

Data-driven Active Rules for the Maintenance of Derived Data and Integrity Constraints in Interactive Database Applications

Abstract

An active rule or trigger is an event-condition-action rule with the following meaning: if the event occurs, evaluate the condition and, in case it evaluates to true, execute the action. Active rules have appeared in active database systems, where they are used for the enforcement of static or dynamic integrity constraints, the automatic calculation of derived data, the enforcement of business rules, etc. In some systems, it is possible to specify absolute or relative priorities among rules that condition the order by which multiple triggered rules are executed.

An interactive database application is a computer program that allows a user to query and manipulate data stored in a database through a user interface based on forms, reports, graphics, etc. In rapid application development tools, triggers with more limited capabilities are also used for data validation, for the automatic calculation of derived data, for the dynamic control of properties of data items, to extend or override automatic behavior embedded in those tools, etc.

A data-driven active rule is an active rule with implicit triggering events of a restricted type - data modification events - which can be inferred from the condition and/or the action of a rule, according to some assumptions. Data-driven active rules differ from event-driven active rules (i.e., active rules with explicit events) in the way they are defined, but not in the way they are executed; the execution of data-driven active rules is also event-driven for efficiency and integration reasons. Formulae for the calculation of derived data in spreadsheets or in rapid application development tools and generic assertions in SQL are examples of entities that may be treated as data-driven active rules.

The main contribution of the work reported in this thesis is the proposal of a model of data-driven active rules specially suited for the maintenance of derived data (calculated data) and integrity constraints (data validation) in interactive database applications (namely screen forms and reports), more flexible and with better integration than other types of data-driven active rules previously proposed. Two important characteristics of the proposed model are the fixpoint semantics and the sequential execution. The fixpoint semantics means that a rule should be executed (i.e., the condition should be evaluated and, in case it evaluates to true, the action should be executed) whenever a change in the state of the data may result from its execution. Sequential execution is important to guarantee that the changes caused by the execution of a rule are not affected by the existence of other rules. Inference rules are established to obtain the triggering events and priorities of data-driven active rules, for their efficient and safe execution, possibly in combination with other types of active rules. Conditions on sets of rules are determined that guarantee termination and determinism of rule execution. The specialized nature of the proposed rules allows the determination of less conservative conditions than the ones that are known for generic active rules. It is shown how to incorporate into the proposed model the optimizations required to deal efficiently with complex data, by means of the differentiation of set oriented rules and the encapsulation of rules within objects.

Finally, it is described a concrete implementation of an active rule system integrated in a rapid application development tool developed at INESC, supporting both data-driven and event-driven rules, demonstrating the advantages of the proposed approach by comparison to other tools, in spite of some compromises taken in the implementation.