



**DECLARATIVE RECONSTRUCTION
OF AN OBJECT ORIENTED
DATA MODEL WITH ROLES**

Internal Report C95-45

21 Dicembre 1995

**M. Carboni
F. Giannotti
G. Manco
D. Pedreschi**

Declarative Reconstruction of an Object-Oriented Data Model with Roles (Extended Abstract) *

M. Carboni¹, F. Giannotti¹, G. Manco¹ and D. Pedreschi²

¹ CNUCE Institute of CNR
Via S. Maria 36, 56125 Pisa, Italy
e-mail: F.Giannotti@cnuce.cnr.it

² Dipartimento di Informatica, Univ. Pisa
Corso Italia 40, 56125 Pisa, Italy
e-mail: pedre@di.unipi.it

Abstract

This paper is aimed at showing how a simple non-deterministic extension of Datalog is sufficient to declaratively reconstruct the essence of object-oriented data models with roles. The target language is a subset of the logical data language $\mathcal{LDL}++$ [AOTZ93]. The reconstruction is performed by means of a compilation into the extended Datalog language of the basic features of the object model, including the schema definition language, the query language and the basic update operations. The purpose of this compilation is twofold. On one side, it provides a logical semantics of the object model, as the semantics of the target language is assigned in purely logical terms [SZ90, ZAO93]. On the other side, the proposed compilation forms the basis of a realistic implementation, as $\mathcal{LDL}++$ is efficiently executed by means of a fixpoint procedure with "in situ" updates.

Keywords. Logic Database Languages, Object-Oriented Data Models, Updates, Stratified Logic Programs, Non-determinism.

1 Introduction

This paper tries to answer the following question: is it possible to declaratively reconstruct the *essence* of object-oriented data models in simple declarative terms, without resorting to non-standard logical frameworks? If we agree that a natural extension of Datalog with a single unary function and a non-deterministic construct is indeed a simple language, then the answer to the above question is affirmative. In this paper, we consider:

- a small language embodying the essential static and dynamic aspects of a deductive object data model with roles, and
- the mentioned extension of Datalog, which corresponds to a fragment of the Logical Data Language $\mathcal{LDL}++$,

*Work partially supported by the EC-US Cooperative Activity Project ECUS-033 - DEUS EX MACHINA.

and provide a compilation of the former into the latter, with a twofold aim. On one side, such a compilation provides directly a logical semantics of the object model, as the semantics of the target language is assigned in purely logical terms [SZ90, ZAO93]. On the other side, the proposed compilation forms the basis of a realistic implementation, as *LDL++* can be efficiently executed by means of a fixpoint procedure with real side-effects to support updates.

The static component of the considered object model, i.e., its data definition and query language, can be seen as a fragment of *F-logic* [BK93], and provides mechanisms for multiple inheritance, multiple roles, virtual objects, views, and methods (derived attributes). The dynamic component of the considered object model, i.e. the basic update and object-migration operations, is inspired by the analogous features of the Fibonacci language [ABGO93]. The object model is compiled into an extension of deductive databases with

- non determinism, used to realize object identifiers, in a way similar to [Zan89], and
- a form of stratification, used to realize updates and object-migration.

The code obtained as a result of the compilation can be understood in model-theoretic terms, thus providing a formal interpretation of the source object-oriented program. Perhaps more importantly, such a code can be executed by the ordinary deductive, fixpoint-based computation integrated with efficient support for updates by means of side effects. A major advantage of this opportunity is the fact that the architecture of the abstract machine supporting deductive databases is left unaltered and, as a consequence, the available optimization techniques, such as magic sets, are directly applicable.

Many proposals in literature enhance the standard semantics of deductive databases, in order to support the object-oriented paradigm, but is missing a reference model which supports static and dynamic aspects in a uniform and simple framework. There are proposals which tackle only static (structural) aspects of an object-oriented data model. This is the case, for example, of the LOGIN proposal [AKN86], in which a new unification algorithm is defined to deal with structural inheritance, or of its extension F-logic [KLW93], in which a model-theoretic semantics is given in terms of F-structures, and a notion of object-identifier is modeled by way of existential quantification. As far as the dynamic aspects is concerned, there are proposals which provide a procedural semantics, as in the case of [Mon93, BGM95], and other proposals provide ad-hoc model theoretic semantics related to modal logic [BK93, MW86, Man89, Che91]. Both the approaches alterate the original standard semantics of deductive databases. Moreover, it is difficult to semantically combine the approaches to structural and dynamic aspects of the object data model in deductive databases. For instance, the F-logic combined with transaction logic leads to a very complicated semantic model, due to the necessity of combining F-structures with path-structures. Yet, it is not clear how to enhance such approaches to handle more complex dynamic features, such as role-dynamics and object-migration, which have been shown to be very important issues in the object-oriented paradigm [ABGO93, RS90, Su91, Wals95].

In the complete version of this paper, the compilation of all aspects of the object model is presented, along with the results concerning the correctness and efficiency of the compilation.

2 XY-Stratification and Nondeterminism

The basic tool used in this paper is a fragment of the *LDL++* language, consisting of *Datalog* augmented with a non-deterministic *choice* construct [GPSZ91] and a form of stratification called XY-stratification [AOTZ93].

Choice goals are used to non-deterministically select subsets of answers to queries which obey a specified FD constraint. For instance, the rule

$$st_ad(St, Ad) \leftarrow major(St, Area), faculty(Ad, Area), choice((St), (Ad)).$$

assigns to each student a unique advisor from the same area, as the choice goal restricts the *st_ad* relation to obey the FD ($St \rightarrow Ad$). Choice programs have a natural declarative reading, which can be formalized using stable-model semantics [SZ90], and can be efficiently executed by enforcing the FD constraints during the bottom-up computation. [GPSZ91]

XY-stratification has been introduced in [AOTZ93] as a syntactically decidable class of programs extending stratified programs. Recursive predicates use a special argument, called *stage* which is a natural number used to count the stages of the computation, and to enforce local stratification of recursive programs. Two kinds of rules are allowed:

- X-rules, where all stage arguments are the same variable, say J , and
- Y-rules, where a stage $s(J)$ is used in the head goals, and either $s(J)$ or J are used in body goals.

Recursion is allowed only in Y-rules, on decreasing stages.

As an example, consider the problem of computing the maximum of an array $a[0, n - 1]$ of integers, represented as a relation

$$a(0, x_0), \dots, a(n - 1, x_{n-1})$$

An XY-stratified program which solves the problem is the following:

$$\begin{aligned} \text{max}(0, X) &\leftarrow a(0, X). & (m_1) \\ \text{max}(s(J), Y) &\leftarrow \text{max}(J, X), a(s(J), Y), Y > X. & (m_2) \\ \text{del}(s(J), X) &\leftarrow \text{max}(J, X), a(s(J), Y), Y > X. & (m_3) \\ \text{max}(s(J), X) &\leftarrow \text{max}(J, X), \neg \text{del}(s(J), X). & (m_4) \end{aligned}$$

The intuitive meaning of this program is quite obvious: at each stage J , $\text{max}(J, X)$ holds if X is the maximum of the array $[0, J]$. Notice the use of the *del* relation to specify when the current maximum value changes from a stage to the next one. This interpretation is formalized by the notion of perfect model, characterized by an intended fixpoint procedure which computes one stratum at a time starting from stage 0.

Although useful to understand programs, this abstract procedural view is hopelessly inefficient, due to the enormous amount of copying needed at each stage. However, a smarter procedural interpretation allows us to avoid copying, based on the observation that new values can be produced starting from values in the current stage only, and therefore the old stages may be discarded. According to this interpretation, the stage argument is represented by a single variable which is incremented at each execution of the Y-values. The only real computation needed, when the body of rules (m_2) and (m_3) is true, corresponds to inserting a new fact $\text{max}(Y)$ and deleting the new fact $\text{max}(X)$, in order to obtain the correct extension of the relation *max* at the current stage. In other words, the effect of Y-rules can be implemented in terms of real side-effects, those making it viable to use XY-stratified programs for database programming. In our example, the execution of the proposed program boils down to the efficient iterative computation of the following Pascal-like program:

```
max := a[0];
for j := 1 to n - 1 do
  if max < a[j] then max := a[j]
```

For this reasons, XY-stratified programs play a central role in the reconstruction of the dynamics aspects of the object model, as discussed later.

<i>atom</i>	::=	<i>pred_name(term) oid : class_name[attrib_val]</i>
<i>clause</i>	::=	<i>atom ← body. class_name isa class_name. class_name[attrib_def].</i>
<i>body</i>	::=	<i>true atom body ∧ body</i>
<i>attrib_val</i>	::=	<i>attrib_name → term mthd_name@term attrib_val; attrib_val</i>
<i>attrib_def</i>	::=	<i>attrib_name ⇒ attrib_type attrib_def; attrib_def</i>
<i>attrib_type</i>	::=	<i>string int class_name</i>

Figure 1: Abstract syntax for the object-oriented deductive model.

3 A Deductive Object-Oriented Data Model

We consider here a paradigmatic object-oriented data model, with mechanisms for supporting a kernel of relevant aspects of (deductive) object-oriented databases, which are listed below.

- Object identity: objects in any class are identified by a unique *oid*.
- Multiple roles: an object belonging to a class can be viewed as playing a role in that class. The set of classes in which an object lives, at a given stage, specifies the set of roles of that object. Objects can dynamically change their roles.
- Multiple structural inheritance: the structure of objects propagates down from classes to subclasses; this form of inheritance is monotonic in the sense that any additional attribute or method specified for a subclass is added to the structure inherited from the superclass(es).
- Virtual objects and views: new classes can be derived from others by means of (deductive) rules, by specifying their extension on the basis of pre-existing classes.
- Methods: objects may have derived attributes, whose values are computed on the basis of pre-existing objects.
- Basic update operations: primitives for object creation, deletion and role migration are provided, whose effects propagates to the whole schema by means of inheritance and derivation rules.

For limitation of space, we choose *not* to discuss some other relevant aspects, listed below.

- Overriding inheritance and late binding: as a consequence, we assume that attributes and methods are not redefined in subclasses. However, these features can be appropriately accommodated in our approach with minor modifications.
- Structured attribute types: sequence and tuple types are not considered here, although they can be easily modeled using complex objects.
- Static type checking: this is actually an orthogonal issue to our approach, and is outside the scope of this paper.
- Transactions and active rules: this is material of future research which we plan to pursue by extending the techniques in [CFG95, Zan95].

$person[name \Rightarrow string].$	$department[ungr \Rightarrow employee; assistant \Rightarrow person].$
$student[affiliation \Rightarrow school].$	$employee[salary \Rightarrow integer; affiliation \Rightarrow department].$
$researcher[aPaper \Rightarrow article].$	$school[about \Rightarrow string; aStudent \Rightarrow student].$
$employee \text{ isa } person.$	$school \text{ isa } department.$
$researcher \text{ isa } employee.$	$student \text{ isa } person.$
$O : employee[affiliation \rightarrow D; name \rightarrow N] \leftarrow$	$D : department[assistant \rightarrow O] \wedge O : person[name \rightarrow N]. \quad (r_1)$
$O : employee[boss@M] \leftarrow$	$O : employee[affiliation \rightarrow D] \wedge D : department[ungr \rightarrow M]. \quad (r_2)$

Figure 2: An example schema.

The syntax of the reference object-oriented (deductive) database language is shown in Fig. 1. The language is closely related to many proposals such as [KLV93, Abi92, AKN86, BJ94]. Figure 2 shows a schema definition in this language. A schema is composed by definitions of classes and (deductive) rules.

A class can be defined in three different ways:

- by explicitly defining the set of its attributes such as, for instance, in the definition of the class *department*;
- by specialization from another class as, for instance, in the definition of the class *school* as a subclass of *department*;
- by derivation from other classes by means of conditions as, for instance, in the definition of rule (r_1), which can be read as follows: all persons which are assistants of any department can be viewed also as employees.

A few remarks about deductive rules are in order. First, objects are referred to by means of atoms of the form $oid : class[attrib_val]$ whose declarative reading is: the object identified by oid , viewed in the role $class$, has attribute values as specified in $attrib_val$. Second, rules are a flexible mechanism to define virtual objects and views. Rule (r_1), for instance, allows to extend the extension of class *employee* with new derived objects. Similarly, views and derived (sub)classes may be constructed. Third, incomplete information is allowed, in the sense that attribute values may be partially specified in rule heads, as for instance, in rule (r_1) where the *salary* attribute is missing. In our approach, missing attribute values are filled in with *null* values. Finally, creation of virtual objects is subject to the constraint of the *oid* uniqueness, so that no new objects are created with *oid*'s already in use.

Rule (r_2) in Fig. 2 defines a method for the *employee* class. It specifies that *boss* is a derived attribute of any *employee* object O , which can be computed as the manager of the department where O works. To comply with structural inheritance, methods are propagated to subclasses.

4 Compiling the Schema

In this section we use XY-stratified, non-deterministic Datalog programs to compile the reference object model, in a way inspired by [Zan89]. Informally, classes are represented by predicates which specify the class attributes augmented with two extra arguments denoting the stage and the *oid*. Therefore, we

associate with every class p with n attributes, n ternary predicates $p_{class}(j, oid, f_i(x))$, $i = 1, \dots, n$, where j is the stage, oid is the object identifier, f_i is the i -th attribute of the class p , and x is its value. Observe that the stage plays a role in modeling the dynamic aspects of the object model. In the compilation of the static part of the model, only X -rules are used, and therefore the stage is only used for deduction within the current state of the database.

The basis of our approach is to represent objects as instances of the most specialized class they dynamically belong to. In other words, each object is completely specified by its *most specialized version* (*msv* in short) which contains all attributes currently (at each stage) available for the object. To this purpose, we introduce a relation

$$msv(j, oid, q(x))$$

which denotes that, at stage j , the tuple x in class q is the most specialized version of object oid .

In our approach we require that the *msv* of each object is unique, albeit possibly different at different stages. Such a property is achieved, in our model, by requiring that the *ISA* hierarchy is closed under intersection, i.e., for any two classes r and q which are not *ISA*-related, the intersection class $r \cap q$ is present. Clearly, $r \cap q$ is a subclass of both r and q , and its attributes are the union of the attributes of r and q . In a real situation, we envisage that the system automatically completes the schema with intersection classes whenever these are not explicitly specified.

The role of *msv* is to activate the deduction process which populates the classes in the whole hierarchy. For each predicate q_{class} and for each attribute f of q the following clause is defined:

$$q_{class}(J, Oid, f(X)) \leftarrow msv(J, Oid, q(Y)). \quad \{\text{MSV Rule}\}$$

where Y is the tuple of variables corresponding to the attributes of the object Oid in class q , X is the variable corresponding to the attribute f in the tuple Y .

EXAMPLE 1 The code corresponding to the definition of the class *employee* in Fig. 2 is the following.

$$employee_{class}(J, Oid, affiliation(A)) \leftarrow msv(J, Oid, employee(N, A, S)).$$

$$employee_{class}(J, Oid, salary(S)) \leftarrow msv(J, Oid, employee(N, A, S)).$$

$$employee_{class}(J, Oid, name(N)) \leftarrow msv(J, Oid, employee(N, A, S)).$$

□

An *ISA* relation between two classes, p *isa* q , is modeled in the following way. For each attribute f in q the following rule is generated:

$$p_{class}(J, Oid, f(X)) \leftarrow q_{class}(J, Oid, f(X)). \quad \{\text{ISA Rule}\}$$

which naturally states that each object of the subclass is also an object of the superclass.

EXAMPLE 2 The code corresponding to the definition of *employee* as a subclass of *person* in Fig. 2 is the following.

$$person_{class}(J, Oid, name(N)) \leftarrow employee_{class}(J, Oid, employee(N, A, S)).$$

□

By the *MSV*-rule, the *msv* of each object is inserted in the appropriate class; then, by the *ISA*-rules each object is propagated up in the hierarchy to the superclasses. It is worth noting how the *ISA*-rules defining the hierarchy from one side provide the declarative specification of subclasses, and from the operational point of view they allow to populate the database—a dual reading which is typical of logic programs.

The above rules can be systematically generated, by gathering from the schema the following static information:

- for each definition $p[f_1 \Rightarrow t_1; \dots; f_n \Rightarrow t_n]$ we associate to the label p the set of attribute-type pairs $\{\langle f_1, t_1 \rangle, \dots, \langle f_n, t_n \rangle\}$;
- for each definition $p \text{ isa } q$, we extend the set of pairs associated to p with those associated to q , and repeat transitively this operation w.r.t. the ISA hierarchy;
- for any pair of not ISA-related classes p and q the class $p \cap q$ is defined, whose attributes are the (disjoint) union of the attributes of p and q .

We now provide a compilation for deductive rules. We consider rules where all atoms are either of the form $oid : p[f \rightarrow t]$ or of the form $oid : p[f@t]$. Clearly, an atom $oid : p[f_1 \rightarrow t_1; \dots; f_n \rightarrow t_n]$ in the body of a rule is equivalent to a conjunction of atoms $oid : p[f_1 \rightarrow t_1] \wedge \dots \wedge oid : p[f_n \rightarrow t_n]$, as well as such a predicate in the head of a rule corresponds to a set of rules with identical body. The compilation function $\mathcal{T}[\cdot]$ is defined as follows.

$$\mathcal{T}[A_0 \leftarrow A_1 \wedge \dots \wedge A_n] = \{a_0 \leftarrow a, a_1, \dots, a_n \mid \langle a_0, a \rangle \in \mathcal{H}[A_0], \mathcal{B}[A_i] = a_i \ (i = 1, \dots, n)\}$$

where $\mathcal{H}[\cdot]$ and $\mathcal{B}[\cdot]$ are defined as follows:

- $\mathcal{H}[Oid : p[f \rightarrow t]] = \{\langle p_{class}(J, Oid, f(t)), \neg msv(J, Oid, q(-)) \rangle \mid q \text{ is a subclass of } p\}$
- $\mathcal{H}[Oid : p[m@t]] = \{\langle q_{class}(J, Oid, m(t)), q_{class}(J, Oid, -) \rangle \mid q \text{ is a subclass of } p\}$;
- $\mathcal{B}[Oid : p[f \rightarrow t]] = \mathcal{B}[Oid : p[f@t]] = p_{class}(J, Oid, f(t))$.

Intuitively, each method is propagated downwards to the subclasses, as specified by the \mathcal{H} function, which produces a conclusion for any subclass of the given one. Moreover, the \mathcal{H} function yields a check of the existence of the object Oid in class for which the method is defined—in a way which recalls the use of **self** in o-o languages. Conversely, the conclusions of derivation rules are only propagated upwards, by means of the *ISA*-rules. It is needed, however, to check that virtual objects are created only if not already existing: this is performed by the \mathcal{H} function by introducing a negated use of the *msv* relation.

EXAMPLE 3 The code for the rule (r_1) of Fig. 2 is the following.

$$\begin{aligned} employee_{class}(J, O, affiliation(D)) &\leftarrow \neg msv(J, Oid, employee(-)), \\ &\quad department_{class}(J, D, assistant(O)), \\ &\quad person_{class}(J, O, name(N)). \\ \\ employee_{class}(J, O, name(N)) &\leftarrow \neg msv(J, Oid, employee(-)), \\ &\quad department_{class}(J, D, assistant(O)), \\ &\quad person_{class}(J, O, name(N)). \\ \\ employee_{class}(J, O, affiliation(D)) &\leftarrow \neg msv(J, Oid, researcher(-)), \\ &\quad department_{class}(J, D, assistant(O)), \\ &\quad person_{class}(J, O, name(N)). \\ \\ employee_{class}(J, O, name(N)) &\leftarrow \neg msv(J, Oid, researcher(-)), \\ &\quad department_{class}(J, D, assistant(O)), \\ &\quad person_{class}(J, O, name(N)). \end{aligned}$$

□

EXAMPLE 4 The code for the rule (r_2) of Fig. 2 is the following.

$$\begin{aligned}
employee_{class}(J, O, boss(M)) &\leftarrow employee_{class}(J, O, -), \\
&employee_{class}(J, O, affiliation(D)), \\
&department_{class}(J, D, mngr(M)). \\
researcher_{class}(J, O, boss(M)) &\leftarrow researcher_{class}(J, O, -), \\
&employee_{class}(J, O, affiliation(D)), \\
&department_{class}(J, D, mngr(M)).
\end{aligned}$$

□

5 Compiling Updates and Role Migration

We are now ready to compile updates and role migration within our object-oriented model. Notice that the previous translation was independent from the stage argument of each class predicate. Now, in order to modify the current state of the database (i.e., to add, delete or modify the objects), we need to operate on the stage argument, in a way similar to that used in [CFG95, SZ95]. We consider five basic operations:

- *new*($q[a], oid$): creation of a new object in class q with attributes a .
- *extend*($oid, q[a]$): insertion of an existing object oid in class q with attributes a .
- *drop*($oid, q[a]$): deletion of an existing object oid from class q with attributes a .
- *delete*($oid, q[a]$): deletion of an existing object oid .
- *modify*($oid, q[b]$): modification of the attributes of an existing object oid in class q with attributes b .

The basic idea is to perform such operations on the *msv* of the referred object. This is done by using a sort of *frame axiom* in the style of [CFG95, SZ95]. Intuitively, the alive objects at the current stage are the newly created ones, and those of the previous stage which have not been explicitly deleted. The following rule states the above concept: relation msv_{del} records the objects which have been explicitly deleted, updated and/or migrated at the current stage.

$$msv(s(J), Oid, X) \leftarrow msv(J, Oid, X), \neg msv_{del}(s(J), Oid, X). \quad \{\text{COPY Rule}\}$$

Notice that, as already mentioned 2, such a clause does not need to be really implemented as a rewrite rule, but it is amenable of a more efficient implementation using side effects.

As a consequence of the assumption that each object has a unique *msv*, we can model updates by simply updating the *msv* relation, as the *MSV*-rules and the *ISA*-rules accomplish the task of reflecting the updates on the whole database. We now show how the *msv* relation is used to the purpose of modeling the above five update operations.

The code realizing the operation *new*($p[a], Oid$) is composed by the following rule and by the *COPY*-rule. Notice that, if a is a partial specification of the attributes of p then a' is the complete list of the attributes of p , where the attributes unspecified in a are assigned to *null*.

$$\begin{aligned}
msv(s(J), Oid, p(a')) &\leftarrow Q, ide(Oid), \neg msv(J, Oid, -), \\
&choice((Oid), (p(a))), choice((p(a)), (Oid)).
\end{aligned}
\quad \{\text{NEW Rule}\}$$

In the *NEW*-rule, the relation *ide* is the domain of object identifiers. According to this rule, a new *Oid* is chosen which is not already in use at stage J , and associated to tuple a in class p . The *choice* operator allows us to non deterministically select, for each tuple, exactly one of the unused identifiers. As usual, the copy rule passes to the next stage all *msv*'s of objects which have not been dropped. Observe that the *NEW*-rule is applicable also when the database is empty, i.e., when no fact $msv(0, -, -)$ is present—thus reflecting the fact that the database can be loaded by a sequence of *new* operations.

EXAMPLE 5 The operation $new(employee[name \rightarrow smith], Oid)$ generates the following NEW-rule:

$$msv(s(J), Oid, employee(smith, null, null)) \leftarrow ide(Oid), \neg msv(J, Oid, _), \\ choice((employee(smith, null, null)), (Oid)), choice((Oid), (student(smith, null, null))).$$

This rule, together with the *MSV*-rules of example 1 has the effect of inserting the new object in class *employee*. Finally, the *ISA*-rules of example 2 insert the same object in class *person*. \square

The code realizing the operation $delete(oid)$ is the following rule plus the COPY-rule:

$$msv_{del}(s(J), oid, X) \leftarrow msv(J, oid, X) \quad \{\text{DELETE Rule}\}$$

The code realizing the operation $extend(oid, q[a])$ is composed by the following three sets of rules plus the COPY-rule (again, a' is the set of attributes of q specified in a and completed with *null*).

1. For each superclass r of p the following rules are defined:

$$msv(s(J), oid, p(a')) \leftarrow msv(J, oid, r(c)). \quad \{\text{SPECIALIZE Rules}\} \\ msv_{del}(s(J), oid, r(c)) \leftarrow msv(J, oid, r(c)).$$

2. For each subclass r of p :

$$msv(s(J), oid, p(a)) \leftarrow msv(J, oid, r(a, b)). \quad \{\text{GENERALIZE Rules}\} \\ msv_{del}(s(J), oid, r(a, b)) \leftarrow msv(J, oid, r(a, b)).$$

3. For each class r not *ISA*-related with p :

$$msv(s(J), oid, p \cap r(a', b)) \leftarrow msv(J, oid, r(b)). \quad \{\text{EXTEND Rules}\} \\ msv_{del}(s(J), oid, r(b)) \leftarrow msv(J, oid, r(b)).$$

Notice that each set of rules is composed by two different rules: an insertion rule, which inserts a new *msv*, and a deletion rule, which deletes the old *msv*. The *SPECIALIZE*-rules model the specialization of the object *oid* from superclass r to subclass p . The associated deletion rule deletes the previous *msv* of the object *oid*. The *GENERALIZE*-rules remove the *msv* of object *oid* if it belongs to a subclass of p , and generate the new *msv* in p . The *EXTEND*-rules deal with the possibility of inserting the object *oid* in a class p while its current *msv* belongs to class r which is not *ISA*-related with p . In this case the new *msv* of the object belongs to the intersection class of p and r .

EXAMPLE 6 the operation $extend(oid, student[name \rightarrow greene; school \rightarrow oid'])$ generates the following insertion and deletion rules.

$$msv(s(J), oid, student(greene, oid')) \leftarrow msv(J, oid, person(greene)). \\ msv_{del}(s(J), oid, person(greene)) \leftarrow msv(J, oid, person(greene)). \\ msv(s(J), oid, student \cap employee(greene, oid', X, Y)) \leftarrow msv(J, oid, employee(greene, X, Y)). \\ msv(s(J), oid, student \cap researcher(greene, oid', X, Y, Z)) \leftarrow msv(J, oid, researcher(greene, X, Y, Z)). \\ msv_{del}(s(J), oid, researcher(greene, X, Y, Z)) \leftarrow msv(J, oid, researcher(greene, X, Y, Z)). \\ msv_{del}(s(J), oid, employee(greene, X, Y)) \leftarrow msv(J, oid, employee(greene, X, Y)).$$

Notice that, to complete the compilation schema, the *EXTEND*-rules related to *department*, *school* and others should be generated. \square

Observe that, by the uniqueness of the *msv* relation, for each *oid* we have that at most one insertion rule is applicable, and analogously for the deletion rules. In fact, if either the student *oid* is not present or its current *msv* belongs to *student* or to *student* \cap *employee*, then none of the above clauses is applicable and therefore the *extend*(*oid*, *student*[*name* \rightarrow *green*; *school* \rightarrow 4]) is not executed.

The code realizing the operation *drop*(*oid*, *q*[*a*]), is composed by the following rules. For each subclass *r* of *p* (possibly *r* = *p*), and for each class *q* which is either a superclass of *p* or a superclass of *r* not ISA-related with *p*:

$$\begin{aligned} msv(s(J), oid, q(c)) &\leftarrow msv(J, oid, r(a, b)). \\ msv_{del}(s(J), oid, r(a, b)) &\leftarrow msv(J, oid, r(a, b)). \end{aligned} \quad \{\text{DROP Rules}\}$$

The DROP rule in the second clause removes the *msv* of object *oid* if it belongs to a subclass of *q* or to *q* itself. To generate the new *msv* of *oid* three cases are considered. If *q* is the root class of the hierarchy, the *oid* is removed from the database as no insertion rule is defined. If *q* is the specialization of a unique class *p*, then the new *msv* of *oid* belongs to *p*. In any other case, by the hypothesis of closure under intersection of the hierarchy, the current *msv* of *oid* belongs to the intersection of a subclass of *q* and some relation *p* which is not ISA-related with *q*. In this case the new *msv* of *oid* belongs to *p*. Observe that two cases are possible. If the current *msv* of *oid* belongs to *student*, then the new *msv* of *oid* will end up in *person*. Otherwise, if current *msv* of *oid* belongs to *student* \cap *employee*, then the new *msv* will end up in *employee*. It is worth noting that the constraint that deletion cannot be applied to intersection classes is needed to guarantee the uniqueness of the *msv*. In the example, it is not allowed to delete a *student* \cap *employee* directly, but it is needed to delete it both as a *student* and as an *employee*.

The code realizing the operation *modify* is similar to that of the *delete* operation, and therefore omitted.

References

- [ABGO93] A. Albano, R. Bergamini, G. Ghelli, and R. Orsini. An Object Data Model with Roles. In *Proceedings of the 19th International Conference on Very Large Data Bases*, 1993.
- [Abi92] S. Abiteboul. Towards a Deductive Object-Oriented Databases. In *Proceedings of the Int. Conf. on Deductive and Object-Oriented Databases*, pages 453-472. W. Kim and Nicolas, J.-M. and S. Nishio, 1992.
- [AKN86] H. Aït Kaci and R. Nasr. LOGIN: A Logic Programming Language with Built-in Inheritance. *Journal of Logic Programming*, 3(3):185-215, 1986.
- [AOTZ93] N. Arni, K. Ong, S. Tsur, and C. Zaniolo. *LDL++: A Second Generation Deductive Databases Systems*. Technical report, MCC Corporation, 1993.
- [BGM95] E. Bertino, G. Guerrini, and D. Montesi. Deductive Object Databases. In *Proceedings of ECOOP'95*, 1995.
- [BJ94] M. Bugliesi and H. M. Jamil. A Logic For Encapsulation in Object-Oriented Languages. In Maria Alpuente, Roberto Barbuti, and Isidro Ramos, editors, *Proceedings of the GULP-PRODE'94 Joint Conference on Declarative Programming*, volume 2, pages 161-175, September 1994.
- [BK93] A. J. Bonner and M. Kifer. Transaction Logic Programming. Technical Report CSRI-270, Computer System Research Institute, University of Toronto, December 1993.

- [CFG95] M. Carboni, V. Foddai, F. Giannotti, and D. Pedreschi. Declarative Reconstruction of Updates in Logic Databases: A Compilative Approach. In M. Alpuente and M. I Sessa, editors, *Proceedings of the GULP-PRODE joint conference on Declarative Programming- GULP-PRODE95*, pages 169–180, september 1995.
- [Che91] W. Chen. Declarative Specification and Evaluation of Database Updates. In *Proceedings of the Second International Conference on Deductive and Object-Oriented Databases (DOOD'91)*, number 566 in Lecture Notes in Computer Science, pages 147–166, 1991.
- [GPSZ91] F. Giannotti, D. Pedreschi, D. Saccà, and C. Zaniolo. Non-Determinism in Deductive Databases. In C. Delobel, M. Kifer, and Y. Masunga, editors, *Proceedings of the 2nd International Conference on Deductive and Object-Oriented Databases (DOOD91)*, Lecture Notes in Computer Science, pages 129–146. Springer-Verlag, Berlin, 1991.
- [KLW93] M. Kifer, G. Lausen, and J. Wu. Logical Foundations of Object-Oriented and Frame-Based Programming. Technical Report 93/06, Department of Computer Science, SUNY at Stony Brook, June 1993. Also appeared in Journal of ACM.
- [Man89] S. Manchanda. Declarative Expression of Deductive Database Updates. In *Proceedings of the ACM Symposium on Principles of Database Systems*, 1989.
- [Mon93] D. Montesi. *A Model for Updates and Transactions in Deductive Databases*. PhD thesis, Dipartimento di Informatica Università di Pisa, 1993.
- [MW86] S. Manchanda and D. S. Warren. A logic-Based Language for Database Updates. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 363–394. Springer-Verlag, Berlin, 1986.
- [RS90] J. Richardson and P. Schwarz. Aspects: Extending Objects to Support Multiple, Independent Roles. Technical Report 7657, IBM Almaden Research center, 1990.
- [Su91] J. Su. Dynamic Constraints and Object Migration. In *Proceedings of the 17th International Conference on Very Large Data Bases*, 1991.
- [SZ90] D. Saccà and C. Zaniolo. Stable Models and Non-determinism in Logic Programs with Negation. In *Proceedings of the ACM Symposium on Principles of Database Systems*, pages 205–217, 1990.
- [SZ95] V. S. Subrahmanian and C. Zaniolo. Relating Stable Models and AI Planning Domains. In *Proceeding of the International Conference on Logic Programming*, 1995.
- [Wds95] R. Wieringa, W. de Jonge, and P. Spruit. Using Dynamic Classes and Role Classes to Model Object Migration. *Theory and Practice of Object Systems*, 1(1):173–196, 1995.
- [Zan89] C. Zaniolo. Object Identity and Inheritance in Deductive Databases - An Evolutionary Approach. Technical Report AST-ST-373-89, MCC Corporation, 1989. Also in the 1st Int. Conf. on Deductive and Object-Oriented Databases.
- [Zan95] C. Zaniolo. Active Database Rules with Transaction Conscious Stable Model Semantics. In *4th International Conference on Deductive and Object-Oriented Databases*, 1995. To appear.
- [ZAO93] C. Zaniolo, N. Arni, and K. Ong. Negation and Aggregates in Recursive Rules: The LDL++ Approach. In *International conference on Deductive and Object-Oriented Databases (DOOD'93)*, 1993.

