Università di Pisa

Dipartimento di Informatica

Technical Report : TR-97-19

Active-U-Datalog: integrating active rules in a deductive database

Vincenzo Gervasi, Alessandra Raffaeta'

July 15, 1997

ADDR: Corso Italia 40,56125 Pisa,Italy. TEL: +39-50-887111. FAX: +39-50-887226

Active-U-Datalog:
integrating active rules in a deductive database

Vincenzo Gervasi and Alessandra Raffaetà
Dipartimento di Informatica
Università di Pisa
{gervasi,raffaeta}@di.unipi.it

Abstract

The marriage between logic programming and databases has given rise to the definition of *deductive databases*. These systems allow the users to express data manipulations and queries in a declarative way, and permit efficient storage and retrieval of intensional knowledge.

Another improvement in the database field has come from the use of *active rules*, linking the occurrence of certain events to a reaction (e.g., updates of some data). This kind of rules has proven to be very useful to ensure integrity constraints and to automatize common procedures.

This paper presents an integration of active rules in U-Datalog, which is an extension to Datalog supporting declarative specification of updates based on a nonimmediate update semantics. The resulting language, called Active-U-Datalog, extends the semantics of U-Datalog in a conservative way, introducing a PARK-like semantics for active rules activation and firing, and for handling conflicting update requests.

Keywords : Deductive databases, Active rules, Semantics.

Contents

1 Introduction

The evolution of database languages in recent years has been driven by two main goals.

The first one has been to avoid any *procedural* description of the operations to perform on the database. The advantages of abstracting from the actual implementation of the data store are obvious: there is no need to know anything about the underlying data representation, thus allowing for easier "program" expression and separation of database engines evolution from linguistic issues. The widespread use of SQL and its derivatives, both as actual query and data manipulation languages and, more recently, as high-level protocols to interact with heterogeneous data repositories, is the most remarkable success of *declarative* languages in this field. The natural evolution of this research has led to the use of *logic-based* languages, like Datalog [1, 28], $\mathcal{LDL}$++ [34] or RDL [17], that add *deductive* power to the simple databases we were used to work with. These deductive databases can efficiently handle vast amounts of derived knowledge, moving most semantic issues from (diffuse) "applicative logic" to (centralized) database design.

The second main goal has been to make the databases *smart* about the data they contain. This need is particularly evident when dealing with consistency issues; one could well ask that a semantic-aware database could *react* correctly to changes to the data it contains, updating other data accordingly, in order to maintain consistency. Besides maintaining integrity constraints, these capabilities are desirable to automatize common procedures (e.g., having an order shipped when the stock is low, or applying a discount when the stock is too high). This need has led to the design of *active* databases, in which *events* of various kinds (both internally- and externally-generated) may cause the firing of so-called "active rules". Also in this case knowledge about the data semantics and the design of the database is moved into the database itself, thus simplifying applicative software and reducing the chance for consistency-breaking data manipulations. The field has produced various results, both commercial [2, 3, 10, 18, 24] and academic [8, 9, 13, 15, 21, 25, 26, 31, 32] (the latter usually being more flexible than the former). However, the experience with these languages has often exposed the lack of a clear, unified and formal foundation for active rules [30].

These two approaches to improve database management systems are not in opposition, but rather they can be composed almost freely to attain the desired functionality. In giving a semantics to these active-deductive languages, two main approaches emerged: to *unify* the different paradigms under a common semantics (often by using compilative techniques) [7, 20, 23, 27, 33], or to *integrate* specific semantics [11].

In this paper, we follow the latter approach. We present an integration of deductive and active databases by extending U-Datalog [5, 22] with a support for active rules in the style of the PARK semantics [14]; we call the resulting language Active-U-Datalog. U-Datalog is an extension to Datalog supporting declarative specification of updates based on a nonimmediate update semantics. U-Datalog is well suited to such an extension thanks to its clean design, that concentrates on deductive issues without polluting the definition with unrelated needs. On the other hand, the PARK semantics is very flexible in its approach to conflicting updates, a subject that U-Datalog does not handle in a satisfactory way. Moreover, the PARK semantics leaves wide freedom in the choice of a conflict resolution policy, a virtue that can prove useful in implementing different strategies.

In section 2, we give an overview of U-Datalog and of Active-U-Datalog. In section 3 we give the formal syntax of Active-U-Datalog, and in section 4 its semantics. Section 5 shows in which sense Active-U-Datalog is an extension of U-Datalog. Some final remarks follow.

2 Overview of U-Datalog and Active-U-Datalog

We present here an overview of U-Datalog and we refer the reader to section 5 or to [22] for a complete description.

Let us consider pairwise disjoint sets of *constants* Σ, *predicate names* Π and *variables* V. *Terms* range over $\Sigma \cup V$ and *predicates* are defined as usual on Π. Π is divided into two disjoint sets: *extensional* (Π^e) and *intensional* (Π^i) predicates. Update operations are expressed by extensional atoms prefixed with $+$ and $-$ to denote respectively insertions or deletions.

A U-Datalog program (or database) consists of an extensional database (EDB), that is a set of ground extensional atoms, and an intensional database (IDB), that is a set of rules of the following form:

$$H \leftarrow U_1, \ldots, U_i, B_{i+1}, \ldots, B_n. \qquad n \geq 1$$

where U_j is an update and B_k is a $(\Pi^e \cup \Pi^i, \Sigma, V)$-atom and H is an intensional atom. [22] considers databases which are safe through invocation by a query. A rule is *safe* if each variable in the head appears in a $(\Pi^i \cup \Pi^e, \Sigma, V)$-atom in the body; it is *safe with respect to a query* if it is safe or rendered safe by bindings from the query.

A *simple transaction* T is a rule with no head. A *complex transaction* is a sequence of transactions, denoted by $T_1; T_2$.

Example 1 *Consider the following U-Datalog program*

```
EDB: r(a).
IDB: p(X) ← r(X),-r(X)
     q(X) ← p(X),+v(X)
     s(X) ← +r(X)
```

The first two rules are safe; on the contrary, the last one is not. But if we want to evaluate the transaction $T = s(b)$ then the database is safe with respect to the query.

The semantics of U-Datalog is essentially given in three steps: in the first one, the *marking phase*, the updates found during the evaluation process are collected, without executing them. Then, in the *update phase*, the requested updates are checked for consistency, that is all the requested updates must be ground and they must not be opposite ($+p(a)$ and $-p(a)$ must not be both requested). If the set of updates is not consistent the transaction is *aborted* and all the updates discarded. The third step is related to the execution of complex transactions.

Example 2 *The transaction $T = q(X)$ evaluated in the U-Datalog program of example 1 computes the binding $X = a$ and collects the updates $+v(a)$, $-r(a)$. Since the requested updates are consistent, the new EDB obtained by executing them is $EDB = v(a)$. If we evaluate the transaction $T = q(X), s(X)$ then the binding is still $X = a$ and the collected updates are $+v(a)$, $-r(a)$, $+r(a)$. The set is not consistent and therefore, the transaction aborts.*

In integrating active rules in U-Datalog, we want to extend its syntax and semantics in a conservative way. In doing so, we augment U-Datalog programs with a set of active (Event-Condition-Action) rules. In this paper, we consider as events only *updates*, but the extension to other kind of events (temporal, system-defined...) is not difficult. Conditions are actually *queries on the deductive part* of the program, and are evaluated with regard to the *intermediate* state of the computation. This is needed to allow active rules to react to the current (possibly inconsistent) state and take appropriate actions to assure desired properties of the final state.

Conflict resolution is handled by a *parametric conflict resolution policy*; this assures that no conflicting updates will still be present during the update of the database. This approach allows us not to abort a transaction in case of conflicts (as happens in U-Datalog): we try to solve conflicts and in the worst case, when the computation requires non-ground updates, we do not modify the database, but we do not need to abort.

Our actions are *sets of updates*. We allow for more than one update in an action so that the database designer can express atomic sets of updates (we could call them sub-transactions), with the intended meaning that, if the conflict resolution policy forbids one of the updates in the set, none of the updates in the set will be requested anymore.

3 Syntax of Active-U-Datalog

We consider a fixed (Π, Σ, V)-language. Π is partitioned into three sets of symbols: Π^e the extensional predicate symbols, Π^i the intensional predicate symbols and Π^u the update predicate symbols defined as $\Pi^u = \{+p, -p \mid p \in \Pi^e\}$.

An Active-U-Datalog program can be seen as a U-Datalog program (EDB $\cup$ IDB) to which we have added a set of active rules AR. We call *deductive* part the U-Datalog program and *active* part the set of active rules.

Definition 1 (Active-U-Datalog syntax) *An* Active-U-Datalog program $P = IDB \cup EDB \cup AR$ *consists of the extensional database EDB, of the intensional database IDB and of a set of active rules AR. The EDB is a set of ground extensional atoms. The IDB is a set of rules of the form*

$$H \leftarrow U_1, \ldots, U_n, B_1, \ldots, B_m$$

where $B_1, \ldots, B_m$ is the query part (B_i is a $(\Pi^i \cup \Pi^e, \Sigma, V)$-atom), $U_1, \ldots, U_n$ is the update part (U_i is a (Π^u, Σ, V)-atom), and the head H is a (Π^i, Σ, V)-atom. The update and query part cannot be both empty.
The AR is a set of rules of the form

$$E_1, \ldots, E_n, B_1, \ldots, B_m \rightarrow U_1, \ldots, U_k$$

where $E_1, \ldots, E_n$ is the event part (E_i is a (Π^u, Σ, V)-atom), $B_1, \ldots, B_m$ is the condition part (B_i is a positive or negative $(\Pi^i \cup \Pi^e, \Sigma, V)$-atom) and $U_1, \ldots, U_k$ is the action part (U_i is a (Π^u, Σ, V)-atom) which cannot be empty.

A *simple transaction* or *query* T is a rule of the deductive part with no head and with a non-empty query part; a *complex transaction* is a sequence of transactions $T_1; T_2$.

We will always assume that our programs are safe with respect to a query.

Example 3 *In a corporation, employees are assigned to a single department; each employee is paid according to his or her level. When a department is deleted, all its "chieves" (that is, all employees of certain levels) are moved to a special "available personnel" list, while the others are simply fired. We want to express this policy by using active rules.*

```
EDB: dep(toys).      chieflevel(3).     chieflevel(4).
     emp(john).      level(john,2).     assign(john,toys).
     emp(mary).      level(mary,4).     assign(mary,toys).
IDB: chief(E) ← level(E,L), chieflevel(L).
AR:  −dep(D), assign(E,D) → −emp(E).
     −emp(E), chief(E) → +avail(E).
     −emp(E) → −assign(E,_), −level(E,_).
```

In response to the query "−dep(toys)", the first active rule asks for the deletion of every employee working in "toys". This request activates the third rule, that deletes the assignment and level of john and mary, and the second one, that puts mary in the available list.

4 Semantics

We present the semantics of Active-U-Datalog in three steps. In the first step, we compute the model of the IDB given the initial EDB and collect the set of bindings that satisfies the query and the requested updates. This step corresponds to the marking phase in U-Datalog. In the second step, we compute the semantics of the active part of the program, according to the model and the updates collected in the first step. The result of this step is the set of updates requested either from the deductive and/or the active part, in which any conflict has been solved by a parametric policy. Finally, we describe how the two semantics fit together and how we apply the computed updates to the EDB, thus obtaining the new database state.

The observable properties of the transaction we want to model are the set of bindings (the answer), the new database state and the indication of success (commit/abort) of the transaction itself.

4.1 Deductive part semantics

The semantics for the deductive part is given by defining a bottom-up operator T analogous to the immediate consequences operator of standard logic programming. This operator is defined over an extended Herbrand Base $\mathcal{B}^{\leftarrow}$containing atoms of the form $H \leftarrow \overline{U}$ where H is a $(\Pi^i \cup \Pi^e)$-atom and $\overline{U}$ is a set of updates. The meaning of such an atom is that H is true, and its evaluation requests the updates in $\overline{U}$.

Definition 2 (Immediate consequence operator) *Given an Active-U-Datalog program P, we define the* immediate consequence operator $T_P : 2^{\mathcal{B}^{\leftarrow}} \to 2^{\mathcal{B}^{\leftarrow}}$ *as follows:*

$$
\begin{aligned}
T_P(I) \;=\{\; & A \leftarrow \overline{U} \mid A \leftarrow \overline{U} \in \mathcal{B}^{\leftarrow}, \\
& H \leftarrow U_1, \ldots, U_n, B_1, \ldots, B_m, \; \textit{is a renamed apart clause of } P \\
& B'_r \leftarrow \overline{U}_r \in I, \quad 1 \le r \le m \\
& \theta = \; mgu\ ((B_1, \ldots, B_m), (B'_1, \ldots, B'_m)) \\
& A = H\theta \\
& \overline{U} = (U_1, \ldots, U_n, \overline{U}_1, \ldots, \overline{U}_m)\theta^1 \\
\}\; &
\end{aligned}
$$

where the function mgu returns an idempotent mgu.

Notice that the rules involved in this definition are only of the deductive part because the head H is a $(\Pi^i \cup \Pi^e)$-atom. The operator T is continuous on the lattice $(2^{\mathcal{B}^{\leftarrow}}, \subseteq)$, and we use it to define the fixpoint semantics for the deductive part.

Definition 3 (Fixpoint semantics) *Given an Active-U-Datalog program P, we define the* fixpoint semantics $\mathcal{F}$ *of P as follows:*

$$\mathcal{F}(P) = T_P^n(\emptyset)$$

Since the domain is finite, the least fixpoint is reached in a finite number of steps (n) [6].

Example 4 *Let us consider the following Active-U-Datalog program P without active rules:*

```
EDB: t(a).    r(b).
IDB: p(X) ← q(X),+r(X)
     q(X) ← s(X)
     q(X) ← r(X)
     s(X) ← t(X),−r(X)
```

The least fixpoint computation is as follows:

$$
\begin{aligned}
T_P^0(\emptyset) &= \emptyset \\
T_P^1(\emptyset) &= \{t(a),\ r(b)\} \\
T_P^2(\emptyset) &= \{t(a),\ r(b),\ q(b), \quad s(a) \leftarrow -r(a)\}
\end{aligned}
$$

[1] As a shorthand, we write $(U_1, \ldots, U_n, \overline{U}_1, \ldots, \overline{U}_m)\theta$ for $\{U_1\theta, \ldots, U_n\theta\} \cup \overline{U}_1\theta \cup \ldots \cup \overline{U}_m\theta$.

$$
\begin{array}{rcl}
T_P^3(\emptyset) & = & \{t(a),\ r(b),\ q(b),\quad s(a) \leftarrow -r(a),\quad q(a) \leftarrow -r(a)\} \\
T_P^4(\emptyset) & = & \{t(a),\ r(b),\ q(b),\quad s(a) \leftarrow -r(a),\quad q(a) \leftarrow -r(a),\quad p(a) \leftarrow -r(a), +r(a), \\
& & p(b) \leftarrow +r(b)\}
\end{array}
$$

$T_P^5(\emptyset) = T_P^4(\emptyset)$, so we reached the least fixpoint $\mathcal{F}(P)$.

A set of updates $\overline{U}$ is *consistent* if it contains no opposite updates, i.e. if $+r(\tilde{t})$ and $-r(\tilde{t})$ are not both in $\overline{U}$.

It is worth remarking that the evaluation of an atom can lead to inconsistent updates as shown in the previous example, where (for instance) the evaluation of $p(a)$ requests the updates $-r(a), +r(a)$. Differing from [22], where inconsistent updates are never introduced in the model, we include them and resolve the conflicts during the active rules evaluation (section 4.2).

Before introducing the set-oriented semantics we give two auxiliary definitions:

Definition 4 *Given a set of bindings b and a transaction T, we define*

$$b_{|T} = \{(X = t) \in b \mid X \textit{ occurs in } T\}$$

Definition 5 *Given a substitution $\theta = \{V_1 \leftarrow t_1, \ldots, V_n \leftarrow t_n\}$ we define*

$$eqn(\theta) = \{V_1 = t_1, \ldots, V_n = t_n\}$$

We define the (set-oriented) semantics of a simple transaction T with regard to a program P, using its model built by $\mathcal{F}$.

Definition 6 (Query answers) *Given an Active-U-Datalog program P and a simple transaction $T = (\overline{U}_0, B_1, \ldots, B_n)$, we define the operator* Set *as follows:*

$$
\begin{array}{rll}
\mathsf{Set}((\overline{U}_0, B_1, \ldots, B_n), P) & = \{ & \langle b, \overline{U}\rangle \mid A_i \leftarrow \overline{U}_i \in \mathcal{F}(P), \quad 1 \le i \le n \\
& & \theta = mgu((B_1, \ldots, B_n), (A_1, \ldots, A_n)) \\
& & b = eqn(\theta)_{|T} \\
& & \overline{U} = (\overline{U}_0, \overline{U}_1, \ldots, \overline{U}_n)\theta \quad \}
\end{array}
$$

In U-Datalog, only pairs with consistent updates are inserted in Set. On the contrary, we release this restriction, deferring again the conflict resolution at the next step.

Example 5 *Let $p(X), q(X)$ be a query to the program P of example 4. Then,*

$$\mathsf{Set}((p(X), q(X)), P) \;\; = \{ \;\; \langle\{X = a\}, \{-r(a), +r(a)\}\rangle,\ \langle\{X = b\}, \{+r(b)\}\rangle \;\; \}$$

4.2 Active part semantics

During this phase, we follow the PARK semantics proposed in [14]. This semantics is particularly suited to a deferred-update approach, like the one we used in the previous step, and adds considerable flexibility in that it uses a parametric policy to resolve conflicts.

This semantics builds an auxiliary model, in which we are interested only in the update atoms needed to trigger active rules and to obtain the new EDB. To this end, we define a bottom-up operator whose domain is $\mathcal{B}^{\pm} = \{p(\tilde{t}) \mid p(\tilde{t}) \in \mathcal{B}\} \cup \{+p(\tilde{t}), -p(\tilde{t}) \mid p(\tilde{t}) \in \mathcal{B}, p \in \Pi^e\}$ where $\mathcal{B}$ is the standard Herbrand Base. A subset of $\mathcal{B}^{\pm}$ is an *i-interpretation* (where the "i" stands for intermediate). An i-interpretation is *consistent* if it does not contain any opposite atoms, i.e. $+a$ and $-a$.

Definition 7 (Validity) *The validity of a ground literal a in an i-interpretation I is defined as follows:*

$$\mathsf{valid}(a, I) = \begin{cases} I \cap \{p(\tilde{t}), +p(\tilde{t})\} \neq \emptyset & \text{if } a = p(\tilde{t}) \\ I \cap \{p(\tilde{t}), +p(\tilde{t})\} = \emptyset \text{ or } -p(\tilde{t}) \in I & \text{if } a = \neg p(\tilde{t}) \\ a \in I & \text{otherwise} \end{cases}$$

A positive $(\Pi^e \cup \Pi^i, \Sigma, V)$-atom is valid if it belongs to I or if an update inserting it belongs to I. A negative $(\Pi^e \cup \Pi^i, \Sigma, V)$-atom is valid if an update deleting it is in I, or if its positive atom is not valid. A (Π^u, Σ, V)-atom is valid if it belongs to I. Notice that both $p(\tilde{t})$ and $\neg p(\tilde{t})$ can be valid according to this definition.

Definition 8 (Purification) *Let IDB be the intensional database of an Active-U-Datalog program P; we define its* purified *version $\widehat{IDB}$ as follows:*

$$H \leftarrow U_1, \ldots, U_n, B_1, \ldots, B_m \in \mathit{IDB} \implies B_1, \ldots, B_m \rightarrow H \in \widehat{\mathit{IDB}}$$

Example 6 *The purified form of the IDB of the program in example 4 is the following:*

```
IDB^:  q(X) → p(X)
       s(X) → q(X)
       r(X) → q(X)
       t(X) → s(X)
```

It is worth noting that a query is provable in $\widehat{\mathrm{IDB}} \cup \mathrm{EDB}$ if and only if it is provable in $\mathrm{IDB} \cup \mathrm{EDB}$, with the same computed answers. The purification only avoids the side effects of the query evaluation. Also notice that we reversed the direction of the arrow in order to have a uniform notation with active rules. No ambiguity can arise since active rules have only update atoms in their heads, while purified ones have (Π^i, Σ, V)-atoms. In the following, we will generically refer to active and purified rules as "rules". Notice that definitions 9, 10, and 11 only consider active rules, while subsequent ones work on both kinds of rules.

Definition 9 (Conflicts) *A pair (r, θ), where r is a rule and θ is a ground substitution for r is called a* rule grounding.

Let P be a set of rules and I an i-interpretation for P. Then conflicts*(P,I) is a set of maximal triples of the form $(a, \mathsf{ins}, \mathsf{del})$ such that a is a ground atom and* ins *and* del *are sets of rule groundings. For each such triple the following conditions must hold:*

1. *$\exists\, r = l_1, \ldots, l_n \rightarrow u_1, \ldots, u_k, \quad r' = l'_1, \ldots, l'_m \rightarrow u'_1, \ldots, u'_s, \quad r, r' \in P$, and $\exists\, \theta, \theta'$ ground substitutions such that*
   - $\forall\, 1 \leq i \leq n,\ \mathsf{valid}(l_i\theta, I)$,
   - $\forall\, 1 \leq i \leq m,\ \mathsf{valid}(l'_i\theta', I)$,
   - $\exists\, i, j.\ 1 \leq i \leq k.\ u_i = +l_0,\ 1 \leq j \leq s.\ u'_j = -l'_0$ *and* $a = l_0\theta = l'_0\theta'$.
2. *For all possible r, r' and θ, θ', satisfying condition 1 above, $(r, \theta) \in \mathsf{ins}$ and $(r', \theta') \in \mathsf{del}$.*

A triple $(a, \mathsf{ins}, \mathsf{del}) \in \mathsf{conflicts}(P, I)$ is called a *conflict.*

The previous step (definitions 2 and 6) could have produced conflicts; we are now going to resolve them via a parametric policy.

Definition 10 (Conflict resolution policy) *Given an extensional database EDB, a set of rules P, an i-interpretation I and a conflict c, we define $\mathsf{sel}(\mathit{EDB}, P, I, c)$ as a total function with codomain $\{\mathsf{insert}, \mathsf{delete}\}$.*

The intended meaning of $\mathsf{sel}(\mathrm{EDB}, P, I, (a, \mathsf{ins}, \mathsf{del}))$ is to choose if the atom a, object of the conflict, should be inserted in or deleted from I, thus effectively choosing which of the conflicting updates requests should prevail. This is obtained by *preventing* the rule instances in one of the two sets from firing.

[14] presents a number of common policies, and discusses their advantages and disadvantages. We point out here just two of the most common ones. The *principle of inertia* states that *both* of the conflicting updates should be discarded, effectively leaving EDB in the same state as before with regard to a (in our framework, this can be obtained by returning insert if a was already in EDB, delete otherwise). The *rule priority* policy, found in systems such as Ariel [15], Postgres [25] and Starburst [32], assumes that each rule has a (static or dynamic) priority associated with it; sel returns insert or delete as needed to preserve the update requested by the highest-priority rule. Other policies, like voting schemes and user queries, are also reasonable, but the final choice is better left to the particular application.

Definition 11 (Blocked rule instances) *Given an extensional database EDB, a set of rules P, a conflict resolution policy* sel*, and an i-interpretation I, let*

$$\begin{aligned} X &= \{\mathsf{del} \mid (a, \mathsf{ins}, \mathsf{del}) \in \mathsf{conflicts}(P, I) \text{ and } \mathsf{sel}(EDB, P, I, (a, \mathsf{ins}, \mathsf{del})) = \mathsf{insert}\} \\ Y &= \{\mathsf{ins} \mid (a, \mathsf{ins}, \mathsf{del}) \in \mathsf{conflicts}(P, I) \text{ and } \mathsf{sel}(EDB, P, I, (a, \mathsf{ins}, \mathsf{del})) = \mathsf{delete}\} \end{aligned}$$

We define then

$$\mathsf{blocked}(EDB, P, I, \mathsf{sel}) = \left(\bigcup_{x \in X} x\right) \bigcup \left(\bigcup_{y \in Y} y\right)$$

Example 7 *Let us consider the following extensional database EDB and the set of rules P:*

```
EDB:    t(a,a).  t(a,b).  r(a).
P:   r1 t(a,X) → p(X)
     r2 s(X), r(X) → p(X)
     r3 +s(X),t(X,Y) → +q(X)
     r4 +s(X),t(X,Y) → −q(Y)
     r5 +s(X) → +q(b)
     r6 +q(X),p(X) → −t(X,X)
```

Given an i-interpretation $I = \{t(a,a), t(a,b), r(a), +s(a)\}$, we have that

$$\begin{aligned} \mathsf{conflicts}(P, I) = \{ \quad & (q(a), \{(r_3, \{X \leftarrow a, Y \leftarrow a\})\}, \{(r_4, \{X \leftarrow a, Y \leftarrow a\})\}), \\ & (q(b), \{(r_5, \{X \leftarrow a\})\}, \{(r_4, \{X \leftarrow a, Y \leftarrow b\})\}) \quad \} \end{aligned}$$

Now, let us suppose that the conflict resolution policy sel *is such that*

$$\mathsf{sel}(EDB, P, I, (q(a), \{(r_3, \{X \leftarrow a, Y \leftarrow a\})\}, \{(r_4, \{X \leftarrow a, Y \leftarrow a\})\})) = \mathsf{insert}$$
$$\mathsf{sel}(EDB, P, I, (q(b), \{(r_5, \{X \leftarrow a\})\}, \{(r_4, \{X \leftarrow a, Y \leftarrow b\})\})) = \mathsf{delete}$$

Then, we have that

$$\mathsf{blocked}(EDB, P, I, \mathsf{sel}) = \{(r_4, \{X \leftarrow a, Y \leftarrow a\}), (r_5, \{X \leftarrow a\})\}$$

To solve the condition part of an active rule, we want to use the knowledge in the deductive part. However, in exploiting this knowledge we work under different conditions with respect to what we did in section 4.1:

- the resolution of a condition should not affect the state of the EDB. To obtain this, we forbid any side effect (i.e., updates) during the resolution process of a condition;

- conditions should be checked against the *current* state of the EDB, in order to take into account the updates requested by the active rules.

Definition 12 (Immediate consequence operator) *Given a set of rules P, a set of blocked rule instances B, and an i-interpretation I, we define $\Gamma_{P,B}(I)$ as the smallest set U satisfying the following conditions:*

- $I \subseteq U$;
- *If* $r = l_1, \ldots, l_n \rightarrow u_1, \ldots, u_k \in P$ *and* θ *is a ground substitution for* r *such that*
  - $(r, \theta) \notin B$
  - $\forall\ 1 \leq i \leq n,\ \mathsf{valid}(l_i\theta, I)$

  then $\{u_1\theta, \ldots, u_k\theta\} \subseteq U$

$\Gamma_{P,B}$ is continuous on the lattice $(\mathcal{B}^{\pm}, \subseteq)$.

Example 8 *Given the definitions of example 7, with* $B = \mathsf{blocked}(EDB, P, I, \mathsf{sel})$*, we obtain*

$$\Gamma_{P,B}(I) \quad = \quad \{t(a,a), t(a,b), r(a), +s(a), p(a), p(b), +q(a), -q(b)\}$$

Notice how the use of blocked *has prevented the insertion of any conflict in* $\Gamma_{P,B}(I)$.

Definition 13 (Bi-structures) *A bi-structure $\langle B, I\rangle$ consists of a set B of blocked rule instances and an i-interpretation I. We define an order relation on bi-structures as follows:*

$$\langle B, I\rangle \prec \langle B', I'\rangle \quad \Leftrightarrow \quad \begin{cases} B \subset B' & \textit{or} \\ B = B' \textit{ and } I \subset I' \end{cases}$$

Moreover, given $\mathcal{A}$ and $\mathcal{B}$ bi-structures, $\mathcal{A} \preceq \mathcal{B} \quad \equiv \quad (\mathcal{A} = \mathcal{B} \vee \mathcal{A} \prec \mathcal{B})$

Definition 14 (Δ operator) *Given a set of rules P, an extensional database EDB, a bi-structure $\mathcal{A} = \langle B, I\rangle$, a set of ground $(\Pi^e \cup \Pi^i, \Sigma, V)$-atoms $\overline{I} \subseteq I$, and a conflict resolution policy* sel*, we define*

$$\Delta_{P,\mathsf{sel},\overline{I}}(\mathcal{A}) = \begin{cases} \langle B, \Gamma_{P,B}(I)\rangle & \textit{if } \Gamma_{P,B}(I) \textit{ is consistent} \\ \langle B \cup \mathsf{blocked}(EDB, P, I, \mathsf{sel}), \overline{I}\rangle & \textit{otherwise} \end{cases}$$

The definition of Δ we give here differs from the original in [14] in two points:

- we use a set of active and purified rules in P, while [14] only consider active rules. This extension does not affect the consistency of i-interpretations, since the purified rules can only add (Π^i, Σ, V)-atoms to the i-interpretation;
- we use an explicit constant parameter $\overline{I}$ to maintain the initial interpretation, instead of obtaining it from the current one as the $I^{\perp}$ operator in [14] does. However, the meaning is not changed, since the property $\Gamma_{P,B} \uparrow n(I)^{\perp} = I^{\perp}$ assures us that $I^{\perp}$ is indeed a constant, despite its definition as an operator.

Notice that the poset of bi-structures, ordered by $\preceq$, is a cpo, and that $\Delta_{P,\mathsf{sel},\overline{I}}$ is continuous on this domain. Moreover, the least fixed point can be obtained in a finite number of steps because the cpo itself is finite.

Theorem 1 *Given a set of rules P, a conflict resolution policy* sel*, a bi-structure $\mathcal{A} = \langle B, I\rangle$, and $\overline{I} = I \setminus \{\pm a \mid \pm a \in I\}$, the following statements hold:*

1. $\mathcal{A} \preceq \Delta_{P,\mathsf{sel},\overline{I}}(\mathcal{A})$

2. $\Delta^{\omega}_{P,\mathsf{sel},\overline{T}}(\mathcal{A})$ *is a least fixed point of* $\Delta_{P,\mathsf{sel},\overline{T}}$

3. *if* $\Delta^{\omega}_{P,\mathsf{sel},\overline{T}}(\mathcal{A}) = \langle B', I'\rangle$, *then* $I' = \mathsf{lfp}(\Gamma_{P,B'})$

Proof. We proceed along the lines of theorem 4.1 in [14].

1. Let $\Delta_{P,\mathsf{sel},\overline{T}}(\mathcal{A}) = \langle B', I'\rangle$. If I' is consistent, then $B' = B$ and $I' = \Gamma_{P,B}(I) \supseteq I$; hence $\langle B, I\rangle \preceq \langle B', I'\rangle$. If instead I' is not consistent, then $B' = B \cup \mathsf{blocked}(\mathrm{EDB}, P, I, \mathsf{sel}) \supseteq B$ and so we have again $\langle B, I\rangle \preceq \langle B', I'\rangle$.

2. Immediate consequence of 1.

3. Since $\Delta^{\omega}_{P,\mathsf{sel},\overline{T}}(\mathcal{A})$ is a fixpoint for $\Delta_{P,\mathsf{sel},\overline{T}}$, we have that $\Delta_{P,\mathsf{sel},\overline{T}}\langle B', I'\rangle = \langle B', I'\rangle$. By definition 14 it follows that $I' = \Gamma_{P,B'}(I)$, and thus I' is a fixpoint for $\Gamma_{P,B'}$. To prove that it is a least fixpoint, let us suppose that there exists $J \subset I'$ fixpoint for $\Gamma_{P,B'}$. By monotonicity of $\Gamma_{P,B'}$ we have that $\Gamma_{P,B'}(J) \subseteq \Gamma_{P,B'}(I')$, hence $\Gamma_{P,B'}(J)$ is consistent and thus

$$\Delta_{P,\mathsf{sel},\overline{T}}\langle B', J\rangle = \langle B', \Gamma_{P,B'}(J)\rangle = \langle B', J\rangle$$

but then $\langle B', J\rangle$ would be a fixpoint for $\Delta_{P,\mathsf{sel},\overline{T}}$ smaller than $\langle B', I'\rangle$, contradicting 2.

□

4.3 Integrating deductive and active semantics

In this section we show how the two semantics presented above fit together and how the result of a transaction is computed.

We are interested in modeling as observable property of a transaction the following information: the set of answers, the database state, and the result of the transaction itself (i.e. Commit and Abort).

Definition 15 (Observables) *An* observable *Oss is a triple* $\langle \mathsf{Ans}, EDB, \mathsf{Res}\rangle$ *where* Ans *is a set of bindings, EDB is an extensional database and* Res $\in${Commit, Abort}. *The set of observables is* Oss.

We now define a function which takes an i-interpretation and the current extensional database and returns the new extensional database obtained by executing the updates in the i-interpretation.

Definition 16 (Updates incorporation) *Given an i-interpretation I and an extensional database EDB, we define*

$$\mathsf{incorp}(I, EDB) = (EDB \setminus \{a \mid -a \in I\}) \cup \{a \mid +a \in I\}$$

Finally, we define an auxiliary operator that transforms the requested updates collected during the first step into active rules.

Definition 17 (Updates as rules) *Given a set of rules P and a set of updates U, we define*

$$P_U = P \cup \{\ \rightarrow \pm a \mid \pm a \in U\}$$

Now, we can give the Active-U-Datalog semantics.

Definition 18 (Active-U-Datalog semantics) *Given an Active-U-Datalog program P (with EDB extensional database, IDB intensional database, and AR active rules), a transaction T and a conflict resolution policy* sel*, the semantics of a transaction T is denoted by the function Sem defined as*

$$Sem_{P,\mathsf{sel}}(T) = \mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T)(\langle \emptyset, EDB, \mathsf{Commit}\rangle)$$

where the function $\mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T) : \mathsf{Oss} \rightarrow \mathsf{Oss}$ *is defined as follows:*
If T *is a simple transaction, then*

$$\mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T)(Oss_i) = \begin{cases} \langle \mathsf{Ans}, \mathsf{incorp}(I, EDB), \mathsf{Commit}\rangle & \textit{if OK} \\ \langle \emptyset, EDB, \mathsf{Commit}\rangle & \textit{otherwise} \end{cases}$$

where

$$\begin{aligned} \mathsf{Ans} &= \{b_j \mid \langle b_j, \overline{u}_j\rangle \in \mathsf{Set}(T,P)\} \\ \overline{U} &= \bigcup \overline{u}_j \\ \langle B, I\rangle &= \Delta^{\omega}_{(\widehat{IDB \cup AR})_{\overline{U}},\mathsf{sel},EDB}(\langle \emptyset, EDB\rangle) \\ OK &\equiv \overline{U} \textit{ is ground} \end{aligned}$$

If T *is a complex transaction* $T_1;T_2$*, then*

$$\mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T_1;T_2)(Oss_i) = \mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T_2)(Oss_{i+1})$$

where

$$\mathsf{Oss}_{i+1} = \mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T_1)(Oss_i)$$

It is worth remarking that our semantics never produces aborts caused by conflicts, differing from [4], augmenting, therefore, the successful computations. This is possible due to the presence of the active part, and especially of the conflict resolution policy. Moreover, note that the "otherwise" case discards all changes to the EDB if a non-ground update is requested. The semantics of a complex transaction is simply given by the sequential composition of the updates of its components, since none of them can abort. Besides, this semantics discards the answer to the first transaction, to stay close to the approach in [4].

Example 9 *Let us return to the corporation we met in example 3. We want that, when a department is closed, all its chieves are moved to a different department, while the remaining personnel is fired. This policy is expressed by the following Active-U-Datalog program* P*:*

```
EDB: dep(toys).     chieflevel(3).     chieflevel(4)
     emp(john).     level(john,2).     assign(john,toys).
     emp(mary).     level(mary,4).     assign(mary,toys).
IDB: chief(X) ← level(X,L), chieflevel(L).
     chief_of(X,D) ← assign(X,D), chief(X).
     close_move(D,B) ← −dep(D), +emp(X), +assign(X,B), chief_of(X,D).
AR:  −dep(D), assign(E,D) → −emp(E).
     −emp(E), assign(E,D), level(E,L) → −assign(E,D), −level(E,L).
     +assign(X,D), assign(X,B), ¬(B=D) → −assign(X,B)
     +assign(X,D), ¬dep(D) → +dep(D)
```

Let us compute the semantics of the query $T = \texttt{close\_move}(\texttt{toys}, \texttt{shoes})$*, using the* inertial *conflict resolution policy* sel *(that is, always leave the extensional database untouched when a conflict arises).*

First, we compute the fixed-point semantics of the deductive part of P *(definition 3):*

$\mathcal{F}(P) = \{$ *dep(toys), chieflevel(3), chieflevel(4), emp(john), level(john,2), assign(john,toys), emp(mary), level(mary,4), assign(mary,toys), chief(mary), chief_of(mary,toys), close_move(toys,B)← −dep(toys), +emp(mary), +assign(mary,B)* $\}$

Then, we compute the query answer (definition 6):

$\mathsf{Set}(T,P) = \{\ \langle\ \emptyset,\ \{-dep(toys),\ +emp(mary),\ +assign(mary,shoes)\}\ \rangle\quad\}$

The set of requested updates is thus

$\overline{U} = \{\ -dep(toys),\ +emp(mary),\ +assign(mary,shoes)\quad\}$

and, since it is ground, we are in the "OK" case of the $\mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}$ *definition and so we can compute the active part semantics.*

The purified IDB (definition 8) is the following:

$\widehat{IDB}$:
```
level(X,L), chieflevel(L) → chief(X)
assign(X,D), chief(X) → chief_of(X,D)
chief_of(X,D) → close_move(D,B)
```

so the set of rules $P' = (\widehat{IDB} \cup AR)_{\overline{U}}$ *we use in the computation of* Δ^ω *is:*

$\widehat{IDB}$:	r_1	`level(X,L), chieflevel(L) → chief(X)`
	r_2	`assign(X,D), chief(X) → chief_of(X,D)`
	r_3	`chief_of(X,D) → close_move(D,B)`
AR:	r_4	`−dep(D), assign(E,D) → −emp(E).`
	r_5	`−emp(E), assign(E,D), level(E,L) → −assign(E,D), −level(E,L).`
	r_6	`+assign(X,D), assign(X,B), ¬(B=D) → −assign(X,B)`
	r_7	`+assign(X,D), ¬dep(D) → +dep(D)`
from U:	r_8	`→ −dep(toys)`
	r_9	`→ +emp(mary)`
	r_{10}	`→ +assign(mary,shoes)`

The computation of $\Delta^\omega_{P',\mathsf{sel},EDB}$ *proceeds as follows (we denote with* $\langle B^i, I^i\rangle$ *the bi-structure computed at the i-th step):*

$$\begin{aligned}
\Delta^0_{P',\mathsf{sel},EDB}(\langle\emptyset, EDB\rangle) &= \langle\emptyset, EDB\rangle\\
\Delta^1_{P',\mathsf{sel},EDB}(\langle\emptyset, EDB\rangle) &= \langle\emptyset, I^0 \cup \{+assign(mary,shoes), +emp(mary),\\
&\qquad -dep(toys), chief(mary), chief_of(mary,toys)\}\rangle
\end{aligned}$$

Now, in the computation of $\Gamma_{P',\emptyset}(I^1)$ *a conflict arises:*

$\Gamma_{P',\emptyset}(I^1) = \{$ *dep(toys), chieflevel(3), chieflevel(4), emp(john), emp(mary), level(john,2), level(mary,4), assign(john,toys), assign(mary,toys), +assign(mary,shoes),* **+emp(mary)**, *−dep(toys), chief(mary), chief_of(mary,toys), close_move(toys,...), −emp(john),* **−emp(mary)**, *−assign(mary,toys), +dep(shoes)* $\}$

The conflict is $c = (emp(mary), \{\langle r_9, \emptyset\rangle\}, \langle r_4, \{D = toys, E = mary\}\rangle)$*; since we are using the inertial conflict resolution policy,* $\mathsf{sel}(EDB, P', I^1, c) = \mathsf{insert}$*, thus blocking* $(r_4, \{D \leftarrow toys, E \leftarrow mary\})$*. The computation proceeds without further conflicts as follows:*

$$\begin{aligned}
\Delta^2_{P',\mathsf{sel},EDB}(\emptyset, EDB) &= \langle\{\langle r_4, \{D = toys, E = mary\}\rangle\}, EDB\rangle)\\
\Delta^3_{P',\mathsf{sel},EDB}(\emptyset, EDB) &= \langle\{\langle r_4, \{D = toys, E = mary\}\rangle\}, I^1\rangle\\
\Delta^4_{P',\mathsf{sel},EDB}(\emptyset, EDB) &= \langle\{\langle r_4, \{D = toys, E = mary\}\rangle\}, I^3 \cup\\
&\qquad \{close_move(toys,\ldots), -emp(john),\\
&\qquad -assign(mary,toys), +dep(shoes)\}\rangle\\
\Delta^5_{P',\mathsf{sel},EDB}(\emptyset, EDB) &= \langle\{\langle r_4, \{D = toys, E = mary\}\rangle\}, I^4 \cup\\
&\qquad \{-assign(john,toys), -level(john,2)\}\rangle\\
\Delta^6_{P',\mathsf{sel},EDB}(\emptyset, EDB) &= \Delta^5_{P',\mathsf{sel},EDB}(\emptyset, EDB)
\end{aligned}$$

We can now incorporate the updates in the EDB, thus obtaining EDB′:

$$EDB' = \mathsf{incorp}(I, EDB) = \{ \ \textit{chieflevel(3), chieflevel(4), emp(mary), level(mary,4),} \\ \textit{assign(mary,shoes), dep(shoes)} \ \}$$

Finally,

$$Sem_{P,\mathsf{sel}}(T) = \mathcal{S}^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}(T)(\langle\emptyset, EDB, \mathsf{Commit}\rangle) = \langle\emptyset, EDB', \mathsf{Commit}\rangle$$

5 Comparison with U-Datalog semantics

The addition of active rules makes Active-U-Datalog semantics quite different from U-Datalog one. In this section, we study the conditions under which our approach is conservative with respect to the original U-Datalog semantics. The main condition is, obviously, the absence of active rules in the Active-U-Datalog program. To this end, we first give a formal description of U-Datalog syntax and semantics [22].

5.1 U-Datalog as CLP($\mathcal{AD}$)

An U-Datalog program can be viewed as an instance of a CLP schema [16], called CLP($\mathcal{AD}$) [5]. We consider a fixed (Π, Σ, V)-language. Π is partitioned into four sets of symbols: Π^e, Π^i, Π^u defined as in section 3, and the binding predicate symbol $\Pi^b = \{=\}$.

Definition 19 (U-Datalog syntax) *An* U-Datalog program *(or database) $P = IDB \cup EDB$ consists of the extensional database EDB and of the intensional database IDB. The EDB is a set of ground extensional atoms. The IDB is a set of rules of the form*

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

where $B_1, \ldots, B_t$ is the query part (B_i is a $(\Pi^i \cup \Pi^e, \Sigma, V)$-atom), $u_1, \ldots, u_s$ is the update part (u_i is a (Π^u, Σ, V)-atom), $b_1, \ldots, b_k$ is the binding part (b_i is a (Π^b, Σ, V)-atom), and the head H is a (Π^i, Σ, V)-atom. The update and query part cannot be both empty.

A simple transaction *or* query *T is a rule with no head and with a non-empty query part; a* complex transaction *is a sequence of transactions $T_1; T_2$.*

We will always consider databases that are safe with respect to a query.

The constraints we deal with are of the form b, u where b is a set of bindings and u is a set of updates. To define the solvability of these constraints we consider the structure $\Re^w$ as defined in [22], definition 7.1.3. We recall that $= \in \Pi^b$ is interpreted as syntactic equality over the Herbrand Base $\mathcal{H}$, and b, u is $\Re^w$-solvable iff there exists a grounding substitution θ such that θ is a solution for b and $u\theta$ is consistent.

The semantics is given in three steps as described in section 2. During the marking phase the model of a program can be computed in a bottom-up style, by using the following immediate consequence operator $T^U_P : 2^{\mathcal{B}^U} \rightarrow 2^{\mathcal{B}^U}$ where $2^{\mathcal{B}^U}$ is the power-set of the constrained atoms: $\mathcal{B}^U = \{p(\tilde{X}) \leftarrow b, u \mid p \in \Pi^i \cup \Pi^e\}$. We call a *CLP-interpretation* a subset of $\mathcal{B}^U$.

Definition 20 (Model) *([22], definition 7.2.3) Given an U-Datalog program P and a CLP-interpretation J, we define*

$$\begin{aligned} T^U_P(J) \ = \{ \ & p(\tilde{X}) \leftarrow b, u \mid \\ & \exists \textit{ a renamed clause} \\ & p(\tilde{t}) \leftarrow b_0, u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \textit{ in } P, \\ & \exists \ p_i(\tilde{X}_i) \leftarrow b_i, u_i \in J, \quad 1 \leq i \leq n \end{aligned}$$

which share no variables
$b = b_0 \cup \{\tilde{X}_1 = \tilde{t}_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\} \cup \{b_1, \ldots, b_n\} \cup \{\tilde{X} = \tilde{t}\}$
$u = u_0 \cup \{u_1, \ldots, u_n\}$
b, u *is* $\Re^w$*-solvable* $\}$

The fixpoint semantics $\mathcal{F}ix(P)$ *is defined as* $\mathcal{F}ix(P) = T_P^U \uparrow n$.

The set of answers to a query T on a program P is computed by the Set^U operator defined as

$$\mathsf{Set}^U(T, P) = \{\langle b_j, u_j\rangle \mid T \mapsto^* \langle b_j, u_j\rangle\}$$

Such answers can be computed by means of the fix-point semantics above, as can be proved by using the following theorem [12].

Theorem 2 *Let* P *be a program and* $G = (c_0, A_1, \ldots, A_n)$ *be a goal. Then* $G \mapsto^* c_{ans}$ *iff there exist* n *constrained atoms* $B_i \leftarrow c_i \in \mathcal{F}ix(P), i = 1 \ldots n$ *which share no variable with* G *and with each other, such that* $c_0 \cup c_1 \cup \ldots \cup c_n \cup \{A_1 = B_1\} \cup \ldots \cup \{A_n = B_n\}$ *and* c_{ans} *have the same solution for the variables in* G.

To apply the updates computed by the marking phase to the current EDB, the function Λ is used:

$$\Lambda(\mathrm{EDB}, u) = (\mathrm{EDB} \setminus \{p(\tilde{t}) \mid -p(\tilde{t}) \in u\}) \cup \{p(\tilde{t}) \mid +p(\tilde{t}) \in u\}$$

Finally, the semantics of a transaction is given by the following definition.

Definition 21 (U-Datalog semantics) *([22], definitions 7.2.6 and 7.2.7) Let* $DB_i = IDB \cup EDB_i$ *be the database. The semantics of a transaction* T *is given by the function*

$$Sem^U_{EDB_i, IDB}(T) = \mathcal{S}^U_{IDB}(T)(\langle\emptyset, EDB_i, \mathsf{Commit}\rangle)$$

where the function $\mathcal{S}^U_{IDB}(T) : \mathsf{Oss} \to \mathsf{Oss}$ *is defined as follows:*

If T *is a simple transaction then*

$$\mathcal{S}^U_{IDB}(T)(Oss_i) = \begin{cases} Oss_{i+1} & \textit{if OK} \\ \langle\emptyset, Oss_i.2, \mathsf{Abort}\rangle & \textit{otherwise} \end{cases}$$

where $Oss_{i+1} = \langle\{b_j \mid \langle b_j, u_j\rangle \in \mathsf{Set}^U(T, DB_i)\}, EDB_{i+1}, \mathsf{Commit}\rangle$, $EDB_{i+1} = \Lambda(EDB_i, \overline{u})$. *The condition OK expresses the fact that the set* $\overline{u} = \bigcup_j u_j\hat{b}_j$ *is consistent.* $u_j\hat{b}_j$ *denotes the ground updates obtained by substituting the variables in* u_j *with the ground terms associated with the variable in* b_j. $Oss_i.n$ *denotes the* n*-th component of the tuple* Oss_i.

If T *is a complex transaction* $(T_1; T_2)$*, then*

$$\mathcal{S}^U_{IDB}(T_1; T_2)(Oss_i) = \begin{cases} Oss_{i+2} & \textit{if OK} \\ \langle\emptyset, Oss_i.2, \mathsf{Abort}\rangle & \textit{otherwise} \end{cases}$$

where $Oss_{i+2} = \mathcal{S}^U_{IDB}(T_2)(Oss_{i+1})$ *and* $Oss_{i+1} = \mathcal{S}^U_{IDB}(T_1)(Oss_i)$*. The condition OK expresses the fact that* $Oss_{i+1}.3 = Oss_{i+2}.3 = \mathsf{Commit}$.

5.2 Equivalence in absence of conflicts.

Active-U-Datalog allows conflicting updates in the evaluation of the deductive part (definition 6), while U-Datalog forbids them. We want to prove that in absence of conflicts the two semantics coincide.

First, we define a transformation from an Active-U-Datalog program without active rules to a U-Datalog program (Definition 19).

Definition 22 *Let P be an Active-U-Datalog program without active rules and T a transaction. Then we call*

- P_{CLP} *the CLP($\mathcal{AD}$) version of P. $P_{CLP} = IDB_{CLP} \cup EDB_{CLP}$ and is defined as follows*
  - *if* $p(t_1, \ldots, t_n) \in EDB$ *then* $p(X_1, \ldots, X_n) \leftarrow X_1 = t_1, \ldots, X_n = t_n \in EDB_{CLP}$ *where $X_1, \ldots, X_n$ are distinct variables not occurring in* $\{t_1, \ldots, t_n\}$
  - *if* $p(t_1, \ldots, t_n) \leftarrow u_1, \ldots, u_n, B_1, \ldots, B_m \in IDB$ *then* $p(X_1, \ldots, X_n) \leftarrow X_1 = t_1, \ldots, X_n = t_n, u_1, \ldots, u_n, B_1, \ldots, B_m \in IDB_{CLP}$ *where $X_1, \ldots, X_n$ are distinct variables not occurring in the rule.*
- T_{CLP} *the CLP($\mathcal{AD}$) version of T*
  - *if T is a simple transaction, $T = u, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)$ then*

$$T_{CLP} \quad = \quad \tilde{X}_1 = \tilde{t}_1, \ldots, \tilde{X}_n = \tilde{t}_n, u, p_1(\tilde{X}_1), \ldots, p_n(\tilde{X}_n)$$

    where $\tilde{X}_1, \ldots, \tilde{X}_n$ are distinct variables not occurring in T.
  - *if $T = (T'; T'')$ then $T_{CLP} = (T'_{CLP}; T''_{CLP})$ with the further condition that the variables introduced in T'_{CLP} and in T''_{CLP} are all distinct and not occurring in T.*

We recall some definitions and notations of the theory of unification [19].

- Let E be a set of equations $\{s_1 = t_1, \ldots, s_n = t_n\}$, a *solution* θ for E is a grounding substitution for V (the set of variables in E) such that

$$s_1\theta \equiv t_1\theta, \ldots, s_n\theta \equiv t_n\theta.$$

  where $\equiv$ is the syntactic identity.
- $soln(E)$ denotes the (possibly empty) set of solutions of the set of equations E.
- Two sets of equations E_1 and E_2 are *equivalent* $E_1 \approx E_2$ if $soln(E_1) = soln(E_2)$.
- A set of equations $E = \{V_1 = t_1, \ldots, V_n = t_n\}$ is *solved* if V_i are distinct variables not occurring in any t_i.
- $vars(t)$ is the set of variables occurring in the term t. This definition is extended to sets of terms and of equations in the obvious way.
- Let $\theta = \{V_1 \leftarrow t_1, \ldots, V_n \leftarrow t_n\}$ be a substitution. $dom(\theta) = \{V_1, \ldots, V_n\}$ and $range(\theta) = vars(\{t_1, \ldots, t_n\})$.

Moreover, we state some theorems [19] that will be useful in the following.

Theorem 3 *Given a set of equations E, if E is solvable, then an equivalent solved set E' can be effectively computed.*

Theorem 4 *θ is an idempotent mgu of a set of equations E iff $eqn(\theta)$ is a solved set equivalent to E.*

Definition 23 *Let E be a solvable set of equations and E' an equivalent solved set. We define θ_E the substitution such that $eqn(\theta_E) = E'$.*

This definition is well formed because theorem 3 ensures the existence of a solved set E' equivalent to E.

Remark 1 *Let E be a set of equations, then θ_E is an idempotent mgu of E (theorem 4).*

We do not consider in the following general sets of equations, but only sets of bindings, i.e. equations of the form

$$X = Y \qquad \text{or} \qquad X = t$$

where X and Y are variables and t is a ground term.

Lemma 1 states that for every constrained atom derived from the CLP program, a corresponding constrained atom can be derived from P.

Lemma 1 *Let P be an Active-U-Datalog program without active rules. Then*

$$p(\tilde{X}) \leftarrow b, u \in Fix(P_{CLP}) \implies p(\tilde{X})\theta_b \leftarrow u\theta_b \in F(P)$$

Proof. We show that for all i

$$p(\tilde{X}) \leftarrow b, u \in T^U_{P_{CLP}} \uparrow i \implies p(\tilde{X})\theta_b \leftarrow u\theta_b \in T_P \uparrow i$$

and by continuity of $T^U_{P_{CLP}}$ and of T_P we can conclude. The proof is by induction on i.
(Base case). Trivial since $T^U_{P_{CLP}} \uparrow 0 = \emptyset$.
(Inductive case). Assume that

$$p(\tilde{X}) \leftarrow b, u \in T^U_{P_{CLP}} \uparrow i \implies p(\tilde{X})\theta_b \leftarrow u\theta_b \in T_P \uparrow i$$

Then:

$p(\tilde{X}) \leftarrow b, u \in T^U_{P_{CLP}} \uparrow i+1$
$\iff$ $\{$definition of $T^U_{P_{CLP}} \uparrow \alpha\}$
$\exists\, p(\tilde{t}) \leftarrow b_0, u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)$ in P_{CLP},
$\exists\, p_i(\tilde{X}_i) \leftarrow b_i, u_i \in T^U_{P_{CLP}} \uparrow i, \quad 1 \leq i \leq n$
which share no variables
$b = b_0 \cup \{\tilde{X}_1 = \tilde{t}_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\} \cup \{b_1, \ldots, b_n\} \cup \{\tilde{X} = \tilde{t}\}$
$u = u_0 \cup \{u_1, \ldots, u_n\}$
b, u is $\Re^w$-solvable
$\iff$ $\{$definition of P_{CLP} (def. 22) and b_0 is solvable, hence θ_{b_0} exists$\}$
$\exists\, p(\tilde{t})\theta_{b_0} \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)$ in P,
$\exists\, p_i(\tilde{X}_i) \leftarrow b_i, u_i \in T^U_{P_{CLP}} \uparrow i, \quad 1 \leq i \leq n$
which share no variables
$b = b_0 \cup \{\tilde{X}_1 = \tilde{t}_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\} \cup \{b_1, \ldots, b_n\} \cup \{\tilde{X} = \tilde{t}\}$
$u = u_0 \cup \{u_1, \ldots, u_n\}$
b, u is $\Re^w$-solvable
$\implies$ $\{$inductive hypothesis$\}$
$\exists\, p(\tilde{t})\theta_{b_0} \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)$ in P,
$\exists\, p_i(\tilde{X}_i)\theta_{b_i} \leftarrow u_i\theta_{b_i} \in T_P \uparrow i, \quad 1 \leq i \leq n$
which share no variables
$b = b_0 \cup \{\tilde{X}_1 = \tilde{t}_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\} \cup \{b_1, \ldots, b_n\} \cup \{\tilde{X} = \tilde{t}\}$
$u = u_0 \cup \{u_1, \ldots, u_n\}$
b, u is $\Re^w$-solvable
$\implies$ $\{\theta_b$ is an idempotent mgu of $\{(p_1(\tilde{X}_1)\theta_{b_1}, \ldots, p_n(\tilde{X}_n)\theta_{b_n}),$
$(p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n))\}\}$
$\exists\, p(\tilde{t})\theta_{b_0} \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)$ in P,
$\exists\, p_i(\tilde{X}_i)\theta_{b_i} \leftarrow u_i\theta_{b_i} \in T_P \uparrow i, \quad 1 \leq i \leq n$
which share no variables
$\theta_b = mgu((p_1(\tilde{X}_1)\theta_{b_1}, \ldots, p_n(\tilde{X}_n)\theta_{b_n}), (p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)))$

$$
\begin{array}{ll}
& u = u_0 \cup \{u_1, \ldots, u_n\} \\
& u\theta_b \text{ is consistent} \\
\Longrightarrow & \quad \{\theta_{b_0} \text{ is an idempotent mgu of } b_0,\ \theta_b \text{ is a unifier of } b_0, \\
& \quad \text{i.e. } \exists \theta'.\theta_b = \theta_{b_0} \circ \theta' \text{ and we apply } \theta_b \text{ to } p(\tilde{t})\theta_{b_0}\} \\
& \exists\ p(\tilde{t})\theta_{b_0} \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \text{ in P}, \\
& \exists\ p_i(\tilde{X}_i)\theta_{b_i} \leftarrow u_i\theta_{b_i} \in T_P \uparrow i, \quad 1 \leq i \leq n \\
& \text{which share no variables} \\
& \theta_b = mgu((p_1(\tilde{X}_1)\theta_{b_1}, \ldots, p_n(\tilde{X}_n)\theta_{b_n}), (p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n))) \\
& p(\tilde{t})\theta_{b_0}\theta_b \equiv p(\tilde{t})\theta_{b_0}\theta' \equiv p(\tilde{t})\theta_b \\
& u = u_0 \cup \{u_1, \ldots, u_n\} \\
& u\theta_b \text{ is consistent} \\
\Longrightarrow & \quad \{\theta_b \text{ is a unifier of } \{\tilde{X} = \tilde{t}\}, b_1, \ldots, b_n \text{ and} \\
& \quad \theta_{b_1}, \ldots, \theta_{b_n} \text{ are idempotent mgus }\} \\
& \exists\ p(\tilde{t})\theta_{b_0} \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \text{ in P}, \\
& \exists\ p_i(\tilde{X}_i)\theta_{b_i} \leftarrow u_i\theta_{b_i} \in T_P \uparrow i, \quad 1 \leq i \leq n \\
& \text{which share no variables} \\
& \theta_b = mgu((p_1(\tilde{X}_1)\theta_{b_1}, \ldots, p_n(\tilde{X}_n)\theta_{b_n}), (p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n))) \\
& p(\tilde{X})\theta_b \equiv p(\tilde{t})\theta_b \equiv p(\tilde{t})\theta_{b_0}\theta_b \\
& u\theta_b = u_0\theta_b, u_1\theta_b, \ldots, u_n\theta_b \equiv u_0\theta_b, u_1\theta_{b_1}\theta_b, \ldots, u_n\theta_{b_n}\theta_b \\
& u\theta_b \text{ is consistent} \\
\Longrightarrow & \quad \{\text{definition of } T_P\} \\
& p(\tilde{X})\theta_b \leftarrow u\theta_b \in T_P \uparrow i+1
\end{array}
$$

□

Lemma 2 shows the reverse inclusion with respect to the previous lemma. We do not prove that $\mathcal{F}(P)$ is included in $Fix(P_{CLP})$ because we insert in $\mathcal{F}(P)$ even constrained atom with inconsistent updates, but if we consider only the constrained atoms with consistent updates the set of derived atoms is *isomorphic* to $Fix(P_{CLP})$.

Lemma 2 *Let P be an Active-U-Datalog program without active rules. Then*

$$
p(\tilde{t}) \leftarrow u \in \mathcal{F}(P) \ \wedge \ u \textit{ consistent} \quad \Longrightarrow \quad \begin{array}{l} p(\tilde{X}) \leftarrow b, u' \in Fix(P_{CLP}) \ \wedge \\ \tilde{X}\theta_b \equiv \tilde{t} \wedge \ u \equiv u'\theta_b \end{array}
$$

Proof. We show that for all i

$$
p(\tilde{t}) \leftarrow u \in T_P \uparrow i \quad \wedge \quad u \text{ consistent} \quad \Longrightarrow \quad \begin{array}{l} p(\tilde{X}) \leftarrow b, u' \in T^U_{P_{CLP}} \uparrow i \ \wedge \\ \tilde{X}\theta_b \equiv \tilde{t} \wedge \ u \equiv u'\theta_b \end{array}
$$

and we conclude by continuity of T_P and of $T^U_{P_{CLP}}$.
The proof is by induction on i.
(Base case). Trivial since $T_P \uparrow 0 = \emptyset$.
(Inductive case). Assume that

$$
p(\tilde{t}) \leftarrow u \in T_P \uparrow i \quad \wedge \quad u \text{ consistent} \quad \Longrightarrow \quad \begin{array}{l} p(\tilde{X}) \leftarrow b, u' \in T^U_{P_{CLP}} \uparrow i \ \wedge \\ \tilde{X}\theta_b \equiv \tilde{t} \wedge \ u \equiv u'\theta_b \end{array}
$$

Then:

$$
\begin{array}{ll}
& p(\tilde{t}) \leftarrow u \in T_P \uparrow i+1 \quad \wedge \quad u \text{ consistent} \\
\Longleftrightarrow & \quad \{\text{definition of } T_P \uparrow \alpha\} \\
& \exists\ p(\tilde{t'}) \leftarrow u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \in P \\
& \exists\ p_r(\tilde{t'_r}) \leftarrow u_k \in T_P \uparrow i \quad 1 \leq r \leq n \\
& \text{which share no variables}
\end{array}
$$

$\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}'_1), \ldots, p_n(\tilde{t}'_n)))$
$\tilde{t} = \tilde{t}'\theta$
$u = (u_0, u_1, \ldots, u_n)\theta$ and u is consistent
$\Longrightarrow$ {definition of P_{CLP}}
$\exists\ p(\tilde{Y}) \leftarrow b_0, u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \in P_{CLP} \ \wedge\ b_0 = \{\tilde{Y} = \tilde{t}'\}$
$\exists\ p_r(\tilde{t}'_r) \leftarrow u_k \in T_P \uparrow i \quad 1 \leq r \leq n$
which share no variables
$\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}'_1), \ldots, p_n(\tilde{t}'_n)))$
$\tilde{t} = \tilde{t}'\theta$
$u = (u_0, u_1, \ldots, u_n)\theta$ and u is consistent
$\Longrightarrow$ {inductive hypothesis and for all $r = 1..n$, u_r is consistent since u is consistent}
$\exists\ p(\tilde{Y}) \leftarrow b_0, u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \in P_{CLP} \ \wedge\ b_0 = \{\tilde{Y} = \tilde{t}'\}$
$\exists\ p_r(\tilde{X}_r) \leftarrow b_r, u'_k \in T^U_{P_{CLP}} \uparrow i \quad 1 \leq r \leq n$
which share no variables
$X_r\theta_{b_r} \equiv \tilde{t}'_r$ and $u_k \equiv u'_k\theta_{b_r}$
$\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}'_1), \ldots, p_n(\tilde{t}'_n)))$
$\tilde{t} = \tilde{t}'\theta$
$u = (u_0, u_1, \ldots, u_n)\theta$ and u is consistent

At this point, we want to build a set of equations b and a set of updates u' which are $\Re^w$-solvable.

Let us consider the following set of equations:

$$b = b_0 \cup \{b_1, \ldots, b_n\} \cup \{\tilde{X}_1 = \tilde{t}_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\} \cup \{\tilde{X} = \tilde{Y}\}$$

where $\tilde{X}$ are new variables. b is solvable and in order to prove it, we show that

$$\overline{\theta} = \theta_{b_0} \circ \theta_{b_1} \circ \ldots \circ \theta_{b_n} \circ \theta \circ \theta_{\{\tilde{X} = \tilde{t}\}}$$

is a unifier of b. To obtain a solution of b, it is sufficient to consider a grounding instance of $\overline{\theta}$.

First we observe that by construction

- $dom(\theta_{b_j}) \cap vars(b_k) = \emptyset$ for all $j, k = 0..n \ \wedge\ j \neq k$
- $dom(\theta_{b_j}) \cap vars(\{\tilde{X}_k = \tilde{t}_k\}) = \emptyset$ for all $j, k.\ j = 0..n,\ k = 1..n,\ j \neq k$
- $dom(\theta_{b_j}) \cap vars(\{\tilde{X} = \tilde{Y}\}) = \emptyset$ for all $j = 1..n$

Therefore if we apply $\overline{\theta}$ to b we have:

$$\begin{aligned}
& b(\theta_{b_0} \circ \theta_{b_1} \circ \ldots \theta_{b_n} \circ \theta \circ \theta_{\{\tilde{X} = \tilde{Y}\}}) \\
& \equiv (b_0\theta_{b_0} \cup b_1\theta_{b_1} \cup \ldots \cup b_n\theta_{b_n} \cup \{\tilde{X}_1 = \tilde{t}_1\}\theta_{b_1} \cup \ldots \cup \{\tilde{X}_n = \tilde{t}_n\}\theta_{b_n} \cup \{\tilde{X} = \tilde{Y}\theta_{b_0}\})\theta \circ \theta_{\{\tilde{X} = \tilde{t}\}}
\end{aligned}$$

The following remarks allow us to conclude that $\overline{\theta}$ is a unifier of b.

- θ_{b_j} is a mgu of b_j for all $j = 0..n$,
- $dom(\theta_{b_j}) \cap vars(\tilde{t}_j) = \emptyset$ and $(\tilde{X}_j\theta_{b_j})\theta \equiv \tilde{t}_j\theta$ for all $j = 1..n$ and
- $vars(\tilde{X}) \cap dom(\theta) = \emptyset$, $vars(\tilde{Y}\theta_{b_0}) \cap dom(\theta_{\{\tilde{X} = \tilde{t}\}}) = \emptyset$ and $\tilde{t} \equiv (\tilde{Y}\theta_{b_0})\theta$

Actually, $\overline{\theta}$ is an idempotent mgu of b and unless renaming it is exactly θ_b. Therefore, $\tilde{X}\theta_b \equiv \tilde{t}$ trivially holds.

Let us consider, now, the following set of updates:

$$u' = u_0, u'_1, \ldots, u'_n$$

We must prove that $u'\theta_b$ is consistent. We first prove that the second syntactic identity of the thesis holds, that is $u \equiv u'\theta_b$. The consistency of $u'\theta_b$ follows from this syntactic identity.
By construction

- $dom(\theta_{b_j}) \cap vars(u'_k) = \emptyset$ for all $j, k = 1..n \ \wedge \ j \neq k$
- $dom(\theta_{b_j}) \cap vars(u_0) = \emptyset$ for all $j = 0..n$
- $dom(\theta_{\{\tilde{X}=\tilde{t}\}}) \cap (vars(u_0, u'_1, \ldots, u'_n)) = \emptyset$ and $vars(\tilde{X}) \cap dom(\theta) = \emptyset$

then

$$\begin{aligned}
& u'\theta_b \\
& \equiv u'(\theta_{b_0} \circ \theta_{b_1} \circ \ldots \theta_{b_n} \circ \theta \circ \theta_{\{\tilde{X}=\tilde{Y}\}}) \\
& \equiv u_0\theta \cup (u'_1\theta_{b_1})\theta \cup \ldots \cup (u'_n\theta_{b_n})\theta
\end{aligned}$$

By inductive hypothesis

$$u'_j\theta_{b_j} \equiv u_j$$

therefore

$$u'\theta_b \equiv u_0\theta, u_1\theta, \ldots, u_n\theta$$

We conclude by recalling that $u = (u_0, u_1, \ldots, u_n)\theta$
Collecting all the above results, we have

$\exists\ p(\tilde{Y}) \leftarrow b_0, u_0, p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n) \in P_{CLP}$
$\exists\ p_r(\tilde{X}_r) \leftarrow b_r, u'_k \in T^U_{P_{CLP}} \uparrow i$ for all $1 \leq r \leq n$
which share no variables
$b = b_0 \cup \{b_1, \ldots, b_n\} \cup \{X_1 = \tilde{t}_1\} \cup \ldots \cup \{X_n = \tilde{t}_n\} \cup \{\tilde{X} = \tilde{Y}\}$
$u' = (u_0, u'_1, \ldots, u'_n)$
b, u are $\Re^w$-solvable
$\tilde{X}\theta_b \equiv \tilde{t} \ \wedge \ u \equiv u'\theta_b$

hence

$$p(\tilde{X}) \leftarrow b, u' \in T^U_{P_{CLP}} \uparrow i+1 \ \wedge \ \tilde{X}\theta_b \equiv \tilde{t} \ \wedge \ u \equiv u'\theta_b$$

□

Definition 24 *Let T be a simple transaction, b and b' sets of bindings, and let $\langle \mathsf{Ans}, EDB, \mathsf{Res}\rangle$ and $\langle \mathsf{Ans}', EDB', \mathsf{Res}'\rangle$ be observables. We define*

$$\begin{aligned}
b \approx_T b' \quad & \textit{iff} \quad b \approx b' \textit{ restricted to the variables in } T \\
\langle \mathsf{Ans}, EDB, \mathsf{Res}\rangle \approx_T \langle \mathsf{Ans}', EDB', \mathsf{Res}'\rangle \quad & \textit{iff} \quad \mathsf{Ans} \approx_T \mathsf{Ans}',\ EDB = EDB' \textit{ and } \mathsf{Res} = \mathsf{Res}'
\end{aligned}$$

Remark 2 *In our semantics any transaction succeeds, therefore*

- *if T is a simple transaction*

$$Sem_{P,\mathsf{sel}}(T) \dot{\approx}_T\ Oss \quad \Longrightarrow \quad Oss.3 = \mathsf{Commit}$$

- *if $T = T_1; \ldots; T_n$*

$$Sem_{P,\mathsf{sel}}(T) \dot{\approx}_{T_n}\ Oss \quad \Longrightarrow \quad Oss.3 = \mathsf{Commit}$$

The same property holds for $S^{\mathcal{DA}}_{IDB,AR,\mathsf{sel}}$.

Lemmas 3 and 4 show that the inclusion relations holding for $\mathcal{F}(P)$ and $Fix(P_{CLP})$ carry over to the respective sets of computed answers and requested updates obtained by Set and the CLP version of Set (Set^U). Corollary 1 presents a sufficient condition that proves that the computed answers are equivalent ($\approx_T$) and the requested updates are equal.

Lemma 3 *Let P be an Active-U-Datalog program without active rules and T a simple transaction. Then*

$$\langle b, u\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP}) \implies \langle eqn(\theta_b)_{|T}, u\theta_b\rangle \in \mathsf{Set}(T, P)$$

Proof. Suppose that $T = \overline{u}_0, p_1(\tilde{t_1}), \ldots, p_n(\tilde{t_n})$.
By definition 22, $T_{CLP} = (\tilde{X_1} = \tilde{t_1}, \ldots, \tilde{X_n} = \tilde{t_n}, \overline{u}_0, p_1(\tilde{X_1}), \ldots, p_n(\tilde{X_n}))$. Let b_0 be $\{\tilde{X_1} = \tilde{t_1}, \ldots, \tilde{X_n} = \tilde{t_n}\}$.

$\langle b, u\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP})$
$\iff$ {Theorem 2}
$\exists\ p_i(\tilde{X_i}') \leftarrow b_i, \overline{u}_i \in \mathcal{F}ix(P_{CLP}) \quad 1 \le i \le n$
$b \equiv b_0 \cup b_1 \cup \ldots \cup b_n \cup \{p_1(\tilde{X_1}) = p_1(\tilde{X_1}')\} \cup \ldots \cup \{p_n(\tilde{X_n}) = p_n(\tilde{X_n}')\}$
$u = \overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$
b, u is $\Re^\omega$-solvable
$\implies$ {lemma 1}
$\exists\ p_i(\tilde{X_i}')\theta_{b_i} \leftarrow \overline{u}_i\theta_{b_i} \in \mathcal{F}(P) \quad 1 \le i \le n$
$b \equiv b_0 \cup b_1 \cup \ldots \cup b_n \cup \{p_1(\tilde{X_1}) = p_1(\tilde{X_1}')\} \cup \ldots \cup \{p_n(\tilde{X_n}) = p_n(\tilde{X_n}')\}$
$u = \overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$
b, u is $\Re^\omega$-solvable
$\implies$ {θ_b idempotent mgu of b and $p_i(\tilde{X}_i)\theta_{b_0} \equiv p_i(\tilde{t_i})$ for all $i = 1..n$}
$\exists\ p_i(\tilde{X_i}')\theta_{b_i} \leftarrow \overline{u}_i\theta_{b_i} \in \mathcal{F}(P) \quad 1 \le i \le n$
$\theta_b = \mathrm{mgu}((p_1(\tilde{X}_1)\theta_{b_0}, \ldots, p_n(\tilde{X}_n)\theta_{b_0}), (p_1(\tilde{X_1}')\theta_{b_1}, \ldots, p_n(\tilde{X_n}')\theta_{b_n})$
$u = \overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n$
$u\theta_b$ is consistent
$\implies$ {θ_b unifier of $b_1, \ldots, b_n$ and $\theta_{b_1}, \ldots, \theta_{b_n}$ are idempotent mgu respectively of $b_1, \ldots, b_n$}
$\exists\ p_i(\tilde{X_i}')\theta_{b_i} \leftarrow \overline{u}_i\theta_{b_i} \in \mathcal{F}(P) \quad 1 \le i \le n$
$\theta_b = \mathrm{mgu}((p_1(\tilde{X}_1)\theta_{b_0}, \ldots, p_n(\tilde{X}_n)\theta_{b_0}), (p_1(\tilde{X_1}')\theta_{b_1}, \ldots, p_n(\tilde{X_n}')\theta_{b_n})$
$u\theta_b = \overline{u}_0\theta_b, \overline{u}_1\theta_b, \ldots, \overline{u}_n\theta_b = (\overline{u}_0, \overline{u}_1\theta_{b_1}, \ldots, \overline{u}_n\theta_{b_n})\theta_b$
$\implies$ {Definition of Set}
$\langle eqn(\theta_b)_{|T}, u\theta_b\rangle \in \mathsf{Set}(T, P)$

□

Lemma 4 *Let P be an Active-U-Datalog program without active rules and T a simple transaction. Then*

$$\langle b, \overline{u}\rangle \in \mathsf{Set}(T, P) \ \wedge\ \overline{u} \text{ is consistent} \implies \langle b', \overline{u}'\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP}) \ \wedge\ b \approx_T b' \ \wedge\ \overline{u}'\theta_{b'} \equiv \overline{u}$$

Proof. Suppose that $T = \overline{u}_0, p_1(\tilde{t_1}), \ldots, p_n(\tilde{t_n})$. By definition 22, $T_{CLP} = (\tilde{X_1} = \tilde{t_1}, \ldots, \tilde{X_n} = \tilde{t_n}, \overline{u}_0, p_1(\tilde{X_1}), \ldots, p_n(\tilde{X_n}))$. Let b_0 be $\{\tilde{X_1} = \tilde{t_1}, \ldots, \tilde{X_n} = \tilde{t_n}\}$.

$\langle b, \overline{u}\rangle \in \mathsf{Set}(T, P) \ \wedge\ \overline{u}$ is consistent
$\iff$ {Definition 6}
$p_i(\tilde{t}_i') \leftarrow \overline{u}_i \in \mathcal{F}(P), \quad 1 \le i \le n$
$\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}_1'), \ldots, p_n(\tilde{t}_n'))$
$b = eqn(\theta)_{|T}$
$\overline{u} = (\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n)\theta$
$\overline{u}$ is consistent
$\implies$ {Lemma 2 and $\overline{u}_1, \ldots, \overline{u}_n$ are consistent since $\overline{u}$ is consistent }
$p_i(\tilde{X}_i') \leftarrow b_i, \overline{u}_i' \in Fix(P_{CLP}) \ \wedge\ \tilde{X}_i'\theta_{b_i} \equiv \tilde{t}_i' \ \wedge\ \overline{u}_i \equiv \overline{u}_i'\theta_{b_i}, \quad 1 \le i \le n$
$\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}_1'), \ldots, p_n(\tilde{t}_n')))$
$b = eqn(\theta)_{|T}$
$\overline{u} = (\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n)\theta$

$\overline{u}$ is consistent
$\Longrightarrow \quad \{\ \tilde{X}_i\theta_{b_0} \equiv \tilde{t}_i,\ \tilde{t}'_i \equiv \tilde{X}'_i\theta_{b_i}$ and $\theta = mgu((p_1(\tilde{t}_1), \ldots, p_n(\tilde{t}_n)), (p_1(\tilde{t}'_1), \ldots, p_n(\tilde{t}'_n)))\}$
$p_i(\tilde{X}'_i) \leftarrow b_i, \overline{u}'_i \in Fix(P_{CLP}) \ \wedge\ \overline{u}_i \equiv \overline{u}'_i\theta_{b_i}, \quad 1 \leq i \leq n$
$b' = b_0 \cup b_1 \cup \ldots \cup b_n \cup \{\tilde{X}_1 = \tilde{X}'_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{X}'_n\}$
$eqn(\theta) \approx b'$
$\overline{u} = (\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n)\theta$
$\overline{u}$ is consistent
$\Longrightarrow \quad \{\theta_{b'}$ is a unifier of $b_1, \ldots, b_n$ and $\overline{u}_i \equiv \overline{u}'_i\theta_{b_i}$ for all $i = 1..n\}$
$p_i(\tilde{X}'_i) \leftarrow b_i, \overline{u}'_i \in Fix(P_{CLP}), \quad 1 \leq i \leq n$
$b' = b_0 \cup b_1 \cup \ldots \cup b_n \cup \{\tilde{X}_1 = \tilde{X}'_1\} \cup \ldots \cup \{\tilde{X}_n = \tilde{X}'_n\}$
$\overline{u}' = \overline{u}_0, \overline{u}'_1, \ldots, \overline{u}'_n$
$\overline{u}'\theta_{b'} \equiv \overline{u}_0\theta_{b'}, (\overline{u}'_1\theta_{b_1})\theta_{b'}, \ldots, (\overline{u}'_n\theta_{b_n})\theta_{b'} \equiv (\overline{u}_0, \overline{u}_1, \ldots, \overline{u}_n)\theta_{b'}$
$eqn(\theta) \approx b'$
$\Longrightarrow \quad \{$definition Set^U and $eqn(\theta) \approx b'\}$
$\langle b', \overline{u}' \rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP}) \ \wedge\ b \approx_T b' \ \wedge\ \overline{u}'\theta_{b'} \equiv \overline{u}$

□

Corollary 1 *Let P be an active-U-Datalog program without active rules and T be a simple transaction. Then*

$$\left(\bigcup_{\langle b,\overline{u}\rangle \in \mathsf{Set}(T,P)} \overline{u}\right) \textit{ is consistent } \Longrightarrow$$
$$\{\overline{u} \mid \langle b, \overline{u}\rangle \in \mathsf{Set}(T,P)\} = \{\overline{u}'\hat{b}' \mid \langle b', \overline{u}'\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP})\}$$
$$\{b \mid \langle b, \overline{u}\rangle \in \mathsf{Set}(T,P)\} \approx_T \{b' \mid \langle b', \overline{u}'\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP})\}$$

Proof. Immediate consequence of lemmas 1 and 2. □

Theorem 5 states that the CLP semantics of a simple transaction is equivalent to the semantics we have given. To prove this theorem, we will use the following lemma.

Lemma 5 *Given an Active-U-Datalog program P without active rules (with intensional database IDB), a simple transaction T, a conflict resolution policy* sel*, and an observable Oss then*

$$\overline{U} \textit{ consistent and ground } \Longrightarrow \mathcal{S}^{\mathcal{DA}}_{IDB,\emptyset,\mathsf{sel}}(T)(Oss) \mathrel{\dot{\approx}_T} \mathcal{S}^U_{IDB_{CLP}}(T_{CLP})(Oss)$$

where

$$\overline{U} \ = \ \bigcup \overline{u}_j \textit{ such that } \langle b_j, \overline{u}_j\rangle \in \mathsf{Set}(T,P)$$

Proof. By hypothesis $\overline{U}$ is consistent and ground, then

$$\mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}(T)(Oss) = \langle \mathsf{Ans}, \mathsf{incorp}(I, \mathrm{EDB}), \mathsf{Commit}\rangle$$

By definition

$$\mathsf{Ans} = \{b_j \mid \langle b_j, \overline{u}_j\rangle \in \mathsf{Set}(T,P)\}$$

and by corollary 1

$$\mathsf{Ans} \approx_T \{b' \mid \langle b', \overline{u}'\rangle \in \mathsf{Set}^U(T_{CLP}, P_{CLP})\}$$

Now we compute the i-interpretation I obtained by

$$\langle B, I\rangle = \Delta^{\omega}_{\widehat{(\mathrm{IDB})}_{\overline{U}},\mathsf{sel},\mathrm{EDB}}(\langle\emptyset, \mathrm{EDB}\rangle)$$

The set of rules used by the Δ operator during the computation is the union of the set of purified rules $\widehat{\mathrm{IDB}}$ and the set of updates as rules, i.e. $\{\ \rightarrow \pm a \mid \pm a \in \overline{U}\}$. Since $\overline{U}$ is consistent and ground, all the heads of updates as rules belong to $\Delta^{\omega}_{\widehat{(\mathrm{IDB})_{\overline{U}}},\mathsf{sel},\mathrm{EDB}}(\langle\emptyset,\mathrm{EDB}\rangle)$. Indeed, such rules are the only ones that add (Π^u,Σ,V)-atoms to the i-interpretation. Therefore,

$$\mathsf{incorp}(I,\mathrm{EDB}) = (\mathrm{EDB}\setminus\{a \mid -a\in\overline{U}\})\cup\{a \mid +a\in\overline{U}\}$$

Let us consider the U-Datalog semantics. Since by corollary 1

$$\{\overline{u}_j \mid \langle b_j,\overline{u}_j\rangle\in\mathsf{Set}(T,P)\} = \{\overline{u}'_j\hat{b}'_j \mid \langle b'_j,\overline{u}'_j\rangle\in\mathsf{Set}^U(T_{CLP},P_{CLP})\}$$

then $\bigcup_j \overline{u}'_j\hat{b}'_j = \bigcup_j \overline{u}_j = \overline{U}$, hence $\bigcup_j \overline{u}'_j\hat{b}'_j$ is consistent and ground and

$$\mathcal{S}^U_{\mathrm{IDB}_{CLP}}(T_{CLP})(Oss) = \langle\{b_j \mid \langle b_j,u_j\rangle\in\mathsf{Set}^U(T,P_{CLP})\},\mathrm{EDB}',\mathsf{Commit}\rangle$$

where $\mathrm{EDB}'=\Lambda(\mathrm{EDB},\overline{U})$.
By definition

$$\Lambda(\mathrm{EDB},\overline{U}) = (\mathrm{EDB}\setminus\{p(\tilde{t}) \mid -p(\tilde{t})\in\overline{U}\})\cup\{p(\tilde{t}) \mid +p(\tilde{t})\in\overline{U}\}$$

This allows us to conclude. □

Theorem 5 *Given an Active-U-Datalog program P without active rules (with extensional database EDB and intensional database IDB), a simple transaction T, and a conflict resolution policy* sel, *then*

$$\overline{U}\ \textit{consistent and ground} \implies Sem_{P,\mathsf{sel}}(T)\ \dot{\approx}_T\ Sem^U_{EDB,IDB_{CLP}}(T_{CLP})$$

where

$$\overline{U} = \bigcup \overline{u}_j \textit{ such that } \langle b_j,\overline{u}_j\rangle\in\mathsf{Set}(T,P)$$

Proof. By definitions 18 and 21

$$Sem_{P,\mathsf{sel}}(T)\ \dot{\approx}_T\ Sem^U_{\mathrm{EDB},\mathrm{IDB}_{CLP}}(T_{CLP}) \iff$$

$$\mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}(T)(\langle\emptyset,\mathrm{EDB},\mathsf{Commit}\rangle)\ \dot{\approx}_T\ \mathcal{S}^U_{\mathrm{IDB}_{CLP}}(T_{CLP})(\langle\emptyset,\mathrm{EDB},\mathsf{Commit}\rangle)$$

We can conclude since the latter equivalence trivially holds by lemma 5. □

The above theorem proves that the two semantics are equivalent with respect to simple transactions; let us now turn our attention to complex ones. By noting that the sequential composition operator ";" is associative, we can always consider a complex transaction as a sequence of simple transactions $T=T_1;T_2;\ldots;T_n$. We denote with $P_i=\mathrm{EDB}_i\cup\mathrm{IDB}$ the Active-U-Datalog program (without active rules) obtained after the execution of $T_1;\ldots;T_i$.

Theorem 6 *Given $P_0 = EDB_0\cup IDB$, $T=T_1;\ldots;T_n$, and a conflict resolution policy* sel, *then*

$$\textit{for all } 1\le i\le n,\ \overline{U}_i\ \textit{consistent and ground} \implies Sem_{P_0,\mathsf{sel}}(T)\ \dot{\approx}_{T_n}\ Sem^U_{EDB_0,IDB_{CLP}}(T_{CLP})$$

where

$$\overline{U}_i = \bigcup \overline{u}_j \textit{ such that } \langle b_j,\overline{u}_j\rangle\in\mathsf{Set}(T_i,P_{i-1})$$

Proof. Let $Oss_0=\langle\emptyset,\mathrm{EDB}_0,\mathsf{Commit}\rangle$. The proof is by induction on the length of the complex transaction.
(Base case). $n=1$. The thesis follows directly from theorem 5.
(Inductive case). Assume that *for all $k<n$*

$$\overline{U}_i \text{ consistent and ground for all } 1\le i\le k \implies$$
$$Sem_{P_0,\mathsf{sel}}(T_1;\ldots;T_k)\ \dot{\approx}_{T_k}\ Sem^U_{\mathrm{EDB}_0,\mathrm{IDB}_{CLP}}((T_1;\ldots;T_k)_{CLP})$$

If $\overline{U}_i$ is consistent and ground for all $1 \leq i \leq n$
then

$$
\begin{array}{ll}
& Sem_{P_0,\mathsf{sel}}(T_1;\ldots;T_n) \\
= & \quad \{\text{definition 18 and associativity of ;}\} \\
& \mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}((T_1;\ldots;T_{n-1});T_n)(Oss_0) \\
= & \quad \{\text{definition 18}\} \\
& \mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}(T_n)(\mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}(T_1;\ldots;T_{n-1})(Oss_0)) \\
\dot{\approx}_{T_n} & \quad \{\text{Lemma 5 and } \overline{U}_n \text{ is consistent and ground}\} \\
& \mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_n)_{CLP})(\mathcal{S}^{\mathcal{DA}}_{\mathrm{IDB},\emptyset,\mathsf{sel}}(T_1;\ldots;T_{n-1})(Oss_0)) \\
= & \quad \{\text{definition 18}\} \\
& \mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_n)_{CLP})(Sem_{P_0,\mathsf{sel}}(T_1;\ldots;T_{n-1})) \\
\dot{\approx}_{T_n} & \quad \{\text{Inductive hypothesis and } \overline{U}_i \text{ are consistent and ground for } 1 \leq i \leq n-1\} \\
& \mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_n)_{CLP})(Sem^{U}_{\mathrm{EDB}_0,\mathrm{IDB}_{CLP}}((T_1;\ldots;T_{n-1})_{CLP})) \\
= & \quad \{\text{definition 21}\} \\
& \mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_n)_{CLP})(\mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_1;\ldots;T_{n-1})_{CLP})(Oss_0)) \\
= & \quad \{\text{definition 21 and by remark 2, } (\mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_1;\ldots;T_{n-1})_{CLP})(Oss_0)).3 = \mathsf{Commit} \\
& \quad \text{and } (\mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_n)_{CLP})(\mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_1;\ldots;T_{n-1})_{CLP})(Oss_0))).3 = \mathsf{Commit}\} \\
& \mathcal{S}^{U}_{\mathrm{IDB}_{CLP}}((T_1;\ldots;T_n)_{CLP})(Oss_0) \\
= & \quad \{\text{definition 21}\} \\
& Sem^{U}_{\mathrm{EDB}_0,\mathrm{IDB}_{CLP}}((T_1;\ldots;T_n)_{CLP})
\end{array}
$$

□

6 Conclusions

In this paper we have shown a conservative extension to U-Datalog with a support for active rules. While the full deductive power of U-Datalog is still intact, active rules allow for clear and efficient expression of integrity constraints and other active policies, moving most logic from applications into the database.

The fact that this extension does not compromise the deductive aspects of U-Datalog let us think that the active and deductive paradigms are more like orthogonal axes rather than to ends of a spectrum as proposed in [29], and that the two approaches can be mixed in a careful design without hindering each other.

More work is needed to enrich the set of events to which an Active-U-Datalog program can react (HiPAC [8, 21], for example, can react to database operations, methods invocations, transaction outcomes, temporal events, external messages and to almost arbitrarily complex combinations of these events) but this can be best done in designing a complete DBMS rather than in giving semantics to a language. Another promising extension is the integration of Active-U-Datalog with Object-U-Datalog [4]. Such a language would offer clean modularity in addition to the advantages of deductive and active languages.

Acknowledgments: We thank P. Baldan for his comments on earlier versions of this paper and E. Bertino for pointing us to the integration of active and deductive databases.

References

[1] S. Abiteboul and E. Simon. Fundamental properties of deterministic and nondeterministic extensions of Datalog. *Theoretical Computer Science*, (78):137–158, 1991.

[2] ANSI TC X3H2 and ISO/IEC JTC 1/SC 21/WG 3. Master index for SQL and all its parts, Mar. 1966. Document X3H2-96-066 DBL:MCI-011.

[3] ASK Computer Co. *INGRES/SQL Reference Manual*.

[4] E. Bertino, G. Guerrini, and D. Montesi. Toward deductive object databases. *Theory and Practice of Object Systems*, 1(1), 1995.

[5] E. Bertino and D. Montesi. Modeling database updates with constraint logic programming. In U. W. Lipeck and B. Thalheim, editors, *Proc. of the Fourth International Workshop on Foundations of Models and Languages for Data and Objects*, pages 120–132, 1992.

[6] S. Ceri, G. Gottlob, and L. Tanca. *Logic Programming and Databases.* Springer-Verlag, Berlin, 1990.

[7] S. Ceri and J. Widom. Deriving incremental production rules for deductive data. *Information Systems*, 19(6):467–490, 1994.

[8] S. Chakravarthy et al. HiPAC: A research project in active, time-constrained database management (final report). Technical Report XAIT-89-02, Xerox Advanced Information Technology, Cambridge, Massachussetts, Aug. 1990.

[9] L. M. L. Delcambre and J. N. Etheredge. The relational procedural language: A production language for relational databases. In *Proc. of the Second International Conference on Expert Database Systems*, Apr. 1988.

[10] Digital Equipment Corporation. *Rdb/VMS - SQL Reference Manual*, Nov. 1991.

[11] A. A. A. Fernandes, M. H. Williams, and N. W. Patton. A logic-based integration of active and deductive databases. *New Generation Computing*, 1996.

[12] M. Gabbrielli and G. Levi. Modeling answer constraints in constraint logic programs. In K. Furukawa, editor, *Proc. of the Eighth International Conference on Logic Programming*, pages 238–252, Cambridge, Mass., 1991. The MIT Press.

[13] N. Gehani and H. V. Jagadish. Ode as an active database: Constraints and triggers. In *Proc. of the Seventeenth International Conference on Very Large Data Bases*, Sept. 1991.

[14] G. Gottlob, G. Moerkotte, and V. S. Subrahmanian. The PARK semantics for active rules. In *Proceedings of the fifth International Conference on Extending Database Technology*, Lecture Notes in Computer Science Nr. 1057, pages 35–55. Springer–Verlag, Mar. 1996.

[15] E. Hanson. Rule condition testing and action execution in Ariel. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 49–58, San Diego, CA, June 1992.

[16] J. Jaffar and J. Lassez. Constraint logic programming. In *Proc. of the Fourteenth Annual ACM Symp. on Principles of Programming Languages*, pages 111–119, New York, USA, 1987. ACM.

[17] J. Kiernan, C. de Maindreville, and E. Simon. Making deductive databases a practical technology: A step forward. In *Proc. of the ACM SIGMOD International Conference on the Management of Data*, pages 237–246, Atlantic City, New Jersey, May 1990.

[18] Kirkwood. *Sybase Architecture and Administration*. Prentice-Hall, 1993.

[19] J. L. Lassez, M. J. Maher, and K. Marriott. Unification revisited. In M. Boscarol, L. C. Aiello, and G. Levi, editors, *Foundations of Logic and Functional Programming*, number 306 in LNCS, pages 67–113. Springer-Verlag, 1987.

[20] B. Ludäscher, W. May, and G. Lausen. Nested Transactions in a Logical Languages for Active Rules. In *Proceedings of the International Workshop on Logic in Databases*, Lecture Notes in Computer Science Nr. 1154, pages 197–222. Springer–Verlag, 1996.

[21] D. R. McCarhy and U. Dayal. The architecture of an active database management system. In *Proc. of the ACM SIGMOD International Conference on Extending Data Base Technology*, Mar. 1989.

[22] D. Montesi. *A Model for Updates and Transactions in Deductive Databases.* PhD thesis, Dipartimento di Informatica – Università di Pisa, 1993.

[23] D. Montesi and R. Torlone. A transaction transformation approach to active rule processing. In *Proc. of the Eleventh International Conference on Data Engineering*, pages 109–116, Mar. 1995.

[24] Oracle Corp. *Oracle 7 Server Concepts Manual, 7.3.* Oracle Corp., Feb. 1996.

[25] M. Stonebraker, A. Jhingran, J. Goh, and S. Potamianos. On rules, procedures, caching and views in data base systems. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 281–290, 1990.

[26] M. Stonebraker and G. Kemnitz. The POSTGRES next-generation database management system. *Communications of the ACM*, 34(10), Oct. 1991.

[27] L. Tanca. (re-)action in deductive databases. In *Proc. of the Second International Workshop on Intelligent and Cooperative Information Systems*, pages 55–61, Como, Italy, May 1990.

[28] J. D. Ullman. *Principles of Database and Knowledge-Base Systems*, volume I and II. Computer Science Press, Rockville, Maryland, 1989.

[29] J. Widom. Deductive and active databases: Two paradigms or ends of a spectrum? In *Proc. of the First International Workshop on Rules in Database Systems*, pages 306–315, Edimburgh, Scotland, Aug. 1993.

[30] J. Widom and S. Ceri, editors. *Active Database Systems - Triggers and rules for advanced database processing*. Morgan Kaupfmann, 1996.

[31] J. Widom, R. J. Cochrane, and B. G. Lindsay. Implementing set-oriented production rules as an extension to Starburst. In *Proc. of the Seventeenth International Conference on Very Large Data Bases*, Sept. 1991.

[32] J. Widom and S. Finkelstein. Set-oriented production rules in relational databases. In *Proc. of the ACM SIGMOD Conference on Management of Data*, pages 259–270, 1990.

[33] C. Zaniolo. A unified semantics for active and deductive databases. In *Proc. of the First International Workshop on Rules in Database Systems*, pages 271–287, Edimburgh, Scotland, Aug. 1993.

[34] C. Zaniolo, N. Arni, and K. Ong. Negation and aggregates in recursive rules. In *Proceedings of the third International Conference on Deductive and O-O DBs*, Phoenix, Arizona, 1993.