

An Approach for Active Deductive Database Management

Thodeti Srikanth¹

¹Research Scholar, Department of Computer Science, Dravidian University, A.P., India

Abstract—After briefly reviewing the simple notions and terminology of effective rules and relating them to creation rules as well as deductive rules, respectively, we survey a selection of formal approaches to effective rules. Subsequently, we present our very own state-oriented rational approach to energetic rules which fuses the declarative semantics of deductive rules with the chance to done revisions in the color of creation rules and active rules. The primary issue is actually that timeliness, i.e., predictability as well as efficiency is actually of supreme importance in realtime databases while reactive conduct might contribute to the unpredictability of the database. This survey addresses reactive behavior as well as timing restrictions from a database perspective

Index Terms—(DL)

I. INTRODUCTION

As the quantity of data in a method grows larger, it has to be saved someplace and saved intelligently, with regard to room requirements as well as time for retrieval. A database serves as being a repository. Database technology has been extensively examined with regards to query languages, database versions, storage strategies, buffer managing, concurrency management, multi-user access and update, etcetera. However, several apps which feature a database method function in a real-time atmosphere, and conventional databases haven't been created with genuine time constraints in mind, particularly predictability. At most, efficiency is actually considered but that's not positive. Moreover traditional databases aren't capable of taking actions on their very own in phrases of consistency maintenance view revisions. Since the first of computer systems there have been methods which have operated in realtime environments present examples of that technology areas are actually chemical processes nuclear power plants as well as space shuttle management. The common denominator for these methods is actually that all of them respond to inputs from the planet and should generate a little output within a specified time. The information connected with and utilized by such

programs has grown steadily inserting a database into such an atmosphere can't be accomplished without concern for keeping timing constraints. In the database activities should be efficient as well as predictable so the website doesn't stop the user from conference deadlines. Suggestions have been created on how to guarantee the timeliness of databases several of which will be discussed in the department on realtime databases.

In a standard database system nothing takes place unless a database activity is explicitly requested. In simple systems that could be good but in larger methods considerable effort needs to be expended on the improvement of applications with the database in order that they preserve data source consistency. Furthermore materialized opinions must be updated whenever they're affected by repository revisions. This may be done incrementally by the database supervisor. The phrase busy databases was first utilized to denote such databases incrementally updating views after energetic databases have come to denote much much more than that for instance triggering facilities constraint checking and so on. Today we do need an energetic website as a database being in a position to respond to events as well as conditions in the database or perhaps perhaps the surrounding programs and environment.

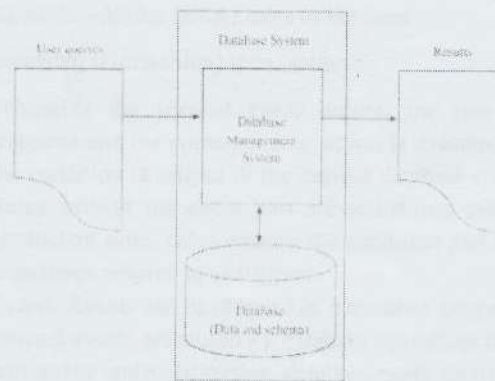


Figure 1: Traditional database

Within traditional databases there has always been a requirement that the data in the database be consistent with the environment. Moreover, enforcing constraints on data has also been desirable, e.g. a certain value must not grow above a given limit. It has also been beneficial to be able to activate alarms on detecting certain events in the database.

II. RELATED WORKS

Traditionally all it has been carried out in 2 methods on passive databases [Cha89]:

1. The database is often polled to determine if something of value has happened.
2. All constraints as well as consistency checks are actually applied in the apps and not in the database management process.

The former strategy has the disadvantage of extremely big overhead if the polling is actually done way too frequently, as the data source has to process queries that frequently return nothing. However, if the polling is done too occasionally, several events could be missed. The latter strategy is better from an efficiency use of perspective however letting the data source management process conduct all checks rather than every software has the following advantages Each and every update of a constraint just has to be accomplished centrally in the database supervisor and not in every application. Since constraints are actually defined centrally they're the same for all apps. No two applications may have inconsistent constraints. A new program doesn't have to be programmed with all inspections before startup, since all details pretty much is kept in a centrally maintained repository supervisor. We will find essentially 2 methods to specify situations to trigger rules, pattern based and event based [HW91]. In a pattern based method a rule is actually fired according to a problem or perhaps a predicate, and an event will need not be explicitly stated. A good example of such a system is actually PRS-I [Sto86], which uses the keyword 'always'. The semantics of such a clause is the fact that a website query is logically consistently executing. A good example of the rule is actually range of E is EMP replace Always EMP (salary = E.salary) where EMP.name = "Mike" and E.name = "Bill" in which many updates on Bill's salary in table EMP instantly triggers an action to replace Mike's salary to be the just like Bill's. Although simple to understand,

this particular technique suffered from a few downsides [SHP89]. It's tough to carry out and it's difficult to figure out on which database actions to perform the query as there may be multiple actions triggering the law some even not appealing or perhaps instance if a principle says that a worker must be removed if no department is actually specified for the worker also inserts of new employees are actually refused if no department is actually specified defining guidelines for inserts & delete independently are therefore not likely furthermore the semantics of the execution isn't clear that's it's not clear which coupling function is used see action or perhaps whether a rule evaluation is actually excited or perhaps idle ie choosing the example above the principle is actually evaluated on update of Mike's salary or perhaps on each and every request for Mike's salary finally there's no chance of defining timing constraints to overcome several of these shortcomings occasion based rules had been released one realization being such guidelines an explicit occasion as opposed to a pattern or perhaps predicate triggers a rule or perhaps example the rule above may be rewritten as
 range of E is EMP
 on replace to E.salary
 where EMP.name = "Mike"
 do replace EMP (salary = new.salary)
 and E.name = "Bill"

Here the replacement is done only on updates of Mike's salary. Moreover, it is possible, although not shown here, to compare and use both the old and the new salary in the rule evaluation.

III. THE PROPOSED APPROACHES

Active rules are typically expressed as Event-Condition-Action (ECA) rules of the form

on(event) if(condition) then(action).

Whenever the specified event occurs, the rule is triggered and the corresponding action is executed if the condition is satisfied in the current database state. Rules without the event part are sometimes called production rules, rules without the condition part are sometimes referred to as triggers.

Events. Events can be classified as internal or external. Internal events are caused by database operations like retrieval or update (insertion, deletion, modification) of tuples, or transactional events like commit or abort. In

object-oriented systems, such internal events may take place through method invocations. External events occurring outside the database system may also be declared and have to be monitored by the ADB.

A question arising from the use of composite events is which of the constituent events "take part" in the composite event and how they are "consumed" by the composite event. This event consumption policy is elaborated using parameter contexts, which were introduced for the SNOOP algebra in [CKAK94, CM94]. In order to illustrate the different parameter contexts, consider the composite event $E := ((F; G); H)$, which occurs if H occurs after both F and G have occurred. Assume the following event history is given:

Different parameter contexts are motivated by applications where constituent events should be consumed by the composite event in a certain way. The following parameter contexts have been proposed [CKAK94]:

Recent: In this context, only the most recent occurrences of constituent events are used; all previous occurrences are lost. In the above event history, E will be raised twice: for the constituent events $\{G2; F3; H1\}$ and for $\{G2; F3; H2\}$.

Chronicle: In this context, events are consumed in their chronological order. In a sense, this corresponds to a first-in-first-out strategy: E will be reported for $\{G1; F1; H1\}$ and for $\{F2; G2; H2\}$.

Continuous: In this context, each event which may initiate a composite event starts a new instance of the composite event. A constituent event can be shared by several simultaneous occurrences of the composite event. In the example, each G_i and each F_j starts a new instance. Thus, the composite occurrences $\{G1; F1; H1\}$, $\{F1; G2; H1\}$, $\{F2; G2; H1\}$, and $\{G2; F3; H1\}$ are reported. The composite event initiated at $F3$ is still to be completed.

Cumulative: In this context, all occurrences of constituent events are accumulated until (and consumed when) the composite event is detected. In the example, E is raised once for the constituent events $\{G1; F1; F2; G2; F3; H1\}$.

Unrestricted: In this context, constituent events are not consumed, but are reused arbitrarily many times. For the above history, E is reported for all twelve possible combinations of $\{F_j; G_k; H_l\}$.

If the condition of the triggered rule is satisfied, the action is executed. Internal actions are database updates (insert, delete, modify) and transactional commands (commit, abort), external actions are executed by procedure calls to application programs and can cause application-specific actions outside the database system (e.g., send-mail, turn-on-sensor). Usually, it is necessary to pass parameters between the different parts of ECA-rules, i.e., from the triggering event to the condition and to the action part. In logic-based approaches this can be modeled very naturally using logical variables, while this issue may be more involved under the intricacies of certain execution models.

Whereas the meaning of deductive rules is based on solid logical foundations, the meaning of the more low-level and operationally intricate active rules is often hard to understand and predict (especially, if the semantics is only given informally). This has led to numerous research towards formal foundations of active rules. In the sequel, we discuss some of these approaches; due to lack of space and the focus on logic-based approaches, we can only provide a rough and necessarily incomplete summary.

Although there is a great variety of execution models for active rules, certain fundamental properties like termination and complexity come up repeatedly and have been studied in the context of the respective execution models:

Whereas the above-mentioned works focus on analysing rule properties in some specific execution model, a lot of research aims at formalizing and characterizing the semantics of active rules in the first place. Once a formal model has been established, abstract properties like termination or expressiveness can be studied. By extending Datalog with a notion of state, (re)active production rules and deductive rules can be handled in a unified way, thereby combining the advantages of active and deductive rules.

IV. CONCLUSIONS

Active rules extend the conventional passive database technology and are actually a potent programming paradigm with a lot of program areas. While an increasing number of methods becomes active and available rule programming is carried out in real life applications, the theoretical foundations of energetic rules continue to be uncommon. In

probably the first portion of the paper, we've introduced the basic principles of active rules and associated them to deductive rules and production rules, respectively. After talking about a selection of formal techniques to energetic rules, we've elaborated on a state oriented logical framework which includes deductive and active rules.

REFERENCES

- [1] [ABW88] K. R. Apt, H. Blair, and A. Walker. Towards a Theory of Declarative Knowledge. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pp. 89-148. Morgan Kaufmann, 1988.
- [2] [AHV95] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison Wesley, 1995.
- [3] [AS91] S. Abiteboul and E. Simon. Fundamental Properties of Deterministic and Nondeterministic Extensions of Datalog. *Theoretical Computer Science*, 78(1):137-158, 1991.
- [4] [AV91] S. Abiteboul and V. Vianu. Datalog Extensions for Database Queries and Updates. *Journal of Computer and System Sciences*, 43(1):62-124, 1991.
- [5] [AWH95] A. Aiken, J. Widom, and J. M. Hellerstein. Static Analysis Techniques for Predicting the Behavior of Active Database Rules. *ACM Transactions on Database Systems (TODS)*, 20(1):3-41, March 1995.
- [6] [BCP95] E. Baralis, S. Ceri, and S. Paraboschi. Improving Rule Analysis by Means of Triggering and Activation Graphs. In Sellis [Sel95], pp. 165-181.
- [7] [BFKM85] L. Brownston, R. Farrel, E. Kant, and N. Martin. *Programming Expert Systems in OPS5: An Introduction to Rule-Based Programming*. Addison Wesley, 1985.
- [8] [BFP+ 95] M. L. Barja, A. A. A. Fernandes, N. W. Paton, M. H. Williams, A. Dinn, and A. I. Abdelmoty. Design and implementation of ROCK & ROLL: a deductive object-oriented database system. *Information Systems*, 20(3):185-211, 1995.
- [9] [BGP97] C. Baral, M. Gelfond, and A. Proveti. Representing Actions: Laws, Observations and Hypotheses. *Journal of Logic Programming*, 31(1(3)):201-243, 1997.
- [10] [BH95] M. Berndtsson and J. Hansson, editors. 1st Intl. Workshop on Active and Real-Time Database Systems (ARTDB), *Workshops in Computing*, Skövde, 1995. Springer.
- [11] [BK94] A. J. Bonner and M. Kifer. An Overview of Transaction Logic. *Theoretical Computer Science*, 133(2):205-265, 1994.
- [12] [BL96] C. Baral and J. Lobo. Formal Characterization of Active Databases. In Pedreschi and Zaniolo [PZ96], pp. 175-195.

Author:



Thodeti Srikanth, Research Scholar, Department of Computer Science in Dravidian University and working as HoD, Department of Computer Science in N.S.V. Degree College, Jagityal, Telangana, India. I received Master of Computer Applications (M.C.A.) from Kakatiya University in April 2004.