Appeared in: Conrad, S.; Höding, M.; Janssen, S. and Saake, G. (Eds.), Kurzfassungen des Workshops "Föderierte Datenbanken", Magdeburg, 22.04–23.04.96, Institut für Technische Informationssysteme, Universität Magdeburg, Pages 51–65, April 1996.

Using Active Mechanisms for Global Integrity Maintenance in Federated Database Systems*

Can Türker and Stefan Conrad

Institut für Technische Informationssysteme
Otto-von-Guericke-Universität Magdeburg
Postfach 4120, D–39016 Magdeburg
E-mail: {tuerker|conrad}@iti.cs.uni-magdeburg.de
Tel.: ++49/391/67-{12994|18066} Fax: ++49-391-67-12020

Abstract

Federated database systems have special requirements w.r.t. integrity maintenance. For instance, it must be possible to express inter-database dependencies in an adequate language and the federated system must be able to enforce such constraints. In order to support these demands, we present an approach where we integrate the active rule paradigm into a federated database framework. Thereby, a very flexible way of expressing integrity constraints involving multiple component database systems is provided. Furthermore, active rules (ECA-rules) are a powerful mechanism for specifying enforcement of constraints in case of violation. ECA-rules can be seen as a fundamental communication mechanism which can be used by different transaction models. For enforcing constraints the existing local mechanisms (e.g. triggers) of the component database systems can be employed. In case this is not possible, an additional layer has to be introduced which realizes the active rule mechanism for the component database system.

Keywords: active integrity maintenance, ECA-rules, federated database systems, global integrity constraints, heterogeneous databases and interoperability.

1 Introduction

Following the needs of present and future database applications, interoperability of database systems becomes more and more important. Today, nearly every enterprise has to face the situation that different database systems are used in its departments. Problems caused by this situation are redundant data, inter-database dependencies and data-exchange between different databases. For each database system there often exists

*This research was partially supported by the ESPRIT Basic Research Working Group No. 8319 ModelAge (A Common Formal **Model** of Cooperating Intelligent **Age**nts) and by the German Country Sachsen-Anhalt under FKZ: 1987A/0025 (Federating Heterogeneous Database Systems and Local Data Management Components for Global Integrity Maintenance).

a large number of applications which have to be preserved. On the other hand, new multidatabase applications must be supported, for instance for realizing workflow systems.

The idea of *federated database systems* has been developed to bring together these requirements (cf. [SL90, LMR90, BHP92, RS94]). A federated database system consists of a number of possibly heterogeneous database systems which are considered to be (partly) autonomous. There are local applications which are working on one of the local database systems without knowing about the other ones. In addition, there can be global applications which access the federated database system through a global interface without knowing by which local database system the corresponding data is managed.

A major aspect in such environments concerns the (global) consistency of federated data. However, global integrity maintenance and enforcement, respectively, seems to contradict the autonomy of the local database systems involved. But on the other hand, there is a clear demand for having global integrity constraints which are needed for instance to maintain the consistency of replicated data and inter-database dependencies. Therefore, an approach for dealing with global integrity constraints in federated database systems is required respecting as much as possible the autonomy of the local database systems.

The approach we describe in this paper is based on mechanisms as they are used for integrity enforcement in active database systems (e.g. SAMOS [GD94b], REACH [BZBW95], Sentinel [CKTB95]). From our point of view, active rules (ECA-rules) are a promising principle not only for integrity enforcement in single, centralized database systems, but also for federated heterogeneous systems. This is due to the fact that the active rule paradigm is in principle system-independent. The active rule mechanism can be considered as a communication mechanism between the component database systems and the federation layer. Therefore, it is rather straightforward to use active rules on the global (federation) level of the system to specify integrity constraints and actions which have to be executed in case of potential integrity violations.

A main idea of our proposal is that we use as much as possible existing local mechanisms for integrity maintenance. Some commercial database systems, e.g. several relational database systems like SYBASE, already provide triggers which may be considered as a subset of active mechanisms. Thereby, an efficient realization is obviously possible by simply using these existing mechanisms. Other (commercial) database systems only offer very restricted mechanisms for integrity enforcement. Furthermore, there are systems which do not support any kind of integrity maintenance at all. For these systems an additional layer has to be introduced into the architecture of a federated database system. We have to emphasize that we are considering tightly coupled federated database systems (cf. [SL90]) in this paper. This is because in our minds tightly coupled federated database systems are better suited for global constraint enforcement than loosely coupled ones in which global integrity constraints do not play a major part anyway.

The remainder of this paper is organized as follows: In Sect. 2 we recall the architecture of federated database systems. In Sect. 3 we briefly survey the concepts and notions of active database systems. In Sect. 4 we consider active rules in the context of a federated database environment. In the main part of this paper (Sect. 5), we propose the application of active mechanisms for integrity enforcement in federated systems. In particular, we show how active mechanisms can be integrated into federated database systems. We illustrate our approach in Sect. 6 by presenting some examples for specifying global integrity constraints by means of ECA-rules. Finally, we conclude by giving an outlook on future work.

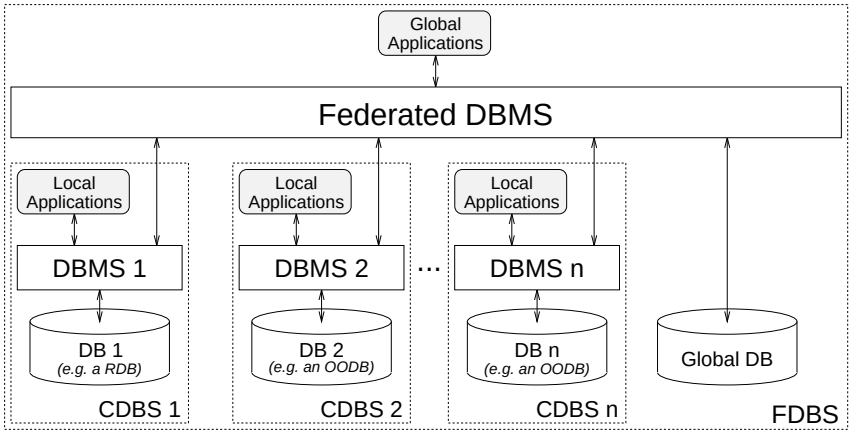

Figure 1: Architecture of federated database systems.

2 Federated Database Systems

In this section, we give a brief overview of concepts and notions of federated database systems. Due to the fact that there are slightly different understandings of what a federated database system is, we have to clearly state our view which goes along with [SL90].

Federated database systems (FDBSs) are a particular kind of multidatabase systems. A multidatabase system consists of several, possibly heterogeneous database systems which cooperate or interoperate in some way. Multidatabase systems can be split up into federated and nonfederated database systems. The latter ones are characterized by the property that their *component database systems* (CDBSs) are not autonomous. Hence, for a multidatabase system to be an FDBS we require a certain degree of autonomy for the CDBSs. In conclusion, an FDBS is a system which consists of a collection of cooperating CDBSs which are to a certain degree *autonomous* and possibly *heterogeneous*.

Besides, FDBSs are categorized as loosely coupled or tightly coupled. In loosely coupled FDBSs the user is responsible for creating and maintaining the federation. In contrast, a federation is tightly coupled if the FDBS and its administrator(s) have the responsibility for creating, managing, and actively controlling the federation. In the following, we focus on tightly coupled FDBSs, because we are interested in global constraint maintenance in FDBSs and it is quite obvious that the possibility of establishing a global constraint enforcement in loosely coupled FDBSs are rather limited. Moreover, in typical application scenarios for loosely coupled FDBSs the aspect of establishing and enforcing global integrity constraints is not significant.

An architecture for tightly coupled FDBSs must provide an extensible federation layer (see Fig. 1). This layer provides a global interface for the *uniform* and *transparent* access to federated data. At the federation layer, there exists a global *common data model* (CDM) onto which the schemata of the specific data models of the CDBSs can be mapped. Because of the requirements that an FDBS should support complex objects and has to be flexible (easy to extend), an object-oriented data model seems to be adequate as common data model. [RS94, BFN94], for instance, propose a model conforming to ODMG-93.

A central issue of FDBSs is that the independent CDBSs can continue performing their local operations while they are participating in a federation. In other words, CDBSs retain their interfaces so that existing applications can continue without any modifications. In FDBSs two kinds of applications are distinguished: local and global (federated)

applications. According to [SL90], global applications access federated data by sending requests to the FDBS while local applications submit their operations directly to a specific CDBS (cf. Fig. 1). Hence, local operations are performed outside the control of the federation layer and may violate global integrity constraints.

Obviously, global integrity maintenance and enforcement, respectively, in federated database systems always requires a certain reduction of autonomy of the local database systems. Therefore, a way has to be found for introducing global integrity maintenance without restricting local autonomy too much. In Sect. 5 we present a pragmatic approach which seems to preserve a high degree of local autonomy by applying mechanisms of active databases on the global level for integrity maintenance and by making as much as possible use of existing local mechanisms for integrity maintenance. For that, in the following section we recall the basic notions and issues of active databases.

3 Active Database Systems

Recently, *active database systems* (ADBSs) have become a very important research field (see for example [Buc94, DHW95]). Initially, active database systems have been proposed as a new data management paradigm to capture more real-world semantics already in the database rather than in applications programs. In fact, active database systems provide a powerful and flexible platform, especially, for application domains in which timely or time-constrained response (which cannot be foreseen in advance) to critical situations is required. Examples for such kinds of applications are air-traffic control, real-time plant control, network management, or workflow management. Common to all these application domains is the demand for handling dynamically changing behavior and environment in an efficient and appropriate way.

Basically, an ADBS is characterized as an extended full-fledged database system which supports *(re)active behavior* beside the conventional database features like transaction management, concurrency control or recovery control. In other words, an ADBS is an extended conventional database system which has the capability to monitor predefined situations (situations of interest) and to react to them under given time constraints with likewise predefined action. Such (re)active behavior is generally expressed by the so-called *event-condition-action rules (or ECA-rules)* which define *what to do* if a certain situation occurs in the database.

In the sequel, we survey the essential features of active database systems. We begin with recalling the ECA-rule concept, thereafter we present the common ingredients of rule definition, rule management and rule execution models, and finally we conclude this section by very briefly presenting two well-known architecture approaches of active database systems.

The ECA-Rule Concept. Today, the ECA-rule concept as introduced in the HiPAC project [CBB+89] is a widely accepted means for describing and realizing active capabilities of active database systems. An ECA-rule consists of three parts: an event, a condition, and an action. The event part of an ECA-rule specifies *when* the rule is to be *triggered (or fired)*, the *condition* determines *if* the action shall be executed, and the *action* defines *what to do* in this case. An ECA-rule can be degraded to an *EA-rule* if the

condition part is missing. In that case the action is executed whenever the corresponding event occurs.

Rule Definition. An active database system must provide a rule definition language as means for specifying (re)active behavior in forms of ECA-rules. In the last years, several languages were proposed for *active relational database systems (ARDBSs)* as well as for *active object-oriented database systems (AOODBSs)*. In the following, we briefly discuss the rule definition which consists of the specification of the three components event, condition, and action.

Event Specification. An event is related to a given point in time. In the literature, e.g. [GJS92a, CM93, GD94b], the following event types[1] are identified:

- *database event*: insertion, deletion, update, and selection of a database object.
- *transaction event*: begin, commit, and abort of a transaction.
- *method call event*: call of an arbitrary method.
- *temporal event*: absolute, relative, and periodic time events.
- *external event*: arbitrary event induced by external devices, users, etc.
- *composite event*: complex event constructed by operators of an event algebra.

Events are classified into primitive and composite events. *Primitive* events in general refer to database or transaction operations, or arbitrary method calls (in case of an AOODBS). Besides, absolute time events and external (or abstract) events are also considered as primitive events. In general, most of the rule definition languages support additional *before* and *after* specifiers to refer to the begin or the end of the execution of an operation (or method). *Composite* events [GD94a, GJS92b], on the other hand, are events which are composed by several "smaller" events using event operators (e.g. *disjunction, conjunction, sequence, history, closure*). Of course, rule definition languages vary in the complexity of specifiable events because of different data models. Whereas AOODBSs support in general most of the events mentioned above, ARDBSs provide only a subset of these events, for instance only database events.

Condition Specification. Whereas AOODBSs support arbitrary predicates and queries over the database state as condition, ARDBSs usually allow only to formulate predicates. In case of a predicate the condition is satisfied if the predicate is true. In the other case the condition is satisfied if the result of the query is non-empty.

Action Specification. The action part of an ECA-rule specifies the operations to be performed when the ECA-rule is triggered and its condition is satisfied. Whereas in AOODBSs arbitrary sequences of arbitrary operations (i.e. methods) can be defined in the action part of an ECA-rule, ARDBSs allow only arbitrary sequences of retrieval and modification operations (incl. rollbacks).

[1] From now on, we simply talk about events whenever the distinction to event instances which occur at run-time is clear from the context.

Rule Management. Usually, a large number of rules exist in a database. In order to be efficient, an ADBS must organize these rules in an appropriate way, for example in collections of enabled and disabled rules. Thus, the number of rules that must be monitored at any given time can be controlled by activating and deactivating rules.

In ARDBS, rules are defined as metadata in the database schema and thus refer to a particular table. Similar to other metadata, operations are provided to add, drop, or modify rules. In contrast, rules in AOODBS are first class objects. Thus, rules can be organized in rule hierarchies, can have attributes (e.g. condition, and action), and can be created, deleted, modified, included into collections, related to other objects etc.

Rule Execution. The rule execution consists of two phases. In the first phase, which is triggered by the occurrence of the corresponding event, the condition of the rule is evaluated. If the evaluation yields true, the second phase which is the execution of the action part of the rule is started.

Both, condition evaluation as well as action execution are performed in transaction boundaries. These transactions are called *triggered transactions* whereas the transaction in which the triggering event occurs is called *triggering transaction*. The relationships between triggering and triggered transactions are defined in the rule execution model. One of these relationships is the commit and abort dependency, e.g. the triggering transaction may only commit if the triggered transaction commits. In addition, there are temporal dependencies concerning the triggering time of the triggered transaction, e.g. the condition must be evaluated immediately after the event occurrence. *Coupling modes*, as introduced in HiPAC [CBB+89], determine how the rule parts are connected together. Up to now, the following coupling modes are proposed: *immediate*, *deferred*, *detached independent*, *detached dependent*, *detached sequential*, and *detached exclusive* (for details see e.g. [BBKZ94]). In the immediate case, the rule execution is started immediately after the event occurrence (in a subtransaction). In the deferred mode, the rule execution is delayed until the end of the transaction. In the detached modes, the rule execution is performed in separate (detached or decoupled) transactions.

Architectures of Active Database Systems. Basically, there are two main approaches for building active database systems: the *layered (or on-top)* approach and the *integrated (or built-in)* approach.

In the layered approach, all active components are built on top of a closed, passive database system. Due to the fact that the underlying database system is closed, the active components use the public interfaces as usual application programs do. SAMOS [GD94b], for example, provides active functionality on top of the object-oriented database system ObjectStore.

In contrast, in the built-in approach all active capabilities are integrated into the kernel of a database system. This approach is used, for example, by the prototype systems Sentinel [CKTB95] and REACH [BZBW95] which integrate active mechanisms into the object-oriented database system OpenOODB.

So far, we have presented the essential notions and issues of active database systems needed to understand how active mechanisms can be used in a federated database environment for integrity constraint enforcement.

4 ECA-Rules in Federated Database Environments

As we have seen in Sect. 3, the ECA-rule concept provides a uniform and powerful mechanism which can be used for integrity constraint enforcement, derived data maintenance, version control, and several other database features. In this section, we show how far ECA-rules can be adapted to a federated database environment. Here, we especially discuss how the specific rule parts have to look like in this case.

Necessary event types for (active) federated database systems. According to [CM93, GD94b], an active system must support the specification and detection of various event types. Starting out from a common object-oriented data model at the federation level, only two primitive event types, method call and (absolute) time events, have to be distinguished (cf. [BZBW95]). This is due to the fact that in an object-oriented database system all operations, e.g. database and transaction operations, are realized in terms of methods.

In addition to primitive events, an appropriate event algebra [CM93, GD94a, BZBW95] is needed for composing events by applying operators on primitive events. Here, the question is whether the proposed operators (see Sect. 3) are usable in federated database environments. Obviously, the conjunction and disjunction operators can be applied in the same way as usual. But, for the other proposed operators (sequence, closure, and history) there are semantic problems which are due to possible communication delays. Please note that in a federated database environment events may occur at two levels, i.e. in the component databases (CDBs) or in the global database (GDB). Therefore, we need active components, e.g. event detectors, in the CDBSs as well as in the federation layer.

Possible conditions. In federated database environments we need complex conditions which may refer to states of various CDBSs, for instance

if (object **in** DB1) **and** (object **not in** DB2).

Such complex conditions have to be subdivided into several (sub-) queries for the related CDBSs and the results of these queries must be merged at the federation level. In order to formulate complex conditions across several CDBSs, an algebra is required in which operators (e.g. and, or, not) for combining the results of subqueries are defined.

Admissible actions. Here, user-defined methods as well as system-provided operations (like insert, update or delete) should be usable in the action part of an ECA-rule. Furthermore, predefined transactions or applications of a specific CDBS should be used as valid actions. It must also be possible to start global actions by calling methods or operations of different CDBSs. To realize such global actions, an FDBS must provide appropriate mechanisms for synchronizing actions.

Coupling modes. The coupling modes of an ECA-rule determine when and in which transaction boundaries the condition and action parts will be executed. Therefore, the power and flexibility of an active system heavily depends on the coupling modes supported by the system. According to [BBKZ94], active systems have to support various coupling modes, especially in case of open environments. From our point of view, all proposed coupling modes are useful and needed for integrity enforcement purposes in federated database environments. In order to support these coupling modes, a FDBS must provide

transaction mechanisms at the federation level. These mechanisms should contain at least open nested transactions to execute parts of global actions in subtransactions on CDBSs. Furthermore, in order to realize detached coupling modes transactions are needed which are able to create new top-level transactions.

5 Integrating Active Mechanisms into a Federated Database System

According to the federated database system architecture presented in Sect. 2, we aim at providing a global mechanism to guarantee the integrity of federated data. In this section, we propose an active federated database framework to provide a powerful and flexible platform for global integrity constraint enforcement in (tightly coupled) federated database systems. Thereby, we pursue the approach to integrate active mechanisms (ECA-rules) into a federated database system for supporting global constraint management. Besides, a main idea is to include existing integrity mechanisms of the CDBSs (see also [CT95]).

We use ECA-rules for specifying events, which potentially violate the integrity of the database, and (re)actions on these events, which should restore a consistent database state by aborting the operation/transaction (which is causing the event) or starting compensating (trans)actions. Generally, global integrity constraints and (compensating) actions in case of violation can be formulated as ECA-rules in the following way:

define rule Global_Integrity_Rule
on event which potentially violates the integrity of the FDB
if a global condition formulated over the state of the FDB is true
do global action(s) which restore(s) a consistent state of the FDB

As we have seen in Sect. 2, federated database systems consists of two levels: the component system level and the federation level. The integration of active mechanisms into such an environment can be done in the following way (cf. Fig. 2). At the federation level, global active mechanisms are offered in form of *global ECA-rules*. These ECA-rules are stored in the global database which also contains all data of the federation layer (e.g. metadata for schema integration). According to the event part, every global ECA-rule is attached to a *global rule manager* which is responsible for the management and processing of this ECA-rule (cf. [BZBW95]). These active components are integrated into the federation layer similar to the REACH approach.

By using global ECA-rules we are able to specify (re)actions on situations which are caused by either global applications/users or the (global) system clock (e.g. absolute time events). For global integrity enforcement global ECA-rules alone are not sufficient. Indeed, we may specify global ECA-rules in order to avoid that global applications violate the integrity of the federated database. But, global ECA-rules cannot prevent local applications from performing operations which may have an effect on the global integrity of the federated database.

Therefore, we introduce the notion of *local ECA-rules* and *local rule managers*, respectively. Local ECA-rules are used to specify (re)actions which are performed upon the occurrence of a local event. Like global ECA-rules, local ECA-rules are managed and

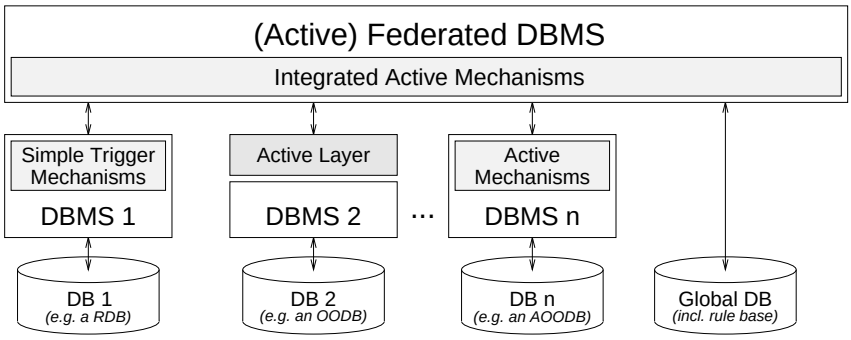

Figure 2: Active federated database framework

processed by a (local) rule manager. A local rule manager detects a local event in the CDBS and informs the appropriate global rule manager about the event occurrence. In contrast to global ECA-rules, local ECA-rules (as the name suggests) are stored locally at the component system level. In our approach, the only purpose of these local ECA-rules is to signal the occurrence of a specific local event to the federation layer. Hence, the format of local E(C)A-rules is as follows:

define rule Local_Rule
on local event
do inform the corresponding global rule manager

A general question is how can we adequately represent local ECA-rules as well as their local rule managers inside an FDBS? Obviously, for that we need some active mechanisms at the CDBS level. In principle, there are two possible scenarios which we discuss in the following:

1. *Using existing active capabilities of the CDBS.* If the underlying CDBS supports some active mechanisms (e.g. triggers), we might use them to signal the occurrence of a local event to the corresponding global rule manager.

   Today, most of the commercial relational database systems, for instance SYBASE or ORACLE, support some kinds of trigger mechanisms which are already very powerful. Triggers, which can be seen as a special kind of active mechanisms, can be used to signal the request of a SQL operation to the federation layer. The term trigger is often used in the context of relational database systems where triggers may be defined on database operations like insert, update, delete, or select as follows:

   create trigger Local_Rule_For_UpdateXXX
   update of <attribute-name>
   on <table-name>
   when <predicate>
   <operations-to-propagate-event>

   In contrast, if the underlying CDBS is an active one which supports ECA-rules the case is trivial. Here, the existing local ECA-rule mechanism can be used as described above.

2. *Providing an additional active layer with active mechanisms.* In case a CDBS does not support any active mechanisms, we have to provide an active layer á la SAMOS for that CDBS (see DBMS2 in Fig. 2). So, we are able to define and handle local ECA-rules in arbitrary CDBSs.

   In general, a technique which is called method wrapping (i.e. putting additional code into an application program) can be applied to equip a database system with active capabilities. Method wrapping is successfully used in active object-oriented database systems (see for instance [GD94b, BZBW95]). However, we have to notice that method wrapping implies the recompilation of existing, *not with active capabilities equipped* application programs.

   In case of relational database systems we may use a similar technique. When a SQL-operation (e.g. an update) is sent to the system it has to pass the active layer. Here, we are able to detect events before and after, respectively, passing the operation to the original (passive) system.

Another very important aspect concerns the derivation of local ECA-rules from global ECA-rules. A global ECA-rule may base on a composite event which consists of several local events (in different databases), for example

define rule RuleX
on DB1::e1(p1) **or** DB2::e2(p2)
if ...
do ...

Such a global ECA-rule has to be subdivided into several local ECA-rules, i.e. in the example above two local ECA-rules have to be generated from the rule RuleX.

define rule Local_Integrity_RuleX1	**define rule** Local_Integrity_RuleX2
on DB1::e1(p1)	**on** DB2::e2(p2)
do GDB::notify(RuleX, DB1::e1(p1))	**do** GDB::notify(RuleX, DB2::e2(p2))

Because these local ECA-rules are only used to trigger the corresponding global ECA-rule (here: RuleX), they have no condition part. Thus, such ECA-rules are actually EA-rules which contain only a notification operation in the action part. Therefore, the mapping process from global to local ECA-rules primarily consists of the subdivision of a composite event in several primitive events, and can be executed completely automatically (by using the metadata of the federated database).

In conclusion, the combination of global and local ECA-rules allows to realize global integrity constraints which must be satisfied by global as well as local applications. These rules are managed and processed by corresponding rule managers. On the one hand, there are the local rule managers which are responsible for detecting local events in the CDBSs, and notifying the responsible global rule manager about the occurrence of that event. On the other hand, the global rule managers are responsible for

- detecting global events which refer to either operations on the global database or absolute time events,
- collecting all incoming messages (from the local rule managers) and, thus, detecting global composite events,

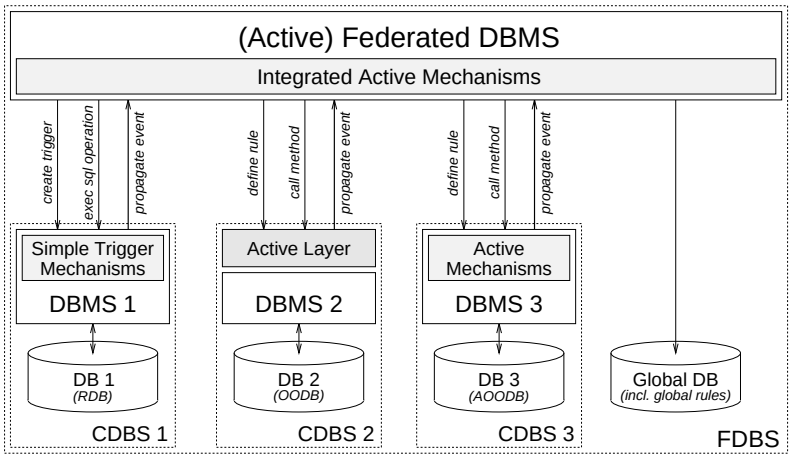

Figure 3: Rule definition and processing in an active FDBS

- evaluating the condition after the occurrence of an event, and
- starting the execution of action(s) when the condition is satisfied.

6 Examples for Active Global Integrity Maintenance in an FDBS

In this section, we exemplarily show how global integrity maintenance in an active federated database environment can work. Especially, we present examples how global integrity constraints such as global key constraints can be specified in the proposed active federated database framework.

For that, we assume that we have the following application scenario: An enterprise wants to couple all (or parts of) the database systems of its departments to a global federated database system. In general, the resulting FBDS has multiple heterogeneous CDBSs as depicted in Fig. 3. Let CDBS1 be a relational database system which supports simple trigger mechanisms, CDBS2 an object-oriented database system with no active capabilities, and CDBS3 an AOODBS. Further, let us assume that there exist several heterogeneous databases in the different database systems:

- a relational database DB1 which contains several relations (i.e. tables), e.g. a relation Employee where all data of the employees of the enterprise are recorded;
- an object-oriented database DB2 which consists of several classes, e.g. a class Project in which all data concerning the projects of that department are recorded;
- an (active) object-oriented database DB3 which belongs to another department and contains similar classes as DB2.

Further, let us assume that all these local database schemata are integrated into a single federated schema, on which we can define global integrity constraints.

We now investigate how integrity constraints can be represented in the presented environment. Here, we have chosen a set of different kinds of integrity constraints to

elucidate how active integrity maintenance can work in a federated database environment (see also Fig. 3).

1. Attribute constraints. Attribute constraints restrict the values an attribute can have. Such constraints are represented, for instance, in relational database systems by the well-known **check** or **not null** clauses.

In our sample scenario, we may have a global attribute constraint saying that no project in the enterprise is allowed to exceed a given budget. Such kinds of constraints can easily be enforced by existing local integrity mechanisms, e.g. by an additional **check** clause, if schema modification is allowed in the CDBSs. Otherwise, we may realize it by a global ECA-rules as proposed in the previous section.

2. Key constraints. In relational database systems, key constraints are used to model the uniqueness of a real-world entity. Although in object-oriented database systems objects are identified by an internal identifier mechanisms, in general additional key attributes are explicitly used to make it possible for the user to distinguish between different objects.

Key constraints are very important for FDBSs, too. Each object of an FDBS has to be identifiable not only in the CDBS but also in the whole FDBS. In order to guarantee the global uniqueness of an FDBS object, global key constraints may be defined, e.g. each project in the enterprise is identified by a unique project number. In order to realize such a constraint adequately, we can use active mechanisms which are represented in the active federated framework as follows: For each CDBS which is affected by the integrity constraint a local rule manager for supervising manipulations on the corresponding attribute(s) is created. Moreover, at the federation level there is a global rule manager which is triggered by these local rule managers when such a manipulation operation is requested locally. In this case, the global rule manager has to check whether this operation violates the global integrity. In case of a violation the local operation has to be rejected. For that, a global ECA-rule can look like as follows:

define rule GlobalKeyConstraintRule
on manipulation of key attribute(s) **or**creation of a new object
if global key constraint is violated
do abort the corresponding transaction

3. Aggregate constraints. Such constraints concern generally several objects, e.g. the sum of the budgets of all projects in the enterprise must not exceed a certain amount. Aggregate constraints can be realized similar to the key constraints:

define rule GlobalAggregateConstraintRule
on budget manipulation operations
if sum of budgets exceeds a given amount
do compensate the corresponding transaction

4. Referential integrity constraints. Referential integrity constraints specify that a database object refers to existing objects. In our sample scenario, we may have the constraint that each participant of a project must be an employee of the enterprise at the same time. The following ECA-rule, for example, guarantees that only employees can be involved in the projects of the enterprise.

define rule GlobalReferentialIntegrityRule
on deletion of an employee
if this employee is involved in a project
do remove this participant from the project

5. Further constraints. By using ECA-rules we are able to realize arbitrary constraints. Especially, we can specify temporal constraints, e.g. in order to model admissible traces of evolution of the federated database. For instance, in our scenario we may specify that the enterprise forbids to start more than five new projects per year.

define rule GlobalTemporalConstraintRule
on start of new project
if already five projects have been started in the same year
do abort transaction

7 Conclusions

Our proposal for integrity enforcement in a federated, heterogeneous framework is based on the idea to bring active mechanisms into federated database systems. We have shown in which way active rules can be applied in FDBSs for global constraint maintenance. Our approach respects the local autonomy of the component database systems as much as possible, especially by using integrity mechanisms offered by them. Obviously, preserving local autonomy and enforcing global constraints represent two contrary requirements.

We emphasize that as foundation for global integrity enforcement, an FDBS must support an appropriate transaction model, for instance an open nested transaction model, which allows the co-existence of local and global (federated) transactions. Another important issue is that the architecture of the FDBS has to be open so that it is possible to add/remove CDBSs into/from an FDBS or to include new transaction managers which realize more sophisticated transaction paradigms.

Beside the development of a prototype system for validating our approach in practice, there are several other questions and problems we are interested in. One problem is the derivation of local ECA-rules from global ECA-rules. Obviously, local events can be composite events themselves. Therefore, the derivation of local ECA-rules from global ones must be done very carefully.

Another very important question is related to the design of global integrity constraints. The definition of global integrity constraints has to be done very carefully. The rule designer must be informed about the existing local integrity constraints in order to avoid unnecessarily redundantly defined constraints, and thus avoiding multiple checks of one and the same constraint. Further, the rule designer has to consider the collaboration of global and local integrity constraints in order to avoid contradictory actions.

These are only two of several important problems we are investigating. For this we have just started a project in which we will put our ideas of global integrity maintenance for federated database systems into practice. This interdisciplinary project deals with the problem of cooperative, distributed planning of factories. Global constraint management and integrity maintenance are absolutely necessary for this application area in order to adequately manage the data which is produced by several designers (engineers) at the

same time and which is highly interrelated. For a prototype implementation of our active federated database system framework we are cooperating with Cadlab Paderborn (a joint venture of University of Paderborn and Siemens-Nixdorf). Cadlab is currently developing a middleware software called OpenDM (partly described in [RS94]) which provides all necessary functionalities for federating heterogeneous database systems.

References

[BBKZ94] H. Branding, A. Buchmann, T. Kudrass, and J. Zimmermann. Rules in an Open System: The REACH Rule System. In N. W. Paton and M. H. Williams, editors, *Proc. of the 1st Int. Workshop on Rules in Database Systems (RIDS'93), Edinburgh, Scotland, August 1993*, pages 111–126. Workshops in Computing, Springer-Verlag, 1994.

[BFN94] R. Busse, P. Fankhauser, and E. J. Neuhold. Federated Schemata in ODMG. In J. Eder and L. A. Kalinichenko, editors, *Extending Information Systems Technology – Proc. of the 2nd Int. East/West Database Workshop, Klagenfurt, Austria*, pages 356–379. Workshops in Computing, Springer-Verlag, September 1994.

[BHP92] M. Bright, A. Hurson, and S. Pakzad. A Taxonomy and Current Issues in Multidatabase Systems. *IEEE Computer*, 25(3):50–56, March 1992.

[Buc94] A. P. Buchmann. Active Object Systems. In A. Dogac, M. T. Öszu, A. Biliris, and T. Selis, editors, *Advances on Object-Oriented Database Systems*, pages 201–224. Springer-Verlag, 1994.

[BZBW95] A. P. Buchmann, J. Zimmermann, J. A. Blakeley, and D. L. Wells. Building an Integrated Active OODBMS: Requirements, Architecture, and Design Decisions. In P. S. Yu and A. L. P. Chen, editors, *Proc. of the 11th IEEE Int. Conf. on Data Engineering (ICDE'95), Taipei, Taiwan*, pages 117–128. IEEE Computer Society Press, March 1995.

[CBB+89] S. Chakravathy, B. Blaustein, A. Buchmann, M. Carey, U. Dayal, D. Goldhirsch, M. Hsu, R. Jauhari, R. Ladin, M. Livny, D. McCarthy, R. McKee, and A. Rosenthal. HiPAC: A Research Project in Active, Time-Contrained Database Management. Technical Report XAIT-89-02, Xerox Advanced Information Technology, 1989.

[CKTB95] S. Chakravarthy, V. Krishnaprasad, Z. Tamizuddin, and R. H. Badani. ECA Rule Integration into an OODBMS: Architecture and Implementation. In P. S. Yu and A. L. P. Chen, editors, *Proc. of the 11th IEEE Int. Conf. on Data Engineering (ICDE'95), Taipei, Taiwan*, pages 341–348. IEEE Computer Society Press, March 1995.

[CM93] S. Chakravarthy and D. Mishra. Snoop: An Expressive Event Specification Language for Active Databases. Technical Report UF-CIS-TR-93-007, Computer Science Department, University of Florida, March 1993.

[CT95] S. Conrad and C. Türker. Active Integrity Maintenance in Federated Database Systems. Preprint Nr. 9, Fakultät für Informatik, Universität Magdeburg, November 1995.

[DHW95] U. Dayal, E. Hanson, and J. Widom. Active Database Systems. In W. Kim, editor, *Modern Database Systems*, pages 434–456. ACM Press, 1995.

[GD94a] S. Gatziu and K. R. Dittrich. Detecting Composite Events in Active Database Systems Using Petri Nets. In J. Widom and S. Chakravarthy, editors, *Proc. of the 4th Int. Workshop on Research Issues in Data Engineering: Active Database Systems (RIDE-ADS'94), Houston, Texas, USA*, pages 2–9. IEEE Computer Society Press, February 1994.

[GD94b] S. Gatziu and K. R. Dittrich. Events in an Active Object-Oriented Database System. In N. W. Paton and M. H. Williams, editors, *Proc. of the 1st Int. Workshop on Rules in Database Systems (RIDS'93), Edinburgh, Scotland, August 1993*, pages 23–39. Workshops in Computing, Springer-Verlag, 1994.

[GJS92a] N. Gehani, H. V. Jagadish, and O. Shmueli. Event Specification in an Active Object-Oriented Database. In M. Stonebraker, editor, *Proc. of the 1992 ACM SIGMOD Int. Conf. on Management of Data, San Diego, CA, SIGMOD RECORD 21(2)*, pages 81–90. ACM Press, June 1992.

[GJS92b] N. H. Gehani, H. V. Jagadish, and O. Shmueli. Composite Event Specification in Active Databases. In L.-Y. Yuan, editor, *Proc. of the 18th Int. Conf. on Very Large Data Bases (VLDB'92), Vancouver, Canada*, pages 265–276. Morgan Kaufmann Publishers, August 1992.

[LMR90] W. Litwin, L. Mark, and N. Roussopoulos. Interoperability of Multiple Autonomous Databases. *ACM Computing Surveys*, 22(3):267–293, September 1990.

[RS94] E. Radeke and M. H. Scholl. Federation and Stepwise Reduction of Database Systems. In W. Litwin and T. Risch, editors, *Proc. of the 1st Int. Conf. on Applications of Databases (ADB'94), Vadstena, Sweden*, pages 381–399. LNCS 819, Springer-Verlag, June 1994.

[SL90] A. P. Sheth and J. A. Larson. Federated Database Systems for Managing Distributed, Heterogeneous, and Autonomous Databases. *ACM Computing Surveys*, 22(3):183–236, September 1990.