

Expanding Database Systems into self-verifying Entities

Kåre J. Kristoffersen and Yvonne Dittrich

IT University of Copenhagen*
DK-2300 S, Denmark

Abstract. The paper presents *work-in-progress* aiming at deploying runtime verification techniques to check whether the state changes in a database system conform with temporal business rules given as expressions in temporal logic. A framework for tailoring enterprise database systems with temporal business rules is defined and an algorithmic framework for checking temporal business rules at runtime is presented. A prototypical implementation of a runtime verifier (called Verification Server) based on this algorithmic framework is presented and discussed.

1 Introduction

Runtime verification is a branch of verification in which a running program is supervised by a concurrently running *verifier*. In this paper we shall employ *timed runtime verification*, in which time will be an important parameter in the task of the verifier. Our idea is to use such a mechanism to monitor a running database system and hereby at runtime check whether a sequence of states of a traditional database obey a set of temporal business rules. In [5] an interesting framework for proving temporal properties of a database *prior* to execution is presented. That task is much harder and the proof of correctness cannot be automated. In our approach, the proof (or dis-proof) is established along with the execution of the database.

Temporal Business Rules and corresponding mechanism to check their success or failure might be hard coded in business systems. Many such systems exists, like e.g. in a library where customers (automatically) get a reclaim of their borrowed material after one month. However, this approach restricts the flexibility to re-define temporal rules at all or at least to anticipated areas. Using runtime verification techniques provides the possibility to formulate and change general temporal business rules and to check them without changing the business application.

Section 2 introduces the temporal logic we use as basis for temporal business rules and the algorithmic base we developed for the checking these at runtime. Section 3 presents the design of an early prototype of the Verification Server. Finally, in Section 4, we sum up our findings so far and present our future line of research on the subject.

* This work is supported by the project NEXT which is a joint effort between Microsoft Business Solutions and the IT University of Copenhagen.

2 Temporal Business Rules

A state of a database is a pair (s, t) where s is a discrete component representing a snapshot of all the data (or more precisely those data that are *relevant* for the temporal business rules) and t is a time stamp. Using s , boolean constraints may be decided, such as if the balance of a bank account is below a certain threshold. A temporal business rule for a database is a specification on how the internal *state* of the database may evolve over time, and what should happen when a rule is satisfied (or the opposite) by the data in the database: IF TC THEN Action₁ ELSE Action₂, where TC is a temporal condition and Actions 1 and 2 are some action to be performed. The core of our temporal conditions are given by *Timed LTL*, LTL_t , in the following abstract syntax (see also [1]):

$$\begin{aligned} \phi ::= & p \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2 \mid \neg\phi \mid \text{ALWAYS } \phi \mid \text{ALWAYS}_c \phi \\ & \mid \text{EVENTUALLY } \phi \mid \text{EVENTUALLY}_c \phi, \end{aligned}$$

where $p \in AP$ and $c \in \mathbb{N}$.

The syntactic elements are: Atomic propositions, AP, which can be the occurrence of insert, update, delete or that an attribute in the database is above/below a certain threshold. Further, logical connectives and then temporal operators ALWAYS and EVENTUALLY both of which may be equipped with a time bound c . Intuitively they mean the following: Where ALWAYS ϕ denotes that the formula ϕ must hold in all timepoints, ALWAYS _{c} ϕ only requires ϕ to hold in the coming c time units. Conversely, the formula EVENTUALLY _{c} ϕ requires that formula ϕ is satisfied *before* c time units have passed, and thus it is a more restrictive operator than EVENTUALLY which only requires the sub-formula to hold at some point arbitrarily long away in the future. We shall use the standard abbreviation such as expressing implication etc. as logical connectives, e.g. $\phi_1 \Rightarrow \phi_2$ for $\neg\phi_1 \vee \phi_2$ and to use *true* instead of $\neg p \vee p$. Using this language we may formulate a temporal business rule for a vehicle assistance company in which the temporal condition is the following, stating that assistances must not be given too frequent:

$$\text{ALWAYS } (\text{new}(C.\text{Assistance}) \Rightarrow \text{ALWAYS}_{30} \neg \text{new}(C.\text{Assistance})).$$

Checking Temporal Business Rules The algorithm in the verification server works by a rewriting principle. For each new state encountered from the database, the algorithm *rewrites* the temporal constraints to a new formula which states what should hold from now on. Such a new formula is called a residual formula. In the algorithm below the residual of formula ϕ with respect to a timed state σ_i is denoted by ϕ/σ_i .

Algorithm: Runtime Verification procedure

Let $\sigma = \sigma_0\sigma_1 \dots$ be a timed trace, let ϕ be a formula, let $\text{exists}(\sigma_i, \text{time})$ be a predicate which is true exactly when σ contains σ_i and $t_i = \text{time}$ and let $\text{forceEvaluation}()$ be a method which returns the systems current state, at the current time.

```

 $\psi := \phi/\sigma_0; i := 1; \text{time} := 1; \text{sit} := \text{SIT}(\psi)$ 
while  $\psi \neq \text{true} \wedge \psi \neq \text{false}$  do
  if  $\text{exists}(\sigma_i, \text{time}) \wedge s_i \neq s_{i-1}$  then
     $\psi := \psi/\sigma_i$ 
     $\text{sit} := \text{SIT}(\psi)$ 

```

```

    i ++
  end if
  if time = sit then
     $\sigma_i := forceEvaluation()$ 
     $\psi := \psi / \sigma_i$ 
    sit := SIT( $\psi$ )
    i ++
  end if
  time ++
end while
return  $\psi$ 

```

Occasionally, the satisfaction (or falsification) of a formula is not triggered by a change of state in the database, but rather by the elapse of time alone. To ensure that the algorithm is timely complete we let the verification server compute the coming smallest interesting timepoint (SIT) and insert an artificial state at such a time point¹. See [4] for a complete description of the algorithm. The rewrite principle is a timed extension of the one in [3].

3 The prototypes

Based on the notation and algorithm introduced in the last section a second prototype of the verification server was implemented. The verification server is designed as an independent program running parallel to the business systems of the company. A first prototype parsed a rule in LTL notation and verified trace of state changes to a simple (one-table) database complied with the rule. Recently the first prototype was extended to handle more than one rule with a more general way to interact with the database of the business system. For each rule, an independent thread is started that takes care of the verification of that rule. This also allows the different rules to be based on different time units. In the momentary version, the verification server queries the database for relevant state changes once each time unit specified as the basic time unit of the given rule. In parallel a prototype for a rule builder is developed. We explore different alternatives for the user interface to find a notation for the temporal business rules and a way to construct them that both leads to correctly formulated expressions in LTL and allows for a meaningful interpretation from a business point of view.

The system so far looks as depicted in figure 1. When the business rule expert wants to insert a new rule, he opens the *rule builder* (1)². The system provides him with an overview of the attributes he can base his rules on by querying the database for its data model (2). He selects a table. When he is satisfied and presses the save button the rule builder constructs an XML file (3) containing the LTL representation of the rule, the time unit the rule is based on and sql statements to access the data defining the states the rule is observing. He then has to define what should happen in case of the violation

¹ The discrete component of the artificial state is the one appearing in the most recent state emitted by the database.

² The numbers refer to the numbers in figure 1.

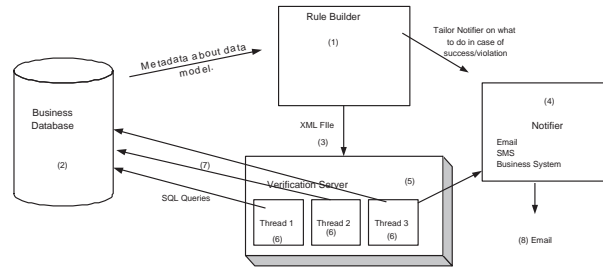


Fig. 1. The system architecture.

of this rule in the administrator interface of the *notifier*. (4) He decides that an e-mail should be sent to the responsible person. The *verification server* (5) parses the rule and creates an internal representation of the rule for the transformation algorithm. A new thread is started for the rule and parameterised with the respective sql queries and the time unit for the rule (6). Each thread independently accesses the database to obtain the information on state changes that are of interest when verifying the fulfillment of the rule (7). When a rule is violated a notification is sent to the notifier. (4) The notifier sends an e-mail (8) to the one responsible for expensive customers as specified by the business rule expert.

4 Conclusions and future research

The prototypes described in the previous section provide a proof of concept for this general way of checking temporal business rules. Our prototype system can handle several temporal rules based on the state of different items in the database. The thread based design allows to define different time units for the different rules. It is possible to define different reactions, in case a temporal rule is broken. Besides sending a notification the Verification Server continues checking that the database changes comply with the different rules. However there are several issues that still are subject to future exploration. We plan to set up the verification server with a database to administer student projects here at our university to explore the following issues.

Using the production Database as is, or defining a specific view The solution the database prototypes implement today offers the whole database to the business expert to define temporal rules. Also the verification server accesses the database without requiring any specific functionality on the database side. This solution is very general and requires no changes at the database side. However, in a normalised database, the data belonging together from the business point of view is often distributed over several tables. Data might be stored in a specific way due to requirements in the business systems. Production databases often contain a huge amount of tables of which only few might be interesting when defining temporal business rules. Defining a specific view for the verification server would on the one hand allow to provide the user with a concise

overview of the kind of data relevant for the formulation of temporal business rules. It would have the additional advantage that it would allow to change the database without interfering with the already defined temporal rules. On the other hand it would constrain the formulation of rules to what is provided in the view. As long as changes in the temporal business rules do not go beyond that scope, the changes can be implemented by the business expert. In case other parts of the database have to be accessed, a database expert has to extend the view.

Checking the whole data of only accessing changes In the momentary implementation of the communication between the database and the verification server, the verification server queries the production table for data that fulfills the different boolean expressions the rules are based on. This way relevant changes will automatically be recognized. For large databases this might lead to performance problems, especially when short time intervals are used as a base for the rules.

Another solution would be to only access changes to the data under observation. This would require to create a buffer and specific triggers as part of the database based on the definition of the rules in the rule builder. This should be possible but would require to change the business database itself. The rule builder module would have to keep track of the changes it implemented in the database and would have to undo them in a controlled way when a temporal rule is change or erased. A third possibility would be to implement an asynchronous observer pattern between the database and the verification server. Here again, the database would have to be adapted based on the rules defined by the user to create the events the verification server needs for checking the temporal business rules.

Finding a suitable user interface Research on End-User Development has shown that users are able to handle complex computations defined in a formal way when the formalisms are presented in a professionally meaningful way. [6][8] We plan to experiment with different kinds of user interfaces in order to find a suitable way to represent the data temporal rules can be formulated about and to support the user in deploying the possibilities that formalism offers. From a technical point of view the notifier and the rule builder are two very different parts of the system. From a business perspective the definition of temporal rules and the definition of what should happen in case of a violation belong together. Here the two tailoring interfaces must be integrated. Another issue that needs more exploration is the requirements such a system poses on the organization: when business rules can be (re-)defined, the organization has to implement procedures to decide what rules to implement.[7][2]

Performance So far, we only tested the verification server based on toy examples. Implementing a sharp version together with the project database of our university will give us the opportunity to evaluate the design and performance of the verification server when interacting with a small but realistic database. The performance of the concrete verification algorithm as well as the influence of the different design possibilities for the interaction with the database can then be evaluated under more realistic conditions.

Acknowledgments

Thanks to Bente B. Petersen, Divya R. Malemane, Bue Pedersen, Ulrik D. Olsen who together developed the different prototypes the article is based on as part of their Master Theses.

References

1. Alur, R., Henzinger, T.: Logics and Models of Real Time: A Survey. Real Time: Theory in Practice, Lecture Notes in Computer Science 600, Springer-Verlag, (1992), pp. 74-106.
2. Dittrich, Y., Lindeberg, O.: Designing for Changing Work and Business Practices. In: N. Patel (ed.). Evolutionary and Adaptive Information Systems. IDEA group publishing, (2002).
3. Havelund, K., Ruso, G.: Monitoring Programs using Rewriting. Automated Software Engineering (ASE'01), San Diego, California, (2001), IEEE Computer Society.
4. Kristoffersen K., Pedersen, C., and Andersen, H. R.: Runtime Verification of Timed LTL using Disjunctive Normalized Equation Systems. Appears in Issue 89.2 of Electronic Notes in Theoretical Computer Science (2003).
5. Kung, C. H.: A Temporal Framework for Database Specification and Verification. Proceedings of the Tenth International Conference on Very Large Data Bases, Singapore, August, (1984).
6. Patern, F., Klann, M., Wulf, V.: Research Agenda and Roadmap for EUD. Deliverables of the Network of Excellence on End-User Development, December, (2003). (<http://giove.cnuce.cnr.it/eud-net.htm>)
7. Trigg, R., Bødker, S. From Implementation to Design: Tailoring and the Emergence of Systematization in CSCW. Proceedings of the CSCW 94, ACM-Press, New York, (1994), pp. 45-55.
8. Stiemerling, O., Kahler, H., Wulf, V.: How to make software softer- Designing tailorable applications. Proceedings of the Designing Interactive Systems (DIS) (1997).

