{Y}_n) \in DB \\
& \exists L_i(\tilde{X}_i) \leftarrow \tilde{b}_i, \tilde{u}_i \in I \ (i = 1, ..., n), \text{ which share no variables} \\
& \tilde{b}' \equiv \textstyle\bigwedge_i (\tilde{b}_i \wedge (\tilde{X}_i = \tilde{Y}_i)) \wedge \tilde{b} \\
& \tilde{u}' \equiv \textstyle\bigwedge_i u_i \wedge \tilde{u} \\
& \tilde{b}' \wedge \tilde{u}' \text{ is } \mathcal{H}\text{-solvable } \}. \qquad \square
\end{aligned}
$$

Before introducing the fixpoint semantics, we define an operator Neg which performs the negation of a constraint belonging to $\mathcal{C}$.

Definition 5
Let $c = c_1 \wedge ... \wedge c_n$. $Neg : \mathcal{C} \to 2^{\mathcal{C}}$ is defined as follows:

$$
Neg(c) = \begin{cases}
\{c'_1, \ldots, c'_m\} & c'_1 \vee ... \vee c'_m \text{ is equivalent to } \neg c_1 \vee ... \vee \neg c_n \\
& c'_i \text{ is } \mathcal{H}\text{-solvable } (i = 1, ..., m) \\
& \text{and } \forall j, j = 1, ..., m, \\
& \ c'_1 \vee ... \vee c'_m \text{ is not equivalent to} \\
& c'_1 \vee ... \vee c'_{j-1} \vee c'_{j+1} \vee ... \vee c'_m, \\
\{false\} & \text{otherwise}
\end{cases} \qquad \square
$$

For example, if $c \equiv X = 2 \wedge Y = 3$, then $Neg(c) = \{X \neq 2, Y \neq 3\}$. Operator Neg is used to define an additional operator $Comp$, which takes a set S of constrained positive literals and returns the set H of constrained negative literals, belonging to the complement of S. This operator is used to make explicit the constraints for negative literals at the end of the computation of the positive literals of each stratum.

Definition 6
$Comp : 2^{B^{\neg}} \rightarrow 2^{B^{\neg}}$ is defined as follows:

- $\neg p(\tilde{X}) \leftarrow \ \in Comp(S)$ if there does not exist any $p(\tilde{Y}) \leftarrow \tilde{b}, \tilde{u} \in S$.
- $\neg p(\tilde{X}) \leftarrow \tilde{b'} \in Comp(S)$ iff $p(\tilde{X}) \leftarrow \tilde{b}_1, \tilde{u}_1, \ldots\ldots, p(\tilde{X}) \leftarrow \tilde{b}_n, \tilde{u}_n$ are the only (renamed apart) constrained atoms for p in S, $\overline{b'_i} \in Neg((\tilde{b_i} \wedge \tilde{bu})|_{\tilde{X}})$,[11] $\tilde{bu} \in Sol(\tilde{u_i}\tilde{b_i}^{12})$ $(i = 1, ..., n)$, $\tilde{b'} \equiv \bigwedge_i \overline{b'_i}$, and $\tilde{b'}$ is $\mathcal{H}$-solvable. □

In the previous definition, $Sol(\tilde{u})$ is the set of *minimal* [13] constraints which implies that $\tilde{u}$ is a consistent set of updates. For example, $Sol(+p(a,Y), -p(X,Z), -p(X,b), -p(b,c)) = \{X \neq a, Y \neq Z \wedge Y \neq b\}$. Of course, if $\tilde{u}$ is an inconsistent set of update atoms, no solution is generated and $Sol(\tilde{u}) = false$. We also assume that all redundant constrained literals contained in $Comp(S)$ are removed.

In the previous definition, operator Neg is applied to computed constraints and to the constraints which make satisfiable the updates generated by the corresponding positive atom. Such solutions have been restricted to the head variables since all the other variables are not needed to define the solutions for the negated atom.

The fixpoint of $T^{\neg}_{DB}$ is computed as follows. First, the fixpoint of a given stratum is computed. Then, all the facts that have not been derived are made explicitly false. This corresponds to locally apply the CWA. Note that, due to stratification, this approach is correct since each predicate is completely defined in one stratum.

Definition 7
Let $DB = IDB \cup EDB$ be a stratified U-Datalog$^{\neg}$ database. Let $(P_i)_{(1 \leq i \leq n)}$ be a stratification for DB. The bottom-up semantics of DB is defined as $\mathcal{FIX}^{\neg}_{DB} = M_n$ where the sequence $M_1, \ldots, M_n$ is computed as follows:
$M_1 = T^{\neg}_{P_1} \uparrow \omega \cup Comp(T^{\neg}_{P_1} \uparrow \omega)$
$M_{i+1} = T^{\neg}_{P_{i+1} \cup M_i} \uparrow \omega \cup Comp(T^{\neg}_{P_{i+1} \cup M_i} \uparrow \omega), 1 < i \leq n.$ □

Theorem 4
Let DB be a stratified U-Datalog$^{\neg}$ database. $\mathcal{FIX}^{\neg}_{DB}$ can be computed in a finite number of steps. □

Due to some basic results presented in (Chandra & Harel, 1985), the bottom-up semantics of a stratified U-Datalog$^{\neg}$ database is independent from the chosen stratification.

[11] $c_{|X}$ denotes the projection of constraint c onto the variables in X (thus, all the other variables are eliminated by applying a variable elimination algorithm (Chang & Keisler, 1973)).
[12] See note 9.
[13] Minimality is defined with respect to the order $\preceq$ defined as follows: $c \preceq c'$ if $\mathcal{H} \models c' \rightarrow c$

The answers to a given U-Datalog$^\neg$ goal are computed as described for U-Datalog programs in Subsection 2.2.2, by replacing $\mathcal{FIX}_{DB}$ with $\mathcal{FIX}^{\neg}_{\mathcal{DB}}$.

Example 4
Consider the extensional database EDB of Example 2 and the U-Datalog$^\neg$ program IDB obtained by adding the following rules to the ones presented in Example 1:

$r4 : change\_man(X) \leftarrow -emp\_man(X,Y), dep\_B(X), dep\_A(Y)$
$r5 : change\_man(X) \leftarrow X = Y, +emp\_man(X,Y), dep\_B(X), \neg ins\_man(X)$

An atom $change\_man(a)$ is now true if 'a' belongs to department B and if there exists at least one employee in department A. In this case, it removes all managers of 'a' belonging to department A. It is also true if 'a' belongs to department B and it has no manager. In this case, the evaluation removes all managers of 'a' belonging to department A and makes 'a' manager of itself.

A possible stratification for $DB = IDB \cup EDB$ is the following: $P_1 = EDB \cup \{r1, r2\}$, $P_2 = \{r3\}$, $P_3 = \{r4, r5\}$. $\mathcal{FIX}^{\neg}_{\mathcal{DB}}$ is computed as follows:

$$
\begin{aligned}
T^{\neg}_{P_1} \uparrow \omega = \{ & emp\_man(X,Y) \leftarrow X = b, Y = b;\\
& emp\_man(X,Y) \leftarrow X = b, Y = c;\\
& dep\_A(X) \leftarrow X = b;\ dep\_A(X) \leftarrow X = c;\ dep\_B(X) \leftarrow X = b;\\
& rem\_man(X,Y) \leftarrow X = b, Y = b, -dep\_A(Y);\\
& rem\_man(X,Y) \leftarrow X = b, Y = c, -dep\_A(Y); \}
\end{aligned}
$$

$$
\begin{aligned}
Comp(T^{\neg}_{P_1} \uparrow \omega) = \{ & \neg emp\_man(X,Y) \leftarrow X \neq b;\\
& \neg emp\_man(X,Y) \leftarrow Y \neq b, Y \neq c;\\
& \neg dep\_A(X) \leftarrow X \neq b, X \neq c;\ \neg dep\_B(X) \leftarrow X \neq b;\\
& \neg rem\_man(X,Y) \leftarrow X \neq b;\\
& \neg rem\_man(X,Y) \leftarrow Y \neq b, Y \neq c; \}
\end{aligned}
$$

$$M_1 = T^{\neg}_{P_1} \uparrow \omega \cup Comp(T^{\neg}_{P_1} \uparrow \omega)$$

$$
\begin{aligned}
& T^{\neg}_{P_2 \cup M_1} \uparrow \omega = \{ins\_man(X) \leftarrow X = b, Y = c, -dep\_A(Y), +dep\_A(X)\} \cup M_1\\
& Comp(T^{\neg}_{P_2 \cup M_1} \uparrow \omega) = \{\neg ins\_man(X) \leftarrow X \neq b\}\\
& M_2 = T^{\neg}_{P_2 \cup M_1} \uparrow \omega \cup Comp(T^{\neg}_{P_2 \cup M_1} \uparrow \omega)
\end{aligned}
$$

$$
\begin{aligned}
T^{\neg}_{P_3 \cup M_2} \uparrow \omega = \{ & change\_man(X) \leftarrow X = b, Y = c, -emp\_man(X,Y);\\
& change\_man(X) \leftarrow X = b, Y = b, -emp\_man(X,Y)\} \cup M_2\\
Comp(T^{\neg}_{P_3 \cup M_2} \uparrow \omega) & = \{\neg change\_man(X) \leftarrow X \neq b\}\\
\mathcal{FIX}^{\neg}_{P} = M_3 & = T^{\neg}_{P_3 \cup M_2} \uparrow \omega \cup Comp(T^{\neg}_{P_3 \cup M_2} \uparrow \omega).
\end{aligned}
$$

Note that rule r5 does not provide any additional answer for predicate $change\_man$. Indeed, $\neg ins\_man$ generates the constraint $X \neq b$ and dep_B generates the constraint $X = b$. Thus, the whole constraint is inconsistent. $\diamond$

3.2 Compositional semantics

The compositional semantics for U-Datalog programs (Section 2.2.1) was defined by using an unfolding operator which replaces the atom $p(\tilde{X})$ in the body of a rule with the body of a rule defining p. Problems arise when unfolding negated atoms. Suppose we want to unfold $\neg p(\tilde{X})$, then the disjunction of the bodies of *all* the rules that define predicate p in the compositional semantics has to be negated. Suppose the following rules represent the compositional semantics of a predicate p:

$$p(\tilde{X}) \leftarrow \tilde{b}_1, \tilde{u}_1, L_1,$$
$$\vdots$$
$$p(\tilde{X}) \leftarrow \tilde{b}_n, \tilde{u}_n, L_n.$$

Since $p(\tilde{X})$ is true (due to the CWA) if and only if $\tilde{b}_1, \tilde{u}_1, L_1 \vee \ldots \vee \tilde{b}_n, \tilde{u}_n, L_n$ is true, $\neg p(\tilde{X})$ has to be unfolded with the negation of $\tilde{b}_1, \tilde{u}_1, L_1 \vee \ldots \vee \tilde{b}_n, \tilde{u}_n, L_n$ (since $\neg p(\tilde{X}) \leftrightarrow \neg(b_1, u_1, L_1 \vee \ldots \vee b_n, u_n, L_n)$).[14]

A problem arises when there exists an infinite set of rules defining $p(\tilde{X})$ in the compositional semantics. In this case, the unfolding operator cannot be applied since it is not effective. In order to solve this problem, a weaker notion of compositionality can be introduced, based on the restriction of the set of extensional databases with respect to which the intensional database can be composed. The additional information available on the considered extensional database has to guarantee that the result of the unfolding operator, which unfolds all the positive literals in the rule bodies, is finite. Gabbrielli et al. in (Gabbrielli *et al.*, 1993) showed that, when the Herbrand Universe $\mathcal{H}$ is finite, it is possible to compute a T-stable semantics of a logic program IDB, which is finite and gives the same answer constraints of IDB when composed with any extensional database defined on $\mathcal{H}$. Intuitively, the T-stable semantics iterates the unfolding operator as many times as the new unfolded rules may give different results on the finite domain $\mathcal{H}$.[15]

Under these hypothesis, the compositional semantics for a stratified U-Datalog$^\neg$ program corresponds to an unfolding semantics computed in two steps:

1. In the first step, all positive literals in the rule bodies are unfolded, by computing the T-stable semantics according to the algorithm given in (Gabbrielli *et al.*, 1993) and the unfolding operator presented in Section 2.2.1. At this stage, negative literals are left unchanged.
2. In the second step, negative literals are unfolded. Due to the finite domain assumption and results presented in (Gabbrielli *et al.*, 1993), the set of rules required to unfold negative literals is finite.

The result is a recursion free program written in an *extended* U-Datalog$^\neg$ language which characterizes the semantics of the intensional database w.r.t. the extensional one.

[14] The way we unfold a negated atom $\neg p(X)$ corresponds to the syntactic transformation performed in the Clark's completion approach (Clark, 1987). However, while Clark's completion is used as a logical theory, our resulting unfolded program is evaluated by using a bottom-up stratified semantics. Therefore, we can prove that P and its unfolded version are equivalent w.r.t answer constraints (see Theorem 7). Moreover, since we deal with stratified programs, no L_i in the formula $f \equiv p(\tilde{X}) \leftrightarrow (b_1, u_1, L_1 \vee \ldots \vee b_n, u_n, L_n)$ can be equal to $\neg p(\tilde{X})$, since no cycle through negation arises in predicate definition. Thus, f is always *consistent* (Clark, 1987).

[15] Note that the finite domain assumption can be guaranteed only by executing updates which do not insert new values inside the database.

3.2.1 Unfolding of positive literals

In order to compute the compositional semantics of a program IDB, we first unfold positive literals by dealing with negative literals as if they were extensional predicates. This means that negative literal are not unfolded. To this purpose, the techniques presented in (Gabbrielli *et al.*, 1993) are applied to obtain a finite set of rules, denoted by $\mathcal{U}^{pos}_{IDB}$. The basic idea of the T-stable semantics is illustrated by the following example.

Example 5
Consider rules r1 and r2 presented in Example 1 and suppose that $\mathcal{H} = \{a, b\}$. After two iterations of the unfolding operator presented in Section 2.2.1, we obtain the following rules (call them T):

$$rem\_man(X,Y) \leftarrow -dep\_A(Y), emp\_man(X,Y)$$
$$rem\_man(X,Y) \leftarrow -dep\_A(Y), emp\_man(X,Z), emp\_man(Z,Y).$$

At the third iteration of the unfolding operator we also obtain the rule

$$\begin{aligned} rem\_man(X,Y) \leftarrow & -dep\_A(Y), emp\_man(X,Z), \\ & emp\_man(Z,W), emp\_man(W,Y) \end{aligned}$$

which cannot infer different results on any database defined on two elements, since $rem\_man(X,Y)$ computes the transitive closure of relation emp_man. Thus, T corresponds to the T-stable semantics of the previous rules. Now suppose that $\mathcal{H} = \{a, b, c\}$. The T-stable semantics $\mathcal{U}^{pos}_{IDB}$ is computed as follows (in the following, $\mathcal{U}^{pos}(P_j)$ denotes the set of rules contained in stratum j of $\mathcal{U}^{pos}_{IDB}$):

$$\begin{aligned} \mathcal{U}^{pos}_{IDB} = \mathcal{U}^{pos}(P_1) :& rem\_man(X,Y) \leftarrow -dep\_A(Y), emp\_man(X,Y) \\ & rem\_man(X,Y) \leftarrow -dep\_A(Y), emp\_man(X,Z), emp\_man(Z,Y) \\ & rem\_man(X,Y) \leftarrow -dep\_A(Y), emp\_man(X,Z), \\ & \qquad emp\_man(Z,W), emp\_man(W,Y) \\ \mathcal{U}^{pos}(P_2) :& ins\_man(X) \leftarrow +dep\_A(X), -dep\_A(Y), emp\_man(X,Y) \\ & ins\_man(X) \leftarrow +dep\_A(X), -dep\_A(Y), emp\_man(X,Z), \\ & \qquad emp\_man(Z,Y) \\ & ins\_man(X) \leftarrow +dep\_A(X), -dep\_A(Y), emp\_man(X,Z), \\ & \qquad emp\_man(Z,W), emp\_man(W,Y) \\ \mathcal{U}^{pos}(P_3) :& change\_man(X) \leftarrow -emp\_man(X,Y), dep\_B(X), dep\_A(Y) \\ & change\_man(X) \leftarrow X = Y, +emp\_man(X,Y), dep\_B(X), \\ & \qquad \neg ins\_man(X). \end{aligned}$$ $\diamond$

By results presented in (Gabbrielli *et al.*, 1993) and (Maher, 1993), we can state the following results.

Theorem 5
Let $\mathcal{H}$ be the fixed and finite Herbrand Universe. Let $\mathcal{U}^{pos}_{IDB}$ the T-stable semantics computed as described in (Gabbrielli *et al.*, 1993), depending on the cardinality of $\mathcal{H}$. For any extensional database EDB, $\mathcal{U}^{pos}_{IDB} \cup EDB$ is equivalent to $IDB \cup EDB$. Moreover, $\mathcal{U}^{pos}_{IDB}$ admits the same stratification of IDB and preserves goal admissibility. $\square$

3.2.2 Unfolding of negative literals

After constructing $\mathcal{U}^{pos}_{IDB}$, negative literals have to be unfolded. In order to unfold a negated literal $\neg p(\tilde{X})$, the disjunction of the bodies of *all* the rules defining predicate p in $\mathcal{U}^{pos}_{IDB}$ has to be negated. This approach should be applied stratum by stratum, generating in a finite number of steps a set of rules not containing negative literals. Note that, due to stratification conditions, the unfolding of $\neg p$ is required only in rules belonging to levels higher than the level where p is defined. The resulting set of rules corresponds to the compositional semantics of IDB. However, unfortunately, the negated disjunction of the bodies defining p is, in general, a first order formula, which cannot be represented in U-Datalog$^\neg$, as the following example shows.

Example 6
Consider the intensional predicate p defined by the rule $r : p(X) \leftarrow X = a, f(X,Y), q(X,Y)$, where f and q are extensional predicates. The previous rule is logically equivalent to the following first order formula: $p(X) \leftarrow X = a \wedge \exists Y \ (f(X,Y) \wedge q(X,Y))$. By assuming that r is the only rule defining p, by CWA, we obtain that $\neg p(X) \leftrightarrow \neg(X = a \wedge \exists Y \ (f(X,Y) \wedge q(X,Y)))$. But $\neg(X = a \wedge \exists Y \ (f(X,Y) \wedge q(X,Y)))$ is logically equivalent to $(X \neq a) \vee (\forall Y \ (\neg f(X,Y) \vee \neg q(X,Y)))$, which can always be transformed in prenex disjunctive normal form (Maher, 1988) obtaining $\forall Y (X \neq a \vee \neg f(X,Y) \vee \neg q(X,Y))$. $\diamond$

From the previous example it follows that, in order to unfold negative literals, the U-Datalog$^\neg$ syntax has to be extended to deal with first order formulas. As shown in the example, the variables which become quantified after negation, correspond to *local variables* of the original rule, i.e., body variables not appearing in the rule head. After this extension, the syntax of a U-Datalog$^\neg$ rule, hereafter called *extended U-Datalog*$^\neg$ rule, becomes the following:

$$H \leftarrow \tilde{b}, \tilde{u}, \tilde{L} \diamond \tilde{Q}(b_1, \tilde{H}_1 \vee ... \vee b_n, \tilde{H}_n)$$

where H is an atom, $\tilde{u}$ is an update constraint, $\tilde{b}$ is a conjunction of equality constraints, $\tilde{L}$ is a conjunction of positive literals, $\tilde{Q}(b_1, \tilde{H}_1 \vee ... \vee b_n, \tilde{H}_n)$ is a first order formula in prenex disjunctive normal form, where $\tilde{Q}$ is a sequence of quantified variables, not appearing in H or in $\tilde{L}$, each b_i is a conjunction of equality and inequality constraints, each $\tilde{H}_i$ is a conjunction of literals.[16] Intuitively, $\tilde{b}, \tilde{u}, \tilde{L}$ is generated by the unfolding of positive literals whereas the first order formula $\tilde{Q}(b_1, \tilde{H}_1 \vee ... \vee b_n, \tilde{H}_n)$ is generated by the unfolding of negative literals. Of course, we still assume that rules are stratified.

An extended U-Datalog rule body is true in a given interpretation if there exist some bindings for the positive literals and some bindings for the free variables of the quantified formula which make the rule body true in the given interpretation. Formally, the truth of an extended U-Datalog rule body can be defined as follows.

Definition 8

[16] Note that the proposed extensions are performed at the body rule level. No change to the Herbrand Base is performed.

Let $R \equiv \tilde{b}, \tilde{u}, L_1(Y_1), \ldots, L_m(Y_m) \diamond F$ where $F \equiv Q(b_1, \tilde{H}_1 \vee \ldots \vee b_n, \tilde{H}_n)$. Let $X_1.,,,.X_n$ be the free variables of F. Let $I \subseteq \mathcal{B}^{\neg}$. R is true in I with answer constraint $\overline{b}, \overline{u}$ if there exist $L_i(\tilde{U}_i) \leftarrow b_i, u_i \in I$ $(i = 1, \ldots, m)$, and $c \equiv (X_1 = t_1 \wedge \ldots \wedge X_n = t_n)$ $(t_j \in \mathcal{H},\ j = 1, \ldots, n)$, such that $\overline{b} \equiv \bigwedge_i b_i \wedge \tilde{b} \wedge c\ \bigwedge_i(\tilde{Y}_i = \tilde{U}_i)$, $\overline{u} \equiv \tilde{u}\ \wedge (u_1 \wedge \ldots \wedge u_m)$, $\overline{b} \wedge \overline{u}$ is $\mathcal{H}$-solvable, and $I \models \overline{b} \wedge F$.[17] □

The bottom-up operator of an extended U-Datalog$^{\neg}$ program can be now defined as follows (in the following, *body*(r) denotes the body of a rule r).

Definition 9
Let $DB = IDB \cup EDB$ be an extended U-Datalog$^{\neg}$ database. The bottom-up operator $T^e_{DB} : 2^{\mathcal{B}^{\neg}} \rightarrow 2^{\mathcal{B}^{\neg}}$ is defined as follows:

$$T^e_{DB}(I) = \{p(\tilde{X}) \leftarrow \overline{b}, \overline{u} | \exists \text{ a renamed rule}$$
$$r : p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, \tilde{L} \diamond \tilde{Q}(b_1, \tilde{H}_1 \vee \ldots \vee b_n, \tilde{H}_n) \in DB$$
$$\overline{b}, \overline{u} \text{ is an answer constraint for } body(r) \text{ in } I\}$$
□

Theorem 6
Let $\mathcal{H}$ be a finite domain. T^e_{DB} is a continuous operator. □

Given an extensional database EDB and an extended U-Datalog$^{\neg}$ program IDB, the semantics of $DB = IDB \cup EDB$ is obtained as the least fixpoint of T^e_{DB}, denoted by $\mathcal{FIX}^e_{DB}$. Due to the presence of a first order formula in rule bodies, a new safeness through query invocation property has to be stated.

Definition 10
Let P be an extended Datalog$^{\neg}$ program. P is *safe through query invocation* if each non-quantified variable, appearing in a rule head, in an update atom, or in a negated atom, also appears in a positive literal of the rule body or is bound by a constant present in the goal. □

The unfolding operator we are going to define works stratum by stratum. First, positive literals in the rule bodies of each stratum are unfolded by using operator $\mathcal{U}^{pos}_{IDB}$. Then, the negative literals contained in the i-th stratum of $\mathcal{U}^{pos}_{IDB}$ are unfolded by using an operator $\mathcal{U}^{neg}_{IDB}$ and the rules resulting from the completed unfolding of strata $1, \ldots, i-1$. The result is an extended U-Datalog$^{\neg}$ program equivalent to IDB but not containing positive or negative intensional literals.

Before presenting the unfolding operator, we define the operator $\mathcal{U}^{neg}_{IDB}$ which unfolds the negative literals of a U-Datalog$^{\neg}$ program IDB, by using the rules of an extended U-Datalog$^{\neg}$ program U. In order to define function $\mathcal{U}^{neg}_{IDB}$, we need an operator, Neg_c, which takes the disjunction of a set of (extended) U-Datalog$^{\neg}$ rule bodies defining a predicate p, performs its logical negation and returns the resulting first order formula in prenex disjunctive normal form. Such formula is then used to construct the unfolded rule. Information on the head variables of each rule defining p is useful to understand which are the local variables, that is, the variables which

[17] $I \models \overline{b} \wedge F$ means that for any assignment of values to quantified variables, satisfying $\overline{b}$, it is possible to find some constrained literals in I unifying with those in F, such that the resulting constraint is $\mathcal{H}$-solvable.

have to be quantified. In performing negation, also constraints which make the update atoms satisfiable have to be considered, similarly to what has been done for operator $Comp$ (see Definition 6).

Definition 11
Let $P^e = \{IDB | IDB$ is an extended U-Datalog$^{\neg}$ program$\}$. Let IDB be a stratified U-Datalog$^{\neg}$ intensional database. The unfolding operator $\mathcal{U}^{neg}_{IDB} : P^e \rightarrow P^e$ is defined as follows:

$$
\begin{array}{l}
\mathcal{U}^{neg}_{IDB}(U) = \{p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, A_1, \ldots, A_n, \diamond F | \exists \text{ a renamed rule} \\
\quad p(\tilde{X}) \leftarrow \tilde{b}, \tilde{u}, A_1, \ldots, A_n, \neg p_1(\tilde{Z}_1), \ldots, \neg p_m(\tilde{Z}_m) \in IDB \\
\quad \text{for all } p_i \ (i = 1, ..., m), \text{ consider the body of all the rules } r^i_j \ (j = 1, ..., l_i) \\
\quad\quad \text{defining } p_i \text{ in } U \\
\quad r^i_1 : p_i(\tilde{V}_{i,1}) \leftarrow \tilde{b}_{i,1}, \tilde{u}_{i,1}, \tilde{L}_{i,1} \diamond F_{i,1} \in U \\
\quad \vdots \\
\quad r^i_{l_i} : p_i(\tilde{V}_{i,l_i}) \leftarrow \tilde{b}_{i,l_i}, \tilde{u}_{i,l_i}, \tilde{L}_{i,l_i} \diamond F_{i,l_i} \in U \\
\quad F^{18} \equiv \tilde{Q}(c_1, \tilde{L_1} \vee \ldots \vee c_e, \tilde{L_e}) \in \\
\quad\quad\quad Neg_c((\tilde{Z}_1 = \tilde{V}_{1,1} \wedge body(r^1_1) \vee ... \vee \tilde{Z}_1 = \tilde{V}_{1,l_1} \wedge body(r^1_{l_1})) \wedge ... \wedge \\
\quad\quad\quad\quad (\tilde{Z}_m = \tilde{V}_{m,1} \wedge body(r^m_1) \vee ... \vee \tilde{Z}_m = \tilde{V}_{m,l_m} \wedge body(r^m_{l_m}))) \\
\quad \tilde{b} \wedge \tilde{u} \wedge \tilde{Q}(c_1 \vee \ldots \vee c_e) \text{ is } \mathcal{H}\text{-solvable } \}
\end{array}
$$
□

By using operator $\mathcal{U}^{neg}_{IDB}$, the compositional semantics is defined as follows.

Definition 12
Let IDB be a stratified U-Datalog$^{\neg}$ program. Suppose that IDB, and therefore $\mathcal{U}^{pos}_{IDB}$, admits a stratification with k strata $P_1, ..., P_k$. Let $\mathcal{U}^{pos}(P_j)$ be the set of rules contained in stratum j of $\mathcal{U}^{pos}_{IDB}$. The compositional semantics $\mathcal{U}_{IDB}$ is defined as follows (see Subsection 2.2.1 for the definition of ID_{EDB}):

$$
\begin{array}{l}
\overline{P}_1 = \mathcal{U}^{neg}_{\mathcal{U}^{pos}(P_1)}(ID_{EDB}) \\
\overline{P}_2 = \mathcal{U}^{neg}_{\mathcal{U}^{pos}(P_2)}(ID_{EDB} \cup \overline{P}_1) \\
\quad \vdots \\
\overline{P}_i = U^{neg}_{\mathcal{U}^{pos}(P_i)}(ID_{EDB} \cup \overline{P}_{i-1}) \\
\mathcal{U}_{IDB} = \bigcup_{1 \leq i \leq k} \overline{P}_i.
\end{array}
$$
□

Theorem 7
$\mathcal{U}_{IDB}$ is safe through query invocation and for each admissible goal G and for each extensional database EDB, the answer constraints for G in $\mathcal{FIX}^{\neg}_{IDB} \cup EDB$ are the same than the ones in $\mathcal{FIX}^{e}_{\mathcal{U}_{IDB}} \cup EDB$. □

Example 7
Consider the program IDB and its stratification, as presented in Example 4, and its positive unfolding $\mathcal{U}^{pos}_{IDB}$, as presented in Example 5. The compositional semantics of IDB is constructed as follows:

$$
\begin{array}{l}
\overline{P}_1 = P_1 \qquad \overline{P}_2 = P_2 \\
\overline{P}_3 = change\_man(X) \leftarrow -emp\_man(X, Y), dep\_B(X), dep\_A(Y)
\end{array}
$$

[18] Note that no updates are generated.

$$change\_man(X) \leftarrow X = Y, +emp\_man(X,Y), dep\_B(X) \diamond$$
$$\forall Z(X = Z \vee \neg emp\_man(X,Z))$$
$$\mathcal{U}_{IDB} = \overline{P}_1 \cup \overline{P}_2 \cup \overline{P}_3.$$

The second rule for predicate *change_man* derives from the unfolding of predicate $\neg ins\_man(X)$. The formula $\forall Z(X = Z \vee \neg emp\_man(X,Z))$ is the simplified result of the application of the Neg_c operator to the disjunction of the bodies of the rules defining *ins_man*. $\diamond$

It is important to remark that the compositional semantics has not to be considered as an alternative semantics w.r.t. the marking phase semantics. Indeed, the computation of the compositional semantics can be quite expensive. However, since the compositional semantics is recursion free and has to be computed just once (unless the Herbrand domain changes) it can be meaningfully used in some cases as a precompilation technique.

4 Concluding remarks

In this paper we have introduced negation inside U-Datalog rules and proposed a stratification-based approach to assign a semantics to such programs. We have also introduced a weaker concept of compositionality and presented a finite and effectively computable compositional semantics for U-Datalog$^{\neg}$ programs. By results presented in (Bertino & Catania, 2000), it is quite immediate to prove that, with respect to the returned answers, U-Datalog$^{\neg}$ is equivalent to Stratified Datalog$^{\neg}$, open with respect to a subset of extensional predicates (Bertino *et al.*, 1999). The presented results can be extended to deal with other Datalog-like update languages safe through query invocation. Future work includes the introduction of negation in other U-Datalog extensions (Bertino *et al.*, 2000; Bertino *et al.*, 1998a) and the definition of static analysis techniques for U-Datalog$^{\neg}$, similarly to those proposed for U-Datalog (Bertino & Catania, 1996).

References

Abiteboul, S., & Vianu, V. (1991). Datalog Extensions for Database Queries and Updates. *Journal of Computer and System Sciences*, **43**(1), 62–124.

Bertino, E., & Catania, B. (1996). Static Analysis of Intensional Databases in U-Datalog. *Pages 201–212 of: Proc. of the ACM SIGACT-SIGMOD-SIGART Int. Symp. on Principles of Database Systems.*

Bertino, E., & Catania, B. (2000). *Analysis of the U-Datalog Expressive Power.* In preparation.

Bertino, E., Catania, B., Gervasi, V., & Raffaetà, A. (1998a). Active-U-Datalog: Integrating Active Rules in a Logical Update Language. *Pages 107–133 of:* Decker, H., Freitag, B., Kifer, M., & Voronkov, A. (eds), *LNCS 1472: Transactions and Change in Logic Databases.* Springer Verlag.

Bertino, E., Martelli, M., & Montesi, D. (1998b). Transactions and Updates in Deductive Databases. *IEEE Transactions on Knowledge and Data Engineering*, **9**(5), 784–797.

Bertino, E., Catania, B., & Gori, R. (1999). *Enhancing the Expressive Power of the U-Datalog Language (Extended version).* Tech. rept. University of Milano, Italy.

Bertino, E., Catania, B., Gervasi, V., & Raffaetà, A. (2000). A Logical Approach to Cooperative Information Systems. *Journal of Logic Programming*, **43**(1), 15–48.

Bonner, A.J., & M.Kifer. (1994). An Overview of Transaction Logic. *Theoretical Computer Science*, **133**(2), 205–265.

Bossi, A., Gabbrielli, M., Levi, G., & Meo, M.C. (1994). A Compositional Semantics for Logic Programs. *Theoretical Computer Science*, **122**(1-2), 3–47.

Ceri, S., Gottlob, G., & Tanca, L. (1990). *Logic Programming and Databases*. Spinger Verlag.

Chandra, A.K., & Harel, D. (1985). Horn Clause Queries and Generalizations. *Journal of Logic Programming*, **2**(1), 1–15.

Chang, C.C., & Keisler, H.J. (1973). *Model Theory*. North-Holland.

Chen, W. (1995). Declarative Updates of Relational Databases. *ACM Transactions on Database Systems*, **20**(1), 42–70.

Chen, W. (1997). Programming with Logical Queries, Bulk Updates, and Hypothetical Reasoning. *IEEE Transactions on Knowledge and Data Engineering*, **9**(4), 587–599.

Clark, K.L. (1987). Negation as Failure. *Pages 293–322 of:* Gallaire, H., & Minker, J. (eds), *Logic and Data Bases*. Plenum Press, New York.

Gabbrielli, M., Giacobazzi, R., & Montesi, D. (1993). Modular Logic Programs over Finite Domains. *Pages 663–678 of:* Saccà, D. (ed), *Proc. Eight Italian Conference on Logic Programming*.

Jaffar, Joxan, Maher, Michael J., Marriott, Kim, & Stuckey, Peter J. (1998). The Semantics of Constraint Logic Programs. *Journal of Logic Programming*, **37**(1-3), 1–46.

Lausen, G., Ludascher, B., & May, W. (1998). On Active Deductive Databases: The Statelog Approach. *Pages 69–106 of:* Decker, H., Freitag, B., Kifer, M., & Voronkov, A. (eds), *LNCS 1472: Transactions and Change in Logic Databases*. Springer Verlag.

Lloyd, J.W. (1987). *Foundations of Logic Programming*. Springer Verlag.

Maher, M.J. (1988). Complete Axiomatizations of the Algebra of Finite, Rational and Infinite trees. *Pages 348–357 of: Proc. of the Third Symp. on Logic in Computer Science*.

Maher, M.J. (1993). A Trasformation System for Deductive Databases with Perfect Model Semantics. *Theoretical Computer Science*, **110**(2), 377–402.

Manchanda, S., & Warren, D.S. (1988). A Logic-Based Language for Database Updates. *Pages 363–394 of:* Minker, J. (ed), *Foundations of Deductive Databases and Logic Programming*. Plenum Press, New York.

Naqvi, S., & Tsur, S. (1989). *A Logical Language for Data and Knowledge Bases*. Computer Science Press.

Wichert, C.-A., & Freitag, B. (1997). Capturing Database Dynamics by Deferred Updates. *Pages 226–240 of: Proc. of the Int. Conf. on Logic Programming*.

Wichert, C.-A., Freitag, B., & Fent, A. (1998). Logical Transactions and Serializability. *Pages 134–165 of:* Decker, H., Freitag, B., Kifer, M., & Voronkov, A. (eds), *LNCS 1472: Transactions and Change in Logic Databases*. Springer Verlag.