

Understanding the Global Semantics of Referential Actions using Logic Rules

WOLFGANG MAY

Institut für Informatik, Universität Freiburg, Germany
and

BERTRAM LUDÄSCHER

San Diego Supercomputer Center, University of California San Diego, USA

Referential actions are specialized triggers for automatically maintaining referential integrity in databases. While the *local effects* of referential actions can be grasped easily, it is far from obvious what the *global semantics* of a set of interacting referential actions should be. In particular, when using procedural execution models, ambiguities due to the execution ordering can occur. No *global, declarative semantics* of referential actions has yet been defined.

We show that the well-known logic programming semantics provide a natural *global semantics* of referential actions that is based on their *local* characterization: To capture the global meaning of a set RA of referential actions, we first define their abstract (but non-constructive) *intended semantics*. Next, we *formalize* RA as a *logic program* P_{RA} . The declarative, logic programming semantics of P_{RA} then provide the constructive, *global semantics* of the referential actions. So, we do not *define* a semantics for referential actions, but we show that there *exists* a unique *natural semantics* if one is ready to accept (i) the intuitive local semantics of local referential actions, (ii) the formalization of those and of the local “effect-propagating” rules, and (iii) the well-founded or stable model semantics from logic programming as “reasonable” global semantics for local rules.

We first focus on the subset of referential actions for deletions only. We prove the equivalence of the logic programming semantics and the abstract semantics via a game-theoretic characterization, which provides additional insight into the meaning of interacting referential actions. In this case a *unique maximal admissible solution exists*, computable by a P_{TIME} algorithm.

Second, we investigate the general case—including modifications. We show that in this case there can be *multiple maximal admissible subsets* and that all maximal admissible subsets can be characterized as *3-valued stable models* of P_{RA} . We show that for a given set of user requests, in the presence of referential actions of the form ON UPDATE CASCADE, the admissibility check and the computation of the subsequent database state, and (for non-admissible updates) the derivation of debugging hints all are in P_{TIME} . Thus, full referential actions can be implemented efficiently.

Categories and Subject Descriptors: H.2.1 [Database Management]: Logical Design—Data Models; H.2.4 [Database Management]: Systems—Relational Databases; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Logic and Constraint Programming

Authors' addresses: W. May, Institut für Informatik, Universität Freiburg, D-79110 Freiburg, Germany; email: may@informatik.uni-freiburg.de; B. Ludäscher, San Diego Supercomputer Center, UCSD, La Jolla, CA 92093-0505; email: ludaesch@sdsc.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2002 ACM 0362-5915/02/1200-0001 \$5.00

2 † W. May and B. Ludäscher

General Terms: Theory, Algorithms

Additional Key Words and Phrases: Database theory, game theory, logic programming, referential integrity, referential actions, relational databases

1. INTRODUCTION

The notion of integrity constraints and their automated maintenance has been an important research issue since the early days of relational databases [Codd 1970; Hammer and McLeod 1975; Eswaran 1976]. Integrity constraints in general, and *referential integrity constraints (rics)* in particular, are central concepts of database models and they are frequently used in real world applications. Many approaches use *ECA (event-condition-action)-rules* for monitoring and enforcing integrity constraints: if some event (here, an update) occurs, a set of actions is executed internally. Rules for integrity maintenance [Ceri and Widom 1990] have been a starting point for the area of *active databases* and their impact is documented by a recent 10-Year Paper Award [Ceri et al. 2000]. Triggers, a special kind of ECA-rules, have been part of database systems from the beginning [Eswaran 1976] and are included in the SQL2 and SQL3 standards [ANSI/ISO 1992a, 1999]. However, even today, researchers still complain about the difficulties in understanding the “subtle behavior” of multiple triggers acting together [Ceri et al. 2000] and the limited progress that has been made [Cochrane et al. 1996].

For enforcing referential integrity, each referential integrity constraint can be associated with *referential actions (racs)* that provide a *declarative, local* specification of how to automatically enforce referential integrity, thereby relieving the user from the burden of enumerating all induced updates that arise from an initial user request U_{\triangleright} . While ECA-rules and triggers are *procedural* means to enforce integrity by locally reacting on an event, the idea behind referential integrity and referential actions is a global one: the semantics of referential actions is given *declaratively* in terms of *local* actions, but with a *global* notion of “original state” (before the update) and “final state” (after the update and all induced changes) in mind.

Date and Darwen [1994] and Date [1990] report the problem of unpredictable behavior when realizing *racs* based on SQL triggers, that is, ambiguities in determining the set of updates on the database and the final database state in certain situations. The solution of the SQL2 standard [ANSI/ISO 1992a] (for a more complete overview of related work, see Section 6) described a *procedural* semantics that was subject to anomalies: since referential actions were executed at the same time as the parent was updated, the outcome depends on the order in which rows are modified or constraints are applied [Cochrane et al. 1996]. Markowitz [1994] presents safeness conditions that aim at avoiding ambiguities at the schema level. However, as shown in Reinert [1996], it is in general undecidable whether a database schema with *racs* is ambiguous.

Horowitz [1992] proposes a *marking algorithm* in the style of a fixpoint computation to define a *global* semantics that avoids these anomalies. An extension of this semantics was later incorporated into the SQL3 standard

[ANSI/ISO 1999]. Nevertheless, the *global* semantics is still—about 30 years after the definition of the idea of referential integrity—given by a complex, not very intuitive *procedural* algorithm, and only few commercial database systems support referential actions in their full extent.

In contrast to the majority of the work on this topic, we present a framework for maintenance of referential integrity based on *logic rules*. The user's original updates, together with the induced updates, yield a set of updates to the database that must be applied *instead* of (only) the original updates. Thus, the notions of “before” or “after” should be understood in a global, all-or-nothing manner without considering intermediate states for *defining* the meaning of updates.

The logic programming characterization given here demonstrates that the problem of *racs* can be solved by specifying *local* behavior in manageable parts, and exploiting the fact that the well-known logic programming semantics define an unambiguous, reasonable¹ *global* semantics of the “puzzle” that results from the interaction of multiple *rics* and *racs*. Moreover, this semantics can be computed efficiently and thus can be implemented in actual database systems. In contrast, the process of developing a procedural characterization as has been done in ANSI/ISO [1992a]; Horowitz [1992]; Markowitz [1994]; Cochrane et al. [1996] and finally ANSI/ISO [1999] required about 20 years, including intermediate solutions that have been proven to be incomplete and/or incorrect.

The Problem. We consider the following problem: Given a database instance D , a set of user-defined update requests U_{\triangleright} , and a set RA of *racs*, find the set of updates Δ that (i) is complete with respect to U_{\triangleright} , (ii) preserves referential integrity in the new database state D^0 , and (iii) reflects the intended meaning of RA , that is, *how* referential integrity should be enforced.

U_{\triangleright} can be given as a single (set-oriented) statement, or as a sequence of statements (including activated triggers) if such behavior is supported by the underlying transaction model. In some examples we construct specific sets U_{\triangleright} in order to illustrate certain interferences. Sometimes, the updates in U_{\triangleright} could be induced via cascading from a single statement, sometimes not. Our investigations not only provide the solution to a practically relevant problem, but also address the basic research problem of interacting *rics* and *racs*—the search for the above set Δ .

In case that no such Δ exists and U_{\triangleright} is rejected, we investigate *maximal admissible* subsets of U_{\triangleright} , and derive hints as to where the problems are located and how they can possibly be solved. Assume that U_{\triangleright} has been collected by several statements (e.g., a subtransaction) that are intended to do a certain amount of work. In case that it is rejected, something in the database (or its specification) is obviously inconsistent with the intended behavior. This points to problems in the design either of the database schema with its *rics* and *racs*, or in the programming of the subtransaction, or the contents of the database in the current situation is not as intended. Here, the additional information from

¹Dix [1995] formally defines this notion using very general principles.

4 † W. May and B. Ludäscher

the analysis of the *rics* and *racs* can be helpful for identifying the problem:

- the schema may be flawed, that is, there are missing *racs* (e.g., a forgotten CASCADE, or a RESTRICT where a NO ACTION would have been correct),
- the definition of the subtransaction is incomplete (e.g., it should generate some more update requests),
- the schema and the subtransaction are correct, but the database state is incorrect due to an incomplete definition of an earlier transaction.

Our semantics can give useful hints about where *exactly* the problem is located, that is, which *rics* are violated, and which tuples cause the problem.

Contributions. From a theoretical perspective, we aim at providing a better understanding of referential actions: We formalize the semantics of referential integrity constraints and referential actions as a logic program P_{RA} where

- (1) the *local behavior* of an individual *rac* $ra \in RA$ is precisely specified, and can be understood by solely looking at the corresponding rules $P_{ra} \sqcup P_{RA}$,
- (2) the (local) *interaction* between different update requests is precisely defined by certain other rules,
- (3) the *global behavior* is precisely specified and understandable from the declarative logic programming semantics. So, we do not *define* a semantics for *racs*, but we show that there *exists* a unique *natural* semantics if one is ready to accept the local semantics (1) and (2), and the logic programming semantics, that is, the well-founded model and the stable model as “reasonable” semantics.

The logic-based characterization not only provides a natural semantics for referential actions, but also leads to efficient procedures for handling referential actions in actual database systems.

From a practical perspective, we give polynomial time constructive characterizations for the following tasks:

- checking if the set U_{\triangleright} is admissible, and
- in case that it is, computing the set of updates to be accomplished (implying that ON UPDATE CASCADE which is currently not supported in most commercial database systems can be implemented efficiently), and
- in case that it is *not*, giving hints what updates, *rics*, *racs*, and tuples caused the problem.

The (complex) rule systems are *not* intended to be used by the designers of the *rics* and *racs*; they encode the *local* semantic conditions that are naturally induced by the application domain, and that the application designer has (correctly) in mind during the development process—so he implicitly relies on a “correct” global semantics that is ensured by our characterization. The use of the logic programming characterization in our approach is (i) in case of deletions for deriving a procedural algorithm and proving its correctness, and (ii) in case of modifications it can serve as a declarative *internal* implementation of referential actions (we do not derive a procedural algorithm for this case

as it would be excessively complex without providing any insight beyond the logic programming semantics). In both cases, if U_{Δ} is admissible, the database silently executes it. Otherwise the problems are presented to the application developer or the user in terms of tuples and foreign key constraints. Thus, the user is not bothered with the actual formalization, presented here as a “black box”—in the same way as he is not required to know about the details of the algorithm given in Horowitz [1992] and ANSI/ISO [1999].

Audience. Thus, the audience is not the typical SQL application programmer from whose point of view the *local* effects of *racs* should be enough to design an application, provided that the underlying DBMS assigns the correct *global* semantics to his specification. The relevance of our results for him is the formal characterization of the “built-in” correctness of lifting his local specification to the global behavior of the database system: A database that acts according to the described global semantics implements the application programmer’s intensions to the largest extent possible, based on his database schema and referential actions. Then, the programmer can rely on the correct interaction of referential actions. In case that an application raises non-admissible updates, something in this specification must be wrong, and the database system—supported by the semantics—can give hints as to where the problems come from. Thus, from the point of view of application programmers, the possible features of an implementation (especially, the conclusions that are drawn in case of rejected updates) based on our approach can be useful for improving their database design.

For implementors of database systems, the presented algorithms and methods for computing the induced set Δ of internal updates could be of interest. Moreover, the possibility of reacting on rejected updates by deriving hints for the application designer as to how to cure the problems can be useful for providing enhanced error reporting messages to the user (see Sections 3.7 and 5). The logical basis provided by three-valued logic and stable models facilitates more flexible investigations than the procedural fixed-point algorithms that are given in the SQL standards (and still only incompletely implemented by database systems).

The theoretical body of the paper—the details of the logic-based specifications and the game-theoretic analysis that lead to a *declarative, model-theoretic* characterization of the *global* semantics of referential actions—is directed at researchers studying database fundamentals and theory. For this audience, the paper provides (i) an elegant formal and “natural” (i.e., declarative) semantics of referential actions, and (ii) an application of theoretical concepts to a practically relevant problem that exhibits several levels of complexity that have to be handled by appropriate theoretical means.

Scope. Although our models and terminology are based on the relational model, the underlying issues of a “justification-based,” declarative semantics as proposed in this paper, are independent of the particular database model chosen. For example, extensions to the (very limited) notion of referential integrity in XML (ID/IDREF) have been proposed [Fan and Siméon 2000], or are

included as integral parts of new XML standards, like XML SCHEMA [2000]. It should be clear that rule-based maintenance of referential integrity in XML will face the same fundamental issues that are explained and resolved by our global semantics.

Relationship with earlier publications. We provide a comprehensive and uniform treatment of our previous work on declarative semantics for referential actions [Ludäscher et al. 1997; Ludäscher and May 1998]. First preliminary results have been reported in Ludäscher et al. [1996b]. In Ludäscher et al. [1997], it is shown that for referential actions (*racs*) with modifications, it may be intractable to compute *all* maximal admissible solutions (since the interactions may lead to an exponential blow up in the number of solutions). In Ludäscher and May [1998], we restricted the investigation to *racs without modifications*—deletions only. This guarantees the existence of a unique optimal solution, which can be efficiently computed. The present paper not only provides a complete and uniform treatment of our previous results, but extends them in various ways: We present a novel game-theoretic characterization of *racs* that gives a more abstract account of referential actions for modifications and shows that important aspects of this problem (admissibility of U_{\triangleright} , computation the actual set of update operations, and deriving debugging hints) are also in PTIME. Additionally, we explore the *practical* implications for actual relational DBMS that result from the theoretical investigations.

Structure of the paper. The paper is organized as follows: In Section 2, we introduce the basics of referential integrity. Then, we illustrate the problem of ambiguity that arises from the *local* specification of referential actions, and describe the disambiguation strategies of the SQL standard.

In Section 3, we investigate the class of *racs* without modifications (i.e., deletions only). In Section 3.1, we identify and formalize desirable abstract properties of updates which lead to the intended (albeit non-constructive) global semantics of *racs*. A constructive definition of this global semantics is obtained by formalizing a set of referential actions RA as a logic program P_{RA} (Section 3.2). The correctness of this characterization is proven via an equivalent game-theoretic characterization (Section 3.3) which allows intelligible proofs on a less technical level (Section 3.4). An algorithm for computing the maximal admissible solution is derived from the logic programming characterization (Section 3.5). So far, Section 3 is based on and extends the previous work [Ludäscher and May 1998]. The correctness of our characterization(s) and of the derived algorithm with respect to the “intended” ECA-style semantics, and the relationship with the SQL3 semantics is shown in Section 3.6. There, we can completely rely on the correctness of the logic programming semantics. The practical consequences of how to debug an application in case a set of updates is rejected are described in Section 3.7.

In Section 4, we extend the investigations to include modifications. We again start by giving an abstract characterization (Section 4.1). In Section 4.2 we associate with every set RA of *racs* a logic program P_{RA} whose rules capture the *local* semantics of modifications with referential actions, and show that the *global* declarative semantics of P_{RA} captures the abstract semantics, and thus

solves the problem in an unambiguous and comprehensive way. In contrast to the restricted deletions-only case, the characterization cannot be reformulated in an efficient algorithm since *stable model* semantics is required for the logical characterization. Sections 4.1 and 4.2 provide a comprehensive treatment of the results of the extended abstract [Ludäscher et al. 1997]. An equivalent game-theoretic characterization that abstracts from some details of the logical characterization is described in Section 4.3. Its details and the proof of the equivalence of all three characterizations can be found in Section B of the *electronic appendix*.

Further results showing the practicability of our approach are developed in Section 5: we show that the following tasks are computable in PTIME and derivable from the well-founded model: (i) a check on whether a user request is admissible, (ii) if so, computation of the subsequent database state, and (iii) for non-admissible user requests, an approximation of a maximal admissible subset, together with debugging hints. Section 6 reviews related work in the area and concluding remarks can be found in Section 7.

2. REFERENTIAL INTEGRITY

2.1 Notation and Preliminaries

In the following, we introduce the necessary notions of the relational model and calculus, which provide the basic formalism of our paper. We use a *positional* and *unnamed* relational calculus/Datalog-style notation—unlike the relational model with named attributes [Abiteboul et al. 1995]. In the unnamed Datalog-style notation, each argument position of a predicate is (implicitly) associated with an attribute. We follow this convention and regard attributes to be ordered according to their argument positions. This is used when correlating foreign keys with candidate keys.

Definition 2.1 Relational Schema, Keys. A relation schema $R(\bar{A})$ consists of a relation name R and a sequence of attributes $\bar{A} \triangleright (A_1, \dots, A_n)$. We identify attribute names A_i of R with the integers $1, \dots, n$. Given \bar{A} , a (possibly reordered) subsequence of \bar{A} (e.g., a key) is a vector $\bar{K} \triangleright (A_{i_1}, \dots, A_{i_k})$ such that $k \cdot n$ and $i_{j_1} \triangleright i_{j_2}$ for $j_1 \triangleright j_2$. Note that we have to allow that the attributes in \bar{K} may have a different order than in \bar{A} .

A relation R consists of *tuples*: Tuples of R are denoted by first-order atoms $R(\bar{X})$ with an n -ary relation symbol R , and a vector \bar{X} of variables or constants from the underlying domain. To emphasize that such a vector is ground, that is, comprises only constants, we write \bar{x} instead of \bar{X} . The *projection* of tuples \bar{X} to an attribute vector \bar{A} is denoted by $\bar{X}[\bar{A}]$: for example, if $\bar{x} \triangleright (a, b, c)$, $\bar{A} \triangleright (1, 3)$, then $\bar{x}[\bar{A}] \triangleright (a, c)$.

For a relation schema R with attributes \bar{A} , a minimal subset \bar{K} of \bar{A} whose values uniquely identify each tuple in R is a *candidate key*. In general, the database schema specifies which attribute vectors are keys. A candidate key $R.\bar{K}$ has to satisfy the first-order formula φ_{key} for every database instance D :

$$\exists \bar{X}_1, \bar{X}_2 (R(\bar{X}_1) \wedge R(\bar{X}_2) \wedge \bar{X}_1[\bar{K}] \triangleright \bar{X}_2[\bar{K}] \wedge \bar{X}_1 \triangleright \bar{X}_2). \quad (\varphi_{key})$$

Usually, in database design, for every relation one candidate key is selected to be *the primary key* of the relation. Since the key values uniquely identify a tuple of the corresponding *parent* relation, they can be used in other *child* relations for referring to the parent tuple. The respective attributes of the child relation are then called a *foreign key* of the *child relation*.

In this work, we assume that candidate and foreign keys do not contain null values (considering null values would add much technical effort and problems that are specific to null values, without giving additional insight).

Example 1 Primary Keys and Foreign Keys. Consider a database that describes countries and cities as depicted below. There, Name and Code are candidate keys of country (we chose Code to be the primary key). The attribute tuple City(Name, Country) is the primary key of City; the attribute City.Country references the key Country.Code and thus is a foreign key in City. Similar, Country(Capital, Code) references a city (identified by City(Name, Country)), thus the attribute tuple (Capital, Code) is a foreign key of Country (note the change of the order in the foreign key with respect to the original attribute list of Country).

Country				City		
1:Name	2:Code	3:Capital	4:Area	1:Name	2:Country	3:Pop.
Germany	D	Berlin	356910	Berlin	D	3472009
Austria	A	Vienna	83850	Munich	D	1244676
Utd. Kingdom	GB	London	244820	Vienna	A	1583000
London	GB			London	GB	6967500
⋮	⋮	⋮	⋮	⋮	⋮	⋮

Definition 2.2 Referential Integrity Constraints. A *referential integrity constraint (ric)* is an expression of the form

$$R_C.\bar{K} \uparrow R_P.\bar{K},$$

where \bar{K} is a *foreign key* of the *child relation* R_C , referencing a candidate key \bar{K} of the *parent relation* R_P . A *ric* $R_C.\bar{K} \uparrow R_P.\bar{K}$ is *satisfied* by a given database D , if for every child tuple $R_C(\bar{x})$ with foreign key values $\bar{x}[\bar{K}]$, there exists a tuple $R_P(\bar{y})$ with matching key value, that is, $\bar{x}[\bar{K}] \sqsubseteq \bar{y}[\bar{K}]$. Thus, for a database instance D , a *ric* is satisfied if $D \models \varphi_{ric}$:

$$\exists \bar{X} (R_C(\bar{X}) \wedge \exists \bar{Y} (R_P(\bar{Y}) \wedge \bar{X}[\bar{K}] \sqsubseteq \bar{Y}[\bar{K}])). \quad (\varphi_{ric})$$

A *ric* is *violated* by D if it is not satisfied by D .

Example 2 Primary Keys and Foreign Keys (Cont'd). Consider again Example 1. There, we have the *rics* City.Country \uparrow Country.Code and Country.(Capital,Code) \uparrow City.(Name,Country) or, in the numerical encoding, City.2 \uparrow Country.2 and Country.(3,2) \uparrow City.(1,2).

Table I. Operations and Possible Repairs

	R_P			R_C		
	ins	del	mod	ins	del	mod
propagate	ok	†	†	-	ok	-
restrict	ok	†	†	†	ok	†
wait	ok	†	†	†	ok	†

ok = *ric* remains satisfied
 † = *ric* may be violated, *rac* applicable
 - = *ric* may be violated, *rac* not applicable

For example, changing the code GB into UK in Country would violate both *rics*; this can be remedied by *propagating* the change, that is, applying the same renaming in City. Changing the capital of Germany to Munich would be allowed; changing it to Hamburg would violate the second *ric* (assuming that Hamburg is not stored in the City table). This could be fixed by either inserting a tuple for (Hamburg, Germany) into City, or by renaming, for example, Berlin to Hamburg. This shows that propagation of a modification is not always desired.

Definition 2.3 Updates. Update requests (*updates*) to a relation R are represented by auxiliary relations $\text{ins}_R(\bar{X})$, $\text{del}_R(\bar{X})$, and $\text{mod}_R(M, \bar{X})$. Here, M is a set of pairs i/c meaning that the i -th attribute of $R(\bar{X})$ should be set to the constant c . We say that a modification $\text{mod}_R(M_1, \bar{X})$ *subsumes* a modification $\text{mod}_R(M_2, \bar{X})$ if $M_1 \uparrow M_2$. As a shorthand for $\text{mod}_R([1/d, 3/e], (a, b, c))$, we sometimes write $\text{mod}_R(a/d, b, c/e)$.

2.2 Referential Actions

Rule-based approaches to referential integrity maintenance are attractive since they describe how *rics* should be enforced using “local repairs”: Given a *ric* $R_C \stackrel{F}{\vdash} R_P \stackrel{K}{\vdash}$ and an update operation on R_P or R_C , a *referential action* (*rac*) defines a *local* operation to be applied to R_C or R_P , respectively. We call this the *locality principle*. The problem with the locality principle is that the intuitive local repairs can lead to complex, “subtle behavior” [Ceri et al. 2000] with different, more or less “reasonable” outcome. Thus, an unambiguous *global* semantics for these local specifications is needed. As mentioned in the introduction, such semantics have been developed in the history of SQL [ANSI/ISO 1992a; Horowitz 1992; Markowitz 1994; Cochrane et al. 1996] over the years, including incomplete and incorrect intermediate solutions, now being specified by a procedural, fixpoint-style algorithm in ANSI/ISO [1999]. In the sequel, we start with a generic, abstract version of *racs* which is then related to the SQL version in Section 2.4.

The updates insert, delete, and modify can be applied to R_P or R_C , leading to six basic cases. It is easy to see from the logical implication in (φ_{ric}) above that insert R_P and delete R_C cannot introduce a referential integrity violation, while the other four operations can. There are in general three possible strategies as to how problems may be resolved; not all of them are applicable for all operations (cf. Table I):

10 † W. May and B. Ludäscher

- propagate: propagate (“cascade”) the update along the current *ric* by executing actions on the *other* tuple. This means, to propagate an update at the parent to all children (with respect to the *current* database state), or to propagate an update at the child to the parent.
- restrict: reject an update if it may cause a problem: (i) reject an update on the parent if there exists a child referencing the parent in the *current* (that means, at the moment when the update is executed) database state, and (ii) reject an update on the child if the referenced parent does not exist in the current state.
- wait: no local action is executed. Instead, the referential integrity constraints are checked—together with the other integrity constraints—after the end of a certain unit of work (note that in contrast to propagate and restrict, this can be seen as a kind of *deferred restrict* which is already a *global* strategy).

Each *rac* consists of the *ric* which should be maintained, the triggering update on either the parent R_P or the child R_C , and the “local repair.” We use the following notation, which should be self-explanatory:²

$$R_C \overset{F}{\leftarrow} R_P \overset{K}{\rightarrow} \text{ on } \{ \text{del} \mid \text{ins} \mid \text{mod} \} \{ \text{parent} \mid \text{child} \} \{ \text{propagate} \mid \text{restrict} \mid \text{wait} \}$$

2.3 The Problem of Ambiguity

With this *local* specification of behavior where each *rac* triggers an action on every child tuple (with respect to the respective *ric*) when an update to a parent tuple is executed (see, e.g., Dayal [1988]; Eswaran [1976]), some nondeterminism with respect to the outcome of a user operation may occur. If there are different possible final states of a database instance D (depending on the execution order of referential actions), D is called *ambiguous* with respect to the given referential actions.

There are several types of ambiguities leading to potentially different final states that are described in the following. In Section 4.2.6, we show that these ambiguities have a very natural and elegant representation in our framework: “controversial” updates are *undefined* in the well-founded model; the different possible results are characterized by certain stable models.

Example 3 Diamond. Consider the database with *racs* as depicted in Figure 1. Solid arcs point from R_C to R_P , *racs* are denoted by dashed (propagate) or double (restrict) arcs. Let $U_{\triangleright} \triangleq \{ \triangleright \text{del} _ R_1(a) \}$ be a user request to delete the tuple $R_1(a)$. Depending on the order of execution of *racs*, one of two different final states may be reached:

- (1) If execution follows the path $R_1 \rightsquigarrow R_3 \rightsquigarrow R_4$, the tuple $R_3(a, c)$ cannot be deleted: Since $R_4(a, b, c)$ references $R_3(a, c)$, the *rac* for R_4 restricts the deletion of $R_3(a, c)$. This in turn also blocks the deletion of $R_1(a)$. The user request $\triangleright \text{del} _ R_1(a)$ is rejected, and the database state remains unchanged, that is, $D^0 \sqsupseteq D$.

²As can be seen from Table I, not all combinations are meaningful: e.g., it is perfectly reasonable to propagate (cascade) a modification from the parent to the referencing child, but not vice versa.

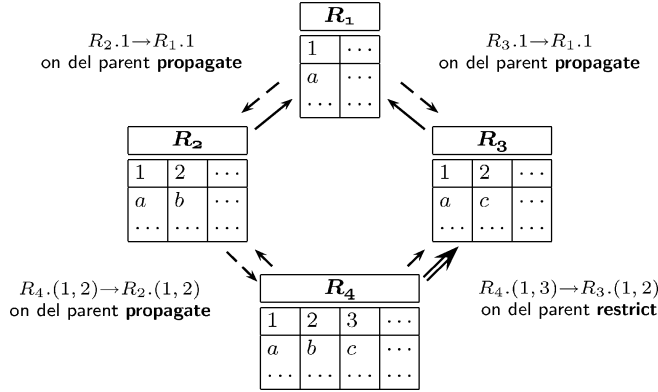


Fig. 1. Database with referential actions.

- (2) If execution follows the path $R_1 \rightsquigarrow R_2 \rightsquigarrow R_4$, the tuples $R_2(a, b)$ and $R_4(a, b, c)$ are requested for deletion. Hence, the *rac* for $R_4 \cdot (1, 3) \mid R_3 \cdot (1, 2)$ can assume that $R_4(a, b, c)$ is deleted, thus no referencing tuple exists in R_4 . Therefore, all deletions can be executed, resulting in a new database state $D^0 \bowtie D$.

In the above “diamond”, when choosing the “right” order of execution, the update is possible, whereas when going the “wrong” way, it is impossible. The reason is that the restrict action looks at the “current” database, and this depends on the order of execution.

This type of ambiguity can be eliminated by specifying that restrictions are always evaluated with respect to the *original* database state instead of the current one (as it is done in SQL, see the following section). However, the situation is more complex for *racs* of the type wait which have to look at the *final* database state. As it turns out, in the presence of modifications, in general, there are still several “equally justified” final states, each of which has to be considered:

Example 4 Mutex. Consider modifications $\triangleright \text{mod}_R(a/b)$ and $\triangleright \text{mod}_R(a/c)$.³ They are mutually exclusive, since they cannot be executed simultaneously. In our logical formalization, both will be *undefined* in the well-founded model. Moreover, there will be *two* stable models, each of which makes one modify request true, and the other false.

Another type of ambiguity may arise due to “self-attacking” requests:

Example 5 Self-Attack. Assume a database with *racs* (again a “diamond” as in Example 3) such that $\triangleright \text{mod}_R([1/b, 2/c](a, a))$ triggers $\text{mod}_{R_1}([1/b], (a))$ and $\text{mod}_{R_2}([1/c], (a))$; $\text{mod}_{R_1}([1/b], (a))$ triggers $\text{mod}_{R_3}([1/b], (a))$, and $\text{mod}_{R_2}([1/c], (a))$ triggers $\text{mod}_{R_3}([1/c], (a))$. Since the original request $\triangleright \text{mod}_R([1/b, 2/c], (a, a))$ causes a conflict at R_3 , it cannot be executed. On the other hand, no *other* request is in conflict with it, so there is no independent

³Here, (as the “ \triangleright ” shows) the modifications come directly from an (already contradictory) user request. However, the Mutex scenario can also occur *indirectly* from a non-contradictory user request.

justification *not* to execute it. Thus, the original request “attacks” itself. In our formalization, there is *no total stable model*.

Example 6 Virgin Birth. This example shows that not every update which does *not* violate any *ric* is also reasonable with respect to the intended semantics:

Consider the following database schema: $R_P(1)$ and $R_C(1, 2)$ with a *ric* $R_C.1 \mid R_P.1$ on del parent wait and a database instance $R_P(a, b)$, $R_P(d, e)$, $R_C(a, h)$. Suppose the user requests $U_{\triangleright} \sqsupset \text{del}_{\triangleright} R_P(a, b)$, $\triangleright \text{mod}_{\triangleright} R_P([1/a], (d, e))g$. Then, deletion of $R_P(a, b)$ is blocked due to the tuple $R_C(a, h)$.

On the other hand, one can argue that deleting $R_P(a, b)$ is possible, since after modifying $R_P(d, e)$ to $R_P(a, e)$, the child tuple $R_C(a, h)$ gets a new parent, so the *ric* $R_C.1 \mid R_P.1$ remains satisfied.

Here, the *semantical* connection which is encoded in $R_C.1 \mid R_P.1$ would be broken: The child tuple $R_C(a, h)$ gets a new parent *although* it is not modified, and the new parent tuple $R_P(a, e)$ “finds” a new child.

In this situation, $U_{\triangleright} \sqsupset \text{del}_{\triangleright} R_P(a, b)$, $\triangleright \text{mod}_{\triangleright} R_P([1/a], (d, e))g$ is not *feasible* according to our abstract semantics.

The underlying idea for treating this (and related) cases is based on the intended semantics of a database and referential integrity: Each database can be seen as a reference network (akin to the network database model), inducing a reference graph. Thus, each set of updates also defines a mapping between reference graphs. From the semantical point of view, references should only be created when a *child* tuple is inserted or modified. Changes on parent tuples are intended to either preserve references (by updating the child tuple accordingly) or delete references.

The above examples showed that the *local* ECA-style characterization considered in Section 2.2 is ambiguous. This ambiguity is caused by considering the *current* database state for applying referential actions. In the following section, we describe the SQL specification of referential actions that solves this problem on the specification level, but whose implementation aspects still suffered from these problems as long as SQL’s trigger functionality was used.

2.4 Referential Actions in SQL and Global Disambiguation Strategies

In SQL [ANSI/ISO 1992a, 1999], referential actions for a referential integrity constraint $R_C.\bar{K} \mid R_P.\bar{K}$ are specified with the definition of the child table. SQL allows referential actions only for modifications at the parent tuple:

```
fCREATE j ALTERg TABLE  $R_C$ 
  ccc
  FOREIGN KEY  $\bar{K}$  REFERENCES  $R_P \bar{K}$ 
  [ON UPDATE fCASCADE j RESTRICT j SET NULL j SET DEFAULT j NO ACTIONg]
  [ON DELETE fCASCADE j RESTRICT j SET NULL j SET DEFAULT j NO ACTIONg]
  ccc
```

Insertions and modifications on child tuples are handled in a straightforward way by rejecting updates which aim to generate a child tuple whose

corresponding parent does not exist. In our work, we deliberately exclude SET NULL/DEFAULT actions, since they are a special case of modifications.

In their *abstract* specification in SQL, these strategies correspond to the abstract local strategies described in Section 2.2, solving the problems that are caused by considering the *current* database state in our localized ECA-like characterization:

- CASCADE is the same as our propagate and propagates the update from the parent to the referencing tuples (evaluation is with respect to the *original* database state).
- RESTRICT is similar to restrict, but refers to the database state *before* the beginning of evaluation (instead of the current database state at the time when the update actually occurs): reject an update on the parent if there exists a child referencing it in the *original* database state,
- NO ACTION is the same as wait: no local action is executed. Instead, the referential integrity constraints are checked—together with the other integrity constraints—after the end of a certain unit of work, thus, referring to the database state *after* completing the updates.

Although this semantics is easy to understand on first sight, it is not a direct, *local* semantics that can be implemented straightforwardly by ECA-rules or triggers in the style of Dayal [1988] and Eswaran [1976]. Thus, its detailed specification in the SQL standard and even more, its realization in actual database systems has proven to be problematic.

Since the final state depends on the updates to be executed, and these may in turn depend on the final state via NO ACTION, there is a (negative) cyclic dependency in the *global* strategy. Thus, any straightforward implementation via ECA-rules/triggers is bound to fail. On the other hand, logic programming semantics provide a natural solution to this kind of problem and will be used in Sections 3 and 4 for an unambiguous, elegant characterization of the *global* semantics of *racs*.

The SQL2 standard [ANSI/ISO 1992a]—where RESTRICT did not yet exist—used the “localized,” immediate specification for cascading updates: the foreign key values in all referencing tuples were immediately updated when the parent was updated. Date and Darwen [1994] and Date [1990] already report the problem of unpredictable behavior, that is, that the outcome depends on the order in which tuples are updated. The same characterization was given in the upcoming SQL3 drafts (e.g., the 1991 version [ANSI/ISO 1991] cited by Horowitz [1992] and ANSI/ISO [1994]). Concerning these specifications, Horowitz [1992] and Cochrane et al. [1996] complain about unpredictable behavior since the outcome depends on the order in which rows are modified or constraints are applied [Cochrane et al. 1996]. In Horowitz [1992], a marking algorithm for a runtime execution model for referential integrity maintenance is presented for unary keys that does not exhibit these problems.

Markowitz [1994] presents safeness conditions which aim at avoiding ambiguities at the schema level. However, as shown in Reinert [1996], it is in general

undecidable whether a database schema with *rac*s is ambiguous (if D and U_{\triangleright} are given, the problem becomes decidable).

This was the state-of-the-art when we started to investigate a logic-based specification of referential actions in order to provide an unambiguous, natural, and “correct” semantics that mirrors the SQL intension described above for CASCADE, RESTRICT, and NO ACTION.

In the meantime, the SQL3 standard [ANSI/ISO 1999] solved the problem of ambiguous semantics of *rac*s by fixing an operational semantics using a marking algorithm based on Horowitz [1992]:

- effectively perform integrity checking at the end of the statement [includes NO ACTION],
- effectively determine and fix matching rows [concerning keys/foreign keys] at the beginning of the statement [(“static matching”)],
- effectively determine update values at the beginning of the statement (includes marking tuples that will be deleted),
- rollback any attempt to update the same data item [in [Horowitz 1992], single attributes] to representationally different values in the same statement,
- perform an update referential action iff the referenced column [Horowitz [1992] was restricted to single-column keys] is updated to a representationally different value,
- effectively perform all deletes at the end of the statement,
- perform a delete referential action if and only if the referenced row has not already been marked for deletion.

Here, “effectively” means that this description specifies the intended semantics; nevertheless, actual implementations can perform actions “on the fly”—if it is guaranteed that this does not violate the above semantics. The actual specification is given in terms of a complex encoding into BEFORE triggers whose interactions are hard to understand. The algorithm associates a *global* semantics with referential actions that will be shown to be equivalent to ours. Horowitz [1992] proves correctness and termination of the algorithm.

Such a procedural specification in form of an algorithm deviates far from the localized, ECA-style specification “ON event action” above, and gives no declarative, easily accessible, semantics of referential actions. Our work shows that the ECA-style specification has an immediate, declarative, “natural” global semantics that is given by the logic programming meta-semantics of rule-based specifications—and that the procedural semantics (developed over about 20 years) coincides with that semantics.

Moreover, in case a set of updates causes referential problems, the transaction is simply aborted. Often in these cases, most of the requested updates are unproblematic, and only one or two are not allowed. Thus, it can be useful to return hints on how to prepare a revised request that realizes the intended changes *and* is accepted by the system, or for debugging the application. We show how to derive such hints from our semantics, and we also sketch how it can be used for deriving suggestions how to correct the behavior of an application in that case.

In Sections 3 (Deletions) and 4 (Modifications), we show how to characterize and qualify the induced semantic problems using different (logical and game-theoretic) characterizations of *racs*.

3. SEMANTICS OF REFERENTIAL ACTIONS WITH DELETIONS

We have shown that in order to avoid ambiguities and nondeterminism as in Example 3, it is necessary to specify the intended *global* semantics of *racs*. In this section, we investigate *racs* $R_C.\vec{K}! R_P.\vec{K}$ with corresponding *racs* of the form

$$R_C.\vec{K}! R_P.\vec{K} \text{ ON DELETE } \{ \text{CASCADE} \mid \text{RESTRICT} \mid \text{NO ACTION} \}$$

according to the above *global* strategies in the SQL sense. For these we present an efficient (PTIME) algorithm that computes the unique solution.

First, we define an abstract, non-constructive semantics that formalizes the SQL notions described in Section 2.4. This semantics then serves as the basis for a notion of correctness. Next, we show how to translate a set of *racs* into a logic program, whose declarative semantics then provides a constructive definition. An equivalent game-theoretic characterization is developed which will be used to prove the correctness of the logic programming semantics with respect to the abstract semantics.

3.1 Abstract Semantics

Let D be a database represented as a set of ground atoms, RA a set of *racs*, and $U_\triangleright \triangleq \{ \triangleright \text{del}_R(x_1), \dots, \triangleright \text{del}_R(x_n) \}$ a set of (external) *user delete requests* which are passed to the system. D and RA define three graphs \mathcal{DC} (ON DELETE CASCADE), \mathcal{DR} (ON DELETE RESTRICT), and \mathcal{DN} (ON DELETE NO ACTION) corresponding to the different types of references:

$$\begin{aligned} \mathcal{DC} : & \{ \triangleright \text{del}_R(x), R_P(y) \mid \triangleright \text{del}_R(x) \in U_\triangleright \} \\ & R_C.\vec{K}! R_P.\vec{K} \text{ ON DELETE CASCADE} \triangleq RA \text{ and } \vec{x}[\vec{K}] \triangleright \vec{y}[\vec{K}], \end{aligned}$$

\mathcal{DR} and \mathcal{DN} are defined analogously. \mathcal{DC}' denotes the reflexive transitive closure of \mathcal{DC} . Note that the graphs describe *potential* interactions due to *racs*, independent of the given user requests U_\triangleright .

Definition 3.1 Abstract Properties. Given RA , D , and U_\triangleright as above, a set $\Delta \triangleq \{ \triangleright \text{del}_R(x_1), \dots, \triangleright \text{del}_R(x_n) \}$ of delete requests is called

- founded*, if for all $\triangleright \text{del}_R(x) \in \Delta$, there is a $\triangleright \text{del}_R(x^0) \in U_\triangleright$ s.t. $(R(x), R^0(x^0)) \in \mathcal{DC}'$ (note that here, we need reflexivity for covering $R^0(x^0)$ itself),
- complete*, if $\triangleright \text{del}_R(y) \in \Delta$ and $(R_C(x), R_P(y)) \in \mathcal{DC}$ implies $\triangleright \text{del}_R(x) \in \Delta$,
- feasible*, if (i) $(R_C(x), R_P(y)) \in \mathcal{DR}$ implies $\triangleright \text{del}_R(y) \notin \Delta$, and (ii) $\triangleright \text{del}_R(y) \in \Delta$ and $(R_C(x), R_P(y)) \in \mathcal{DN}$ implies $\triangleright \text{del}_R(x) \in \Delta$,
- admissible*, if it is founded, complete, and feasible.

(for individual updates, we also write “ $\triangleright \text{del}_R(x)$ is admissible” instead of “ $\triangleright \text{del}_R(x)$ is admissible”; analogous for “founded.”)

Foundedness guarantees that all deletions are “justified” by some user request, completeness guarantees that no cascading deletions are “forgotten” (see the next lemma), and feasibility ensures that RESTRICT/NO ACTION *rics* are “obeyed.”

For given external updates, the induced internal updates are characterized by the transitive closure of \mathcal{DC} :

Definition 3.2 Induced Updates. For given $RA, D, U_{\triangleright}$, and $U \mu U_{\triangleright}$

$\Delta(U) :D \text{ fdel_}R(\bar{x}) \text{ j there is a } \triangleright\text{del_}R^0(\bar{y}) \triangleright U \text{ s.t. } (R(\bar{x}), R^0(\bar{y})) \triangleright \mathcal{DC}'g$

is called the set of *induced updates of U*.

LEMMA 3.3 INDUCED UPDATES. For given $RA, D, U_{\triangleright}$, and $U \mu U_{\triangleright}$, $\Delta(U)$ is the least set⁴ Δ which contains the updates given by U and is complete.

PROOF. \mathcal{DC}' is the reflexive, transitive closure of \mathcal{DC} . Hence $\Delta(U)$ contains all user requested updates and all cascaded updates (completeness) and nothing else (least set). \square

Definition 3.4 Admissibility and Application of U_{\triangleright} . Let RA, D , and U_{\triangleright} be given. $U \mu U_{\triangleright}$ is *admissible* if $\Delta(U)$ is admissible, and *maximal admissible* if there is no other admissible U^0 , such that $U \subsetneq U^0 \mu U_{\triangleright}$. For a set Δ of user requests, $D^0 \triangleright D \S \Delta$ denotes the database obtained by applying Δ to D .

This definition provides a precise and elegant characterization of the intended semantics. However, it is non-constructive in the sense that it does not lend itself to a computation of the intended semantics.

From the above fundamental definitions, we derive the following:

PROPOSITION 3.5 CORRECTNESS.

- a) If $U \mu U_{\triangleright}$, then $\Delta(U)$ is founded and complete.
- b) If Δ is complete and feasible, then $D^0 :D D \S \Delta(U)$ satisfies all *rics*.

PROOF. a) $\Delta(U)$ is defined as the *least* complete set. Since $U \mu U_{\triangleright}$, $\Delta(U)$ is founded.

b) Completeness guarantees that all *rics* labeled with ON DELETE CASCADE in RA are satisfied, feasibility guarantees that all *rics* labeled with ON DELETE RESTRICT/NO ACTION are satisfied. \square

PROPOSITION 3.6 UNIQUENESS. For given RA, D , and U_{\triangleright} ,

- (i) if $U_1, U_2 \mu U_{\triangleright}$ are admissible, then $U_1 \sqcup U_2$ is also admissible,
- (ii) thus, there is exactly one maximal admissible $U_{\max} \mu U_{\triangleright}$.

PROOF. (i) is obvious. (ii) follows from (i) together with the fact that μ is always admissible. Thus, the union of all admissible subsets of U_{\triangleright} yields U_{\max} . \square

Note that for an admissible set U , not necessarily each subset is also admissible (there can be updates that “need” each other to be feasible).

⁴i.e., there is no proper subset that satisfies the required properties, and it is the only minimal set.

3.2 Logic Programming Characterization

We show how a set RA of *rac*s can be translated into a logic program P_{RA} whose rules specify their *local* behavior. The advantage of this logical formalization is that the declarative semantics of P_{RA} defines a precise *global* semantics. Moreover, by choosing an appropriate evaluation strategy, this logical specification can be *executed* as well, yielding the desired constructive semantics.

Unblocked requests. The following rule derives for every user request $\triangleright\text{del}_R(\bar{x}) \not\subseteq U_\triangleright$ an *internal delete request* $\text{req_del}_R(\bar{x})$, provided there is no *blocking* $\text{blk_del}_R(\bar{x})$.

$$\text{req_del}_R(\bar{X}) \wedge \triangleright\text{del}_R(\bar{X}), R(\bar{X}), : \text{blk_del}_R(\bar{X}). \quad (I)$$

Referential actions. Each referential action is specified by an appropriate rule:

— $R_C.\bar{F}! R_P.\bar{K}$ ON DELETE CASCADE is encoded into two rules: the first one propagates internal delete requests downwards from the parent to the child:

$$\text{req_del}_{R_C}(\bar{X}) \wedge \text{req_del}_{R_P}(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \supset \bar{Y}[\bar{K}]. \quad (DC_1)$$

Additionally, blockings are propagated upwards, that is, when the deletion of a child is blocked, the deletion of the referenced parent is also blocked:

$$\text{blk_del}_{R_P}(\bar{Y}) \wedge R_P(\bar{Y}), \text{blk_del}_{R_C}(\bar{X}), \bar{X}[\bar{F}] \supset \bar{Y}[\bar{K}]. \quad (DC_2)$$

Note that the atom $R_P(\bar{Y})$ can be added to the body of (DC_1) , and $R_C(\bar{X})$ can be added to the body of (DC_2) , but are redundant since delete requests and blockings are only derived for tuples that actually exist.

— $R_C.\bar{F}! R_P.\bar{K}$ ON DELETE RESTRICT blocks the deletion of a parent tuple if there is a corresponding child tuple:

$$\text{blk_del}_{R_P}(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \supset \bar{Y}[\bar{K}]. \quad (DR)$$

— $R_C.\bar{F}! R_P.\bar{K}$ ON DELETE NO ACTION blocks the deletion of a parent tuple if there is a corresponding child tuple that is not requested for deletion:

$$\text{blk_del}_{R_P}(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), : \text{req_del}_{R_C}(\bar{X}), \bar{X}[\bar{F}] \supset \bar{Y}[\bar{K}]. \quad (DN)$$

Note that (i) the *local semantics* of each individual *rac* is precisely specified by one or two logic rules, and (ii) P_{RA} is in general not stratified due to the negative cyclic dependency $\text{req_del} \rightsquigarrow \text{blk_del} \rightsquigarrow \text{req_del}$. Therefore, the *global semantics* is not necessarily unique. We first consider a “skeptical” global semantics: the unique well-founded model of the generated logic program. The more “brave” stable models are considered in Section 3.4.2.

First, we add two rules that define an auxiliary relation pot_del , which contains all tuples that are potentially deleted when executing U_\triangleright and cascading deletions. This relation is not used for checking the admissibility of U_\triangleright , but, as shown in Section 3.7, is useful to locate the problems in case U_\triangleright is *not* admissible:

$$\begin{aligned} & \text{pot_del}_R(\bar{X}) \wedge \triangleright\text{del}_R(\bar{X}), R(\bar{X}). \\ & \text{for each } R_C.\bar{F}! R_P.\bar{K} \text{ ON DELETE CASCADE (analogous to } (DC_1)\text{):} \quad (P) \\ & \text{pot_del}_{R_C}(\bar{X}) \wedge \text{pot_del}_{R_P}(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \supset \bar{Y}[\bar{K}]. \end{aligned}$$

18 † W. May and B. Ludäscher

(note that the atom $\text{pot_del_}R_P(\bar{Y})$ can be added to the body of (DC_2) , (DR) , and (DN) as an optimization.)

We have the following simple lemma:

LEMMA 3.7. *Given a database D and a set of user requests U_\triangleright , for the minimal model $\mathcal{M} : \triangleright \mathcal{M}(\text{f}P\text{g}, D, U_\triangleright)$ of rule (P) alone,*

$$U_{\text{pot}} : \triangleright \text{f} \text{del_}R(\bar{x}) \text{ j } \mathcal{M}(\text{pot_del_}R(\bar{x})) \triangleright \text{trueg} \\ \triangleright \text{f} \text{del_}R(\bar{x}) \text{ j } \text{there is a } \triangleright \text{del_}R^0(\bar{x}^0) \triangleright U_\triangleright \text{ and } (R(\bar{x}), R^0(\bar{x}^0)) \triangleright \mathcal{DC}'\text{g} \triangleright \Delta(U_\triangleright)$$

contains exactly the deletions that are obtained when cascading all deletions in U_\triangleright .

Well-Founded Semantics. The well-founded model [Van Gelder et al. 1991] is widely accepted as a (skeptical) declarative semantics for logic programs containing negation. Given a database D and a set of user requests U_\triangleright , the well-founded model $\mathcal{W} : \triangleright \mathcal{W}(P_{RA}, D, U_\triangleright)$ assigns truth-values true and false to all uncontroversial update requests—those that are true or false under any *reasonable semantics* of P_{RA} [Dix 1995]. \mathcal{W} assigns a third truth value *undefined* to atoms whose truth cannot be determined using a “well-founded” argumentation. The atoms that are undefined in \mathcal{W} are controversial due to some kind of ambiguity (cf. Section 2.3). In Section 3.4.1, we will prove the following:

THEOREM 3.8 CORRECTNESS.

The logic programming characterization is correct with respect to the abstract semantics (for an atom at , Let $\mathcal{W}(at)$ denote the value of at in the model $\mathcal{W} : \triangleright \mathcal{W}(P_{RA}, D, U_\triangleright)$):

$$\begin{aligned} & \text{— } U_t : \triangleright \text{f} \triangleright \text{del_}R(\bar{x}) \triangleright U_\triangleright \text{ j } \mathcal{W}(\text{req_del_}R(\bar{x})) \triangleright \text{trueg} \text{ and} \\ & \quad U_{t,u} : \triangleright \text{f} \triangleright \text{del_}R(\bar{x}) \triangleright U_\triangleright \text{ j } \mathcal{W}(\text{req_del_}R(\bar{x})) \triangleright \text{ftrue, undefgg} \text{ are admissible,} \\ & \text{— } U_{t,u} \triangleright U_{\text{max}}, \text{ and} \\ & \text{— } \Delta(U_{\text{max}}) \triangleright \Delta(U_{t,u}) \triangleright \text{f} \text{del_}R(\bar{x}) \text{ j } \mathcal{W}(\text{req_del_}R(\bar{x})) \triangleright \text{ftrue, undefgg.} \end{aligned}$$

COROLLARY 3.9. *In case that U_\triangleright is admissible, we have $U_\triangleright \triangleright U_{\text{max}}$ and $U_{\text{pot}} \triangleright \Delta(U_\triangleright)$.*

Often, even if not all requested updates can be accomplished, a subset of them is admissible. Thus, the information as to which tuple or update really causes problems is valuable for preparing a refined update that realizes the intended changes *and* is acceptable. In Section 3.7, we will systematically investigate the information that is available in case U_\triangleright is not admissible.

Example 7. Consider the database depicted in Figure 2 (ignoring R_0 for now) and the user request $U_\triangleright \triangleright \text{f} \triangleright \text{del_}R_1(a), \triangleright \text{del_}R_1(b)\text{g}$. Here, $\text{del_}R_1(b)$ is not admissible since it is blocked by $R_5(b)$. The other request, $\text{del_}R_1(a)$, can be executed without violating any *ric* by deleting $R_1(a)$, $R_2(a, x)$, $R_3(a, y)$, and $R_4(a, x, y)$.

The well-founded semantics reflects the different status of the single updates:

Given the user request $U_{\triangleright_a} \triangleright \text{f} \triangleright \text{del_}R_1(a)\text{g}$, the delete requests req_del for $R_1(a)$, $R_2(a, x)$, $R_3(a, y)$, $R_4(a, x, y)$, as well as the blockings blk_del for $R_1(a)$ and $R_3(a, y)$ will be *undefined* in the well-founded model.

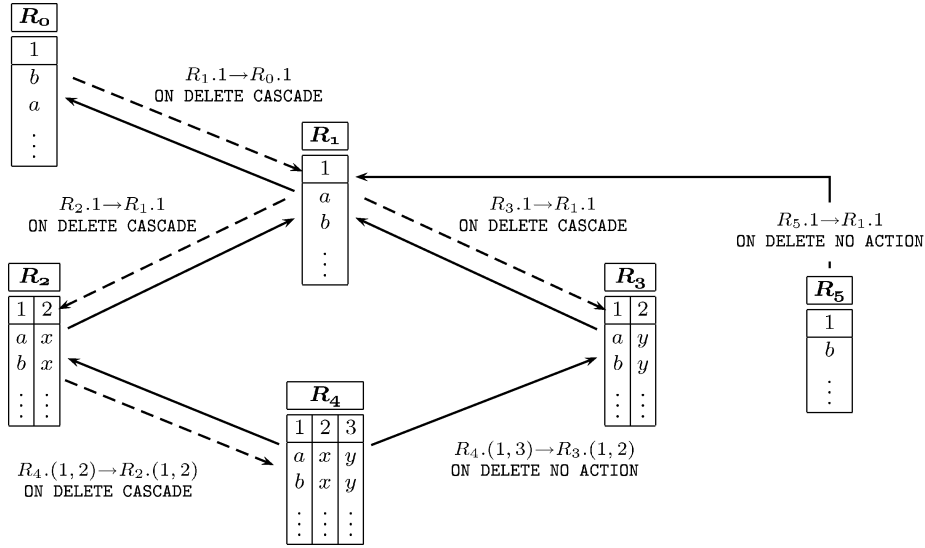


Fig. 2. Extended database with modified *racs*.

For the user request $U_{\triangleright b} \text{ Df } \triangleright \text{del}_{R_1}(b)g$, blk_del is *true* for $R_1(b)$ due to the referencing tuple $R_5(b)$. Thus, $\text{req_del}_{R_1}(b)$ is *false*, and $\text{del}_{R_1}(b)$ is not admissible; hence there are no cascaded delete requests. Due to the referencing tuple $R_4(b, x, y)$, which cannot be deleted in this case, $\text{blk_del}_{R_3}(b, y)$ is also *true*.

Note that the extended set $U_{\triangleright}^0 \text{ Df } \triangleright \text{del}_{R_1}(a), \triangleright \text{del}_{R_1}(b), \triangleright \text{del}_{R_5}(b)g$ is a candidate for a refined request which accomplishes the deletion of $R_1(a)$ and $R_1(b)$.

\mathcal{W} contains some ambiguities that can be interpreted constructively as *degrees of freedom*: The blockings and deletions induced by $U_{\triangleright} \text{ Df } \triangleright \text{del}_{R_1}(a)g$ in Example 7 are undefined due to the dependency $\text{req_del} \rightsquigarrow \text{blk_del} \rightsquigarrow \text{req_del}$. This may be used to define different global policies by giving priority either to deletions or blockings, as will be done in Section 3.4.2.

3.3 Game-Theoretic Characterization

The following game-theoretic formalization provides an elegant characterization of *racs* which yields additional insight into the well-founded model of P_{RA} and the intuitive meaning of *racs*.

The game is played with a pebble by two players, I (the “Deleter”) and II (the “Spoiler”), who argue whether a tuple may be deleted. The players move alternately in *rounds*; each round consists of two *moves*. A player who cannot move loses. The set of *positions* of the game is $D \sqcup U_{\triangleright} \sqcup \text{restricted}g$. The possible moves of I and II are defined below. Note that I moves from D to U_{\triangleright} , while II moves from U_{\triangleright} to $D \sqcup \text{restricted}g$. Initially, the pebble is placed on some tuple in D (or U_{\triangleright}) and I (or II) starts to move. If II begins, the first round only consists of the move by II .




Fig. 3. Possible Moves.

The possible moves are illustrated in Figure 3: By moving the pebble from $R(\bar{x}) \in D$ to some $\triangleright del_R^0(\bar{x}^0) \in U_\triangleright$, which cascades down to $R(\bar{x})$, I claims that the deletion of $R(\bar{x})$ is “justified” (i.e., founded) by $\triangleright del_R^0(\bar{x}^0)$. Conversely, II claims by her moves that $\triangleright del_R^0(\bar{x}^0)$ is not feasible. II can use two different arguments: Assume that the deletion of $R^0(\bar{x}^0)$ cascades down to some tuple $R_P(\bar{x}_P)$. First, if the deletion of $R_P(\bar{x}_P)$ is restricted by a referencing child tuple $R_C(\bar{x}_C)$, then II may force I into a lost position by moving to restricted (since I cannot move from there). Second, assume that the deletion of $R^0(\bar{x}^0)$ cascades down to some other tuple $R_P^0(\bar{x}_P^0)$. Then, II can move to a child tuple $R'_C(\bar{x}'_C)$, which references $R_P^0(\bar{x}_P^0)$ with a NO ACTION *rac*. With this move, II claims that this reference to $R_P^0(\bar{x}_P^0)$ will remain in the database, so $R_P^0(\bar{x}_P^0)$ and, as a consequence, $R^0(\bar{x}^0)$ cannot be deleted. In this case, I may start a new round of the game by finding a justification to delete the referencing child $R'_C(\bar{x}'_C)$. More precisely:

Player I can move from $R(\bar{x})$ to $\triangleright del_R^0(\bar{x}^0) \in U_\triangleright$ if $(R(\bar{x}), R^0(\bar{x}^0)) \in DC'$.

Player II can move from $\triangleright del_R^0(\bar{x}^0)$

- to restricted if there are $R_P(\bar{x}_P)$ and $R_C(\bar{x}_C)$ such that $(R_P(\bar{x}_P), R^0(\bar{x}^0)) \in DC'$ and $(R_C(\bar{x}_C), R_P(\bar{x}_P)) \in DR$.
- to $R'_C(\bar{x}'_C)$, if $(R_P^0(\bar{x}_P^0), R^0(\bar{x}^0)) \in DC'$ and $(R'_C(\bar{x}'_C), R_P^0(\bar{x}_P^0)) \in DN$.

LEMMA 3.10 CLAIMS OF I AND II.

- (1) If I can move from $R(\bar{x})$ to $\triangleright del_R^0(\bar{x}^0)$, then deletion of $R^0(\bar{x}^0)$ is founded by U_\triangleright and induces the deletion of $R(\bar{x})$.
- (2) If II can move from $\triangleright del_R^0(\bar{x}^0)$ to restricted, then deletion of $R(\bar{x})$ is not feasible due to the existence of a referencing tuple.
- (3) If II can move from $\triangleright del_R^0(\bar{x}^0)$ to $R^0(\bar{x}^0)$, then deletion of $R(\bar{x})$ is admissible only if $R^0(\bar{x}^0)$ is also deleted.

PROOF. (1) The move of I implies that $(R(\bar{x}), R^0(\bar{x}^0)) \in DC'$.

The move of II means that either

- (2) there are $R_P(\bar{x}_P)$, $R_C(\bar{x}_C)$ such that $(R_P(\bar{x}_P), R(\bar{x})) \in DC'$ and $(R_C(\bar{x}_C), R_P(\bar{x}_P)) \in DR$. Then, deletion of $R(\bar{x})$ induces the deletion of $R_P(\bar{x}_P)$, but the deletion of $R_P(\bar{x}_P)$ is restricted by $R_C(\bar{x}_C)$, or

(3) $(R_C^0(\bar{x}_C), R(\bar{x})) \not\leq \mathcal{DN} - \mathcal{DC}'$, that is, there is a $R_P^0(\bar{x}_P)$ such that $(R_P^0(\bar{x}_P), R(\bar{x})) \leq \mathcal{DC}'$ and $(R_C^0(\bar{x}_C), R_P^0(\bar{x}_P)) \leq \mathcal{DN}$. Hence, by (1), deletion of $R(\bar{x})$ induces deletion of $R_P^0(\bar{x}_P)$, which is only allowed if $R_C^0(\bar{x}_C)$ is also deleted.⁵ \square

LEMMA 3.11. *The moves are correlated with the logical specification as follows:*

- The moves of I correspond to rule (DC_1) : I can move from $R(\bar{x})$ to $\triangleright del.R^0(\bar{x}^0)$ if, given the fact $req_del.R^0(\bar{x}^0)$, $req_del.R(\bar{x})$ can be derived using (DC_1) .
- The moves by Π are reflected by the rules (DC_2) and $(DR)/(DN)$:
- Π can move from $\triangleright del.R(\bar{x})$ to restricted if $blk_del.R(\bar{x})$ is derivable using (DR) and (DC_2) only, or
- Π can move from $\triangleright del.R(\bar{x})$ to $R_C^0(\bar{x}_C^0)$ if $blk_del.R(\bar{x})$ is derivable using (DC_2) and an instance of (DN) if $req_del.R_C^0(\bar{x}_C^0)$ is assumed to be false.
- The negative dependencies in (I) , $req_del \rightsquigarrow : blk_del$, and (DN) , $blk_del \rightsquigarrow : req_del$, mirror the alternation of moves between I and Π , respectively.

Definition 3.12. A position $R(\bar{x}) \leq D$ is won (for I), if I can win the game starting from $R(\bar{x})$ no matter how Π moves. If p is won (lost) for a player, p is lost (won) for the opponent. A position which is neither lost nor won is drawn. In the sequel, “is won/lost” stands for “is won/lost for I .” An update $\triangleright del.R(\bar{x}) \leq U_\triangleright$ is won if $R(\bar{x}) \leq D$ is won.

Drawn positions can be viewed as ambiguous situations. For the game above, this means that neither can I prove in finitely many moves that $R(\bar{x})$ has to be deleted, nor can Π prove that it is infeasible to delete $R(\bar{x})$.

Example 8. Consider again Figure 2 with $U_\triangleright \sqsupset \triangleright del.R_1(a), \triangleright del.R_1(b)g$. From each of the “ a ”-tuples $R_1(a), R_2(a, x), R_3(a, y), R_4(a, x, y)$, I can move to $\triangleright del.R_1(a)$, while Π can move from $\triangleright del.R_1(a)$ to $R_4(a, x, y)$. Thus, after I has started the game moving to $\triangleright del.R_1(a)$, Π will answer with the move to $R_4(a, x, y)$, so I moves back to $\triangleright del.R_1(a)$ again, and so forth. Hence the game is drawn for each of the “ a ”-tuples.

In contrast, for the “ b ”-tuples, there is an additional move from $\triangleright del.R_1(b)$ to $R_5(b)$ for Π , who now has a winning strategy: by moving to $R_5(b)$, there is no possible answer for I , so I loses.

THEOREM 3.13 GAME SEMANTICS. *For every tuple $R(\bar{x}) \leq D$:*

- $R(\bar{x})$ is lost , it is not possible with the given set of user delete requests to delete $R(\bar{x})$ without violating a ric .
- $R(\bar{x})$ is won or drawn , simultaneous execution of all user delete requests $\triangleright del.R^0(\bar{x}^0)$ that are won or drawn does not violate any ric and deletes $R(\bar{x})$.

PROOF. Note that if $R(\bar{x})$ is won or drawn, then there is no $R_C(\bar{x}_C) \leq D$ such that $(R_C(\bar{x}_C), R(\bar{x})) \leq \mathcal{DR}$ (otherwise, if I moves from $R(\bar{x})$ to some $\triangleright del.R_d(\bar{x}_d)$, Π moves to restricted since $(R_C(\bar{x}_C), R_d(\bar{x}_d)) \leq \mathcal{DR} - \mathcal{DC}'$ and wins). Thus, no ric of the form ON DELETE RESTRICT is violated when deleting a won or drawn tuple.

⁵ $\mathcal{DN} - \mathcal{DC}' : \sqsupset f(x, y) \sqsupset g(z : (x, z) \leq \mathcal{DN}$ and $(z, y) \leq \mathcal{DC}'g$.

22 † W. May and B. Ludäscher

“() ”: A tuple $R(\bar{x})$ is lost in n rounds if either

- ($n \geq 0$) there is no user request $\triangleright\text{del}_d R_d(\bar{x}_d)$ such that $(R(\bar{x}), R_d(\bar{x}_d)) \in \mathcal{DC}'$, i.e., the deletion of $R(\bar{x})$ is unfounded, or
- ($n > 0$) for every user request $\triangleright\text{del}_d R_d(\bar{x}_d)$ such that $(R(\bar{x}), R_d(\bar{x}_d)) \in \mathcal{DC}'$, $\triangleright\text{del}_d R_d(\bar{x}_d)$ is lost in $\bullet n + 1$ rounds, that is, either Π can move from $\triangleright\text{del}_d R_d(\bar{x}_d)$ to restricted (in this case, by Lemma 3.10(2), $\triangleright\text{del}_d R_d(\bar{x}_d)$ is not feasible), or there is some tuple $R^0(\bar{x}^0)$ such as Π can move from $\triangleright\text{del}_d R_d(\bar{x}_d)$ to $R^0(\bar{x}^0)$ and which is lost in $\bullet n + 1$ rounds. By induction hypothesis, $R^0(\bar{x}^0)$ cannot be deleted, but by Lemma 3.10(3), it must be deleted if $R(\bar{x})$ is deleted. Thus, $R(\bar{x})$ cannot be deleted.

“(() ”: If $R(\bar{x})$ cannot be deleted without violating a *ric*, then either,

- deletion of $R(\bar{x})$ is unfounded—therefore it is lost immediately since Π cannot move,
- or deletion of $R(\bar{x})$ is founded, but none of its founding user delete requests $\triangleright\text{del}_d R^0(\bar{x}^0)$ is executable. This can be either due to a $\mathcal{DC}' - \mathcal{DR}$ chain to a tuple $R_C^0(\bar{x}_C^0)$ —then $\triangleright\text{del}_d R^0(\bar{x}^0)$ is lost in one round since Π moves to restricted— or due to a $\mathcal{DC}' - \mathcal{DN}$ chain to a tuple $R_C^0(\bar{x}_C^0)$ that must be deleted, but cannot. Then, Π can move there and will win (the detailed proof would argue with induction over the length of the proof *why* $\triangleright\text{del}_d R^0(\bar{x}^0)$ is not executable, analogous to the proof of “() ”).
- The second statement follows from the first by contraposition. \square

The correspondence between the game semantics and the abstract semantics yields the following:

COROLLARY 3.14 CORRECTNESS. *The game-theoretic characterization is correct with respect to the abstract semantics:*

- $U_w : \text{D } \text{fu} \in U_{\triangleright} \mid u \text{ is won}$ and $U_{w,d} : \text{D } \text{fu} \in U_{\triangleright} \mid u \text{ is won or drawn}$ are admissible,
- $U_{w,d} \supseteq U_{\max}$,
- $\Delta(U_w) \supseteq \text{del}_d R(\bar{x}) \mid R(\bar{x}) \text{ is won}$ and $\Delta(U_{\max}) \supseteq \Delta(U_{w,d}) \supseteq \text{del}_d R(\bar{x}) \mid R(\bar{x}) \text{ is won or drawn}$.

3.4 Equivalence and Correctness

We show that the logical characterization is equivalent to the game-theoretic one. Thus, the correctness of the logical characterization reduces to the correctness of the game-theoretic one proven above.

3.4.1 Well-Founded Semantics. The *alternating fixpoint computation* (AFP) is a method for computing the well-founded model based on successive rounds [Van Gelder 1993]. This characterization finally leads to an algorithm for determining the maximal admissible subset of a given set U_{\triangleright} of user requests. We introduce the AFP by using **Statelog**, a state-oriented extension of Datalog which allows the integration of active and deductive rules Ludäscher et al. [1996a]; Ludäscher [1998]. It can be seen as a restricted class of logic programs

where every intensional predicate contains an additional distinguished argument for *state terms* of the form $[S \subset k]$. EDB predicates and built-in predicates are state-independent. Here, S is the distinguished *state variable* ranging over \mathbb{N}_0 . Statelog rules are of the form

$$[S \subset k_0] H(\bar{X}) \wedge [S \subset k_1] B_1(\bar{X}_1), \dots, [S \subset k_n] B_n(\bar{X}_n),$$

where the head $H(\bar{X})$ is an atom, $B_i(\bar{X}_i)$ are atoms or negated atoms, and k_0, k_i , for all $i \in \{1, \dots, n\}$. A rule is *local* if $k_0 \sqsupset k_i$ for all $i \in \{1, \dots, n\}$.

In Statelog, the AFP is obtained by attaching state terms to the program P such that all positive IDB literals refer to $[S \subset 1]$ and all negative IDB literals refer to $[S]$. The resulting program P_{AFP} computes the alternating fixpoint of P :

$$[S \subset 1] \text{ req_del_}R(\bar{X}) \wedge \triangleright \text{del_}R(\bar{X}), R(\bar{X}), [S] : \text{ blk_del_}R(\bar{X}). \quad (I^A)$$

% R_C. \bar{K} ! R_P. \bar{K} ON DELETE CASCADE :

$$[S \subset 1] \text{ req_del_}R_C(\bar{X}) \wedge R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}], [S \subset 1] \text{ req_del_}R_P(\bar{Y}). \quad (DC_1^A)$$

$$[S \subset 1] \text{ blk_del_}R_P(\bar{Y}) \wedge R_P(\bar{Y}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}], [S \subset 1] \text{ blk_del_}R_C(\bar{X}). \quad (DC_2^A)$$

% R_C. \bar{K} ! R_P. \bar{K} ON DELETE RESTRICT :

$$[S \subset 1] \text{ blk_del_}R_P(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}]. \quad (DR^A)$$

% R_C. \bar{K} ! R_P. \bar{K} ON DELETE NOACTION :

$$[S \subset 1] \text{ blk_del_}R_P(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}], [S] : \text{ req_del_}R_C(\bar{X}). \quad (DN^A)$$

P_{AFP} is *locally stratified*, thus there is a *unique perfect model* [Przymusinski 1988] \mathcal{M}_{AFP} of P_{AFP} [$D \sqcup U \triangleright$. \mathcal{M}_{AFP} mimics the alternating fixpoint computation of \mathcal{W} : even-numbered states $[2n]$ correspond to the increasing sequence of underestimates of true atoms, while odd-numbered states $[2n \subset 1]$ represent the decreasing sequence of overestimates of true or undefined atoms. The *final state* n_f of the computation is reached if $\mathcal{M}[2n_f] \sqsupset \mathcal{M}[2n_f \subset 1]$. Then, the truth value of atoms A in \mathcal{W} can be determined from \mathcal{M}_{AFP} as follows:

$$\mathcal{W}(A) \sqsupset \begin{cases} \text{true} & \text{if } \mathcal{M}_{AFP} \sqsupset [2n_f] A, \\ \text{undef} & \text{if } \mathcal{M}_{AFP} \sqsupset [2n_f] : A \wedge [2n_f \subset 1] A, \\ \text{false} & \text{if } \mathcal{M}_{AFP} \sqsupset [2n_f \subset 1] : A. \end{cases}$$

THEOREM 3.15 EQUIVALENCE. *The well-founded model is equivalent to the game-theoretic characterization:*

- $R(\bar{x})$ is won , $\mathcal{W}(\text{ req_del_}R(\bar{x})) \sqsupset \text{true}$.
- $R(\bar{x})$ is lost , $\mathcal{W}(\text{ req_del_}R(\bar{x})) \sqsupset \text{false}$.
- $R(\bar{x})$ is drawn , $\mathcal{W}(\text{ req_del_}R(\bar{x})) \sqsupset \text{undefined}$.

PROOF. The proof is based on a lemma which follows from the correspondence between moves and reference chains that has been established in Lemma 3.11 (using the same argumentation as in the proof of Theorem 3.13):

LEMMA 3.16.

- $R(\bar{x})$ is won (for I) within $\bullet n$ rounds iff $\mathcal{M}_{AFP} \sqsupset [2n] \text{ req_del_}R(\bar{x})$.
- $R(\bar{x})$ is lost within $\bullet n$ rounds iff $\mathcal{M}_{AFP} \sqsupset [2n \subset 1] : \text{ req_del_}R(\bar{x})$.

From this, Theorem 3.15 follows immediately: The n th overestimate excludes deletions provably non-admissible in n rounds, whereas the n th underestimate contains all deletions which can be proven in n rounds. Thus, there is an n such that $\mathcal{M}_{AFP} \uparrow \mathbb{D} [2n] \text{ req_del_}R(\bar{x})$ iff $\mathcal{W}_{RA}(\text{req_del_}R(\bar{x})) \mathbb{D} \text{true}$, and there is an n such that $\mathcal{M}_{AFP} \uparrow \mathbb{D} [2n \subset 1] : \text{req_del_}R(\bar{x})$ iff $\mathcal{W}(\text{req_del_}R(\bar{x})) \mathbb{D} \text{false}$.

A position $R(\bar{x})$ is drawn if for every user request $\triangleright\text{del_}R^0(\bar{x}^0)$ that I uses for deleting it, II can find a witness against $\triangleright\text{del_}R^0(\bar{x}^0)$, and conversely, I claims to be able to delete the witness. Thus, no player has a “well-founded” proof for or against deleting those tuples (caused by NO ACTION links that introduce cycles into the argumentation). \square

From Corollary 3.14 and Theorem 3.15, the correctness of the logic programming formalization (and thus the proof of Theorem 3.8) follows. In the following section, it is shown that the maximal admissible subset $U_{\max} \uparrow U_{\triangleright}$ (by Theorem 3.8, $U_{\max} \mathbb{D} U_{t,u}$) also corresponds to a *total* (i.e., not involving atoms with an undefined truth value) semantics of P .

3.4.2 Stable Models. The undefined atoms in the well-founded model leave some scope for further interpretation. This “freedom of choice” can often be used to obtain alternative solutions, given by *stable models*:

Definition 3.17 Stable Model [Gelfond and Lifschitz 1988]. Let M_P denote the minimal model of a positive program P . Given a ground-instantiated program P and an interpretation I of the atoms of P , P/I denotes the *reduction* of P with respect to I —the program obtained by replacing every negative literal of P by its truth-value with respect to I . An interpretation I is a *stable model* if $M_{P/I} \mathbb{D} I$.

Every stable model S extends the well-founded model \mathcal{W} with respect to true and false atoms: $S^{\text{true}} \uparrow \mathcal{W}^{\text{true}}$, $S^{\text{false}} \uparrow \mathcal{W}^{\text{false}}$. However, not every program has a two-valued stable model (e.g., the “self attack” in Example 5).

THEOREM 3.18. *Let S_{RA} be defined by*

$$S_{RA} : \mathbb{D} D [U_{\triangleright} [\text{freq_del_}R(\bar{x}) \uparrow \mathcal{W}(\text{req_del_}R(\bar{x})) \mathbb{D} \text{true, undef} \uparrow \uparrow \\ [\text{blk_del_}R(\bar{x}) \uparrow \mathcal{W}(\text{blk_del_}R(\bar{x})) \mathbb{D} \text{true} \uparrow \uparrow]]]$$

Then, S_{RA} is a total stable model of $P_{RA} [D [U_{\triangleright}$.

S_{RA} is the “maximal” stable model in the sense that it contains all delete requests that are true in some stable model. Consequently, deletions have priority over blockings (cf. Example 7). For the diamond example, there are two stable models:

Example 9 Diamond—Stable Models. Consider Example 3 and the “diamond” in Figure 1. Assume the *rac* $R_4.(1, 3) \uparrow R_3.(1, 2)$ ON DELETE RESTRICT to be replaced by $R_4.(1, 3) \uparrow R_3.(1, 2)$ ON DELETE NO ACTION. From the rules of P_{RA} it can be derived that the deletion of $R_1(a)$ is blocked (via $R_4 \rightsquigarrow R_3 \rightsquigarrow R_1$) if $R_4(a, b, c)$ cannot be deleted. $R_4(a, b, c)$ can be deleted (via $R_1 \rightsquigarrow R_2 \rightsquigarrow R_4$) if the deletion of $R_1(a)$ is not blocked. Hence there is a negative cycle of the form

$\{block \wedge : exec, exec \wedge : block\}$. Setting all requests in the diamond to true (as done in S_{RA}) or all to false results each in a stable model.

THEOREM 3.19 CORRECTNESS.

—Let S be a stable model of $P_{RA} [D [U_{\triangleright}$. Then $U_S :D \triangleright del_R(\bar{x}) \triangleright U_{\triangleright} \mid S \mid req_del_R(\bar{x})$ is admissible and $\Delta(U_S) \triangleright \Delta_S :D \triangleright del_R(\bar{x}) \mid S \mid req_del_R(\bar{x})$.

— $U_{max} \triangleright U_{S_{RA}}$ and $\Delta(U_{max}) \triangleright \Delta_{S_{RA}}$.

PROOF. *Foundedness:* follows directly from the fact that S is stable (an unfounded $req_del_R(\bar{x})$ would not be stable).

Completeness: For every $ric R_C.\bar{F} \mid R_P.\bar{K}$ ON DELETE CASCADE, if $S \mid R_C(\bar{x}) \wedge req_del_R_P(\bar{y}) \wedge \bar{x}[\bar{F}] \triangleright \bar{y}[\bar{K}]$, then, due to (DC_1) , $S \mid M_{P/S} \mid req_del_R_C(\bar{x})$.

Feasibility: Suppose a $ric R_C.\bar{F} \mid R_P.\bar{K}$ ON DELETE RESTRICT or $R_C.\bar{F} \mid R_P.\bar{K}$ ON DELETE NO ACTION would be violated—Then $S \mid req_del_R_P(\bar{y}) \wedge R_C(\bar{x}) \wedge \bar{x}[\bar{F}] \triangleright \bar{y}[\bar{K}]$ (for NO ACTION also $S \mid req_del_R_C(\bar{x})$), and thus because of (DR) or (DN) , respectively, $S \mid M_{P/S} \mid blk_del_R_P(\bar{y})$. Thus, by (DC_2) , for the founding delete request $del_R(\bar{z})$, $S \mid blk_del_R(\bar{z})$, and by (I) , $S \mid req_del_R(\bar{z})$, which is a contradiction to the assumption that $del_R(\bar{z})$ is the founding delete request. $\Delta_S \mid \Delta(U_S)$ follows from foundedness, and $\Delta_S \nmid \Delta(U_S)$ follows from completeness. \square

3.5 A Procedural Translation

The declarative semantics of the well-founded model is translated into a more “algorithmic” implementation in Statelog by “cutting” the cyclic dependency at one of the possible points—at the rules (I) and (DN) (cf. the AFP characterization). From Theorem 3.15 and Corollary 3.14, the undefined deletions (which are drawn in the game-theoretic characterization) are also admissible (Theorem 3.19). Cutting in (DN) implements the definition of S_{RA} (giving priority to deletions over blockings), corresponding to the observation that S_{RA} takes exactly the blockings from the underestimate and the internal delete requests from the overestimate.

The rules (DC_1) , (DC_2) , and (DR) are already local rules:

$$\begin{aligned} [S] req_del_R_C(\bar{X}) \wedge R_C(\bar{X}), \bar{X}[\bar{F}] \triangleright \bar{Y}[\bar{K}], [S] req_del_R_P(\bar{Y}). & \quad (DC_1^S) \\ [S] blk_del_R_P(\bar{Y}) \wedge R_P(\bar{Y}), \bar{X}[\bar{F}] \triangleright \bar{Y}[\bar{K}], [S] blk_del_R_C(\bar{X}). & \quad (DC_2^S) \\ [S] blk_del_R_P(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \triangleright \bar{Y}[\bar{K}]. & \quad (DR^S) \end{aligned}$$

The rule (I) is also translated into a local rule:

$$[S] req_del_R(\bar{X}) \wedge \triangleright del_R(\bar{X}), R(\bar{X}), [S] : blk_del_R(\bar{X}). \quad (I^S)$$

(DN) incorporates the state leap and is augmented to a *progressive* rule (DN^S) :

$$[S \subset 1] blk_del_R_P(\bar{Y}) \wedge R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \triangleright \bar{Y}[\bar{K}], [S] : req_del_R_C(\bar{X}).$$

In the following, we refer to this program as P_S .

P_S is *state-stratified*, which implies that it is locally stratified, so there is a unique perfect model \mathcal{M}_S of $P_S [D [U_{\triangleright}$. The state-stratification

$\text{fblk_del_Rg} \setminus \text{freq_del_Rg}$, mirrors the stages of the algorithm: First, only blockings resulting from ON DELETE RESTRICT *rac*s are considered (local rules (DR^S) and (DC_2^S)). Based on these, the first maximal overestimate of internal delete requests is computed by (I^S) and (DC_1^S). Then, the blockings resulting from ON DELETE NO ACTION *rac*s whose child nodes are not deleted (by (DN^S), the only progressive rule; it does not derive anything for the initial state) are derived in the step to the subsequent state. (DR^S) contributes the blockings resulting from ON DELETE RESTRICT *rac*s. Again, the induced blockings are derived by (DC_2^S). The second stratum, consisting of (I^S) and (DC_1^S) determines the remaining non-blocked user delete requests and its induced delete requests. Then, the next iteration is started, calculating a decreasing sequence of overestimates which leads to S_{RA} .

LEMMA 3.20. \mathcal{M}_{AFP} corresponds to \mathcal{M}_S as follows:

1. $\mathcal{M}_{AFP} \text{ jD } [2n] \text{ blk_del_R}(\bar{x})$, $\mathcal{M}_S \text{ jD } [n] \text{ blk_del_R}(\bar{x})$.
2. $\mathcal{M}_{AFP} \text{ jD } [2nC1] \text{ req_del_R}(\bar{x})$, $\mathcal{M}_S \text{ jD } [n] \text{ req_del_R}(\bar{x})$.

PROOF. P_S and P_{AFP} differ in the rules (I^S) and (I^A): In every iteration, P_S takes the blockings from the latest underestimate, and the delete request from the latest overestimate, resulting in S_{RA} . \square

THEOREM 3.21 TERMINATION. For every database D and every set U_{\triangleright} of user delete requests, the program reaches a fixpoint, that is, there is a least $n_f \cdot \text{j}U_{\triangleright}\text{j}$, such that $\mathcal{M}_S[n_f] \text{ jD } \mathcal{M}_S[n_f \subset 1]$.

PROOF. A fixpoint is reached if the set of blocked user delete requests becomes stationary. Since this set is nondecreasing, there are at most $\text{j}U_{\triangleright}\text{j}$ iterations. \square

The correctness of P_S follows from Lemma 3.20 and Theorem 3.18:

THEOREM 3.22 CORRECTNESS. The final state of \mathcal{M}_S , $\mathcal{M}_S[n_f]$, represents U_{\max} and $\Delta(U_{\max})$:

- $\mathcal{M}_S[n_f] \text{ jD } S_{RA}$,
- $U_{\max} \text{ jD } \text{f} \triangleright \text{del_R}(\bar{x}) \text{ j} U_{\triangleright} \text{ j} \mathcal{M}_S[n_f] \text{ jD } \text{req_del_R}(\bar{x}) \text{ j}$, and
- $\Delta(U_{\max}) \text{ jD } \text{f} \triangleright \text{del_R}(\bar{x}) \text{ j} \mathcal{M}_S[n_f] \text{ jD } \text{req_del_R}(\bar{x}) \text{ j}$.

3.5.1 *Implementation in a Procedural Programming Language.* The Statelog formalization P_S above is translated into the algorithm given in Figure 4. Initially, it is assumed that there are only those blockings which result directly from ON DELETE RESTRICT *rac*s. Then, blockings are propagated upwards along the ON DELETE CASCADE chains, finally blocking the triggering user requests. For the remaining unblocked user requests, the cascaded requests are recomputed. Thus, some more tuples will remain in the database, which can block other requests. In the next step, all blockings are computed that are caused by ON DELETE NO ACTION *rac*s from tuples that are not reachable via cascaded deletions. These steps are repeated until a fixpoint is reached. Observe that each iteration corresponds to the evaluation of a query with PTIME data

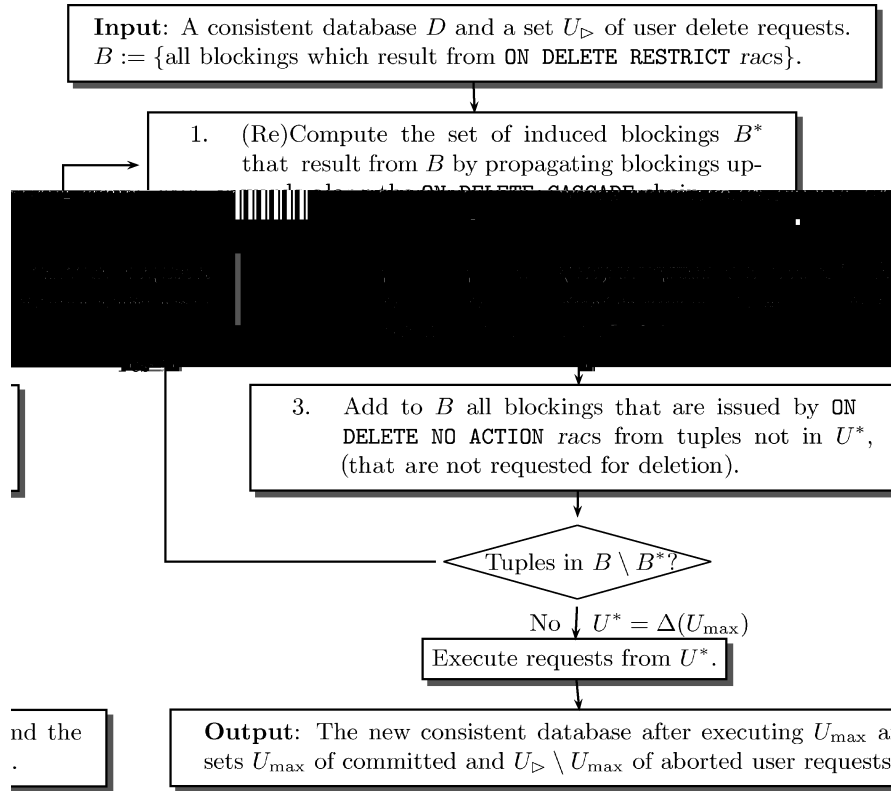


Fig. 4. Algorithm for executing all admissible deletions.

complexity. Since the fixpoint is reached after at most $|U_D|$ iterations (Theorem 3.21), the overall algorithm also has polynomial data complexity.

THEOREM 3.23. *The algorithm in Figure 4 is correct:*

— $U_{max} \sqsubseteq_{\triangleright} \text{del}_R(\bar{x}) \sqsubseteq_{\triangleright} U_D \sqcup \text{req_del}_R(\bar{x}) \sqsubseteq_{\triangleright} U'$, and $\Delta(U_{max}) \sqsubseteq_{\triangleright} U'$.

PROOF. In the n th iteration, $B' \sqsubseteq_{\triangleright} \text{blk_del}_R(\bar{x}) \sqcup \mathcal{M}_S \sqcup_{\triangleright} [n] \text{ blk_del}_R(\bar{x})$, and $U' \sqsubseteq_{\triangleright} \text{del}_R(\bar{x}) \sqcup \mathcal{M}_S \sqcup_{\triangleright} [n] \text{ req_del}_R(\bar{x})$. \square

For given D , U_D , and RA , the algorithm in Figure 4 computes the maximal subset U_{max} of U_D that can be executed without violating any ric , and the set $\Delta(U_{max})$ of “internal deletions” that are induced by it. For cases when U_D is not admissible, troubleshooting is described in Section 3.7.

3.6 Relationship with the SQL Semantics

The SQL semantics as presented in Horowitz [1992] and specified in the SQL3 standard [ANSI/ISO 1999] (see also Section 2.4) coincides with ours:

— matching rows (i.e., $R_C.\bar{F}(\bar{x})$ and $R_P.\bar{K}(\bar{y})$) are selected with respect to the original database state,

- the tuples to be deleted are fixed with respect to the original state, RESTRICT is also evaluated with respect to the original state,
- NOACTION is evaluated against the (prospective) result.
- an update referential action is executed whenever the parent is updated.
- effectively, Δ is applied after the computation is completed.

In the logic programming characterization, these “correctness” requirements are directly encoded into the rules (each individual rule corresponding to a single *ric* with *rac*), thus the application of the logic programming semantics cannot destroy them (provided one accepts the declarative logic programming semantics). Thus, a correctness proof with respect to the intended semantics—except that the specification of the individual rules is correct—is redundant.

Moreover, if a set of updates causes referential problems, the SQL semantics simply rejects the updates. Roughly speaking, it corresponds to the (*P*) rules and a simple RESTRICT test, and then checks the NO ACTION *rics*. In case of a rejection, it is hard to find which update caused the problem. Since the SQL semantics has no notion of the semantics of the *racs*, it cannot return useful details to the user as to how to prepare a revised request, or for debugging the application. Next, we show how the logic-based semantics directly incorporates this information, and how it can be used.

3.7 Troubleshooting in Case of Rejected Deletions

If $\triangleright\text{del_}R_d(\bar{x}_d) \not\leq U_\triangleright$ is not admissible, there are local “situations” that cause the problem. Each such situation consists of a parent tuple that should be deleted, and a child tuple that references it. Thus, the problem is caused by

- a parent tuple $R_P(\bar{y})$ that is reachable by a chain of ON DELETE CASCADE references from $R_d(\bar{x}_d)$. Correctly, $R_P(\bar{y})$ is potentially deleted (and must be deleted when $R_d(\bar{x}_d)$ is deleted),
- a child tuple $R_C(\bar{x})$ that references $R_P(\bar{y})$,
- a *ric* that concerns the reference from $R_C(\bar{x})$ to $R_P(\bar{y})$, associated with a *rac* (either RESTRICT or NO ACTION, and the child is not requested for deletion).

In a correct application—is where the database schema, containing the *rics* and *racs* is correct with respect to the semantics of the application domain, and where the program itself that computes U_\triangleright are correct—such a situation must not occur. Thus, there is a design problem that shows up in this situation:

- the definition of the *rac* is wrong: if it were ON DELETE CASCADE, the child would be deleted. Why did the programmer not specify ON DELETE CASCADE (did he forget it, since the default is ON DELETE NO ACTION)?
- the program is wrong: the blocked $\triangleright\text{del_}R_d(\bar{x}_d)$ should not be derived. Some object of the application must not be deleted as long as there are objects that need it, and the program did not check this correctly. Possibly, the whole transaction should not be executed.

- In case that the *rac* from $R_P(\bar{y})$ to $R_C(\bar{x})$ is ON DELETE NO ACTION, it is implicitly assumed that $R_C(\bar{x})$ is deleted by some other delete request in U_{\triangleright} (possibly via another cascading chain). Thus, there can be two reasons:
1. the program is wrong: it should derive $\triangleright\text{del}_R R_C(\bar{x})$ or it should derive a $\triangleright\text{del}_R R^0(\bar{x}^0)$ such that deleting $R^0(\bar{x}^0)$ cascades to deletion of $R_C(\bar{x})$, or
 2. another *ric* or *rac* is wrong: The program derives such an $\triangleright\text{del}_R R^0(\bar{x}^0)$, but it does not cascade to $R_C(\bar{x})$.
- the original database state is wrong: R

30 † W. May and B. Ludäscher

- the definition of the transaction is wrong. It should yield U_{\triangleright}^0 (delete $R_5(b)$ too, explicitly). Thus add a program line `DELETE . . . FROM R_5` .
- both are correct, but $R_5(b)$ should not be there. Another transaction had a wrong effect (e.g., by not deleting it).

Consider now the complete database in Figure 2 with $U_{\triangleright}^0 \text{ f } \triangleright \text{ del_}R_0(a), \triangleright \text{ del_}R_0(b) \text{ g}$. It induces the update U_{\triangleright} above and is again rejected. A possible conclusion from the detailed description of which tuples cause the problems could be

- there should be a *rac* $R_5(1) ! R_0(1)$ ON DELETE CASCADE that would then induce the updates given in U_{\triangleright}^0 .

Detecting problems during computation. The above problem analysis is based on the *outcome* of computing the well-founded model. However, the problem could easily be located *during* the computation of the well-founded model: the *first* blocking $\text{blk_del_}R_P(\bar{y})$ that is derived for a requested (internal) update in an *underestimate* (even-numbered states) locates the problem. The rule that is applied for deriving it identifies the referencing child tuple $R_C(\bar{x})$ and the reference. Then, the above scenario of “indirect” blockings can be avoided: the detected problem *is* a problem.

If the procedural algorithm given in Section 3.5 is applied, the same applies: the first blocking that is found identifies the problem situation.

Concerning the game, the problem situations correspond directly to the “final moves” of Π , either, to restricted, or via $\mathcal{DC}' - \mathcal{DN}$ to a tuple whose deletion is not founded (if its deletion is founded, but also “lost,” the above indirect situation applies, and the game goes on, leading to another final move of Π).

Solving the Problem. Still, if a problem is located, it is not clear how it must be solved. As described above, there can be several reasons and solutions. For example, a sophisticated reasoning component could be added that investigates the well-founded model in order to analyze the possible causes of the problem:

- try the same U_{\triangleright} for different arguments (e.g., checking what happens for $\triangleright \text{ del_}R_1(c)$ in the above example) to check if there is a general problem, or if the problem exists only for a special value (output: a set of problematic values).
- try extensions of U_{\triangleright} , and check how they are related to the original update, (output: a set of additional external updates; cf. Example 7).
- try possible modifications of the *racs* (e.g., turn NO ACTION into CASCADE), and
- try to extend P_{RA} with additional *rics* and *racs*.

In all cases, this information should be returned to the user. Due to the nature of the well-founded model as a stable (and reproducing) model, it is sufficient to adapt the rules or facts (immediately when a problem is detected during the computation) and to continue the evaluation starting with \mathcal{W} . The same holds for the presented procedural algorithm (which is a “hard-coded” version of the computation of the well-founded model for the given rule schemata).

4. SEMANTICS OF REFERENTIAL ACTIONS WITH MODIFICATIONS

In this section, we extend our approach to include modifications of tuples, resulting in a more involved translation and semantics. Interferences between cascaded modifications can occur especially in the presence of *overlapping* foreign keys. In contrast to deletions, modifications can create new foreign key values where the existence of an appropriate parent tuple has to be checked. Thus, instead of the SQL syntax, we use the more detailed syntax

ON UPDATE OF fPARENTjCHILDg fCASCADEjRESTRICTjNO ACTIONg .

Recall that all updates to a relation R are represented by auxiliary relations $\text{ins}_R(\bar{X})$, $\text{del}_R(\bar{X})$, and $\text{mod}_R(M, \bar{X})$. M is a list of pairs i/c meaning that the i -th attribute of $R(\bar{X})$ should be set to the constant c . As a shorthand for $\text{mod}_R([1/d, 3/e], (a, b, c))$, we may write $\text{mod}_R(a/d, b, c/e)$. The *restriction of a modification M to a key \bar{A}* is denoted by $M[\bar{A}]$; the result of *applying a modification M to \bar{X}* is denoted by $M(\bar{X})$. For example if $M \triangleright [1/d, 3/e]$, $\bar{X} \triangleright (a, b, c)$, then $M[(2, 3)] \triangleright [3/e]$, and $M(\bar{X}) \triangleright (d, b, e)$. The union of two modifications is the union of the lists: $[1/a, 2/b, 3/d] \cup [1/e, 2/b, 4/f] \triangleright [1/a, 1/b, 2/b, 3/d, 4/f]$. A modification is *consistent* if it does not assign two different values to a position (the above union is inconsistent).

4.1 Abstract Semantics

Let D be a database instance, U_\triangleright a set of user requests, and RA a set of *racs*. We define abstract properties that a set of updates Δ may have with respect to D , U_\triangleright , and RA . These allow us to define the intended meaning of a set of *racs* in an abstract (and non-constructive) way. Recall that $D^0 \triangleright D \S \Delta$ denotes the database obtained by applying Δ to D . In contrast to deletions, we cannot simply base our considerations on the transitive closure of references. Every step has to be analyzed individually, taking into consideration possible interferences due to overlapping keys.

Definition 4.1 Abstract Properties. An individual update instruction is called *founded* with respect to U_\triangleright , Δ , D , and RA if it can be justified by the user requests and propagations:

- A delete instruction $\text{del}_R(\bar{x})$ is *founded in n steps* with respect to U_\triangleright , Δ , D , and RA if either $n \triangleright 0$ and $\triangleright \text{del}_R(\bar{x}) \triangleright U_\triangleright$, or there is a delete instruction $\text{del}_{R_i}(\bar{x}_i) \triangleright \Delta$ that is founded in $< n$ steps, and a *rac* $R.\bar{K}_i, ! R_i.\bar{K}_i$ ON DELETE OF PARENT CASCADE such that $\bar{x}[\bar{K}_i] \triangleright \bar{x}_i[\bar{K}_i]$.
- A modify instruction $\text{mod}_R(M, \bar{x})$ is *founded in n steps* with respect to U_\triangleright , Δ , D , and RA if there is a set M_1, \dots, M_n of modifications s.t. $M \triangleright \bigcup_{i \triangleright 1..n} M_i$ (not necessarily disjoint) and for every i either $\triangleright \text{mod}_R(M_i, \bar{x}) \triangleright U_\triangleright$ or there is a modify instruction $\text{mod}_{R_i}(M_i, \bar{x}_i) \triangleright \Delta$ that is founded in $< n$ steps, and a *rac* $R.\bar{K}_i, ! R_i.\bar{K}_i$ ON UPDATE OF PARENT CASCADE such that $\bar{x}_i[\bar{K}_i] \triangleright M_i^0(\bar{x}_i)[\bar{K}_i]$ and $\bar{x}[\bar{K}_i] \triangleright \bar{x}_i[\bar{K}_i]$ and $M_i \triangleright \bar{K}_i / M_i^0(\bar{x}_i)[\bar{K}_i]$.
- An insert instruction $\text{ins}_R(\bar{x})$ is *founded* with respect to U_\triangleright , if $\triangleright \text{ins}_R(\bar{x}) \triangleright U_\triangleright$.

For given D , U_{\triangleright} , and RA , a set Δ of updates is called:

—*founded with respect to U_{\triangleright} , D , and RA* if every update instruction $\text{del}_R(\bar{x})$ or $\text{mod}_R(M, \bar{x}) \triangleright \Delta$ is founded.

—*complete with respect to D and RA* if it is closed with respect to propagations, that is, satisfies the following conditions:

1. if $\text{del}_{R_P}(\bar{y}) \triangleright \Delta$, $R_C(\bar{x}) \triangleright D$, $R_C.\bar{K}, ! R_P.\bar{K}$ ON DELETE OF PARENT CASCADE $\triangleright RA$, and $\bar{x}[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$ then $\text{del}_{R_C}(\bar{x}) \triangleright \Delta$.
2. if $\text{mod}_{R_P}(M, \bar{y}) \triangleright \Delta$, $R_C(\bar{x}) \triangleright D$, $R_C.\bar{K}, ! R_P.\bar{K}$ ON UPDATE OF PARENT CASCADE $\triangleright RA$, $\bar{y}[\bar{K}] \sqsupseteq M(\bar{y})[\bar{K}]$ and $\bar{x}[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$ then there is an M^0 such that $\text{mod}_{R_C}(M^0, \bar{x}) \triangleright \Delta$ and $M^0 \uparrow \bar{K}/M(\bar{y})[\bar{K}]$.

—*feasible with respect to D and RA* if all racs specified as no action and restrict are satisfied:

1. a) Δ contains no delete instruction $\text{del}_{R_P}(\bar{y})$ such that there is a rac $R_C.\bar{K}, ! R_P.\bar{K}$ ON DELETE OF PARENT RESTRICT $\triangleright RA$ and a tuple $R_C(\bar{x}) \triangleright D$ with $\bar{x}[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$.
b) Δ contains no modify instruction $\text{mod}_{R_P}(M, \bar{y})$ such that there is a rac $R_C.\bar{K}, ! R_P.\bar{K}$ ON UPDATE OF PARENT RESTRICT $\triangleright RA$, $\bar{y}[\bar{K}] \sqsupseteq M(\bar{y})[\bar{K}]$ and a tuple $R_C(\bar{x}) \triangleright D$ with $\bar{x}[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$.
2. for every rac $R_C.\bar{K}, ! R_P.\bar{K}$ ON DELETE OF PARENT NO ACTION or $R_C.\bar{K}, ! R_P.\bar{K}$ ON UPDATE OF PARENT NO ACTION $\triangleright RA$, every delete instruction $\text{del}_{R_P}(\bar{y})$ or modify instruction $\text{mod}_{R_P}(M, \bar{y})$ such that $\bar{y}[\bar{K}] \sqsupseteq M(\bar{y})[\bar{K}]$ that is contained in Δ , and every tuple $R_C(\bar{x}) \triangleright D$ such that $\bar{x}[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$, Δ contains either a delete instruction $\text{del}_{R_C}(\bar{x})$ or a modify instruction $\text{mod}_{R_C}(M, \bar{x})$ such that $M(\bar{x})[\bar{K}] \sqsupseteq \bar{y}[\bar{K}]$.
3. for every rac $R_C.\bar{K}, , ! R_P.\bar{K}$ ON ... OF CHILD RESTRICT $\triangleright RA$ and every tuple $R_C(\bar{x}) \triangleright D$ that either results from an insertion $\text{ins}_{R_C}(\bar{x}) \triangleright \Delta$ or a modification $\text{mod}_{R_C}(M, \bar{x}^0) \triangleright \Delta$ (i.e., $\bar{x} \sqsupseteq M(\bar{x}^0)$) such that $\bar{x}^0[\bar{K}] \sqsupseteq M(\bar{x}^0)[\bar{K}]$, D contains a tuple $R_P(\bar{y})$ such that $\bar{y}[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$ and Δ contains neither $\text{del}_{R_P}(\bar{y})$ nor any $\text{mod}_{R_P}(M, \bar{y})$ such that $M(\bar{y})[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$.
4. for every rac $R_C.\bar{K}, ! R_P.\bar{K}$ ON ... OF CHILD NO ACTION $\triangleright RA$ and every tuple $R_C(\bar{x}) \triangleright D$ that either results from an insertion $\text{ins}_{R_C}(\bar{x}) \triangleright \Delta$ or a modification $\text{mod}_{R_C}(M, \bar{x}^0) \triangleright \Delta$ (i.e., $\bar{x} \sqsupseteq M(\bar{x}^0)$) such that $\bar{x}^0[\bar{K}] \sqsupseteq M(\bar{x}^0)[\bar{K}]$, one of the following conditions hold:
 - a) D contains a tuple $R_P(\bar{y})$ such that $\bar{y}[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$ and Δ contains neither $\text{del}_{R_P}(\bar{y})$ nor any $\text{mod}_{R_P}(M, \bar{y})$ such that $M(\bar{y})[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$,
or
 - b) there is an $\text{ins}_{R_P}(\bar{y}) \triangleright \Delta$ such that $\bar{y}[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$, or
 - c) there is a $\text{mod}_{R_P}(M, \bar{y}) \triangleright \Delta$ such that $R_P(\bar{y}) \triangleright D$ and $M(\bar{y})[\bar{K}] \sqsupseteq \bar{x}[\bar{K}]$.

—*coherent* if no contradicting updates are issued on the same tuple—if $\text{upd} \sqsupseteq \text{del}_R(\bar{x}) \triangleright \Delta$, then $\text{ins}_R(\bar{x}) \not\triangleright \Delta$ and there is no M such that $\text{mod}_R(M, \bar{x}) \triangleright \Delta$; similar for other updates upd . Note that if Δ is coherent, $D^0 \sqsupseteq D \sqsubseteq \Delta$ is well-defined.

—*key-preserving* if in $D^0 \sqsupseteq D \sqsubseteq \Delta$ all key dependencies are satisfied.

—*admissible* if Δ is founded, complete, feasible, coherent, and key-preserving.

Again, these abstract properties are used to formalize our *intended semantics*:

Definition 4.2 Induced Updates, Application of U_{\triangleright} . Let U_{\triangleright} , RA , and D be given.

- The set of *induced updates* $\Delta(U)$ of a set of user requests $U \mu U_{\triangleright}$ is the least set Δ which contains U and is complete.
(analogously to the case of deletion, $\Delta(U)$ there is a unique least such set by construction).
- A set of user requests $U \mu U_{\triangleright}$ is *admissible* if $\Delta(U)$ is admissible, and *maximal admissible* if there is no other admissible U^0 , such that $U \subsetneq U^0 \mu U_{\triangleright}$.
- For a set Δ of user requests, $D^0 \text{D} D \S \Delta$ denotes the database obtained by applying Δ to D .

Note that $\Delta(U)$ does not contain any subsumed updates. This semantics reflects the intended behavior of the database system—it neither “invents” nor “forgets” updates and guarantees referential integrity:

THEOREM 4.3 ADEQUACY.

- (1) If $U \mu U_{\triangleright}$, then $\Delta(U)$ is founded and complete.
- (2) If a coherent Δ is complete and feasible, then $D^0 \text{D} D \S \Delta(U)$ satisfies all *racs*.

PROOF.

- (1) $\Delta(U)$ is defined as the *least* complete set, thus it is founded.
- (2) Since Δ is complete, all updates propagated by RA are contained in Δ . Feasibility of Δ guarantees that no *upd* $\geq \Delta$ is restricted, and all **NO ACTION** *racs* are satisfied in D^0 . \square

The abstract semantics specifies the notions of maximal admissible sets U and induced updates $\Delta(U)$, but provides no method as to how to compute them. In contrast to the case of deletions, where the union of admissible sets of updates was again admissible (cf. Proposition 3.6), this does not hold in general due to conflicting updates. Given a set of n user requests, there can be an exponential (in n) number of maximal admissible subsets. Not surprisingly, it is not sufficient to consider the well-founded model, but stable models will be required for the general case.

Moreover, even if it is known that U is admissible, computing $\Delta(U)$ is not straightforward: In contrast to deletions that can be propagated in a “greedy” way without considering parallel updates, in the presence of modifications, parallel updates have to be taken into account:

Example 11. Consider the database schema depicted in Figure 5. Among others there are *racs* of type **ON UPDATE OF PARENT CASCADE** for the *rics* $T.(1, 2)!$ $R.(1, 2)$, $T.(3, 4)!$ $S.(1, 2)$, $U.(1, 2)!$ $T.(2, 3)$, and a *rac* **ON UPDATE OF CHILD RESTRICT** for $T.(1, 4)!$ $V.(1, 2)$. Assume the database contains $R(a, b)$, $S(c, d)$, $T(a, b, c, d, \dots)$, $U(b, c, \dots)$, and $V(a, d, \dots)$, $V(a^0, d, \dots)$, $V(a, d^0, \dots)$.




Fig. 5. Example scenario for modifications (see Example 11).

Given $\text{mod}_R(a/a^0, b/b^0)$ and $\text{mod}_S(c/c^0, d/d^0)$, the *rac*s trigger the modifications $\text{mod}_T(a/a^0, b/b^0, c, d, \dots)$ and $\text{mod}_T(a, b, c/c^0, d/d^0, \dots)$. Since these updates to T are coherent, they can be merged into $\text{mod}_T(a/a^0, b/b^0, c/c^0, d/d^0, \dots)$, which triggers $\text{mod}_U(b/b^0, c/c^0, \dots)$.

On the other hand, the *rac* $T.(1, 4) \uparrow V.(1, 2)$ ON UPDATE OF CHILD RESTRICT restricts this modification since there is no tuple $V(a^0, d^0, \dots)$. So each of the updates is admissible, but they are not admissible together, even though they do not directly contradict each other. Note that it is completely irrelevant, what modifications would be raised on the dotted parts of the tuples.

Example 11 illustrates some of the problems that may arise due to overlapping foreign keys and candidate keys, and gives an idea of the inherent complexity of rule-based referential integrity maintenance in the presence of modifications.

As indicated by Example 11, the propagation of updates cannot be seen tuple-oriented (too coarse: cf. the dotted parts of T, U, V) or attribute-oriented (too fine: $\text{mod}_T(a/a^0, b/b^0, c, d, \dots)$ and $\text{mod}_T(a, b, c/c^0, d/d^0, \dots)$ are both allowed in isolation, but their combination is forbidden), but must be handled *reference-oriented*. Thus, parent keys, foreign keys, references, and overlapping keys play an important role in our logical formalization.

4.2 Logic Programming Characterization

Problems and Solutions. The handling of updates adds several problems in contrast to deletions since tuples (and references) are not only removed, but also new tuples are generated that possibly contain new foreign key values. From the point of view of updating *referenced* relations, updates are not more complicated than deletions: propagate, wait, or restrict, as before. But from the point of view of updating the *referencing* relation (where deletions were trivial), we have to check whether there is an appropriate parent. If not, the update must be blocked. For deletions, we could list all blocked deletions—all tuples of the database that must not be deleted. In contrast, it is not possible to simply list all updates that are not allowed due to the absence of an appropriate parent (there are infinitely many). The considerations must be restricted to those updates that *potentially* arise from U_{\triangleright} . Additionally, there are *overlapping keys* as illustrated above where *interfering* modifications must be considered. As explained above, the solution is based on the analysis of *changes of key values*. Thus, the relations $\text{chg}_R.\bar{A}(M, \bar{X})$ (defined below in (CH_1)) and $\text{prp}_{R_P}.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{X})$

(defined below in (MPC_1)) are the “keys” to the solution:

— $\text{chg}_R.\bar{A}(M, \bar{X})$ denotes that the (candidate or foreign) key value $\bar{X}[\bar{A}]$ in relation R is modified by M , that is, changes to $(M(\bar{X}))[\bar{A}]$.

For candidate keys $R.\bar{A}$, this modification is propagated to all referencing foreign keys:

— $\text{prp}_{R_P}.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{x})$ denotes that a modification M of $R_P.\bar{K}$ is propagated along $R_P.\bar{K} \dashv R_C.\bar{F}$ to $R_C(\bar{x})$ resulting in the modification $M_C \sqsupset \bar{F}/M_P(\bar{y})[\bar{K}]$ of the foreign key value.

Changes of overlapping candidate keys are computed from the incoming foreign key changes (see rule (CH_1) below).

In the logic programming characterization given for deletions in Section 3.2, the logical rules represented the *local* semantics of referential actions, and all global aspects were covered by the logic programming semantics (i.e., by the definition of the well-founded model and the stable model). When modifications are concerned, there are many more local parts in the global puzzle of the semantics of referential actions: Logical rules do not only represent the *local* semantics of referential actions, but also represent local *interferences* of updates. Again, all global aspects are provided by the chosen logic programming semantics.

The meaning of a set RA of *racs* is formalized as a logic program P_{RA} , consisting of the sets P_{ra} that specify the local behavior of every *rac* ra , and a set of rules that specify the meaning of interacting update requests (instantiations for a given set of *rics* can be found in Example 12 in the *electronic appendix*).

4.2.1 Initialization and Bookkeeping

User Requests. The handling of user requests incorporates the selection of admissible updates. User requests that are not blocked, raise an update to the database:

$$\begin{aligned} \text{del}_R(\bar{X}) &\quad \wedge \quad \triangleright \text{del}_R(\bar{X}), \quad : \text{blk_del}_R(\bar{X}). \\ \text{ins}_R(\bar{X}) &\quad \wedge \quad \triangleright \text{ins}_R(\bar{X}), \quad : \text{blk_ins}_R(\bar{X}). \\ \text{mod}_R(M, \bar{X}) &\quad \wedge \quad \triangleright \text{mod}_R(M, \bar{X}), \quad : \text{blk_mod}_R(M, \bar{X}). \end{aligned} \quad (EXT_1)$$

Auxiliary Relations. As stated above, the approach is key-based. Thus, several auxiliary relations are maintained that contain information about referenced and referenceable candidate key values:

- $\text{is_refable}_R.\bar{K}(\bar{x})$: the (key) value $R.\bar{K}(\bar{x})$ is referenceable in the original state.
- $\text{rem_refable}_R.\bar{K}(\bar{x})$: the (key) value $R.\bar{K}(\bar{x})$ remains referenceable.
- $\text{new_refable}_R.\bar{K}(\bar{x})$: the (key) value $R.\bar{K}(\bar{x})$ becomes referenceable.
- $\text{is_refd}_{R_P}.\bar{K} _by_{R_C}.\bar{F}(\bar{v})$: in the current database, the key value $R.\bar{K}(\bar{v})$ appears as foreign key value of \bar{F} in some tuple $R_C(\bar{x})$.
- $\text{rem_refd}_{R_P}.\bar{K} _by_{R_C}.\bar{F}(\bar{v})$: there is a reference to the key value $R.\bar{K}(\bar{v})$ as foreign key value of \bar{F} in some tuple $R_C(\bar{x})$ such that $\bar{x}[\bar{F}]$ does not change.
- $\text{new_refd}_{R_P}.\bar{K} _by_{R_C}.\bar{F}(\bar{v})$: a reference to the key value $R.\bar{K}(\bar{v})$ as foreign key \bar{F} in some tuple $R_C(\bar{x})$ is introduced by some update.

36 † W. May and B. Ludäscher

For every candidate key \bar{K} mentioned in some $ric\ R_C.\bar{K}!\ R_P.\bar{K}$:

$$\begin{aligned}
is_refable_R_P.\bar{K}(\bar{V}) &\hat{\ } R_P(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}]. \\
rem_refable_R_P.\bar{K}(\bar{V}) &\hat{\ } R_P(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}], : del_R_P(\bar{X}), : \exists M: \\
&\quad chg_R_P.\bar{K}(M, \bar{X}). \qquad (RCK) \\
new_refable_R_P.\bar{K}(\bar{V}) &\hat{\ } ins_R_P(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}]. \\
new_refable_R_P.\bar{K}(\bar{V}) &\hat{\ } chg_R_P.\bar{K}(M, \bar{X}), M(\bar{X})[\bar{K}] \sqsupset \bar{V}.
\end{aligned}$$

For every $ric\ R_C.\bar{K}!\ R_P.\bar{K}$:

$$\begin{aligned}
is_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{V}) &\hat{\ } R_C(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}]. \\
rem_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{V}) &\hat{\ } R_C(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}], : del_R_C(\bar{X}), \\
&\quad : \exists M : chg_R_C.\bar{K}(M, \bar{X}). \qquad (Refd) \\
new_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{V}) &\hat{\ } R_C(\bar{X}), chg_R_C.\bar{K}(M, \bar{X}), M(\bar{X})[\bar{K}] \sqsupset \bar{V}. \\
new_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{V}) &\hat{\ } ins_R_C(\bar{X}), \bar{V} \sqsupset \bar{X}[\bar{K}].
\end{aligned}$$

4.2.2 Deletions. We only have to consider *rac*s of the form $R_C.\bar{K}!\ R_P.\bar{K}$ ON DELETE OF PARENT (cf. Table I). Logic rules that again describe their *local* behavior are generated for these *rac*s as given below:

- ON DELETE OF PARENT CASCADE: Deletions of parent tuples are propagated downwards to every child tuple by rule (DC_1). Additionally, blockings are propagated upwards: if the deletion of a child tuple is blocked, the deletion of the parent tuple is also blocked (DC_2).
- ON DELETE OF PARENT RESTRICT: The deletion of a parent tuple is blocked, if there is a referencing child tuple (DR).
- ON DELETE OF PARENT NOACTION: The deletion of a parent tuple is blocked, if there is a corresponding child tuple which is neither requested for deletion nor *modified away* (i.e., modified such that it references another parent) (DN).

$del_R_C(\bar{X}) \hat{\ } del_R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}].$	(DC_1)
$blk_del_R_P(\bar{Y}) \hat{\ } blk_del_R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}].$	(DC_2)
$blk_del_R_P(\bar{Y}) \hat{\ } is_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{Y}[\bar{K}]).$	(DR)
$blk_del_R_P(\bar{Y}) \hat{\ } rem_refd_R_P.\bar{K}_by_R_C.\bar{K}(\bar{Y}[\bar{K}]).$	(DN)

Again, we add the rules for tracing potential deletions (that are only used for analyzing problem situations if an update is rejected):

$$\begin{aligned}
pot_del_R(\bar{X}) &\hat{\ } \triangleright del_R(\bar{X}), R(\bar{X}). \\
\text{for each } R_C.\bar{K}, ! R_P.\bar{K} \text{ ON DELETE OF PARENT CASCADE} &\qquad (P) \\
\text{(analogous to } (DC_1)\text{):} & \\
pot_del_R_C(\bar{X}) &\hat{\ } pot_del_R_P(\bar{Y}), R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}].
\end{aligned}$$

4.2.3 Modifications. The handling of modifications follows the same principle as presented for deletions, but since the propagation of modifications is handled key-oriented, the details are more involved. Here, the predicates

$\text{chg_R}.\bar{A}(M, \bar{X})$ and $\text{prp_R}_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{x})$ are used as described above for describing changing key values and their propagation. In case of a partially modified parent key, the referencing foreign key in the child is regarded as *atomic*: no other update may change parts of it (which would mean to cut the reference to the original parent although it is used for cascading). Thus, with a modification, the whole key value is propagated, even if not all parts of it change. On the other hand, modifications on a tuple trigger a *rac* only if the key referred to in the *rac* is actually changed.

User requests for modifications are decomposed into their effects on keys. An external request is blocked if it causes a forbidden change to a key. For every key (candidate and foreign keys) $R.\bar{A}$:

$$\begin{aligned} \text{pot_prp}_{\triangleright R} &\rightsquigarrow R.\bar{A}(M, \bar{X}) \wedge \triangleright \text{mod_R}(M^0, \bar{X}), \bar{X}[\bar{A}] \text{ } \textcircled{D} \text{ } M^0(\bar{X})[\bar{A}], \\ &M \text{ } \textcircled{D} \text{ } M^0[\bar{A}]. \\ \text{blk_mod_R}(M, \bar{X}) &\wedge \triangleright \text{mod_R}(M, \bar{X}), \text{blk_chg_R}.\bar{A}(M^0, \bar{X}), \quad (\text{EXT}_2) \\ &M^0 \text{ } \textcircled{D} \text{ } M[\bar{A}]. \\ \text{prp}_{\triangleright R} \rightsquigarrow R.\bar{A}(M, \bar{X}) &\wedge \text{mod_R}(M^0, \bar{X}), \bar{X}[\bar{A}] \text{ } \textcircled{D} \text{ } M^0(\bar{X})[\bar{A}], \\ &M \text{ } \textcircled{D} \text{ } M^0[\bar{A}]. \end{aligned}$$

Interaction of Modifications. Assume a modification $\text{mod_R}_P(M_P, \bar{y})$ and a *rac* $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE OF PARENT CASCADE such that the *key value* $R_P.\bar{K}$ of $R_P(\bar{y})$ changes, which is denoted by $\text{chg_R}_P.\bar{K}(M_P, \bar{y})$ (analogously, there are relations pot_chg and blk_chg). Then, for every referencing child $R_C(\bar{x})$, this modification is propagated to the corresponding foreign key— $M_C \text{ } \textcircled{D} \text{ } \bar{F}/M_P(\bar{y})[\bar{K}]$. This is stored in the propagation relation $\text{prop_prp_R}_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{x})$.

The changes of candidate and foreign key values in a (child) tuple are collected: modifications can be founded either by external requests or by propagating modifications from parent relations (e.g., consider a tuple $T(a, b, c, d)$ in Figure 5 where modifications come in to the primary key $T[2, 3]$ from $R[2]$ and $S[1]$).

For a given database schema, (CH) defines a finite set of rules for computing all the possible ways a key can change by collecting the elementary modification requests that are propagated along references.

The *only* restriction in this presentation is, that for every foreign or candidate key, only one user modify request is raised that changes the key (which is satisfied if no parallel modifications of the same tuple are allowed; overcoming this restriction requires no conceptual, but some technical expense).

For a given foreign or candidate key $R.\bar{A}$, define the set of sources of influences on $R.\bar{A}$, $S_{R.\bar{A}} \sqcup \text{Keys}(R) \in (\text{ForeignKeys}(DB) \text{ } [\text{f} \triangleright R])$ as follows:

- $S_{R.\bar{A}}$ contains all pairs $(R_{P_i}.\bar{K}_i, R.\bar{F}_i)$ such that there is a *rac* $R_{P_i}.\bar{K}_i, ! R.\bar{F}_i$ ON UPDATE OF PARENT CASCADE and F_i overlaps \bar{A} (that is, key references whose propagation influences the value of $R.\bar{A}$), and
- $(R.\bar{A}, \triangleright R) \text{ } \textcircled{D} \text{ } S_{R.\bar{A}}$ since external modifications of R also influence the value of $R.\bar{A}$.

38 † W. May and B. Ludäscher

For the following rule (CH_1), let S range over all subsets of $S_{R.\bar{A}}$ ⁶:

$$\begin{aligned} \text{pot_chg_R}.\bar{A}(M, \bar{X}) \wedge & (\bigwedge_{i \in 2S} \text{pot_prp_R}_{P_i}.\bar{K}_i \rightsquigarrow R.\bar{F}_i(M_i, \bar{X})), \\ & M \sqsupset \bigcup_{i \in 2S} M_i[\bar{A}] \text{ consistent, } M(\bar{X})[\bar{A}] \sqsupset \bar{X}[\bar{A}]. \\ \text{chg_R}.\bar{A}(M, \bar{X}) \wedge & (\bigwedge_{i \in 2S} \text{prp_R}_{P_i}.\bar{K}_i \rightsquigarrow R.\bar{F}_i(M_i, \bar{X})), \\ & (\bigwedge_{i \in 2S_{R.\bar{A}} \cap S} \exists M_i : \text{prp_R}_{P_i}.\bar{K}_i \rightsquigarrow R.\bar{F}_i(M_i, \bar{X})), \\ & M \sqsupset \bigcup_{i \in 2S} M_i[\bar{A}] \text{ consistent, } M(\bar{X})[\bar{A}] \sqsupset \bar{X}[\bar{A}]. \end{aligned} \quad (CH_1)$$

Additionally, the interferences between blockings of changes of overlapping keys must be considered: A change on the intersection of two overlapping keys is allowed, if both changes coincide on the intersection. Furthermore, a change of a key is forbidden, if its effect on the intersection with another key is not allowed:

For every foreign key \bar{F} and foreign or candidate key \bar{A} such that \bar{F} and \bar{A} overlap:

$$\begin{aligned} \text{allow_chg_R}.\bar{F} \setminus \bar{A}(M, \bar{X}) \wedge & \text{chg_R}.\bar{F}(M_1, \bar{X}), : \text{blk_chg_R}.\bar{F}(M_1, \bar{X}), \\ & M \sqsupset M_1[\bar{F} \setminus \bar{A}], \text{chg_R}.\bar{A}(M_2, \bar{X}), \\ & : \text{blk_chg_R}.\bar{A}(M_2, \bar{X}), M \sqsupset M_2[\bar{F} \setminus \bar{A}]. \quad (CH_2) \\ \text{blk_chg_R}.\bar{F}(M, \bar{X}) \wedge & \text{pot_chg_R}.\bar{F}(M, \bar{X}), : \text{allow_chg_R}.\bar{F} \setminus \bar{A}(M^0, \bar{X}), \\ & M^0 \sqsupset M[\bar{F} \setminus \bar{A}]. \end{aligned}$$

Remark 4.4. Consider a foreign or candidate key $R.\bar{A}$ such that there is only a single constraint $R_P.\bar{K}, ! R.\bar{F}$ ON UPDATE OF PARENT CASCADE such that \bar{F} and \bar{A} overlap. Then, (CH_1) maps the incoming change on \bar{F} to a change of \bar{A} :

$$\begin{aligned} \text{pot_chg_R}.\bar{A}(M, \bar{X}) \wedge & \text{pot_prp_R}_P.\bar{K} \rightsquigarrow R.\bar{F}(M^0, \bar{X}), M \sqsupset M^0[\bar{A}], \\ & M(\bar{X})[\bar{A}] \sqsupset \bar{X}[\bar{A}]. \\ \text{chg_R}.\bar{A}(M, \bar{X}) \wedge & \text{prp_R}_P.\bar{K} \rightsquigarrow R.\bar{F}(M^0, \bar{X}), M \sqsupset M^0[\bar{A}], \\ & M(\bar{X})[\bar{A}] \sqsupset \bar{X}[\bar{A}]. \end{aligned} \quad (CH_1^0)$$

Modifications of Parent Tuples. When a candidate key in a parent tuple is modified, the usual *local* referential actions apply, yielding the rules given below:

- ON UPDATE OF PARENT CASCADE: Changes of parent keys are propagated downwards to foreign keys (MPC_1). If a propagated modification would change a foreign key in a forbidden way, the propagation of the modification and the change of the parent key are also blocked (MPC_2).
- ON UPDATE OF PARENT RESTRICT: The change of the parent key $R_P.\bar{K}$ is blocked, if there is a referencing child (MPR).

⁶if \bar{A} is a foreign key, the possibilities can further be restricted to either (i) a propagation along its “own” parent-reference, or (ii) only influences from other sources, see Example 12 in the *electronic appendix*.

—ON UPDATE OF PARENT NOACTION: The change of the parent key $R_P.\bar{K}$ is blocked, if there is a referencing child which is neither requested for deletion nor modified away (MPN).

$\text{prp_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{X})$	\wedge	$\text{chg_}R_P.\bar{K}(M_P, \bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \sqsupset \bar{Y}[\bar{K}],$ $M_C \sqsupset \bar{F}/M_P(\bar{Y})[\bar{K}].$	(MPC ₁)
$\text{pot_prp_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{X})$	\wedge	$\text{pot_chg_}R_P.\bar{K}(M_P, \bar{Y}), R_C(\bar{X}), \bar{X}[\bar{F}] \sqsupset \bar{Y}[\bar{K}],$ $M_C \sqsupset \bar{F}/M_P(\bar{Y})[\bar{K}].$	
$\text{blk_chg_}R_P.\bar{K}(M_P, \bar{Y})$	\wedge	$\text{pot_chg_}R_P.\bar{K}(M_P, \bar{Y}), \text{blk_prop_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M_C, \bar{X}),$ $\bar{X}[\bar{F}] \sqsupset \bar{Y}[\bar{K}], M_C \sqsupset \bar{F}/M_P(\bar{Y})[\bar{K}].$	(MPC ₂)
$\text{blk_prop_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M, \bar{X})$	\wedge	$\text{pot_prp_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M, \bar{X}), \text{blk_chg_}R_C.\bar{F}(M, \bar{X}).$	
$\text{blk_chg_}R_P.\bar{K}(M_P, \bar{Y})$	\wedge	$\text{pot_chg_}R_P.\bar{K}(M_P, \bar{Y}), \text{is_refd_}R_P.\bar{K} \text{_by_}R_C.\bar{F}(\bar{Y}[\bar{K}]).$	(MPR)
$\text{blk_chg_}R_P.\bar{K}(M_P, \bar{Y})$	\wedge	$\text{pot_chg_}R_P.\bar{K}(M_P, \bar{Y}), \text{rem_refd_}R_P.\bar{K} \text{_by_}R_C.\bar{F}(\bar{Y}[\bar{K}]).$	(MPN)

Modifications on Child Tuples. When a foreign key in a child tuple is changed, it must be checked whether there is a suitable parent tuple—note that the check against the reference along which the update has been cascaded is redundant. Thus, we have only to check foreign keys where an incoming cascaded update overlaps the foreign key of *another* reference. A change on a foreign key value $R_C.\bar{F}^0$ of a child tuple with respect to a ric $R_C.\bar{F}^0 ! R_P.\bar{K}^0$ is blocked if the change is influenced from a propagation along *another ric* $R_C.\bar{F} ! R_P.\bar{K}$ (i.e., $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE OF PARENT CASCADE and $R_C.\bar{F}$ and $R_C.\bar{F}^0$ overlap) or from an external modification and the resulting tuple violates $R_C.\bar{F}^0, ! R_P.\bar{K}^0$ ON UPDATE OF CHILD By considering only changes which are propagated along *another ric*, the inherent negative cycle of “propagation allowed if result’s reference exists,” “result’s reference exists if parent is modified,” and “parent is modified if propagation is allowed” is avoided.⁷

For every pair $R_C.\bar{F}^0, ! R_P.\bar{K}^0$ ON UPDATE OF PARENT CASCADE and $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE OF CHILD RESTRICT such that $R_C.\bar{F} \sqsupset R_C.\bar{F}^0$ or $R_P.\bar{K} \sqsupset R_P.\bar{K}^0$ and $R_C.\bar{F}$ and $R_C.\bar{F}^0$ overlap:

$$\begin{aligned} \text{blk_chg_}R_C.\bar{F}(M, \bar{X}) \wedge & \text{pot_chg_}R_C.\bar{F}(M, \bar{X}), \text{prp_}R_P.\bar{K}^0 \rightsquigarrow R_C.\bar{F}^0(M^0, \bar{X}), \\ & M[\bar{F} \setminus \bar{F}^0] \sqsupset M^0[\bar{F} \setminus \bar{F}^0], \\ & : \text{is_refable_}R_P.\bar{K}(\bar{F}[M(\bar{X})]). \end{aligned} \quad (\text{MCR}_1)$$

$$\begin{aligned} \text{blk_chg_}R_C.\bar{F}(M, \bar{X}) \wedge & \text{pot_chg_}R_C.\bar{F}(M, \bar{X}), \text{prp_}\triangleright \rightsquigarrow R_C.\bar{F}(M^0, \bar{X}), \\ & M^0 \sqcup M, : \text{is_refable_}R_P.\bar{K}(M(\bar{X})[\bar{F}]). \end{aligned}$$

Additionally to (MCR₁) which guarantees the existence of a (unique) parent tuple to be referenced in the *current* database state, any modification of the attributes $R_P.\bar{K}$ or deletion of this tuple is blocked: For every *rac* $R_C.\bar{F}, ! R_P.\bar{K}$

⁷Note that it still occurs with every “diamond” (Figure 1).

40 † W. May and B. Ludäscher

ON UPDATE OF CHILD RESTRICT:

$$\begin{aligned}
\text{blk_chg_}R_P.\bar{K}(M_P, \bar{Y}) \wedge & \text{pot_chg_}R_P.\bar{K}(M_P, \bar{Y}), R_C(\bar{X}), \\
& \text{chg_}R_C.\bar{F}(M_C, \bar{X}), M(\bar{X})[\bar{F}] \sqsupset \bar{Y}[\bar{K}]. \quad (MCR_2) \\
\text{blk_del_}R_P(\bar{Y}) \wedge & R_C(\bar{X}), \text{chg_}R_C.\bar{F}(M_C, \bar{X}), M(\bar{X})[\bar{F}] \sqsupset \bar{Y}[\bar{K}].
\end{aligned}$$

Analogous to (MCR_1) there is a rule for maintaining $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE OF CHILD NO ACTION which checks if the parent is available *after* execution of the updates. For every pair $R_C.\bar{F}^0, ! R_P.\bar{K}^0$ ON UPDATE OF PARENT CASCADE and $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE OF CHILD NO ACTION such that $R_C.\bar{F} \sqsupset R_C.\bar{F}^0$ or $R_P.\bar{K} \sqsupset R_P.\bar{K}^0$ and $R_C.\bar{F}$ and $R_C.\bar{F}^0$ overlap:

$$\begin{aligned}
\text{blk_chg_}R_C.\bar{F}(M, \bar{X}) \wedge & \text{pot_chg_}R_C.\bar{F}(M, \bar{X}), \text{prp_}R_P.\bar{K}^0 \rightsquigarrow R_C.\bar{F}^0(M^0, \bar{X}), \\
& M[\bar{F} \setminus \bar{F}^0] \sqsupset M^0[\bar{F} \setminus \bar{F}^0], \\
& : \text{rem_refable_}R_P.\bar{K}(M(\bar{X})[\bar{F}]), \\
& : \text{new_refable_}R_P.\bar{K}(M(\bar{X})[\bar{F}]). \\
\text{blk_chg_}R_C.\bar{F}(M, \bar{X}) \wedge & \text{pot_chg_}R_C.\bar{F}(M, \bar{X}), \text{prp_}\triangleright \rightsquigarrow R_C.\bar{F}(M^0, \bar{X}), \\
& M^0 \sqcup M, : \text{rem_refable_}R_P.\bar{K}(M(\bar{X})[\bar{F}]), \\
& : \text{new_refable_}R_P.\bar{K}(M(\bar{X})[\bar{F}]). \quad (MCN)
\end{aligned}$$

Resulting Modifications. Since modifications are handled key-oriented by (CH_1) , the incoming modifications must be collected for every tuple. For a given n -ary relation R , let $S_R \sqsupset \text{ForeignKeys}(R) [f(1, \dots, n)]g$.

$$\text{chg_}R.(1, \dots, n)(M, \bar{X}) \wedge \triangleright \text{mod_}R(M, \bar{X}). \quad (CH_\triangleright)$$

For the following rule (MOD) , let S range over all subsets of S_R :

$$\begin{aligned}
\text{mod_}R(M, \bar{X}) \wedge & (\bigwedge_{\bar{F} \in S} \text{chg_}R.\bar{F}(M_{\bar{F}}, \bar{X})), M \sqsupset \bigcup_{\bar{F} \in S} M_{\bar{F}} \text{ consistent,} \\
& (\bigwedge_{\bar{F} \in S_R \setminus S} \exists M^0 : \text{chg_}R.\bar{F}(M^0, \bar{X})). \quad (MOD)
\end{aligned}$$

4.2.4 Insertions. Since insertions on parent tuples are not critical, only insertions of child tuples have to be handled analogously to (MCR) and (MCN) :

—For every *ric* $R_C.\bar{F}, ! R_P.\bar{K}$ ON INSERT OF CHILD RESTRICT:

$$\begin{aligned}
\text{blk_ins_}R_C(\bar{X}) \wedge & \triangleright \text{ins_}R_C(\bar{C}), : \text{is_refable_}R_P.\bar{K}(\bar{X}[\bar{F}]). \\
\text{blk_chg_}R_P.\bar{K}(M, \bar{Y}) \wedge & \text{pot_chg_}R_P.\bar{K}(M, \bar{Y}), \bar{Y}[\bar{K}] \sqsupset M_P(\bar{Y})[\bar{K}], \quad (ICR) \\
& \text{ins_}R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}]. \\
\text{blk_del_}R_P(\bar{Y}) \wedge & \text{ins_}R_C(\bar{X}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}].
\end{aligned}$$

—For every *ric* $R_C.\bar{F}, ! R_P.\bar{K}$ ON INSERT OF CHILD NO ACTION:

$$\begin{aligned}
\text{blk_ins_}R_C(\bar{X}) \wedge & \triangleright \text{ins_}R_C(\bar{X}), : \text{rem_refable_}R_P.\bar{K}(\bar{X}[\bar{F}]), \\
& : \text{new_refable_}R_P.\bar{K}(\bar{X}[\bar{F}]). \quad (ICN)
\end{aligned}$$

4.2.5 *Coherence and Key-Preservation.* The following rule prevents requests which are directly incoherent:

$$\begin{array}{l}
\text{blk_ins_}R(\bar{X}) \quad \wedge \triangleright \text{ins_}R(\bar{X}), \text{ del_}R(\bar{X}). \\
\text{blk_del_}R(\bar{X}) \quad \wedge \text{ins_}R(\bar{X}). \\
\text{blk_mod_}R(M, \bar{X}) \quad \wedge \triangleright \text{mod_}R(M, \bar{X}), \text{ del_}R(\bar{X}). \\
\text{blk_del_}R(\bar{X}) \quad \wedge \text{mod_}R(M, \bar{X}).
\end{array} \tag{C}$$

For every *rac* $R_P.\bar{K}, ! R_C.\bar{F}$ ON UPDATE OF PARENT CASCADE:

$$\begin{array}{l}
\text{blk_prop_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M, \bar{X}) \wedge \text{pot_prp_}R_P.\bar{K} \rightsquigarrow R_C.\bar{F}(M, \bar{X}), \text{ del_}R(\bar{X}). \\
\text{blk_del_}R(\bar{X}) \quad \wedge \text{prp_}R_P.\bar{K} \rightsquigarrow R.\bar{F}(M, \bar{X}).
\end{array}$$

Since propagated modifications are handled key-oriented as foreign-key-modifications, it is sufficient to handle contradicting modifications at this granularity: For every pair of *rics* $R_{P_1}.\bar{K}_1, ! R.\bar{F}_1$ ON UPDATE OF PARENT CASCADE and $R_{P_2}.\bar{K}_2, ! R.\bar{F}_2$ ON UPDATE OF PARENT CASCADE such that $R.\bar{F}_1$ and $R.\bar{F}_2$ overlap:

$$\begin{array}{l}
\text{blk_prop_}R_{P_1}.\bar{K}_1 \rightsquigarrow R.\bar{F}_1(M_1, \bar{X}) \wedge \text{pot_prp_}R_{P_1}.\bar{K}_1 \rightsquigarrow R.\bar{F}_1(M_1, \bar{X}), \\
\text{prp_}R_{P_2}.\bar{K}_2 \rightsquigarrow R.\bar{F}_2(M_2, \bar{X}), M_1 [M_2 \text{ inconsistent.}
\end{array} \tag{C}$$

The uniqueness of a candidate key $R.\bar{K}$ is guaranteed by the following rules:

$$\begin{array}{l}
\text{blk_ins_}R(\bar{X}) \quad \wedge \triangleright \text{ins_}R(\bar{X}), \text{ rem_refable_}R.\bar{K}(\bar{X}[\bar{K}]). \\
\text{blk_chg_}R.\bar{K}(M, \bar{X}) \wedge \text{pot_chg_}R.\bar{K}(M, \bar{X}), \text{ rem_refable_}R.\bar{K}(M(\bar{X})[\bar{K}]). \\
\text{blk_ins_}R(\bar{X}) \quad \wedge \triangleright \text{ins_}R(\bar{X}), \text{ ins_}R(\bar{Y}), \bar{X}[\bar{K}] \sqsupset \bar{Y}[\bar{K}]. \\
\text{blk_ins_}R(\bar{X}) \quad \wedge \triangleright \text{ins_}R(\bar{X}), \text{ chg_}R.\bar{K}(M, \bar{Y}), \bar{X}[\bar{K}] \sqsupset M(\bar{Y})[\bar{K}]. \tag{K} \\
\text{blk_chg_}R.\bar{K}(M, \bar{Y}) \wedge \text{pot_chg_}R.\bar{K}(M, \bar{Y}), \text{ ins_}R(\bar{X}), \bar{X}[\bar{K}] \sqsupset M(\bar{Y})[\bar{K}]. \\
\text{blk_chg_}R.\bar{K}(M, \bar{X}) \wedge \text{pot_chg_}R.\bar{K}(M, \bar{X}), \text{ chg_}R.\bar{K}(M^0, \bar{Y}), \\
M(\bar{X})[\bar{K}] \sqsupset M^0(\bar{Y})[\bar{K}].
\end{array}$$

On Delete / Update Set Default / Set Null. The additional *racs* $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE/DELETE SET DEFAULT and $R_C.\bar{F}, ! R_P.\bar{K}$ ON UPDATE/DELETE SET NULL can be handled by variants of the rules (MPC), (MPR), and (MPR). Analogously, NOT NULL conditions can be integrated into the framework.

4.2.6 *Declarative Semantics and Results.* The examples in Section 2.3 illustrate different types of ambiguities which can occur for a set RA of *racs*. These ambiguities become apparent by the declarative semantics of our logical formalization P_{RA} . Again, we consider well-founded and stable models:

Mutex. For two mutually exclusive operations (cf. Example 4), if one of them is rejected, the other can be executed: setting some undefined requests to false admits stable models where other updates are true, and the false ones are blocked. This situation is analogous to $\text{fblock}_1 \wedge \text{exec}_2, \text{block}_2 \wedge \text{exec}_1 \text{g} [\text{fexec}_i \wedge : \text{block}_i \mid i \sqsupset 1, 2 \text{g}$.

Self-Attack. For a self-attacking request (cf. Example 5), there is no other support for rejecting it than its “internal contradiction,” thus there is no total

42 † W. May and B. Ludäscher

(i.e., two-valued) stable model making it true or false. This situation corresponds to $fexec \hat{=} : block, block \hat{=} execg$, where no total stable model exists.

Definition 4.5. Every 3-valued model $\mathcal{M} : \mathbb{D} \mathcal{M}(P_{RA}, D, U_{\triangleright})$ defines sets of updates $\Delta_{\mathcal{M}}$ and user requests $U_{\mathcal{M}} \sqcup U_{\triangleright}$ which are true, false, or undefined in \mathcal{M} . Let upd be any of $ins_R(\bar{x})$, $del_R(\bar{x})$, $mod_R(M, \bar{x})$, then:

$$\Delta_{\mathcal{M}}^{true} : \mathbb{D} fupd \text{ j } \mathcal{M}(upd) \mathbb{D} trueg, \text{ and } U_{\mathcal{M}}^{true} : \mathbb{D} f \triangleright upd \text{ } \mathbb{D} U_{\triangleright} \text{ j } \mathcal{M}(upd) \mathbb{D} trueg$$

(analogous for false and undef).

Again, we first examine the *well-founded model* $\mathcal{W} : \mathbb{D} \mathcal{W}(P_{RA}, D, U_{\triangleright})$, which provides a safe, *skeptical* semantics that is computable in *polynomial* time, and the the *stable semantics* of P_{RA} . Here, overlapping and subsuming modifications must be taken into account:

Definition 4.6. For any 3-valued model $\mathcal{M} : \mathbb{D} \mathcal{M}(P_{RA}, D, U_{\triangleright})$, let $\Delta_{\mathcal{M}}^{trueC} \sqcup \Delta_{\mathcal{M}}^{true}$ denote the set of *non-subsumed updates*: the set of all $upd \in \Delta_{\mathcal{M}}^{true}$ such that there is no M^0 which subsumes M and $mod_R(M^0, \bar{x}) \in \Delta_{\mathcal{M}}^{true}$.

The result of applying a set of updates does not change when restricting to non-subsumed updates:

$$\text{LEMMA 4.7. } D \S \Delta_{\mathcal{W}}^{true} \mathbb{D} D \S \Delta_{\mathcal{W}}^{trueC}.$$

THEOREM 4.8 CORRECTNESS: WELL-FOUNDED SEMANTICS.

- (1) $\Delta_{\mathcal{W}}^{trueC}$ is admissible,
- (2) $\Delta_{\mathcal{W}}^{trueC} \mathbb{D} \Delta(U_{\mathcal{W}}^{true})$ is the set of updates that are induced by $U_{\mathcal{W}}^{true}$.
- (3) $U_{\mathcal{W}}^{true}$ is admissible; the new database after submitting $U_{\mathcal{W}}^{true}$ is $D^0 \mathbb{D} D \S \Delta_{\mathcal{W}}^{true}$.

PROOF.

- (1) Foundedness, completeness, and feasibility are proven using the rules of all $racs \text{ } ra \in RA$; coherence and key-preservation is guaranteed by the rules specifying the interaction of updates.
- (2) $\Delta_{\mathcal{W}}^{true} \sqcup \Delta(U_{\mathcal{W}}^{true})$ follows from foundedness, $\Delta_{\mathcal{W}}^{true} \nabla \Delta(U_{\mathcal{W}}^{true})$ from completeness.
- (3) follows from (1) and (2). \square

We see that already $\Delta_{\mathcal{W}}^{trueC}$ abstracts from some intermediate results by considering only non-subsumed updates. The game-theoretic characterization—that corresponds to stable models—given in Section 4.3 uses the same abstraction.

The following corollary states that \mathcal{W} is a skeptical approximation: (i) every maximal admissible $U \sqcup U_{\triangleright}$ extends $U_{\mathcal{W}}^{true}$, and (ii) updates classified as false by \mathcal{W} are definitely not admissible:

COROLLARY 4.9. For every maximal admissible $U \sqcup U_{\triangleright}$:

- (i) $U_{\mathcal{W}}^{true} \sqcup U$,
- (ii) $U_{\mathcal{W}}^{false} \setminus U \mathbb{D} ;$,
- (iii) $\Delta_{\mathcal{W}}^{false} \setminus \Delta(U) \mathbb{D} ;$.

PROOF. Follows from Theorem 4.8 and model-theoretic properties of the well-founded semantics. \square

The different types of undefined updates $upd \not\leq U_{\mathcal{W}}^{undef}$ can be characterized according to the different types of controversial atoms:

- $upd \leq U$ for every maximal admissible $U \sqcup U_{\triangleright}$ (“diamond”), or
- there are maximal admissible sets $U, U^0 \sqcup U_{\triangleright}$ such that $upd \leq U$ and $upd \not\leq U^0$ (“mutex”), or
- $upd \not\leq U$ for any admissible $U \sqcup U_{\triangleright}$ (“self-attack”).

As already mentioned, “diamonds”—even if they do not cause a problem—cause a negative cycle that is undefined in the well-founded model. Such a case is described in Example 12 in Section A of the *electronic appendix*: there, all actual modifications and changes, as well as all blockings are undefined in the well-founded model. Nevertheless, the modification is admissible.

In such cases, similar to the case of deletions only, there are two stable models: the one that executes the user request, and the one that rejects it. In this case, taking the overestimate of an alternating fixpoint as the intended result—similar to the case where only deletions were considered—is correct. On the other hand, this is *not* correct in cases such as *Mutex* (cf. Example 4), or for self-attacking requests (cf. Example 5).

Stable Models. For further investigation of these cases, we use stable models, which provide a more detailed logical semantics for normal logic programs. Since self-attacking updates exclude the possibility of total stable models, we have to consider *P-stable* (partial stable) models:

Definition 4.10 P-, M-Stable Models. [Eiter et al. 1996] Let $I \sqcup I^{true}$, I^{false} be a 3-valued interpretation (I^{true} the set ground of atoms that are true, I^{false} the set ground of atoms that are false, $I^{true} \setminus I^{false} \sqcup ;$). The reduction P/I of a ground-instantiated logic program P is obtained by replacing every negative literal in P by its truth-value with respect to I . Thus, P/I is positive and has a unique minimal (with respect to the truth-order $false <_t undef <_t true$) 3-valued model $\mathcal{M}_{P/I}$.

I is a *P-stable* model, if $\mathcal{M}_{P/I} \sqcup I$. A *P-stable* model I is *M-stable* (*maximal stable*) if there is no *P-stable* model $J \sqcup I$ such that $J^{true} \not\sqcup I^{true}$ and $J^{false} \not\sqcup I^{false}$.

In contrast to the well-founded model, which is the “most skeptical” *P-stable* model, *M-stable* models are “more brave” and handle mutually exclusive requests as expected; in particular, *all* admissible solutions are represented by *P-stable* models. This fact, and the generalization of Theorem 4.8 is expressed by the following theorem (proven analogously to Theorem 4.8).

THEOREM 4.11 CORRECTNESS, COMPLETENESS: STABLE SEMANTICS.

—For every *P-stable* model S :

- (i) $\Delta_S^{true^C}$ is admissible, (ii) $\Delta_S^{true^C} \sqcup \Delta(U_S^{true})$, (iii) U_S^{true} is admissible.

—For every maximal admissible $U \sqcup U_{\triangleright}$, there is an *M-stable* model \mathcal{M}_S such that $U \sqcup U_{\mathcal{M}_S}^{true}$ and $\Delta(U) \sqcup \Delta_{\mathcal{M}_S}^{true^C}$.

M-stable models of P_{RA} almost capture the notion of “optimal” (maximal admissible) solutions. Note that in case of mutual exclusion, there can be

44 † W. May and B. Ludäscher

several M-stable models which describe maximal admissible solutions. In case of a “diamond” $fblock \hat{\ } : exec, exec \hat{\ } : blockg$ (cf. Example 9) there are two M-stable models. However, executing an update should be preferred to blocking it in order to capture the notion of maximal admissibility. Therefore, we define a semantic ordering $<_a$ on P-stable models according to our intended application:

$$\mathcal{S}_1 <_a \mathcal{S}_2 \text{ ;. } U_{\mathcal{S}_1}^{true} \%_{00} U_{\mathcal{S}_2}^{true}.$$

Among M-stable models, $<_a$ prefers those that make more updates true. This holds as well for the user requests as for the resulting updates:

LEMMA 4.12. *For two M-stable models \mathcal{S}_1 and \mathcal{S}_2 , the following is equivalent:*

- (1) $\mathcal{S}_1 <_a \mathcal{S}_2$,
- (2) for every $\triangleright upd \in U_{\triangleright}$, $\mathcal{S}_1(\triangleright upd) \cdot_t \mathcal{S}_2(\triangleright upd)$,
- (3) for every (internal) update upd , $\mathcal{S}_1(upd) \cdot_t \mathcal{S}_2(upd)$, or $upd \in mod_R(M, \bar{x})$ for some R, M, \bar{x} , and there is an M^0 which subsumes M and $\mathcal{S}_1(upd) \cdot_t \mathcal{S}_2(mod_R(M^0, \bar{x}))$ (cf. Definition 4.6 and Theorem 4.8).

The maximal stable models with respect to $<_a$ represent exactly the maximal admissible sets:

THEOREM 4.13 MAXIMALITY. *Let D and U_{\triangleright} be as usual. Then, the following sets coincide:*

- the set of all maximal admissible sets $U \sqcup U_{\triangleright}$,
- the set of all U_{AS}^{true} such that AS is a $<_a$ -maximal M-stable model of P_{RA}, D , and U_{\triangleright} .

4.2.7 *An Upper Bound.* Above, we have shown that the well-founded model provides a lower bound of maximal admissible subsets. Analogously, an upper bound can be derived. For deletions, Theorems 3.18 and 3.19 stated that the set of all true and undefined user requests (i) is admissible, and (ii) corresponds to a stable model. In case of modifications, this, in general, does not hold:

LEMMA 4.14. *The well-founded model induces an upper bound for the set of admissible updates:*

Let $\mathcal{W} \in \mathcal{W}(P_{RA}, D, U_{\triangleright})$ be the well-founded model of P_{RA}, D , and U_{\triangleright} , and

$$U^{ub} := \{ U_{\mathcal{W}}^{true[undef]} \in U_{\triangleright} \mid \exists upd \in U_{\triangleright} \text{ j } \mathcal{W}(upd) \in \{true, undef\} \}.$$

- If U^{ub} is admissible, then there is a unique $<_a$ -maximal stable model ASU of RA, D , and U_{\triangleright} such that $U_{ASU}^{true} \in U^{ub}$. Then $\Delta(U^{ub}) \in \Delta_{ASU}^{trueC}$, and the new database is $D \sqcup \Delta_{ASU}^{trueC}$.
- for every $upd \in U_{\triangleright}$ which is in any admissible subset of U_{\triangleright} , $upd \in U^{ub}$.

Thus, the following necessary condition for admissibility U_{\triangleright} can be decided in PTIME using the well-founded model:

COROLLARY 4.15. *If $U^{ub} \subsetneq U_{\triangleright}$, then U_{\triangleright} is not admissible.*

Due to the various predicates, the above ASU is not as easy to describe based on the well-founded model as it has been for deletions (cf. Theorem 3.18). Nevertheless, for the “positive” statements (updates, propagations, and changes) in P_{RA} , we have the following “monotonicity result”:

LEMMA 4.16. *Let $\mathcal{W} : \mathbb{D} \mathcal{W}(P_{RA}, D, U_{\triangleright})$. If U^{ub} is admissible, then for every $upd \in \Delta_{\mathcal{W}}^{true[undef]}$, $ASU(upd) \sqsupset true$, or $upd \sqsupset mod_R(M, \bar{x})$ for some R, M, \bar{x} , and there is an M^0 which subsumes M and $ASU(mod_R(M^0, \bar{x})) \sqsupset true$ (cf. Lemma 4.12). (analogously for prop and chg predicates.)*

Thus, in case that U^{ub} is admissible, it is “safe” to do all changes in the database which are true or undefined in $\mathcal{W} : \mathbb{D} \mathcal{W}(P_{RA}, D, U_{\triangleright})$. We come back to this issue in Section 5. First, the game theoretic characterization gives more insight into the correctness of the stable model characterization and subsumed updates.

4.3 Game-Theoretic Semantics

The maximal admissible sets of updates can also be characterized in a game-theoretic way (the details of the game can be found in Section B of the *electronic appendix*). This “update game” needs to also consider a *history* of the game that was not required for the simpler game-theoretic characterization in Section 3.3 with deletions only (that was in the famous *win-move*-style).

For given U_{\triangleright} , Player I claims a subset $U \sqsubseteq U_{\triangleright}$ to be maximal and admissible. In her first move, Player II chooses to falsify either the maximality or the admissibility. If Player II challenges the maximality, she chooses a proper superset U^0 such that $U \subsetneq U^0 \sqsubseteq U_{\triangleright}$ which she claims to be maximal and admissible, then the roles are changed. Thus, after finitely many moves, a player challenges the admissibility of a set U suggested by the other player by examining this set with respect to its coherence and feasibility by questions. The other player has to defend U to be admissible by stepwise showing what updates are actually executed. By doing this, he constructs $\Delta(U)$ (both players are assumed to play optimally). The game is an abstraction of the logic programming characterization in the sense that I uses only non-subsumed updates (anticipating the overall result), thus the details of interfering updates can be ignored. This abstraction step is similar to that in Section 5 for deriving the practical results from the construction of the well-founded and stable models.

Setting and Initialization. The positions of the game are all tuples of the database D , and a sufficient number of empty positions for representing insertions. Actually, the “board” is practically a graphical representation of the database. The game is then played by putting *plates* that represent the update operations performed on the database: Each plate consists of a source tuple ($2 D$), an update (with additional detail information concerning the foundedness), and a result tuple, for example, $\boxed{R(a, b, c) \text{ mod_} R(1/x, 2/y], (a, b, c)) \text{ h. . . i } j R(x, y, c)}$. At the beginning, Player I puts all plates that describe a set U of updates that

46 † W. May and B. Ludäscher

he claims to be maximal admissible with respect to given D , RA , and U_{\triangleright} . E.g., for $\text{del}_R(x) \triangleright U$, he puts the plate $\boxed{R(\bar{x}) \mid \text{del}_R \bar{x} \text{h.} \dots \mid \varepsilon}$ the tuple $R(\bar{x})$.

The Moves. During the game, the admissibility (that is, foundedness, completeness, and feasibility; cf. Definition 4.1) of U in the given situation is challenged by Π , and defended by I . Π asks a question by pointing to an instance of one of the above aspects, and I has to show how to guarantee the respective property—if he has no answer, he loses (that is, Π showed that U is actually not admissible). In most answers, I puts a new plate to show that the database changes in a consistent way.

Child tuples. For any plate that describes an update to a tuple $R_P(\bar{y})$, Π can point to a referencing tuple, asking I what happens to it. I answers by putting a plate on the tuple that describes an update (depending whether the reference is maintained via CASCADE, NO ACTION or RESTRICT, different restrictions apply).

Parent tuples. For any plate that describes an update to a tuple $R_C(\bar{x})$ such that a foreign key is changed, Π can ask I what parent node is referenced. I has either to show an unchanged tuple, or to show a modification or insertion that generates a suitable parent.

Foundedness. In case of children whose references are maintained by the NO ACTION policy, I claims in his answer that the tuples are accordingly modified—but this update must be propagated from somewhere else (cf. diamonds). Similar to the “delete game” Π can ask I to prove that these updates are founded. In contrast to the “delete game,” updates can be collected from several parent tuples (cf. the *(CH)* rules of the logic programming characterization). Thus, the problem is much more involved than in the case of deletions. For each atomic component of the update, Π asks a separate question, and I has to show a parent tuple that cascades a suitable modification (by putting a plate whose foundedness is then again attacked by Π).

By putting new plates, I constructs $\Delta(U)$ (since Π asks for all CASCADE references), and U is admissible if and only if I wins the game; see Theorems B.2 B.4 in the *electronic appendix*. For challenging maximality, Π is also allowed to add another plate for $\text{upd} \triangleright U_{\triangleright}$, $\text{upd} \not\triangleright U$ and to change the roles to show that $U \uparrow \{\text{upd}\}$ is also admissible.

Model-theoretic aspects. Note that for given D , RA and U_{\triangleright} —in contrast to the case of deletions only—the question is not to win *individual* positions, but to win a game for an *initial setting* $U \uparrow U_{\triangleright}$. For a given initial setting, I either wins or loses, there is no draw. In case that I wins, the whole $\Delta(U)$ is created explicitly on the board—retaining the whole history of the game. For given D , RA and U_{\triangleright} , there can be several US such that I wins the game. Each of them is a maximal admissible subset of U_{\triangleright} . In contrast, in the “delete game,” the set of all won and drawn positions characterizes the *unique* maximal admissible subset.

5. PRACTICAL ASPECTS

5.1 Admissibility and Execution is Polynomial

In Corollary 4.15, we have given a necessary condition for admissibility of U_{\triangleright} that can be checked in P_{TIME} by using only the well-founded model. In case that this condition is satisfied, we have found in Lemma 4.14 that a certain $<_a$ -maximal stable model \mathcal{AS} has to be considered for the final check, and for computing the actual set Δ of updates to the database. In case U_{\triangleright} is not admissible, information is needed about maximal admissible subsets and about the problem situations. In the general case, there is no unique maximal admissible solution for a set of user requests, and an exponential number of stable models may have to be considered. Thus, the computation of *all* maximal admissible subsets is not feasible in practice. But, it is also not needed. For practical use, the following tasks are relevant:

- check if U_{\triangleright} is admissible, and
- if U_{\triangleright} is not admissible, locate the problem situations.

These tasks are now addressed based on the well-founded model (that can be computed in P_{TIME}), building on the results of Section 4.2.7 where an upper bound for the admissible updates has been characterized by the well-founded model. Recall also that for deriving a procedural algorithm for handling deletions in Section 3.5, the negative cycle in the logic programming characterization has been “cut” by taking some (“positive”—the requested deletions) predicates from the overestimate, and the other (“negative”—the blockings) predicates from the underestimate.

- The following predicates are *fixed* predicates: `pot_del`, `pot_prop`, `pot_chg`.
- The following are *positive* predicates: `del`, `prop`, `mod`, `ins`, `chg`, `allow_chg`.
- blockings are *negative* predicates: `blk_del`, `blk_prop`, `blk_mod`, `blk_ins`, `blk_chg`.

Positive predicates represent the knowledge of Player I of what plates to play to win the game, including that he only plays non-subsumed updates. We will consider a special stable model that is induced by the fixed and positive predicates. In contrast to guessing a stable model, or to the game-theoretic characterization (that relies on guessing the correct plates), the following “model,” \mathcal{W}^C can be generated in polynomial time:

Definition 5.1. Given $\mathcal{W} :D \mathcal{W}(P_{RA}, D, U_{\triangleright})$, let \mathcal{W}^C be defined as follows (note that \mathcal{W}^C is in general not a model of P_{RA}):

- for fixed predicates (evaluate to *true* or to *false* in \mathcal{W}), $\mathcal{W}^C(atom) :D \mathcal{W}(atom)$,
- for positive predicates, $\mathcal{W}^C(atom) \sqsupset true$ if $\mathcal{W}(atom) \not\sqsupset true$, *undefg* and *atom* is not subsumed by another one.
- $\mathcal{W}^C(atom) :D false$ for all other atoms, including those over negative predicates.

Then, \mathcal{W}^C has to be checked as to whether there is a stable model that coincides with \mathcal{W}^C on *positive* and *fixed* predicates. We call \mathcal{W}^C *positive-stable* if this

is the case, which corresponds to a game where Player Π has no successful attacks.

First, we give the relationship to stable models, especially to the $<_a$ -maximal M-stable models that characterize *maximal* admissible sets:

THEOREM 5.2 POSITIVE-STABLE VS. STABLE MODELS. *If \mathcal{W}^C is positive-stable, then there is a (unique) $<_a$ -maximal M-stable model AS that coincides in all atoms over positive and fixed predicates with \mathcal{W}^C . AS is computed by applying the computation of the well-founded model starting with \mathcal{W}^C (instead of \cdot).*

PROOF. As described above, \mathcal{W}^C corresponds to cutting negative cycles in the rules with priority to executing updates. In a stable or positive-stable model, these atoms reproduce themselves. The other atoms (blockings etc.) only “fill” the gaps in the stable model. If no blockings are derived that “kill” updates in \mathcal{W}^C , it is stable (it cannot be attacked by Player Π).

\mathcal{W}^C corresponds to the *unique* $<_a$ -maximal M-stable model since in Corollary 4.9 and Lemma 4.14, it has been shown that updates $upd \supseteq U_{\triangleright}$ that are false in \mathcal{W} are not contained in any admissible $U \%_0 U_{\triangleright}$. \square

In the following, we call the above AS the *positive-stable model to \mathcal{W}* . Note again that such a model does not always exist (e.g., if \mathcal{W}^C contains mutually exclusive or “self-killing” updates). With respect to the upper bound that has been discussed in Section 4.2.7, the above theorem checks whether U^{ub} is admissible. Especially, concerning the admissibility check for U_{\triangleright} , the following result shows that this task can be done in PTIME since computing \mathcal{W} and AS is polynomial:

COROLLARY 5.3 ADMISSIBILITY OF U_{\triangleright} . *If $U^{ub} \supseteq U_{\triangleright}$ and \mathcal{W}^C is positive-stable (that is, there is an AS which is the positive-stable model to \mathcal{W}), then U_{\triangleright} is admissible.*

For computing the set of induced updates, we again consider the game-theoretic characterization with its set of “winning plates”:

LEMMA 5.4 INTERNAL UPDATES IN AS VS. GAME. *If AS is the positive-stable model to \mathcal{W} , then $\Delta_{AS}^{true^C} \supseteq \Delta_{\mathcal{W}^C}^{true^C} \supseteq \Lambda$ where Λ is as defined in Section B.3 based on the plates that are played by I for winning the game.*

Proof Sketch. The non-subsumed updates are exactly those that are rep-

COROLLARY 5.5 MAIN RESULT. *If \mathcal{AS} is the positive-stable model to \mathcal{W} and $U_{\mathcal{AS}}^{true} \sqsupset U_{\triangleright}$, then U_{\triangleright} is admissible and $\Delta_{\mathcal{AS}}^{true^C} \sqsupset \Delta(U_{\triangleright})$. The whole problem can be solved in polynomial time.*

Thus, the steps for executing a set U_{\triangleright} of updates on a database D are as follows:

- (1) compute $\mathcal{W} \sqsupset \mathcal{W}(P_{RA}, D, U_{\triangleright})$,
- (2) check whether $U_{\mathcal{W}}^{true,undef} \sqsupset U_{\triangleright}$,
- (3) if not, reject, otherwise check if there is a positive-stable model \mathcal{AS} to \mathcal{W} ,
- (4) if not, reject, otherwise compute $\Delta_{\mathcal{AS}}^{true^C}$ and apply it to the database.

Note that if U_{\triangleright} is accepted by the SQL semantics as presented in Horowitz [1992] and specified in the SQL3 standard [ANSI/ISO 1999] (see also Section 2.4), the semantics coincides with ours. Similarly to the considerations in Section 3.6, the global correctness is implied by the correct specification of the individual rules and the meta-correctness of the logic programming semantics. Again, in contrast to the procedural semantics for SQL, the logic-based characterization also provides information if U_{\triangleright} is rejected.

5.2 Troubleshooting in Case of Rejected Updates

Similarly to Section 3.7, in case U_{\triangleright} is not admissible, the information contained in the well-founded model \mathcal{W} and the corresponding stable model \mathcal{AS} can be used for deriving debugging hints. Problem situations are defined analogously as in Definition 3.24; additionally there can be problems due to conflicts between updates.

The above considerations also apply for the admissibility of $U^{ub} \sqsupset U_{\mathcal{W}}^{true,undef}$ that provides an upper bound to which subset of U_{\triangleright} can possibly be admissible. Thus, first problems are already identified during the computation of \mathcal{W} . The first time that a blocking for an $upd \in U_{\triangleright}$ or for a propagation along a cascading reference is derived in an underestimate, a problem situation is identified that can be reported.

Later, if $U^{ub} \sqsupset U_{\triangleright}$, but it is still not admissible, the check as to whether \mathcal{W}^C is positive-stable immediately shows where the problems are located: the first rule in the subsequent computation of the well-founded model that derives a blocking against an update in $\Delta_{\mathcal{W}^C}^{true^C}$ identifies the problem situation. The conclusions as to how to cure the problem are the same as described in Section 3.7.

5.3 Refined Analysis

In the above well-founded model \mathcal{W} , in general most updates will be undefined, thus, the subsequent investigation of \mathcal{W}^C is generally necessary. There are several typical situations, where even single, admissible updates yield only undefined updates in \mathcal{W} , for example, diamonds or cyclic dependencies. These can be detected by schema analysis, and then can be handled by refining the first (MCN) rule.

Diamonds. For a given schema, it is possible to detect diamonds a-priori by checking *transitive* dependencies and appropriately modifying the program

(similar to the use of the *referential action graph* in Gallagher [1986] and Horowitz [1989]). For every diamond, the key on top of the diamond, and its children are considered when reasoning about changes. Changes inside the diamond are only checked if they are due to references that come into the diamond from outside.

Cyclic Dependencies in a Single Update. Cyclic dependencies also lead to undefined modifications and blockings. In fact, cyclic dependencies are a special case of the diamond where the top relation is the same as the bottom relation.

Cyclic Dependencies between Updates. NO ACTION references can lead to cyclic dependencies between a set of updates. In this case, the well-founded model yields undefined updates although often all of them as a group are admissible. The “second try,” based on \mathcal{W}^C is then successful in deriving a total model.

In the above cases, undefined atoms in the well-founded model were caused by problems of the logic programming characterization and could be eliminated by either changing to stable models or rewriting the program according to a given schema. Additionally, undefined atoms in the well-founded model can be caused by mutual exclusion or “self-killing” updates.

Self-Killing Updates. An update is self-killing if it causes conflicting cascaded updates. In most cases this points to a severe design error. Again, potential problems can be detected by regarding the transitive closure of referential actions.

Mutual Exclusion: Multiple Admissible Sets. Mutual exclusion is the only case where the expressive power of the well-founded model is not sufficient (it cannot represent the choice between alternatives in the result). It also results in undefined updates and blockings wherever atoms are involved in mutual exclusion. Here, stable models provide the cure to the problem.

6. RELATED WORK

Referential integrity for relational databases was first considered in Codd [1970] with an implicit *or-semantics* between a foreign key and *several* parent keys (in different tables). The definition has been reformulated in Date [1981]; specifying whether a parent key must be present in *all*, *some*, or *exactly one* of the parent tables. Later [Date 1990], the *or-semantics* were cancelled and *referential actions* were defined for enforcing these referential integrity constraints after the execution of updates.

In the SQL2 [ANSI/ISO 1992a] standard, referential integrity constraints and referential actions are specified using the syntax given in Section 2.2, according to the above *all-semantics*. While the syntactic specification of referential actions in SQL is in a declarative style, the given procedural semantics causes ambiguities in some cases of network-like referential structures [Hammer and McLeod 1975; Markowitz 1990, 1991b] due to different execution orderings of referential actions (see also Date [1990] and Date and Darwen [1994]). The same characterization was also used in the early SQL3 working drafts, for example, ANSI/ISO [1991, 1992b, 1994].

Motivated by the problem of ambiguities, several research directions have been followed. *Schema-based* analysis provides conditions that are sufficient, but more restrictive than necessary. In Markowitz [1991b], such a *schema-based* strategy is presented to exclude ambiguous situations. The approach considers only delete operations with referential actions CASCADE, SET NULL, and RESTRICT; neither DELETE NO ACTION nor UPDATE CASCADE are allowed. Thus, already situations similar to the “diamond” from Example 3 (which has an intuitively clear semantics) are not considered. The approach is tuple-oriented, and there is no “natural” extension of the solution for updates. The approach is refined in Markowitz [1994], dropping several simplifying assumptions.

Another schema-based approach using a relation-column-based (instead of tuple-based) *referential action graph* was presented in Gallagher [1986] and Horowitz [1989].

A “*semantic*” approach is followed by Markowitz [1990]: Starting from an object-oriented or EER model of an application, criteria are given as to how to map this model to a relational model with *rics* and *racs*, depending on the semantics of a reference (of “blocking” or “cascading” nature). Here, the problem of cascading modifications is solved by the introduction of *surrogate attributes*, which simulate a kind of object-identity. Then, it is shown that the result does not have the negative properties identified in Date [1990]—that there cannot occur any ambiguities.

At that time, in commercial database systems ON DELETE/UPDATE NO ACTION was used by default, only ON DELETE CASCADE could be specified optionally. Thus, especially cascading modifications were not supported. For an overview of the support for referential integrity in commercial RDBMSs at that time, see Markowitz [1991a].

Compile-time approaches: based on triggers and schema information only, are too restrictive, since their criteria are sufficient but not necessary to avoid ambiguities (cf. the abovementioned undecidability of the problem). In other words, they prohibit a schema if there is the *possibility* of an anomaly in a “wrong-use-worst-case.” Reinert [1996] shows that it is undecidable whether a given *schema* with referential actions can, for some database instances, lead to ambiguous update situations under the SQL2 semantics. In contrast, the problem becomes decidable for a *given* situation: when a database instance and a set of updates is given. Such considerations have led to the investigation of *run-time* approaches:

Horowitz [1992] presented a procedural execution model using a marking algorithm based on bookkeeping about deletions and modifications under *the restriction that keys consist only of a single column*. Thus, this (in practice unrealistic) assumption avoids the problems of overlapping keys and foreign keys. The evaluation model has been extended and accepted for the SQL3 [ANSI/ISO 1999] standard (see also Section 2.4); also commercial DBMS implementations conform to it (as far as they support referential actions). An integration of the semantics of SQL triggers and declarative constraints (including referential integrity constraints and referential actions) is investigated in Cochrane et al. [1996]. Their model is fully compatible with the SQL2 (and also the later SQL3)

requirements. The internal problems of referential actions are not considered. From the aspect of *active databases*, production rules have been considered for investigating triggers and referential actions, for example, Abiteboul and Vianu [1991], Picouet and Vianu [1995]; Ceri and Widom [1990] and Widom et al. [1990].

SQL3 [ANSI/ISO 1999] extends the above-mentioned marking strategy given in Horowitz [1992] for cascading referential actions (dropping the unary-key-constraint). The actual SQL3 specification is then given in terms of a complex characterization via BEFORE triggers. SQL3 also specifies several “levels”: For the *intermediate SQL* specification (which is currently supported by commercial systems), update rules are not allowed (i.e., the only ON UPDATE action is the default ON UPDATE NO ACTION). For *full SQL*, all referential actions are allowed. Note that the approach is *not key-oriented* but tuple-oriented: already *non-interfering* updates on a tuple (concerning disjoint foreign keys) are explicitly forbidden. For updates that are allowed by full SQL3, our semantics coincides with the one specified for SQL3—note that our semantics does not use longwinded procedural or trigger-based algorithms, but only specifies the intuitive *local* behavior whereas the *global* semantics is naturally given by the well-known logic programming semantics.

Commercial DBMSs. For a long time, in commercial DBMSs, ON DELETE/UPDATE NO ACTION was used by default, only ON DELETE CASCADE could be specified optionally. Only recently, Microsoft SQL Server (since 2000) and PostgreSQL support ON UPDATE CASCADE.

7. CONCLUSIONS

We have investigated the semantics of arbitrary sets of user-requested updates to a database in the presence of referential integrity constraints and referential actions. Our results contribute both to database theory, and to the practical implementation of referential actions.

Theoretical Aspects. We have shown how a logic-based specification of referential actions can yield a better understanding of ambiguities and conflicts: the logic programming characterization shows how to obtain a declarative *global* semantics that is computable in polynomial time from a concise and intuitive *local* definition of *rac*s accessing only two tuples at a time. The game-theoretic formalizations abstract from details and provide additional insight into the behavior of *rac*s and are used to establish the correctness of the logic programming formalization.

In contrast to the latest SQL3 standard [ANSI/ISO 1999], where the semantics is specified by a complicated, procedural computation (after previous versions that suffered from ambiguities), our results show that the well-known semantics for general logic programs already unambiguously specify a semantics of *rac*s (that coincides with that of ANSI/ISO [1999]).

Provided one accepts the appropriateness of the well-established Logic Programming semantics, our semantics is the “natural” semantics of referential integrity constraints and referential actions specified in SQL’s ECA-style syntax.

Practical Aspects. For the restricted case of deletions only, we have shown how this specification can be transformed into an efficient algorithm and implementation in an arbitrary procedural language. For modifications *in the presence of referential actions of the form ON UPDATE CASCADE* (which is still only partly supported in commercial database systems), we have shown how the logical characterization can be used for efficiently checking admissibility of a set of updates, and for computing the subsequent database state. In both cases, if the initial set of updates is not admissible, the computation can also be used for detecting the exact location of the problems, and for giving hints as to how to change the application specification.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: www.acm.org/pubs/citations/journals/tods/20YY-V-N/p1-May.

It contains an example for the logic-programming characterization of referential actions in the case of update operations that has been given in Section 4.2. Additionally, it describes the *detailed* game-theoretic characterization of that case that has been sketched in Section 4.3, illustrates it by an example, and shows its equivalence with the logic programming characterization.

ACKNOWLEDGMENTS

The authors thank Joachim Reinert and Georg Lausen for fruitful discussions especially in the early stages of this work, and Jörg Flum for interesting discussions about logic programming and game theory. We also want to thank Richard Snodgrass and three anonymous reviewers for their constructive critique and suggestions as to how to improve the presentation of the different aspects of the paper.

REFERENCES

- ABITEBOUL, S., HULL, R., AND VIANU, V. 1995. *Foundations of Databases*. Addison Wesley.
- ABITEBOUL, S. AND VIANU, V. 1991. Datalog extensions for database queries and updates. *J. Comput. Syst. Sci.* 43, 1, 62–124.
- ANSI/ISO. 1991. (ISO/ANSI working draft) Database Languages—SQL3 ANSI X3H2-91-254.
- ANSI/ISO. 1992a. Information Technology—Database Languages—SQL2 Standard.
- ANSI/ISO. 1992b. (ISO/ANSI working draft) Database Languages—SQL3 ANSI X3H2-92-154.
- ANSI/ISO. 1994. (ISO/ANSI working draft) Database Languages—SQL3 ANSI X3H2-94-080.
- ANSI/ISO. 1999. Information Technology—Database Languages—SQL3 Standard.
- CERI, S., COCHRANE, R., AND WIDOM, J. 2000. Practical applications of triggers and constraints: Success and lingering issues (10-year award for Ceri and Widom [1990]). In *International Conference on Very Large Data Bases (VLDB)*. 254–262.
- CERI, S. AND WIDOM, J. 1990. Deriving production rules for constraint maintenance. In *International Conference on Very Large Data Bases (VLDB)*. 566–577.
- COCHRANE, R., PIRAHESH, H., AND MATTOS, N. 1996. Integrating triggers and declarative constraints in SQL database systems. In *International Conference on Very Large Data Bases (VLDB)*. 567–578.
- CODD, E. 1970. A relational model for large shared data banks. *Comm. ACM* 13, 6, 377–387.

- DATE, C. 1981. Referential Integrity. In *International Conference on Very Large Data Bases (VLDB)*. 2–12.
- DATE, C. 1990. *Relational Database Writings 1985–1989*. Addison-Wesley.
- DATE, C. AND DARWEN, H. 1994. *A Guide to the SQL standard: A User's Guide to the Standard Relational Language SQL*. Addison-Wesley.
- DAYAL, U. 1988. Active database management systems. In *Proceedings of the 3rd International Conference on Data and Knowledge Bases: Improving Usability and Responsiveness*. Morgan Kaufmann, 150–169.
- DIX, J. 1995. Semantics of logic programs: Their intuitions and formal properties. In *Logic, Action and Information*, A. Fuhrmann and H. Rott, Eds. de Gruyter.
- EITER, T., LEONE, N., AND SACCA, D. 1996. The expressive power of partial models for disjunctive deductive databases. In *International Workshop on Logic in Databases (LID)*. Springer LNCS 1154.
- ESWARAN, K. P. 1976. Specification, Implementation and Interactions of a Trigger Subsystem in an Integrated Database System. IBM Research Report RJ-1820(26414).
- FAN, W. AND SIMÉON, J. 2000. Integrity constraints for XML. In *ACM Symposium on Principles of Database Systems (PODS)*. 23–34.
- GALLAGHER, L. 1986. Referential integrity. ANSI X3H2-86-164.
- GELFOND, M. AND LIFSCHITZ, V. 1988. The stable model semantics for logic programming. In *International Conference on Logic Programming (ICLP)*. 1070–1080.
- HAMMER, M. AND MCLEOD, D. 1975. Semantic Integrity in a Relational Data Base System. In *International Conference on Very Large Data Bases (VLDB)*. 25–47.
- HOROWITZ, B. 1989. Relaxing referential integrity syntax rule restrictions. ANSI X3H2-89-260.
- HOROWITZ, B. M. 1992. A run-time execution model for referential integrity maintenance. In *International Conference on Data Engineering (ICDE)*. 548–556.
- LUDÄSCHER, B. 1998. Integration of active and deductive database rules. Ph.D. thesis, Institut für Informatik, Universität Freiburg.
- LUDÄSCHER, B. AND MAY, W. 1998. Referential actions: From logical semantics to implementation. In *6th International Conference on Extending Database Technology (EDBT)*. Springer LNCS 1377, 404–418.
- LUDÄSCHER, B., MAY, W., AND LAUSEN, G. 1996a. Nested transactions in a logical language for active rules. In *International Workshop on Logic in Databases (LID)*. Springer LNCS 1154, 196–222.
- LUDÄSCHER, B., MAY, W., AND LAUSEN, G. 1997. Referential actions as logical rules. In *Proceedings of the 16th ACM Symposium on Principles of Database Systems (PODS)*. 217–227.
- LUDÄSCHER, B., MAY, W., AND REINERT, J. 1996b. Towards a logical semantics for referential actions in sql. In *6th Intl. Workshop on Foundations of Models and Languages for Data and Objects (FMLDO)*. 57–72.
- MARKOWITZ, V. M. 1990. Referential integrity revisited: An object-oriented perspective. In *International Conference on Very Large Data Bases (VLDB)*. 578–589.
- MARKOWITZ, V. M. 1991a. Problems underlying the use of referential integrity in relational database management systems. In *International Conference on Data Engineering (ICDE)*.
- MARKOWITZ, V. M. 1991b. Safe referential integrity structures in relational databases. In *International Conference on Very Large Data Bases (VLDB)*. 123–132.
- MARKOWITZ, V. M. 1994. Safe referential integrity and null constraint structures in relational databases. *Information Systems* 19, 4, 359–378.
- PICOUET, P. AND VIANU, V. 1995. Semantics and expressiveness issues in active databases. In *ACM Symposium on Principles of Database Systems (PODS)*.
- PRZYMUSINSKI, T. C. 1988. On the declarative semantics of deductive databases and logic programs. In *Foundations of Deductive Databases and Logic Programming*, J. Minker, Ed. Morgan Kaufmann, 191–216.
- REINERT, J. 1996. Ambiguity for referential integrity is undecidable. In *Constraint Databases and Applications*, G. Kuper and M. Wallace, Eds. Springer LNCS 1034, 132–147.
- VAN GELDER, A. 1993. The alternating fixpoint of logic programs with negation. *J. Comput. Syst. Sci.* 47, 1, 185–221.

Understanding the Global Semantics of Referential Actions using Logic Rules † 55

VAN GELDER, A., ROSS, K., AND SCHLIPF, J. 1991. The well-founded semantics for general logic programs. *J. ACM* 38, 3 (July), 620–650.

WIDOM, J., COCHRANE, R., AND LINDSAY, B. 1991. Implementing set-oriented production rules as an extension to Starburst. In *International Conference on Very Large Data Bases (VLDB)*. 275–285.

XML Schema 2000. XML Schema Parts 1/2: Structures/Datatypes. W3C Candidate Recommendation, www.w3.org/TR/xmlschema-1,2.

Received January 2001; revised April 2002, July 2002; accepted July 2002