

# Active Database Rules with Transaction Conscious Stable Model Semantics

Carlo Zaniolo

Computer Science Department  
University of California  
Los Angeles, CA 90024  
zaniolo@cs.ucla.edu

## Abstract

Semantics represents a major problem area for active databases inasmuch as (i) there is no formal framework for defining the abstract semantics of active rules, and (ii) the various systems developed so far have ad-hoc operational semantics that are widely different from each other. This situation contributes to the difficulty of predicting the run-time behavior of sets of rules: thus, ensuring the termination of a given set of rules is currently recognized as a major research issue. This situation hampers the applicability of this powerful technology in critical application areas.

In this paper, we introduce a *durable change* semantics for active database rules; this semantics improves Starburst's *deferred activation* notion with concepts taken from Postgres and Heraclitus and the semantic foundations of deductive databases. We provide a formal logic-based model for this transaction-oriented semantics, show that it is amenable to efficient implementation, and prove that it solves the non-termination problem.

## 1 Introduction

Several active database languages and systems have been developed so far — a very incomplete list include [2, 5, 9]. Furthermore, active rules are now becoming a part of several commercial databases and of the SQL3 proposed standards. Indeed, active databases represent a powerful new technology that finds important application in the market place. However, this new technology is faced with several technical challenges; among these the lack of uniform and clear semantics stands out as one of the most and pressing and difficult problems [10]. The lack of formal models for characterizing the abstract semantics of active systems is a first facet of this problem. The second facet is represented by the differences between the many operational semantics proposed and implemented by the various systems in an ad-hoc fashion, with little progress towards unification and convergence. The result is that the behavior of complex rule sets is very difficult to predict, and critical questions such as confluence and termination are extremely hard to answer. These questions must be answered before active rules can be trusted with critical functions in an information system. Improved semantics and formal models for active databases, will solve these problem, and foster generalizations and simplifications that enhance the implementability, generality and intuitive appeal of active rules. The results presented in this paper illustrate the feasibility of achieving many of the desiderata of such an ideal scenario.

Let us consider the typical ECA rules of active databases:

## EVENT, CONDITION $\rightarrow$ ACTION

The basic structure of these rules make them different from those used in expert systems shells such as CLIPS and OPS5 (which follow a  $\text{CONDITION} \rightarrow \text{ACTION}$  pattern) or those used by deductive databases (that follow a  $\text{CONDITION} \rightarrow \text{CONDITION}$  pattern). But in addition to these differences, active rules are also unique inasmuch as their meaning is intertwined with the concept of database transactions. Therefore, an active database system must specify, whether the  $\text{ACTION}$  part of the rule is to be fired in the same transaction as the  $\text{EVENT}$  (*coupled semantics*) or as a separate transaction (*decoupled semantics*) [10]. Most systems adopt the coupled semantics, inasmuch as this is more effective at enforcing integrity constraints via active rules, and this framework will be adopted in this paper as well. Furthermore, while in the *immediate* interpretation of coupled semantics, rules are fired as soon as  $\text{EVENT}$  is detected, the *deferred* semantics used in Starburst is more transaction-conscious [9], inasmuch as it takes into account the fact that individual actions might be of no consequence before the transaction commits. For instance, the insertion of a record  $r$ , followed by the deletion of the same  $r$  within a transaction is *ephemeral*, inasmuch as the net effect of these two opposite actions (i.e., their composition) is null. Inasmuch as ephemeral actions leave no trace once the transaction completes, they should be disregarded by transaction-conscious ECA rules, which should instead be triggered only by actions that are durable, i.e., persist till the end of the transaction.

Thus, a critical contribution of Starburst is the notion that rules should be deferred until a *rule processing point*, which, by default, occurs at the end of the transaction, after regular actions have completed but before the transaction commits. At rule processing point, the net effect of all accumulated actions is computed using *composition* rules; for instance the net effect of an insert followed by an update on the same tuple is a modified insert [9].

While the basic idea of deferred semantics is both elegant and profound, there is the complication that many competing rules are fireable at once at rule processing point. At this time, the many changes requested by the transaction (often on several relations) could trigger several rules with incompatible action requests in their heads. The Starburst designers recognized the complexity of the situation, and the fact that different firing orders might lead to different results or to non-terminating programs [10]. Their proposed solution calls for a very smart designer, who after studying the rules to ensure termination steers the system clear of problems through explicit assignment of inter-rule precedence. This approach has several drawbacks, including the fact that termination analysis is exceedingly difficult, and that rules cannot be added or deleted independent of each other. Therefore, in this paper we take deferred semantics a step further and show that, with a simple extension, the system can solve the rule termination and priority-assignment problems for the designer.

## 2 An Example

Consider the following example. We have three relations:

```
Dept (D#, DName, Div, Loc)
EMP(E#, Ename, JobTitle, SAL, Dept)
HPaid(JobTitle)
```

`HPaid` is actually a derived relation, which stores those job titles for which there are, or have been employees, who make more than \$100,000. This concrete view is maintained via the rules `EMP_INSERT` and `EMP_UPDATE` specified as follows:

Rules EMP\_INSERT and EMP\_UPDATE. Upon an insertion into EMP or an update to EMP, the new SAL is checked, and if it exceeds \$100,000, then the JobTitle of this employee is added to HPaid, assuming that it was not there already.

There is also a foreign key constraint between EMP and Dept. This is supported by a rule DEPT\_DELETE that propagates deletions from Dept to EMP :

Rule DEPT\_DELETE: *When a tuple is deleted from Dept, then delete all employees who were working in the deleted department.*

Now assume that our transaction has executed the following actions (in the order listed):

- Change the location of every department who was in LA (Los Angeles) into SM (Santa Monica).
- Delete the department with D# = 1300 from the database.
- Give a raise of \$4,000 to all employees whose Job\_Title is analyst.

Say that in the initial database there are only two departments with location 'LA', say, one with D# = 1300 and the other with D# = 2500. Then, the Starburst *composition* semantics prescribes that the update on the Dept tuple with D# =1300 followed by the deletion of the same is equivalent to the deletion the original tuple. The update on the department tuple with D# = 2500 instead remains till the end of the transaction . Thus, the update on the tuple with D# =1300, which does not persist until the end of the transaction, will be called *ephemeral* while the second will be called *durable*. Therefore, at rule processing point only two change remain and the following two rules can be activated:

- Rule DEPT\_DELETE can be triggered by the resulting deletion of department with D# = 1300 (and Loc = 'LA').
- Rule EMP\_UPDATE can be triggered by the +\$4,000 salary update gotten by analysts.

The issue of which of the two rules above should be fired first is left by Starburst up to programmer, who can direct the system by assigning explicit precedence to the rules. However, consider the situation in which EMP\_UPDATE is fired before, or even at the same time as, DEPT\_DELETE. Say that there is only one network specialist, Bob White who now makes \$98,000. Then, with the \$4,000 raise the new salary exceeds the \$100,000 threshold and rule EMP\_UPDATE might add a new entry into HPaid. However, if Bob White happens to work for department 1300, then there is a problem. Once the rule DEPT\_DELETE fires the Bob White tuple is deleted and the \$4000 salary-raise becomes ephemeral, and, therefore, the addition of analyst to HPaid becomes totally unjustified— it should never have happened.

Therefore, we propose a semantics whereby only durable-change events can fire rules— ephemeral-change events cannot. Observe that this restriction produces a natural ordering between rules; in our case, the DEPT\_DELETE rule must be fired before the EMP\_UPDATE rule is considered for firing. Once the rules are processed in this order, for the database content previously described, then the update on of Bob White's record will be removed, and second rule will not fire at all.

There is also an obvious implication upon the termination problem, since the composition rules have the following property: every event that is followed by a later event on the same tuple is ephemeral. Therefore, each durable event on a tuple  $t$  is the final event on  $t$ . Since only final events can trigger rules, the computation cannot fall into an infinite loop.

In the rest of the paper we give a logic-based formalization to these simple concepts.

Figure 1: Three entries in the delta tables at rule-processing point

---

```

updDept0(2500, ims, 1000, 'LA',
          2500, ims, 1000, 'SM').
delDept0(1300, media, 1000, 'LA').
updEMP0(E2309, 'Bob White', analyst, 98000, 1300,
        e2309, 'Bob White', analyst, 102000, 1300).

```

---

### 3 Active Rules as Deductive Rules

We will now define a logic-based semantics for active rules with priority assignment. We use the notion of delta relations from Heraclitus [5], which we extend to handle updates along with inserts, deletes (the term "changes" will be used here to denote any of these three). Thus, for each relation  $R(X)$  in the database schema, where  $X$  denotes the attribute vector for  $R$ , we also keep three additional relations (*delta relations*):

$\text{insR}, \text{delR}, \text{updR}$

The delta relations  $\text{insR}(X)$  and  $\text{delR}(X)$  have the same attributes as the original  $R(X)$ . However,  $\text{updR}$  has an old-value and a new-value for each attributes of  $R$ , and can therefore be viewed according to the scheme  $\text{updR}(X_{\text{old}}, X_{\text{new}})$ . A sample of the initial value delta relations just before the rule-processing point is shown in Figure 1. The same tuple cannot appear in more than one delta table of the same relation.

We will use Datalog<sub>1S</sub> to model the state changes occurring in the various relations [4, 3]. In Datalog<sub>1S</sub>, tables and predicates are allowed to have an additional argument or column called the *stage argument*. The values in the stage argument are taken from the domain  $0, 0 + 1, 0 + 1 + 1, \dots$ , i.e., the integers generated by using the postfix successor function  $+1$ ; thus, the integer 3 is represented as  $0 + 1 + 1 + 1$ . Alternatively using the normal functional notation, the successor of  $J$  is denoted  $s(J)$ —this notation is at the root of the name Datalog<sub>1S</sub>. The merits of Datalog<sub>1S</sub> for modeling temporal and dynamic systems have been described in several papers [3, 14, 11]. Therefore, delta predicates with stage argument  $J$  have the form  $\text{insR}(J, X)$ ,  $\text{delR}(J, X)$  and  $\text{updR}(J, X_{\text{old}}, X_{\text{new}})$ . For notational convenience, we will instead write the stage argument as a superscript:  $\text{insR}^J(X)$ ,  $\text{delR}^J(X)$  and  $\text{updR}^J(X_{\text{old}}, X_{\text{new}})$ .

In addition to the delta relations, several auxiliary predicates are needed for each  $R$  in our database schema. In particular we need:

- *Initial Relation*:  $\text{iniR}(X)$ . This stores the value of  $R$  at the beginning of the transaction. It does not have a stage argument since it remains constant during the whole transaction.
- *Delta Relations*:  $\text{insR}^J(X)$ ,  $\text{delR}^J(X)$ ,  $\text{updR}^J(X_{\text{old}}, X_{\text{new}})$ .
- *Current Relations*:  $\text{curR}^J(X)$  represents the current content of relation  $R$ , as seen within the transaction. It is computed from the initial relation and the delta relations.
- *Action Request Relations*:  $\text{rinR}(X)$ ,  $\text{rdeR}^J(X)$ ,  $\text{rupR}^J(X_{\text{old}}, X_{\text{new}})$ . These contain the actions on  $R$  produced by fired active rules. Their union yields the *change request relation*  $\text{chrR}^J(X)$ . The union of these for all  $R$  in the schema produces the request relation  $\text{req}^J$ .

Figure 2: The current state of the database via frame axioms

---


$$\begin{aligned}
\text{curDept}^J(D, N, Dv, L) &\leftarrow \text{iniDept}(D, N, DvL), \text{levl}^J(-), \\
&\quad \neg \text{delDept}^J(D, N, Dv, L), \neg \text{updDept}^J(D, N, Dv, L, -, -, -, -), \\
&\quad \neg \text{ddeDept}^J(D, N, Dv, L), \neg \text{dupDept}^J(D, N, Dv, L, -, -, -, -). \\
\text{curDept}^J(D, N, Dv, L) &\leftarrow \text{insDept}^J(D, N, Dv, L). \\
\text{curDept}^J(D, N, Dv, L) &\leftarrow \text{updDept}^J(-, -, -, -, D, N, Dv, L). \\
\text{curDept}^J(D, N, Dv, L) &\leftarrow \text{ddeDept}^J(D, N, Dv, L). \\
\text{curDept}^J(D, N, Dv, L) &\leftarrow \text{dupDept}^J(-, -, -, -, D, N, Dv, L). \\
\text{curEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}) &\leftarrow \text{iniEMP}(\text{Eno}, N, \text{Jt}, \text{Dno}), \text{levl}^J(-) \\
&\quad \neg \text{delEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}), \neg \text{updEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}, -, -, -, -), \\
&\quad \neg \text{ddeEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}), \neg \text{dupEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}, -, -, -, -). \\
\text{curEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}) &\leftarrow \text{insEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}). \\
\text{curEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}) &\leftarrow \text{updEMP}^J(-, -, -, -, \text{Eno}, N, \text{Jt}, \text{Dno}). \\
\text{curEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}) &\leftarrow \text{dinEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}). \\
\text{curEMP}^J(\text{Eno}, N, \text{Jt}, \text{Dno}) &\leftarrow \text{dupEMP}^J(-, -, -, -, \text{Eno}, N, \text{Jt}, \text{Dno}). \\
\text{curHPaid}^J(\text{Jt}) &\leftarrow \text{iniHPaid}(\text{Jt}), \text{levl}^J(-), \neg \text{delHPaid}^J(\text{Jt}), \neg \text{updHPaid}^J(\text{Jt}, -). \\
&\quad \neg \text{ddeHPaid}^J(\text{Jt}), \neg \text{dupHPaid}^J(\text{Jt}, -). \\
\text{curHPaid}^J(\text{Jt}) &\leftarrow \text{insHPaid}^J(\text{Jt}). \\
\text{curHPaid}^J(\text{Jt}) &\leftarrow \text{updHPaid}^J(-, \text{Jt}). \\
\text{curHPaid}^J(\text{Jt}) &\leftarrow \text{dinHPaid}^J(\text{Jt}). \\
\text{curHPaid}^J(\text{Jt}) &\leftarrow \text{dupHPaid}^J(-, \text{Jt}).
\end{aligned}$$


---

- *Durable-change Relations*:  $\text{dinR}^J(X)$ ,  $\text{ddeR}^J(X)$ ,  $\text{dupR}^J(X_{\text{old}}, X_{\text{new}})$ . These contain all the changes assumed durable at step  $J$ .
- *A Current Level relation*:  $\text{levl}^J(Nr)$  with  $Nr$  the name of a rule. This predicate is used to enforce priorities between rules by denoting the rules that are currently active. The priorities between rules is represented by a binary **prec** relation.

We begin by computing the current value of the database as shown in Figure 2, using frame axioms. The current value of relation for  $R(X)$  is obtained by first subtracting from its initial value  $\text{iniR}$  the tuples deleted and the old values of tuples updated, and then adding the tuples inserted and the new values of tuples updated (for both delta relations and durable-change relations):

$$\begin{aligned}
\text{curR}^J(X) &\leftarrow \text{iniR}(X), \text{levl}^J(-). \\
&\quad \neg \text{delR}^J(X), \neg \text{updR}^J(X, \text{New}), \\
&\quad \neg \text{ddeR}^J(X), \neg \text{dupR}^J(X, \text{New}). \\
\text{curR}^J(X) &\leftarrow \text{insR}^J(X). \\
\text{curR}^J(X) &\leftarrow \text{updR}^J(\text{Old}, X). \\
\text{curR}^J(X) &\leftarrow \text{dinR}^J(X). \\
\text{curR}^J(X) &\leftarrow \text{dupR}^J(\text{Old}, X).
\end{aligned}$$

The next set of rules, called *action request* rules, capture the behavior of the actual active rules in the system.

Figure 3: Translations of Active Rules

DEPT\_DELETE:

$$\text{rdeEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}) \leftarrow \text{delDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}), \text{curEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}), \text{levl}^J(\text{dd}), \\ \neg \text{lchDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}).$$

EMP\_INSERT:

$$\text{rinHPaid}^J(\text{Jt}) \leftarrow \text{inseEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}), \text{S} > 100000, \neg \text{curHPaid}^J(\text{Jt}), \text{levl}^J(\text{ei}), \\ \neg \text{lchEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}).$$

EMP\_UPDATE:

$$\text{rinHPaid}^J(\text{Jt}) \leftarrow \text{updEMP}^J(\text{En}, \_, \text{Jt}, \text{So}, \_, \text{En}, \_, \text{Jt}, \text{Sn}, \_), \text{Sn} > 100000, \neg \text{curHPaid}^J(\text{Jt}), \text{levl}^J(\text{eu}), \\ \neg \text{lchEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}).$$

Figure 4: Changes assumed durable in firing the rules of Figure 3

DEPT\_DELETE:

$$\text{ddeDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}). \leftarrow \text{delDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}), \text{curEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}), \text{levl}^J(\text{dd}), \\ \neg \text{lchDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}).$$

EMP\_INSERT:

$$\text{dinEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}) \leftarrow \text{inseEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}), \text{S} > 100000, \neg \text{curHPaid}^J(\text{Jt}), \text{levl}^J(\text{ei}), \\ \neg \text{lchEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}).$$

EMP\_UPDATE:

$$\text{dupEMP}^J(\text{E}, \text{N}, \text{Jt}, \text{S}, \text{D}, \text{En}, \text{Nn}, \text{Jtn}, \text{Sn}, \text{Dn}) \leftarrow \text{updEMP}^J(\text{E}, \text{N}, \text{Jt}, \text{S}, \text{D}, \text{En}, \text{Nn}, \text{Jtn}, \text{Sn}, \text{Dn}), \text{Sn} > 100000, \\ \neg \text{curHPaid}^J(\text{Jt}), \text{levl}^J(\text{eu}), \neg \text{lchEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}).$$

Figure 5: Change Requests and later changes

$$\begin{aligned} \text{chrDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}) &\leftarrow \text{rinDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}). \\ \text{chrDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}) &\leftarrow \text{rdeDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}). \\ \text{chrDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}) &\leftarrow \text{rupdDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}, \_, \_, \_). \\ \text{lchDept}^I(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}) &\leftarrow \text{chrDept}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{DN}), \text{I} < \text{J}. \\ \\ \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}) &\leftarrow \text{rinEMP}^I(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}). \\ \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}) &\leftarrow \text{rdeEMP}^I(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}). \\ \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}) &\leftarrow \text{rupEMP}^I(\text{En}, \text{N}, \text{Jt}, \text{Dn}, \text{S}, \_, \_, \_). \\ \text{lchEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}) &\leftarrow \text{chrEMP}^I(\text{En}, \text{N}, \text{Jt}, \text{S}, \text{Dn}), \text{I} < \text{J}. \\ \\ \text{chrHPaid}^J(\text{Jt}) &\leftarrow \text{rinHPaid}^I(\text{Jt}). \\ \text{chrHPaid}^J(\text{Jt}) &\leftarrow \text{rdeHPaid}^I(\text{Jt}). \\ \text{chrHPaid}^J(\text{Jt}) &\leftarrow \text{rupHPaid}^I(\text{Jt}, \_). \\ \text{lchHPaid}^J(\text{Jt}) &\leftarrow \text{lchHPaid}^I(\text{Jt}), \text{I} < \text{J}. \end{aligned}$$

Obviously events are represented by the tuples of the delta relations, while conditions must be evaluated against the current relations. Finally the actions in the head of active rules are modeled by action requests. For instance, an immediate translation of rule DEPT\_DELETE is:

$$\begin{aligned} \text{rdeEMP}^{J+1}(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}) \leftarrow & \text{delDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}), \\ & \text{curEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}), \text{levl}^J(\text{dd}). \end{aligned}$$

This rule specifies that all the employees working in a certain department must be deleted once their department is in the delta relation. The rule will fire only if its proper level of precedence, i.e., only if  $\text{levl}^J(\text{dd})$  is true, where  $\text{dd}$  is just a shorter name for DEPT\_DELETE.

To express the durable-change semantics we need to add an additional goal  $\neg\text{lchDept}$  to ensure that the event triggering the rule is a durable one and will not be obliterated by later change requests ( “later” refer to stage values larger than the current ones). Thus our original rule becomes:

$$\begin{aligned} \text{rdeEMP}^{J+1}(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}) \leftarrow & \text{delDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}), \\ & \text{curEMP}^J(\text{En}, \text{E}, \text{JT}, \text{S}, \text{Dn}), \text{levl}^J(\text{dd}), \\ & \neg\text{lchDept}^J(\text{Dn}, \text{N}, \text{V}, \text{L}). \end{aligned}$$

This rule specifies that all the employees working in a certain department must also be deleted if their department is deleted. The translation of active rules for the example at hand is shown in Figure 3.

At each step, the firing of active rules might generate several action requests on  $\mathbf{R}$ . These have the form  $\text{rinR}$ ,  $\text{rdeR}$ ,  $\text{rupR}$ , respectively for tuples inserted, deleted or updated. Thus we have three rules as follows:

$$\begin{aligned} \text{chrR}^J(\mathbf{X}) \leftarrow & \text{rinR}^J(\mathbf{X}). \\ \text{chrR}^J(\mathbf{X}) \leftarrow & \text{rdeR}^J(\mathbf{X}). \\ \text{chrR}^J(\mathbf{X}) \leftarrow & \text{rupR}^J(\mathbf{X}, \text{New}). \end{aligned}$$

From these, we can now derive  $\text{lchR}^I(\mathbf{X})$  for values of  $I$  preceding the current stage value of  $J$ . (Say that the  $<$  relation between stage values is part of  $\text{Datalog}_{1S}$ , or alternatively that we define recursive rules to achieve the same effect.)

$$\text{lchR}^I(\mathbf{X}) \leftarrow \text{chrR}^J(\mathbf{X}), I < J.$$

Now, we have to use the composition rules to compose the action requests with old deltas yielding new deltas (of course new and old deltas are denoted by their respective stage values of  $J+1$  and  $J$ ). Basically there are three cases:

1. The action request  $\text{rinR}(\mathbf{X}), \text{rde}(\mathbf{X}), \text{rup}(\mathbf{X}, -)$ , does not compose with any object in the delta tables. In this case the action request is simply entered in the delta tables. Thus:

$$\begin{aligned} \text{insR}^{J+1}(\mathbf{X}) \leftarrow & \text{rinR}^J(\mathbf{X}), \neg\text{insR}^J(\mathbf{X}), \\ & \neg\text{delR}^J(\mathbf{X}), \neg\text{upd}(-, \mathbf{X}). \\ \text{delR}^{J+1}(\mathbf{X}) \leftarrow & \text{rdeR}^J(\mathbf{X}), \neg\text{insR}^J(\mathbf{X}), \\ & \neg\text{delR}^J(\mathbf{X}), \neg\text{upd}(-, \mathbf{X}). \\ \text{updR}^{J+1}(\mathbf{X}, \mathbf{Y}) \leftarrow & \text{rupR}^J(\mathbf{X}, \mathbf{Y}), \neg\text{insR}^J(\mathbf{X}), \\ & \neg\text{delR}^J(\mathbf{X}), \neg\text{upd}(-, \mathbf{X}). \end{aligned}$$

2. The second case concerns delta tuples that are neither moved to durable-change tables nor affected by the last action requests. These tuples are simply copied into the next-state delta tables. We also have added a  $\text{wt4}^J$  predicate to ensure that these rules do not fire until the current change-requests have been computed:

$$\begin{aligned}
\text{insR}^{J+1}(X) &\leftarrow \text{insR}^J(X), \text{wt4}^J, \\
&\quad \neg \text{dinR}(X), \neg \text{chrR}^J(X). \\
\text{delR}^{J+1}(X) &\leftarrow \text{delR}^J(X), \text{wt4}^J, \\
&\quad \neg \text{ddeR}(X), \neg \text{chrR}^J(X). \\
\text{updR}^{J+1}(X, Y) &\leftarrow \text{updR}^J(X, Y), \text{wt4}^J, \\
&\quad \neg \text{dupR}(X, Y), \neg \text{chrR}^J(X).
\end{aligned}$$

3. This is the situation where an object in the delta tables at stage  $J$  must be composed with action requests to yield an entry in the delta table at stage  $J + 1$ . In this case we have to apply the composition rules as follows:

$$\begin{aligned}
\%null &\leftarrow \text{insR}^J(X), \text{rdeR}^J(X). \\
\text{insR}^{J+1}(X_{\text{new}}) &\leftarrow \text{insR}^J(X), \text{rupR}^J(X, X_{\text{new}}). \\
\text{error} &\leftarrow \text{insR}^J(X), \text{rinR}^J(X). \\
\text{error} &\leftarrow \text{delR}^J(X), \text{rdeR}^J(X). \\
\text{error} &\leftarrow \text{delR}^J(X), \text{rupR}^J(X, Y). \\
\%null &\leftarrow \text{delR}^J(X), \text{rinR}^J(X). \\
\text{del}^{J+1}(X) &\leftarrow \text{updR}^J(X, Y), \text{rdeR}^J(Y). \\
\text{error} &\leftarrow \text{updR}^J(X), \text{rinR}^J(X). \\
\text{updR}^{J+1}(X_{\text{old}}, X_{\text{new}}) &\leftarrow \\
&\quad \text{updR}^J(X_{\text{old}}, X), \text{rupR}^J(X, X_{\text{new}}).
\end{aligned}$$

The set of rules for updating the delta relations for the example at hand is shown in Figures 6, 7 and 8.

Finally, we have the  $\text{prec}$  table that describes the (inverse) priority between rules and ensures that only the rules at the correct precedence level will fire. An entry  $\text{prec}(r_1, r_2)$  denotes that a rule at level  $r_2$  should fire only after all rules at level  $r_1$  have stopped firing (if  $r_2$  is non-recursive, then it can only fire at one stage value; if  $r_2$  is recursive, then it can fire at successive stage values while keeping at the same precedence level). Therefore, the first two rules in Figure 9 specify that, if there has been some action request we keep the same level; otherwise we move to the rules at the next precedence level. Naturally,  $\text{req}^J$  is defined as the disjunction of all possible action requests. When we reach the last level in  $\text{prec}$ , for a stage value of say  $m$ , then,  $\text{levl}^{m+1}$  is never set to true and we thus reached the end of the computation. The third rule in Figure 9 specifies that at the first step of the computation the precedence to be used to select the rules should be the first (bottom) in the  $\text{prec}$  tree.

For each database,  $\mathcal{D}$ , the program containing the rules so generated, augmented with the facts describing the content of the database at the beginning of the transaction, will be called the *durable-delta* program for  $\mathcal{D}$ .

## 4 Declarative Semantics and Operational Semantics

At this point, our reader is probably puzzled by the many rules that our  $\text{Datalog}_{15}$  model has generated starting from a rather simple example. It is therefore important to point out



Figure 6: Action requests that not compose with objects in delta relations (Case 1)

---

$\text{insDept}^{J+1}(D, N, \text{DivL}) \leftarrow$	$\text{rinDept}^J(D, N, \text{Div}, L), \neg \text{insDept}^J(D, N, \text{Div}, L),$ $\neg \text{delDept}^J(D, N, \text{Div}, L), \neg \text{updDept}^J(-, -, -, D, N, \text{Div}, L).$
$\text{delDept}^{J+1}(D, N, \text{Div}, L) \leftarrow$	$\text{rdeDept}^J(D, N, \text{Div}, L), \neg \text{insDept}^J(D, N, \text{Div}, L),$ $\neg \text{delDept}^J(D, N, \text{Div}, L), \neg \text{updDept}^J(-, -, -, D, N, \text{Div}, L).$
$\text{updDept}^{J+1}(\text{Do}, \text{No}, \text{Dvo}, \text{Ln}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}) \leftarrow$	$\text{rupDept}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Ln}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}), \neg \text{insDept}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}),$ $\neg \text{delDept}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}), \neg \text{updDept}^J(-, -, -, \text{Do}, \text{No}, \text{Dvo}, \text{Lo}).$
$\text{insEMP}^{J+1}(\text{En}, \text{N}, \text{Jt}, \text{Dn}) \leftarrow$	$\text{rinEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{insEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}),$ $\neg \text{delEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{updEMP}^J(-, -, -, \text{En}, \text{N}, \text{Jt}, \text{Dn}).$
$\text{delEMP}^{J+1}(\text{En}, \text{N}, \text{Jt}, \text{Dn}) \leftarrow$	$\text{rdeEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{insEMP}^J(\text{Eno}, \text{N}, \text{Jt}, \text{Dn}),$ $\neg \text{delEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{updEMP}^J(-, -, -, \text{En}, \text{N}, \text{Jt}, \text{no}).$
$\text{updEMP}^{J+1}(\text{Eno}, \text{No}, \text{Jto}, \text{Dno}, \text{Enn}, \text{Nn}, \text{Jtn}, \text{Dnn}) \leftarrow$	$\text{rupEMP}^J(\text{Eno}, \text{No}, \text{Jto}, \text{Dno}, \text{Enn}, \text{Nn}, \text{Jtn}, \text{Dnn}), \neg \text{insEMP}^J(\text{Eno}, \text{No}, \text{Jto}, \text{Dno}).$ $\neg \text{updEMP}^J(\text{Eno}, \text{No}, \text{Jto}, \text{Dno}), \neg \text{updEMP}^J(-, -, -, \text{Eno}, \text{No}, \text{Jto}, \text{Dno}).$
$\text{insHPaid}^{J+1}(\text{Jt}) \leftarrow$	$\text{rinHPaid}^J(\text{Jt}), \neg \text{insHPaid}^J(\text{Jt}).$
$\text{delHPaid}^{J+1}(\text{Jt}) \leftarrow$	$\text{rdeHPaid}^J(\text{Jt}), \neg \text{insHPaid}^J(\text{Jt}).$
$\text{updHPaid}^{J+1}(\text{Jto}, \text{Jtn}) \leftarrow$	$\text{rupHPaid}^J(\text{Jto}, \text{Jtn}),$

---

Figure 7: Delta tuples copied into the next-state table without any change (Case 2)

---

$\text{insDept}^{J+1}(D, N, \text{Div}, L) \leftarrow$	$\text{insDept}^J(D, N, \text{Div}, L), \neg \text{dinDept}^J(D, N, \text{Div}, L), \neg \text{chrDept}^J(D, N, \text{Div}, L),$
$\text{delDept}^{J+1}(D, N, \text{Div}, L) \leftarrow$	$\text{delDept}^J(D, N, \text{Div}, L), \neg \text{ddeDept}^J(D, N, \text{Div}, L), \neg \text{chrDept}^J(D, N, \text{Div}, L),$
$\text{updDept}^{J+1}(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}) \leftarrow$	$\text{updDept}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}),$ $\neg \text{dupDept}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}), \neg \text{chrDept}^J(D, N, \text{Div}, L)$
$\text{insEMP}^{J+1}(\text{En}, \text{N}, \text{Jt}, \text{Dn}) \leftarrow$	$\text{insEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{dinEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}),$
$\text{delEMP}^{J+1}(D, N, \text{Div}, L) \leftarrow$	$\text{deleEMP}^J(D, N, \text{Div}, L), \neg \text{ddeEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}), \neg \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}),$
$\text{updEMP}^{J+1}(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}) \leftarrow$	$\text{updEMP}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}),$ $\neg \text{dupEMP}^J(\text{Do}, \text{No}, \text{Dvo}, \text{Lo}, \text{Dn}, \text{Nn}, \text{Dvn}, \text{Ln}), \neg \text{chrEMP}^J(\text{En}, \text{N}, \text{Jt}, \text{Dn}).$
$\text{insHPaid}^{J+1}(\text{Jt}) \leftarrow$	$\text{insDept}^J(\text{Jt}), \neg \text{dinHPaid}^J(\text{Jt}), \neg \text{chrHPaid}^J(\text{Jt}),$
$\text{delHPaid}^{J+1}(\text{Jt}) \leftarrow$	$\text{delDept}^J(\text{Jt}), \neg \text{ddeHPaid}^J(\text{Jt}), \neg \text{rdeHPaid}^J(\text{Jt}).$
$\text{updHPaid}^{J+1}(\text{Jto}, \text{Jtn}) \leftarrow$	$\text{delDept}^J(\text{Jto}, \text{Jtn}), \neg \text{dupHPaid}^J(\text{Jto}, \text{Jtn}), \neg \text{chrHPaid}^J(\text{Jtn}).$

---

Figure 8: Applying the composition rules (Case 3)

---

```

insDeptIJ+1(D, N, Dv, L) ← insDeptJ(X), rupDeptJ(D, N, Dv, L, -, -, -).
error ← insDeptJ(D, N, Dv, L), rinDeptJ(D, N, Dv, L).
error ← delDeptJ(D, N, Dv, L), rdeDeptJ(D, N, Dv, L).
error ← delDeptJ(D, N, Dv, L), rupDeptJ(J, D, N, Dv, L, -, -, -).
delDeptJ+1(Y) ← updDeptJ(X, Y), rdeDeptJ(Y).
error ← updDeptJ(-, -, -, Dn, Nn, Dvn, Ln), rinDeptJ(Dn, Nn, Dvn, Ln).
updDeptJ+1(Do, No, Dvo, Lo, Dn, Nn, Dvn, Ln) ←
    updDeptJ(Do, No, Dvo, Lo, D, N, Dv, L), rupDeptJ(D, N, Dv, L, Dn, Nn, Dvn, Ln)

insEMPJJ+1(En, N, Jt, Dn) ← insEMPJ(X), rupEMPJ(En, N, Jt, Dn, -, -, -).
error ← insEMPJ(En, N, Jt, Dn), rinEMPJ(En, N, Jt, Dn).
error ← deleEMPJ(), rdeEMPJ().
error ← deleEMPJ(En, N, Jt, Dn), rupEMPJ(En, N, Jt, Dn, -, -, -).
deleEMPJ+1(En, N, Jt, Dn) ← updEMPJ(-, -, -, En, N, Jt, Dn), rdeEMPJ(En, N, Jt, Dn).
error ← updEMPJ(-, -, -, En, N, Jt, Dn), rinEMPJ(En, N, Jt, Dn).
updEMPJ+1(Eo, No, Jto, Do, En, Nn, Jtn, Dn) ←
    updEMPJ(Eo, No, Jto, Do, E, N, Jt, D), rupDeptJ(E, N, Jt, D, Dn, Nn, Dvn, Ln).

insHPaidJJ+1(Jt) ← insHPaidJ(Jt), rupHPaidJ(Jt, -).
error ← insHPaidJ(Jt), rinHPaidJ(Jt).
error ← delHPaidJ(Jt), rdeHPaidJ(Jt).
error ← delHPaidJ(Jt), rupHPaidJ(Jt, -).
delHPaidJ+1(Jt) ← updHPaidJ(Jt), rdeHPaidJ(Jt).
error ← updHPaidJ(-, Jt), rinHPaidJ(Jt).
updHPaidJ+1(Jto, Jtn) ← updHPaidJ(Jto, Jt), rupHPaidJ(Jt, Jtn).

```

---

Figure 9: Moving to the next level till no more

---

```

levlJ+1(X) ← levlJ(X), wt1J+1, -reqJ, -error.
levlJ+1(Y) ← levlJ(X), wt1J+1, reqJ, -error,
                prec(X, Y).
levl0(X) ← wt10, prec(nil, X).
reqJ ← chrDdepJ(-, -, -, -).
reqJ ← chrEMPJ(-, -, -, -).
reqJ ← chrHPaidJ(-).
wt4J ← wt3J. %same strtm as chrR
wt3J ← levlJ(-).
wt1J+1 ← wt4J. %a new stage value
wt10. %begin ruleprocessing

```

---

that most of these rules, namely all those in Figures 2, 6, 7, 8 and 9 are needed to express Starburst’s deferred evaluation with composition semantics using Heraclitus’ delta relation approach. Any complete formalization of these complex operations is bound to be a lengthy one. For instance, the usage of relational algebra or relational calculus would lead to even longer formulas—and in fact most topical papers only provide semi-formal English-based descriptions of these. The examples of active rules, normally expressed in SQL or QUEL-like syntax in such papers, however, find a very simple expression in our framework—Figure 3.

Finally, all rules but those that define action requests and durable changes, can be generated directly from the schema, and they obey highly repetitive patterns. Therefore the use of a meta-level notation, where variables represent relation names and their attribute lists, would cut down dramatically in the number of final rules generated. Nevertheless, we have restricted ourselves to the basic Datalog<sub>IS</sub> representation, because we want to apply the standard stable model semantics to the problem at hand.

Observe that all these rules, but the active rules and the durable change rules, can be given a simple operational interpretation. These safe rules can, for instance, be translated into equivalent relational algebra expressions. Then the overall computation proceeds in a bottom-up fashion from level  $J$  to level  $J + 1$  (in fact, if we remove the  $\text{lchr}^J$  goals, the whole program becomes  $XY$ -stratified and thus efficiently computed [14]).

In our durable-changes policy, however, we use the negation of  $\text{lchr}^J$  as a goal to predict the absence of conflicting *future events*. This feature puts us beyond the scope of any operational semantics, and in the realm of declarative semantics based on the notion of stable models (Definition 2). Therefore, we can now define our durable-change semantics for active databases as follows:

**Definition 1** Let  $\mathcal{D} = (\mathcal{S}, \mathcal{C}, \mathcal{A})$  be a database where:

- $\mathcal{S}$  denotes a set of schema relations
- $\mathcal{C}$  denotes the current content of the database
- $\mathcal{A}$  denotes a set of active rules on  $\mathcal{S}$

Let  $P$  denote the durable-delta program for  $\mathcal{D}$ . If  $P$  has a stable model semantics, then  $\mathcal{D}$  is said to obey a durable-change semantics.

For all its superior conceptual benefits this new declarative semantics is bound to remain of little practical consequence until we can translate it into some efficient operational semantics. In general, stable models represent an egregious basis for efficient implementation since computing stable models is NP-hard [12]. Even more restrictive subclasses of programs, such as locally stratified programs or those that have well-founded models, might not yield computation procedures that can be realistically used for active database applications. At this point, therefore, our reader might suspect of having being led to the quagmire of current non-monotonic reasoning research whereby: ‘*The semantics we like cannot be implemented efficiently...*’. Fortunately, in this case, a careful assignment of priorities to rules and events, will take us out that quagmire and to the solid grounds of very efficient operational semantics. As described more formally next, this can be done by reconciling the stable model semantics with an efficient inflationary-fixpoint computation.

Let  $r$  be a rule of a logic program  $P$  and let  $h(r)$ ,  $gp(r)$  and  $gn(r)$ , respectively, denote the head of  $r$ , the set of positive goals of  $r$  and the set of negated goals of  $r$  without the negation sign. For instance, if  $r : a \leftarrow b, \neg c, \neg d$ , then  $h(r) = a$ ,  $gp(r) = \{b\}$  and  $gn(r) = \{c, d\}$ . In the following,  $P$  denotes a logic program with negated goals,  $I$  and  $N$  are subsets of  $P$ 's Herbrand Base  $B_P$  (here,  $I$  represents the set of atoms that are true, and  $N$  represents those that are false);  $ground(P)$  represents the Herbrand instantiation of  $P$ .

**Definition 2** *Let  $P$  be a logic program, and let  $I$  and  $N$  be subsets of  $B_P$ . The immediate positive-consequence operator for  $P$  given  $N$  is defined as:*

$$\Gamma_{P(N)}(I) = \{h(r) \mid r \in ground(P), gp(r) \subseteq I, gn(r) \subseteq N\}$$

While  $\Gamma$  can also be viewed as a two-place function (on  $I$  and  $N$ ), in the following definition, we view it as a function of  $I$  only, inasmuch as  $N$  is kept constant. The following characterization of two-valued stable models follows directly from the one given in [1]:

**Definition 3** *Let  $P$  be a logic program with Herbrand base  $B_P$  and  $\overline{M} = B_P - M$ . Then,  $M$  is a stable model for  $P$  iff:*

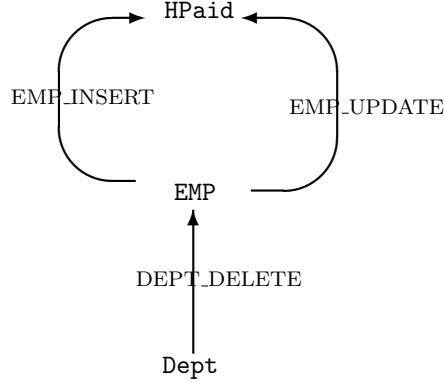
$$\Gamma_{P(\overline{M})}^{\uparrow\omega}(\emptyset) = M$$

Thus  $M$  is a stable model if it can be obtained as the  $\omega$  power of the positive consequence operator, where the set of false atoms is kept constant and equal to the set of atoms not in  $M$ . Using this last definition, it is easy to check whether a model  $M$  is stable in polynomial time, by simply letting the set of false atoms to be  $\overline{M} = B_P - M$ . In actual computations, however, the set of false atoms is not known a priori, and educated guesses must be made in the course of the computation when firing rules with negated goals. For instance, it is customary to use a *naive immediate consequence operator*, defined as follows ( $\overline{I} = B_P - I$ ):

$$T_P(I) = \Gamma_{P(\overline{I})}(I).$$

$T_P^{\uparrow\omega}(\emptyset)$  yields the least model for positive programs where  $T_P$  is continuous. However, for programs with negated goals, this operator makes the naive closed-world assumption that every atom that is currently not in  $I$  is false. However, as successive powers  $T_P$  are computed, larger and larger sets  $I$  are constructed, and the original assumptions about negated facts are frequently contradicted. Therefore, for most programs with negation,  $T_P^{\uparrow\omega}(\emptyset)$  does not yield a

Figure 10: The EPG for the rules of Figure 3



stable model, or not even a minimal model. Fortunately, our durable-delta programs offer a very useful exception to this general rule. Let us begin with the concept of *Event Precedence Graph (EPG)*:

**Definition 4** *Let  $P$  be a durable-delta program. The Event Precedence Graph (EPG) for  $P$  is a directed labeled graph that has as nodes the relation names of the database schema. The graph contains an arc from relation  $R_1$  to relation  $R_2$  with label  $\alpha$  iff there is an active rule  $\alpha$  having as goals either  $\text{ins}R_1, \text{del}R_1$ , or  $\text{upd}R_1$  and having either  $\text{rin}R_2, \text{rde}R_2$ , or  $\text{rup}R_2$  as its head.*

The EPG for the example at hand is shown in Figure 10. We will now discuss the treatment of acyclic EPG graphs: the treatment of graphs with cycles is discussed in the next section. The *Canonical Rule Precedence Assignment* for an EPG graph is defined as follows:

- Nodes with zero in-degree are assigned level 0
- The arcs departing from a node of level  $j \geq 0$  are assigned level  $j$ .
- Every node that is the end-node of one or more arcs, is assigned the maximum level of such arcs, plus 1.

Thus, in our example `Dept` (and the rules triggered by its changes) are at level 0, `EMP` is at level 1 and `HPaid` is at level 2. In order to avoid using integers outside the stage argument, we will represent level through a binary precedence relation `prec`. For the example at hand, for instance, we have

$$\text{prec}(\text{nil}, \text{dd}) \quad \text{prec}(\text{dd}, \text{ei}) \quad \text{prec}(\text{dd}, \text{eu})$$

Thus, `prec` is a graph having as nodes the abbreviated rule names. For each rule  $r$  at level 0 there is an arc from a special node `nil` to  $r$ ; for each rule  $r$  at level  $j$  there must be an arc connecting some rule at level  $j - 1$  to  $r$ . Then we have the following theorem:

**Theorem 1** *Let  $P$  denote the durable-change program, with acyclic EPG graph, and canonical rule precedence assignment. Then,  $P$  has a stable model which is equal to  $T_P^{\uparrow\omega}(\emptyset)$ .*

*Proof.* It suffices to show every `1chRJ` atom assumed false to fire a rule instance  $r$  is not in  $T_P^{\uparrow\omega}(\emptyset)$ . Indeed, durable-change rules can fire only at their canonical level—i.e., at a level where rules that could affect their triggering events have already fired. Also these rules can never fire again since the EPG is acyclic.  $\square$

Observe for example the computation for the example at hand. The computation begins with the exit rules in Figure 9, setting  $\text{wt}1^J$  and then  $\text{lev}1^0(\text{dd})$  to true. Thus, the rules of Figure 2 compute database state at level 0, by combining the database before the transaction with the net effect of all actions till the rule-processing point. The durable changes are also evaluated at this point, assuming as a default that all  $\neg\text{chR}$  goals are true. While this assumption is incorrect, no arm follows from it, since only rules enabled by  $\text{lev}1^J$  can fire. For instance, for stage value of 0, only the  $\text{dd}$  rule can fire, and its firing event is entered in the durable-change table. The action requested by  $\text{DEPT\_DELETE}$  rule is the deletion of the last tuple in Figure 1 (the  $\text{analyst}$  tuple) which is thus removed by the composition rules. As the computation proceeds with a stage value of 1, no rule fires; thus the delta relations are copied unchanged to next stage value of 2, and  $\text{lev}1^2(\text{ei})$  and  $\text{lev}1^2(\text{eu})$  are set to true and the first two rules can fire; but with no change was left in the delta tables for this level, the computation proceeds by setting level  $\text{lev}1^3(\text{ei})$ . There is no candidate triggering event at this level either, and we are now at the top of the EPG graph. Thus  $\text{wt}1^4$  is set to true while  $\text{lev}1^4$  remains false. Thus the computation terminates yielding a stable model for our durable delta program.

Upon successful termination, all remaining entries in the delta relations and all the entries accumulated in the durable-change relations, are written back into stable storage as the transaction commits.

#### 4.1 Recursive Rules

In the previous example, the durable-delta program is recursive, but the EPG is acyclic. Let us now consider the situation where the EPG is cyclic, which corresponds to the situation where the set of active rules alone are recursive. For instance, assume that we have a hierarchy of organizations each identified by a  $D\#$ ; the column  $\text{Div}$  in the  $\text{Dept}$  relation, now denotes the organization to which the department is reporting. Then, we have an active rule  $\text{DEPT\_DEL\_PROP}$  which, once an organization is deleted, deletes all organizations reporting to it. The logical counterpart of such a rule is:

$$\begin{aligned} \text{DEPT\_DEL\_PROP:} \\ \text{rdeDept}^J(Dc, N, Dp, Loc), \leftarrow & \text{delDept}^J(Dp, -, -, -), \\ & \text{curDept}^J(Dc, N, Dp, Loc), \\ & \text{lev}1^J(\text{ddp}), \neg\text{chDept}^J(Dp, -, -, -). \end{aligned}$$

This last rule introduces a loop from  $\text{Dept}$  to  $\text{Dept}$  the EPG graph of Figure 10. While, programs with acyclic EPGs always have stable models, not all cyclic programs have one. Take, for instance, the following rule that reacts to a tuple with  $\text{nil}$  value being inserted in  $\text{HPaid}$  by deleting the same tuple:

$$\begin{aligned} \text{COUNTER\_ACTION:} \\ \text{rdeHPaid}^{J+1}(\text{nil}) \leftarrow & \text{insHPaid}^J(\text{nil}), \\ & \text{lev}1^J(\text{ca}), \neg\text{chHPaid}^J(\text{nil}). \end{aligned}$$

Say now that our delta tables contain  $\text{insHPaid}(\text{nil})$ , and our durable-delta program  $P$  has no active rule, but  $\text{COUNTER\_ACTION}$ , affecting this entry in the delta relation. Then, if we assume the insertion of  $\text{HPaid}(\text{nil})$  to be durable, we must fire  $\text{COUNTER\_ACTION}$ , requesting the deletion of this delta tuple—making the original insertion ephemeral. Conversely, if we consider the initial insertion ephemeral, then we cannot fire the rule, thus making the insertion of  $\text{HPaid}(\text{nil})$  durable. This is contradiction means that our program  $P$  (which

also include the durable change rules not listed above) does not have a stable model, much in the way in which a program containing the rule  $a \leftarrow \neg a$  cannot have a stable model. The COUNTER\_ACTION rule is therefore disallowed in our durable-delta semantics; to provide the same operational effect as this rule we propose the use of ‘instead’ rules from Postgres [2], which are given a modified logical translation [15].

Therefore, durable-delta programs with cyclic EPG might not have stable model semantics; moreover, deciding if a given durable-delta program has a stable model is as complex a problem as deciding whether an arbitrary program has a stable model; <sup>1</sup>

Therefore, it appears that cyclic EPG pose an insuperable obstacle to the implementation of our durable change semantics. Fortunately, we can take advantage of the roll-back mechanism of transactions, whereby a computation that has incurred in errors or semantic constraint violations can be simply aborted, while the database is returned to the initial consistent state. We have already used **error** conditions in composition semantics, where, e.g., an insert followed by another insert on the same tuple produces an error. Once the **error** predicates becomes true then the transaction aborts. For a cyclic EPG, therefore, we can monitor the computation as it takes place, and once we detect that this will not generate a stable model, we can simply abort the computation. As we describe next, this policy can be implemented efficiently—consistently with the fact that checking that a model is stable is PTIME.

Let  $G$  be a directed graph, and  $S$  be a strong component for  $G$ . The *contraction* of  $S$  in  $G$  yields a new graph  $G'$  obtained by (i) eliminating all the arcs of  $S$  and merging the nodes of  $S$  into one node, say  $N_S$ , and (ii) replacing each arc  $A \rightarrow B$  by  $N_S \rightarrow B$  if  $A \in S$ , and by  $A \rightarrow N_S$  if  $B \in S$ . The graph obtained from  $G$  by contracting all its maximal strong components of  $G$  is unique and will be called the *acyclic contraction* of  $G$ .

The canonical rule precedence assignment for a cyclic EPG is then constructed as follows: first compute the canonical assignment for its acyclic contraction, and then set all arcs (rules) in a strong component  $S$  to the same level as  $N_S$ .

For the example at hand, the addition of rule DEPT\_DEL\_PROP to those of Figure 4, adds a loop on Dept; then DEPT\_DEL\_PROP is assigned to level 0 and the levels of the remaining rules does not change, although the computation of  $T_P^{\uparrow\omega}(\emptyset)$  is changed by this rule. Say, for instance, that the database contains the following Dept tuples:

```
iniDept (2500, ims, 1000, 'LA')
iniDept (1300, media, 1000, 'LA').
iniDept (2300, prodc, 1300, 'LA').
```

Then, the computation begin with `levl(dd)` and `lev(ddp)` being set to true and the rules DEPT\_DELETE and DEPT\_DEL\_PROP being triggered by the first tuple in the delta table of Figure 1. The rule DEPT\_DELETE triggers a deletion on EMP which composes with the last entry from the delta table, and removes it as in the non-recursive case. The recursive rule DEPT\_DEL\_PROP instead generates a new request on Dept `rdeDept(2300, prodc, 1300, 'LA')`. This does not compose with any current request, and it is entered as `delDept(2300, prodc, 1300, 'LA')` in the delta relation. Now, the stage value is increased, but the precedence level is not changed, and will remain the same until all the requests at this level have been exhausted. At this point the request `delDept(2300, prodc, 1300, 'LA')` is assumed durable, (the durable

---

<sup>1</sup>It suffices to write active rules for an NP-complete problem—e.g., to decide whether a graph has a Hamiltonian circuit.

change rules have been omitted for brevity). Next, the rule `DEPT_DEL_PROP` can no longer fire since the condition part of the rule fails. Thus, and the computation moves to the next precedence level where it continues as in the non-recursive case.

From the various examples proposed in the literature, it appears that  $T_P^{\uparrow\omega}(\emptyset)$  succeeds in computing a stable model for most durable-delta programs of practical interest. However, precautions must be taken against rules such as the `COUNTER_ACTION` rule where there is no stable model, or even situations where  $T_P^{\uparrow\omega}(\emptyset)$  cannot find it. To this end, we add the following rules:

$$\begin{aligned} \text{fail\_sc} &\leftarrow \text{dinDept}^J(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}), \text{lchDept}^J(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}) \\ \text{fail\_sc} &\leftarrow \text{ddeDept}^J(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}), \text{lchDept}^J\text{t}(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}) \\ \text{fail\_sc} &\leftarrow \text{dupDept}^J(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}), \\ &\quad \text{lchDept}^J(\mathbf{X}, \mathbf{Y}, \mathbf{Z}, \mathbf{W}, -, -, -, -). \end{aligned}$$

We need to add a similar rule for each event in a strongly connected component of the EPG, only. Whenever any such a rule fires  $T_P^{\uparrow\omega}(\emptyset)$  is no longer a stable model. In this case, `fail_sc` the transaction can be aborted using the following rule:

$$\text{error} \leftarrow \text{fail\_sc}$$

Observe that, as per the rules of Figure 9, `error` immediately terminates computation of the model  $M$  and aborts the transaction. Then,  $M$  is a stable mode iff and only iff `fail_sc`  $\notin M$ , i.e., when the error has been produced by the composition rules rather than by a violation of the stability condition. Independent of its cause, `error` always results in an immediate transaction-abort.

## 5 Termination

When `fail_sc` does not occur  $T_P^{\uparrow\omega}(\emptyset)$  produces a stable model. The main question that remains open is whether it terminates after a finite number of steps, or only an infinite computation to the first ordinal can yield the stable model. Using the durable change semantics, and the `Datalog1S` formalism, we can now derive a simple a practical solution to this problem, that is in normally of very intractable nature.

Since in `Datalog1S` functions symbols are confined to an argument,  $T_P^{\uparrow\omega}(\emptyset)$  defines a computation that either terminates or becomes ultimately periodic

**Definition 5** *A function  $f$  on natural numbers is said to be ultimately periodic with period  $(n, k)$ , where  $n$  and  $k$  are non-negative integers, if for all  $j \geq n$  we have  $f(j + k) = f(j)$ .*

Let  $M = T_P^{\uparrow\omega}(\emptyset)$ , and let  $M^J$  denote the set of atoms in  $M$  with stage value equal to  $J$ . For a `Datalog1S` program  $P$ ,  $M^J$  can be viewed as a function that maps an integer  $J$  to the set of atoms in  $T_P^{\uparrow\omega}(\emptyset)$  that have stage argument  $J$ . Then, we have the following theorem [4]:

**Lemma 1** *Let  $P$  be a `Datalog1S` program. Then one of the following two cases must hold:*

1. **[Finite Set of Stage Values]** *There exist an integer  $n$  such that, for  $J > n$ :  $M^J = \emptyset$*
2. **[Periodic Behavior]** *The set of stage values is not finite, but there exist two integers  $n$  and  $k$  such that for every  $J > n$ :  $M^{J+k} = M^J$ .*



We now have the following theorem:

**Theorem 2** *Let  $P$  be a durable-delta program. If  $P$  is Datalog<sub>IS</sub>, then there exists an integer  $n$  such that for every  $J > n$ ,  $M^J = \emptyset$ . Then  $M = \bigcup_{1 \leq j \leq n} M^J$  is the stable model of  $P$  iff `fail_sc`  $\notin M$ .*

*Proof:* It suffices to prove that the computation is not eventually periodic. Indeed, assume that the computation becomes periodic after  $n$  with periodicity  $k$ . Then if  $M$  contains a  $\text{din}R^j(X)$  with  $j > n$  then it must also contain  $\text{din}R^{j+k}(X)$ . Observe that the latter requires a  $\text{ins}R^{j+k}(X)$  to in delta relation—and this requires that  $M$  contains some  $\text{chr}R^{j+h}(X)$  for  $0 < h \leq k$ . Then,  $\text{lch}R^j(X)$  is true, and that is a contradiction, as error is generated and the computation terminates. Similar considerations hold for  $\text{dde}R^{j+k}(X)$  and  $\text{dup}R^{j+k}(X)$ .  $\square$

Therefore, only a finite number of distinct stage values is possible for durable-delta programs where the computation can be stopped at the first  $n$  for which `lev1n` is not set to true. If that occurs at the  $m$ -step of the computation <sup>2</sup> of  $T_P^{\uparrow\omega}(\emptyset)$  then we have that:

$$M = \bigcup_{1 \leq j \leq m} T_P^{\uparrow j}(\emptyset).$$

The rules in Figure 3 are Datalog<sub>IS</sub>, and so is `DEPT_DEL_PROP`. For these rules, and any other Datalog<sub>IS</sub> rule we know a priori that the execution of our recursive rule will terminate, even though the current database viewed as a graph has cycles (in our example this would mean a department reporting to and transitively supervising the same department—a situation that can only be associated with incorrect data).

Durable-change semantics is immune from the problems of recursive rules falling into a cyclic behavior, whereby the active rules cycle forever, periodically repeating the same actions. Even if the rules are not Datalog<sub>IS</sub> this behavior cannot occur; the proof of this follows the footsteps of Theorem 2. Thus, the only active rules where termination is an issue (under durable change semantics) are those that are not Datalog<sub>IS</sub>, and keep generating new change request as the computation moves along (thus the change requests relations in  $T_P^{\uparrow\omega}(\emptyset)$  are infinite).

Practical examples of such behaviors are rare in common applications of active rules; however, we will now seek them out to better understand the problem and correct it.

Active rules that are not Datalog<sub>IS</sub> are those that make use of interpreted functions including arithmetic and aggregates. Rules that keep summary information in materialized views are often of these kind—but these are normally non-recursive and thus they are free of termination concerns.

As a more interesting example, consider a part/subpart database, where active rules contribute in the computation of various Bill-of-Materials applications. Say, for instance that the cost of a part, defined as the sum of the cost of its subparts, is computed using recursive active rules. For instance, one might have a relation `PartCost(P#, Cost)`, where the SQL transaction inserts the costs for the basic parts, purchased from external suppliers, and then the active rules compute the cost of derived parts in a bottom-up fashion. While, a correct part/sub-parts relation should have no cycles, these might be there anyway due to software bugs and human errors. Then, all parts that lay in these cycles, will be visited an infinite number of times, each time with a new and larger value for `Cost`. As our readers might have observed, the non-termination can here be cured by simply ensuring that `P#` is the unique

---

<sup>2</sup>It is easy to show that  $m = 5 \times n + 1$ .

key for `partcost(P#, Cost)`. Then an error occurs once the same part is visited a second time, and the insertion of a second cost is attempted. However, to compute the cost for the parts, the programmer may have chosen to initialize first the relation `partcost(P#, Cost)` by entering a tuple for each available part where `Cost = 0`. Then the SQL statement might have initialized the basic parts to their actual cost; the rules would then be triggered by update events and request that the cost in super-part be updated to the sum of its sub-parts. Under this approach no key constraint is violated, and the presence of cycles in the database will here cause an infinite loop.

This scenario suggests that a simple improvement is desirable for our durable-change policy. Basically, the composition rules must be made more strict, to ensure that the only durable changes are those *for which there is no later request with the same key value*. Under this policy, Under this new policy the fact `PartCost(100, 34.5)`, will be considered ephemeral if a later request, such as `rinPartCost(100, 49.8)` or `rdePartCost(100, 34.5, 100, 55.1)` are generated, inasmuch as these have the same key value as the previous delta tuple. For reason of space, details about these and other improvements related to keys will be discussed in later papers. In a nutshell however, they further restrict the the set of active rules that will not terminate. In fact it can be shown that the only rules for which non-termination is possible are those that generate and store in the database an infinite number of new key values. As long as the generation of resources such as employee-numbers or other logical identifiers is not delegated to the active rules, then termination is ensured in every case. Since this policy is consistent with all active rule applications I have seen so far, and that it is easily enforced by the system, I make the bold claim that durable change semantics provides an efficient, complete and permanent cure to the non-termination problem for active database rules.

## 6 Instead Rules and Refuse Rules

There are basically two approaches to integrity constraint enforcement. One is *change propagation*, the other is *change rejection*. The typical example is a deletion that violates a foreign-key constraint. The change rejection policy simply returns an error; the change propagation policy propagates the deletion to all tuples that refer to the deleted tuple through its key value. The `COUNTER_ACTION` rule is an example of a rejection policy enforced implicitly through a propagation rule whose final effect is to reverse the original action (the mixing of the two approaches leads to a violation of the durable change semantics). The need for having active rules that explicitly reverse their triggering event was recognized by the designers of Postgres [2], who introduced the optional `refuse` and `instead` annotation for this purpose. A rule

$$\text{EVENT, CONDITION} \rightarrow \text{REFUSE}$$

is interpreted as a rule that remove its triggering `EVENT` from the delta relation. Therefore, a rule declared with the `instead` option, specifies that in addition to its head `ACTION` a `REFUSE` action must also be taken. Refuse and Instead rules can be implemented with minor modification in our approach. The observation to be made is that if a regular rule  $r_A$  has the same body `EVENT` as a refuse (or instead) rule  $r_B$ , then every instance of such `EVENT` that triggers  $r_B$  is ephemeral for  $r_A$ . While this observation can be implemented in several ways (e.g., through the re-writing of  $r_B$  to exclude the changes that trigger  $r_A$ ), the simplest solution seems to be changing the precedence of the rules as follows:

*Canonical Precedence Assignment* for an EPG graph with instead or refuse rules:

- Nodes with zero in-degree are assigned level 0
- The arc departing from a node of level  $j \geq 0$  are assigned level  $j$ , if none of these corresponds to an refuse rule or an instead rule. Otherwise the latter arcs are assigned level  $j$  and the remaining arcs (if any) are assigned level  $j + 1$ .
- Every node that is the end-node of one or more arcs, is assigned the maximum levels of such arcs, plus 1.

Say for instance that we have an instead rule that upon the insertion of a nil value in HPaid employee writes the following warning message into some output relation OUT (not part of the schema) "insertion refused.". This is modeled with the following pair of logical rules:

```

INSTEAD_RULE:
rinOUTJ+1("insertion refused.") ← insHPaidJ(nil), lev1J(ir), ¬lchHPaid(nil).
chrHPaidJ+1(nil) ← insHPaidJ(nil), lev1J(ir), ¬lchHPaid(nil).

```

The key difference here is the second rule that blocks the event  $\text{insHPaid}^J(\text{nil})$  from being passed to the next stage (this is the only rule generated for a refuse rule). Furthermore, there is no durable-change rule for an instead or refuse rule; thus, the change that has triggered the instead or refuse rule is obliterated at the next step, and it is not seen by rules having later precedence, including those that have the same EVENT in their bodies.

For the example at hand, the canonical precedence assignment will add the fact  $\text{prec}(\text{ei}, \text{ir})$  (i.e., in priority INSTEAD\_RULE immediately follows EMP-INSERT. (If there were a regular rule  $\text{rr}$  triggered triggered by an event on HPaid and short name, we would add  $\text{prec}(\text{ir}, \text{rr})$ .)

These improvements model the instead and refuse semantics in a natural fashion, and support the typical applications of these rules, including as integrity by rejection, signaling of illegal requests, concrete view and version maintenance. The theorems and formal properties discussed in the previous section remain valid.

## 7 Conclusion

This paper has presented several seminal contributions. A first novelty is represented by the notion of durable-change semantics, which ensures termination of active rule programs by making their behavior more consistent with transaction semantics. This result, obtained using the Datalog<sub>IS</sub> framework, provides a tangible proof that the power of active databases, that was previously considered impervious to formal treatment, can in fact be tamed and put to great uses with the help of the semantics of deductive databases. (Nor the benefits flow in one direction only, since this paper provides a rare example of a successful derivation of efficient operational semantics for a problem characterized by stable-model semantics).

In my recent papers, I have been expousing the thesis that a conceptual unity underlies the areas of active databases, temporal databases and deductive databases [13, 14, 11]. The results of this paper, bring further support to this thesis, and, hopefully, will promote the confluence of these three areas of database research.

## 8 Conclusion

This paper presented several new results. A first novelty is the notion of durable-change semantics, which ensures termination of active rule programs by making their behavior more

consistent with transaction semantics. This result, obtained using the Datalog<sub>IS</sub> framework, provides a tangible proof that the power of active databases, that was previously considered impervious to formal treatment, can in fact be tamed and improved with the help of the semantics of deductive databases. (Nor benefits flow in only one direction, since this paper provides a rare example of a successful derivation of efficient operational semantics for a problem characterized by stable-model semantics).

In my recent research, I have been pursuing the thesis that a conceptual unity underlies the areas of active databases, temporal databases and deductive databases [13, 14, 11]. The results of this paper, bring further support to this thesis, and, hopefully, will promote the confluence of these three areas of database research.

## Acknowledgments

Thanks are due to Antonio Brogi for many improvements.

## References

- [1] M. Gelfond, V. Lifschitz, The stable model semantics for logic programming, *Proc. 5th Int. Conf. on Logic Programming*, MIT Press, 1988.
- [2] M.L. Stonebraker, A. Jhingran, J. Goh, and S. Potamianos. On rules, procedure, cacheing and views in data base systems. In *ACM SIGMOD Int. Conf. on Management of Data*, pages 281–290, 1990.
- [3] J. Chomicki, Temporal deductive databases, *Temporal Databases: Theory, Design and Implementation*, A. Tansel et al. (eds), Benjamin/Cummings, 1993.
- [4] J. Chomicki, “Polynomial-time Computable Queries in Temporal Deductive Database Systems,” *PODS 1990*.
- [5] S. Ghandeharizadeh, R. Hull and D. Jacobs, “On Implementing a Language for Specifying Active Database Execution’ Models, *Procs. Int. Conf. on Very Large Databases*, 1993.
- [6] D. McCarty and U. Dayal. The architecture of an active database management system. In *ACM SIGMOD International Conf. on Management of Data*, pages 215–224, 1989.
- [7] S. Gatzju and K. R. Dittrich. Events in an object-oriented database system. In *Proceedings of the First Intl. Conference on Rules in Database Systems*, pages 23–39, September 1993.
- [8] N.H. Gehani and H.V. Jagadish. Ode as an active database: Constraints and triggers. *Seventeenth Int. Conf. on Very Large Data Bases, Barcelona*, pages 327–336, 1991.
- [9] Widom J., “The Starburst Active Database Rule System”, To appear in *IEEE Trans. On Knowledge and Data Engineering*.
- [10] U. Dayal, E.N. Hanson, and J. Widom *Active Database Systems*, ”Modern Database Systems, W. Kim (ed.), Addison Wesley, 1995.
- [11] Y. Motakis, and C. Zaniolo, *Composite Temporal Events in Active Databases: a Formal Semantics*, submitted for publication.
- [12] J.S. Schlipf, The expressive powers of logic programming semantics, *Proc. ACM-PODS*, 1990, 196-204.
- [13] Zaniolo, C., N. Arni, K. Ong, “Negation and Aggregates in Recursive Rules: the  $\mathcal{LDL}++$  Approach”, *Proc. 3rd Int. Conf. on Deductive and O-O DBs, DOOD-93*, Phoenix, AZ, Dec 6-8, 1993.
- [14] C. Zaniolo, “A unified semantics for active and deductive databases”, In *Procs. 1st Int. Workshop on Rules in Database Systems*, pages 271–287, Springer-Verlag, 1993