

# STEPWISE SPECIFICATION OF DYNAMIC DATABASE BEHAVIOUR

Udo W. Lipeck

Institute für Informatik  
Technische Universität Braunschweig  
Postfach 3329, D-3300 Braunschweig  
West Germany

**ABSTRACT:** This paper presents a methodology for the stepwise specification of dynamic database behaviour. A conceptual schema is described in three levels: data, objects and transactions. To determine which sequences of database states are "admissible", integrity constraints on objects are given in temporal logic. Transactions are specified by pre/postconditions to produce "executable" state sequences. In order to guarantee that executable state sequences already become admissible, integrity constraints are completely transformed into additional pre/postconditions. We introduce general rules for these transformations. Thus, schema specifications can be refined and simplified systematically.

**KEYWORDS:** database design, conceptual schema, database integrity, dynamic constraints, temporal logic, pre/postconditions

## 1. INTRODUCTION

Conceptual database design aims at defining a global information structure independently of any implementation [TF82, Ce83]; its result is the so-called conceptual schema. This phase corresponds to the specification phase in general software design meant to bridge the gap between requirements analysis and implementation. If conceptual design is supported by formal

specification techniques, it may form a basis of rapid prototyping, too [Fu84].

Usually, the conceptual schema characterizes the structure of the database contents at any single instant, i.e., it defines the common structure of all database states. If a database, however, is considered together with its usage, the behaviour over the course of time also has to be described.

Dynamic database behaviour results from a sequence of database states. Whereas "static" integrity constraints define "consistency" of a single state, "dynamic" integrity constraints are concerned with "admissibility" of a state sequence. To specify admissible state sequences formally, temporal logic is an appropriate language. States within a sequence can conveniently be related by means of temporal formulas which provide special quantifiers like "always ... until" or "sometime ... before" [CCF82, Ku84, EL84/LE85].

During runtime of a database, transitions between states are caused by user transactions. To exclude undesirable database manipulations which lead to inconsistent states, a set of basic transactions should be agreed upon during design. Then, all applications have to be composed of them. In the implementation, transactions are realized by programs that access the database system. During design, they can be specified independently of their later implementation by means of pre/postconditions. Thus, the set of possible transaction sequences is restricted to "exe-

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

© 1986 ACM 0-89791-191-1/86/0500/0387 \$00 75

cutable" sequences in which each state between two transactions must satisfy the respective post- and precondition.

This specification of transactions completes the original specification of dynamic database behaviour: Only those state sequences are allowed which are induced by executable transaction sequences and which are admissible with respect to the integrity constraints.

In order to ensure correct database behaviour, it becomes necessary to monitor the constraints during runtime. In contrast to expensive universal monitoring mechanisms, transactions offer a possibility of controlling database usage "locally" and therefore efficiently: Checks may be limited to the data actually affected and may be suited to the respective operation. The same principle is pursued in [N182] to simplify static constraints for standard relational updates.

To incorporate this kind of monitoring, the specification of transactions should be refined in subsequent design steps: Integrity constraints have to be converted into more restrictive pre/postconditions so that executable transaction sequences already guarantee admissible state sequences. A similar procedure is recommended by [VF85], but only for static and transitional constraints (i.e. conditions on single state transitions).

This paper presents a methodology for such a stepwise specification of dynamic database behaviour. Unlike the related literature, we concentrate on dynamic integrity constraints which refer to arbitrary intervals of state sequences and we give rules how to transform constraints into pre/postconditions systematically. In particular, verification of admissibility becomes unnecessary when applying general transformation rules.

The following section introduces conceptual schema specifications and especially explains integrity constraints and transaction specifications. Section 3 describes our method of refining schema specifica-

tions by means of transformations. Finally, related work and some conclusions will be discussed.

## 2. CONCEPTUAL SCHEMA SPECIFICATION

The original specification of a conceptual database schema consists of three levels, each built upon the previous one:

- (1) data types
- (2) object types  
with integrity constraints
- (3) transactions

In the following, we only give a short introduction to schema specifications, but emphasize those aspects which are relevant to the main subject of this paper. More detailed work on levels (1) and (2) can be found in [LEG85].

### 2.1 DATA AND OBJECTS

Data types are a fixed reservoir of values for the database, e.g., bool, int, no, text, year; they may be defined by, for instance, algebraic equational specifications. Objects are things about which information is to be stored in the database. They are partitioned into sorts, and information about them is carried by functions between objects and data. Predicates on objects are represented as bool-valued functions. For example, in a schema for an "Automobile Registration Authority" similar to that of [ISO82], the following object sorts and functions appear:

```
sorts      CAR, MODEL, CAR-OWNER, ...
subsorts  CAR-OWNER =
            MANUFACTURER+GARAGE+PERSON
            ...
functions regno:      CAR    -> no
            model-of:  CAR    -> MODEL
            produced:  CAR    -> bool
            registered: CAR    -> bool
            approved:  MODEL  -> bool
            name:      CAR-OWNER -> text
            owner:     CAR    -> CAR-OWNER
            this-year: -> year
```

Later examples will refer to this syntactic structure.

We assume the reader to be familiar with the construction of predicate calculus formulas from object functions. General temporal formulas can be built from those formulas using

- logical connectives ( $\wedge, \vee, \neg, \Rightarrow, \dots$ )
- quantification over possible objects ( $\forall, \exists$ )
- temporal quantification by always, sometime, next
- bounded temporal quantification by always/sometime ... before/until

In order to express dynamic integrity constraints, only the following special kinds of temporal formulas are used:

$$\forall x_1:s_1 \dots \forall x_n:s_n$$

$$\text{always } (\alpha \Rightarrow \left. \begin{array}{l} \text{always } \beta \\ \text{sometime } \beta \end{array} \right\} \left. \begin{array}{l} \text{before } \tau \\ \text{until } \tau \end{array} \right)$$

Here  $\alpha$ ,  $\beta$ , and  $\tau$  are nontemporal formulas, i.e. formulas without temporal quantifiers, and  $x_1, \dots, x_n$  are all free variables of sorts  $s_1, \dots, s_n$  that occur in these formulas. In the following, a simplified syntax is used for constraints; it omits the standard prefix and allows to omit the  $\alpha$  and  $\tau$  parts:

$$\left[ \text{from } \alpha \right] \left\{ \begin{array}{l} \text{always } \beta \\ \text{sometime } \beta \end{array} \right\} \left[ \begin{array}{l} \text{before } \tau \\ \text{until } \tau \end{array} \right]$$

Examples: The integrity constraint

$$(I1) \text{ from produced}(c) \wedge \text{this-year} = y$$

$$\text{sometime registered}(c)$$

$$\text{before this-year} \geq y + 2$$

(with variables  $c:CAR$  and  $y:year$ ) says that each car must be registered in the year of its production or in the following year. The mentioned database object this-year is assumed to contain the current year.

Two further constraints shall restrict the succession of ownership relations. For this purpose, we consider the subsort to which an object of sort  $CAR-OWNER$  belongs, i.e. either to  $MANUFACTURER$ ,  $GARAGE$ , or  $PERSON$ . Membership in a subsort is expressed by the standard predicate is.

$$(I2) \text{ from owner}(c) \text{ is } MANUFACTURER$$

$$\text{sometime owner}(c) \text{ is } GARAGE$$

$$\text{before owner}(c) \text{ is } PERSON$$

$$(I3) \text{ from owner}(c) = co \wedge co \text{ is } GARAGE$$

$$\text{always owner}(c) = co$$

$$\text{before owner}(c) \text{ is } PERSON$$

(variables:  $c:CAR, co:CAR-OWNER$ )

A car must be passed from a manufacturer to a garage first, before it may be passed to a person. It must remain property of the same garage before it is sold to a person. \*\*\*

Aside from the constraints above one can imagine more complicated temporal formulas or even other kinds of dynamic assertions which might refer to the past [Ku84]; such constraints are not treated in this work. The restricted formulas from above are not only an appropriate basis of monitoring or refining specifications, but also turn out to be sufficient for many applications.

The object level of the schema specification determines how database states are structured and which state sequences are admissible. In a state  $\sigma$ , the sort and function names are interpreted by finite sets of "actual" objects and by functions between object and data sets. In a state sequence  $\underline{\sigma} = \langle \sigma_0, \sigma_1, \sigma_2, \dots \rangle$ , this interpretation may change from state to state.  $\underline{\sigma}$  is called admissible iff it satisfies the integrity constraints.

We briefly define what it means that an infinite state sequence  $\underline{\sigma}$  satisfies a constraint  $\Gamma$ . Let ' $[\sigma, \zeta] \models \psi$ ' denote the fact that a nontemporal formula  $\psi$  holds in a single state  $\sigma$  for a substitution  $\zeta$ .  $\zeta$  indicates the objects and data which are substituted for the free variables in  $\psi$ . For substitutions, all "possible" objects must be considered, i.e. all objects that belong to some database state as actual objects.

$$(1) \underline{\sigma} \text{ satisfies } \Gamma \equiv \text{from } \alpha \text{ always } \beta \text{ before } \tau \text{ iff.}$$

For all substitutions  $\zeta$  of free variables in  $\alpha, \beta, \tau$  and for all states

$\sigma_1, i \geq 0$ , the following holds:  
 If  $[\sigma_1, \zeta] \models \alpha$ , then  $[\sigma_j, \zeta] \models \beta$   
 for each state  $\sigma_j, 1 \leq j < \mu$ , where  
 $\mu = \min(\{k \mid k \geq 1 \wedge [\sigma_k, \zeta] \models \tau\} + \{\infty\})$ .

Starting from any state  $\sigma_1$  in which the "start condition"  $\alpha$  holds, all states following that state (including  $\sigma_1$  itself) must satisfy the condition  $\beta$ , before the "termination condition"  $\tau$  becomes valid for the first time (in  $\sigma_\mu$ ). This situation is required for each substitution  $\zeta$ . If  $\tau$  does not occur at all,  $\beta$  must hold in the entire future of  $\sigma_1$  ( $\mu = \infty$ ).

(11)  $\underline{g}$  satisfies

$\models$  from  $\alpha$  sometime  $\beta$  before  $\tau$  iff:  
 For all substitutions  $\zeta$  and for all states  $\sigma_1, i \geq 0$ , holds:  
 If  $[\sigma_1, \zeta] \models \alpha$ , then there is some state  $\sigma_j, 1 \leq j < \mu$ , such that  $[\sigma_j, \zeta] \models \beta$   
 ( $\mu$  as above).

In contrast to (1), there must exist at least one state following  $\sigma_1$  which satisfies  $\beta$  before the termination condition  $\tau$  occurs.

(111) Analogous definitions apply to the combinations of always/ sometime with until. Here the bound  $\mu$  must be included.

According to the kind of implicit quantification over states, we distinguish between universal constraints in which the keyword always appears and existential ones with sometime.

The above semantics of constraints refers to infinite state sequences which represent the complete behaviour of a database from an initial state ( $\sigma_0$ ) to the future. From a practical point of view, it is more interesting to consider properties of the partial database behaviour up to an actual state ( $\sigma_n$ ). Therefore, we introduce the following notions.

A finite state sequence  $\underline{g} = \langle \sigma_0, \dots, \sigma_n \rangle$  is called (strictly) admissible iff the infinite sequence  $\langle \sigma_0, \dots, \sigma_n, \sigma_n, \sigma_n, \dots \rangle$  (remaining constant from position  $n$  on) is admissible. If some infinite continua-

tion  $\underline{g}$  (not necessarily  $\langle \sigma_n, \sigma_n, \dots \rangle$ ) exists such that  $\underline{g}.\underline{g}$  becomes admissible, then  $\underline{g}$  is possibly admissible. In this case,  $\underline{g}$  is not strictly admissible if and only if the sometime-condition of some existential constraint has not been satisfied up to  $\sigma_n$ , although the start condition had occurred. For example, the condition 'registered(c)' in constraint (II) might still be pending, after 'produced(c)' has become valid.

In order to avoid any checking that looks ahead from the actual state, we proceed on the "optimistic" assumption that all pending conditions are satisfiable in the future. Nevertheless, the user might choose a wrong continuation after the actual state. Formally, we assume that a distinguished (infinite) state sequence  $\underline{\lambda}$  exists in which every formula holds. A finite state sequence  $\underline{g}$  is defined to be provisionally admissible iff  $\underline{g}.\underline{\lambda}$  is admissible. To assure provisional admissibility, it suffices to prove that no constraint has been violated in  $\sigma_n$  or in any state  $\sigma_j$  before  $\sigma_n$ . For instance, an existential constraint is violated in this sense, if its termination condition occurs although the sometime-condition has not been valid before.

The latter kind of admissibility can be controlled by means of an integrity monitor which only refers to the database history, but does not involve any lookahead [LEGB5]. For this reason, that property will be the goal of refining specifications in section 3.

## 2.2 TRANSACTIONS

The third level of a schema specification describes all transactions that may change the database contents. In the style of Hoare's assertions, transactions are specified by means of pre- and postconditions.

Such conditions may be expressed by arbitrary formulas of predicate logic (nontemporal formulas). A precondition contains a

condition on which the transaction may be executed. Any state directly following the transaction must satisfy the corresponding postcondition. Besides, it is implicitly assumed as a "frame rule" that all object functions that are not affected by the postcondition remain unchanged. To specify different cases of state transitions, several pairs of pre/postconditions may be given for one transaction.

Examples: The automobile registration schema of [ISO82] includes transactions to register or sell a car. They are specified here as follows:

```
(T1) REGISTER (c:CAR, m:MODEL,...)
      pre  approved(m)
      post registered(c) [=true]
           ^ model-of(c)=m ^ ...
```

A car  $c$  may only be registered, if the declared model  $m$  has already been approved. During registration, any information passed as parameter is appropriately stored.

```
(T2) SELL (c:CAR, newco:CAR-OWNER)
      pre  --- [true; no restriction]
      post owner(c) = newco
```

A car  $c$  is sold to a new owner  $newco$ . \*\*\*

The frame rule simplifies specifications, because only changes must be mentioned in a postcondition. Since arbitrary formulas are allowed, conditional or alternative effects may be formulated. Thus, the post-state of a transaction may be left indefinite within a certain range, and it will be determined uniquely only by later refinement or implementation.

[ku84/ku85a] and [VF85] also use pre/postconditions to specify transactions, but apart from the frame rule they have other implicit assumptions which seem to be more restrictive than helpful in many applications. Besides, [VF85] only admit conjunctions of ground formulas (like in the examples above), thus excluding implications, existence statements, etc. In [VCFB1] and [CVF84] different formalisms of specifying transactions and operations are compared by means of examples.

Formally, syntax and semantics of transactions are given by the following definitions:

A transaction specification consists of a name  $T$ , (sorted) parameter variables  $X$  and several pre/postconditions  $F_j/Q_j$  ( $j \geq 1$ ); notation:  $\{F_j\} [(X) \{Q_j\}]$ . These conditions are nontemporal formulas in which additional free variables  $Y$  may occur beside  $X$ . A substitution  $\zeta$  of parameter variables is called an actual parameter.

$T$  transfers a state  $\sigma$  into a state  $\sigma'$  with an actual parameter  $\zeta$  iff:

- There is a precondition  $F_j$  and a substitution  $\xi$  of the free variables  $Y$  such that.  $[\sigma, \zeta + \xi] \models F_j$ .
- For all pre/postconditions  $F_j/Q_j$  and for all substitutions  $\xi$  of the free variables  $Y$ , the following holds:  $[\sigma, \zeta + \xi] \models F_j$  implies  $[\sigma', \zeta + \xi] \models Q_j$ .
- $\sigma$  is minimally changed to  $\sigma'$ , i.e., no change of any object function value from  $\sigma$  to  $\sigma'$  can be undone without violating condition (b).

Part (c) of the definition gives a semantic characterization of the frame rule. The so-called frame of  $T$  is constituted by all formulas  $\varphi$  which are invariant in the following sense:

For all triples  $(\sigma, \sigma', \zeta)$  such that  $T$  transfers  $\sigma$  into  $\sigma'$  with  $\zeta$  and for all  $F_j, \xi$  as above, the following holds:  $[\sigma, \zeta + \xi] \models F_j \wedge \varphi$  implies  $[\sigma', \zeta + \xi] \models \varphi$  (We assume that  $\lambda + Y$  already contains the free variables occurring in  $\varphi$ .)

Obviously formulas that do not mention any function used in the postconditions belong to the frame. But there may also be conditions in the frame which involve those functions, trivial examples are the postconditions themselves. (Of course, their negations are not invariant.) To analyze specifications, derivation rules for frame formulas are needed.

Transaction specifications restrict possible sequences of transactions or states to "executable" ones.

A transaction sequence  $\langle T_1, \dots, T_n \rangle$  is called executable iff there is a sequence  $\underline{\sigma} = \sigma_0, \dots, \sigma_n$  of states, such that  $\sigma_{i-1}$  is transferred into  $\sigma_i$  by  $T_i$  with some actual parameter  $\zeta_i$  (for  $i=1, \dots, n$ ). In this case, the state sequence  $\underline{\sigma}$  is also called executable.

Thus, transactions and dynamic integrity constraints can be seen as complementary specifications, together they restrict state sequences to such sequences that are executable and admissible. In the example above, state transitions changing a "car-owner" must be induced by corresponding transactions like SELL (T2), and they must satisfy the constraints (I2) and (I3).

## 2. REFINING THE SCHEMA SPECIFICATION

The possibility of restricting state sequences by specifying transactions shall be utilized to monitor integrity constraints. The constraints are converted into a refined specification of the transactions so that executable state sequences become provisionally admissible up to each actual state. Thus, dynamic database behaviour is determined by transactions only, and temporal constraints need not be considered for implementation any longer.

This refinement offers the opportunity to observe the impacts of "global" integrity constraints on the "local" usage of the database. If new pre/postconditions lead to undesirable restrictions or even to contradictions, the original schema specification should be corrected. Moreover, there is some prospect that integrity monitoring by transactions can be implemented quite efficiently, if checks are suited to the respective kind of change and to the objects affected.

As a preparation for refining a schema specification, the set of object functions is extended by some predicates which keep track of partially satisfied constraints: there may be sometime-conditions which are yet to come or always-conditions which must hold only temporarily. It makes sense

to store such "memos" in the database, since they can indicate the actual degree of admissibility to the user.

Then, the pre/postconditions of transactions are completed by further conditions. Semantically, executability of state sequences becomes stronger restricted. At first we show which completion would be necessary without regarding any specific knowledge about a transaction. This transformation corresponds to the universal integrity monitor presented in [LEG85]. By considering which transactions can change which object functions, those completions can be essentially simplified afterwards.

The key idea of monitoring or transforming a dynamic constraint is to split a temporal formula into a nontemporal formula that must hold in the actual state, and a temporal one that remains to be checked in the future. For this purpose, so-called "recursion rules" of temporal logic [MP81] are helpful:

$$\begin{aligned} \text{always } \varphi & \equiv \varphi \wedge \text{next } (\text{always } \varphi) \\ \text{sometime } \varphi & \equiv \varphi \vee \text{next } (\text{sometime } \varphi) \end{aligned}$$

For  $\Phi \equiv (\alpha \Rightarrow \text{sometime } \beta \text{ before } \tau)$  and  $\Phi \equiv (\text{sometime } \beta \text{ before } \tau)$ , the following laws can be derived:

$$\begin{aligned} (1) \quad & \Phi \equiv (\alpha \Rightarrow \neg \tau) \wedge (\alpha \wedge \neg \beta \Rightarrow \text{next } \Phi) \\ (11) \quad & \Phi \equiv \neg \tau \wedge (\neg \beta \Rightarrow \text{next } \Phi) \end{aligned}$$

If one of the temporal formulas  $\Phi$  or  $\Phi$  is required for a state sequence, part (1) excludes an illegal occurrence of the termination condition  $\tau$  in the first state, and (2) indicates on which premises  $\Phi$  must be satisfied by the rest sequence starting with the next state. According to the interpretation of a constraint "from  $\alpha$  always  $\beta$  before  $\tau$ " (see section 2.1),  $\Phi$  applies to arbitrary (starting) states and substitutions. According to the laws above  $\Phi$  may apply to following states for certain substitutions; it then replaces the requirement  $\Phi$ , since  $\Phi$  implies  $\Phi$ . This temporary requirement  $\Phi$  must be recorded in a special memo.

In this section we discuss those two basic

forms of integrity constraints which were explained in section 2.1 explicitly. Each existential constraint

$\Gamma \equiv \text{from } \alpha \text{ sometime } \beta \text{ before } \tau$   
is transformed in the following way:

Step I:

An additional predicate named sometime  $\beta$  before  $\tau$  is introduced whose arguments correspond to the free variables in  $\beta$  and  $\tau$ ; in the following, it is abbreviated to SB (though depending on  $\Gamma$ ). This predicate must be entered in the database as true for a substitution  $\zeta$ , if the corresponding condition ( $\emptyset$  from above) remains to be checked with that substitution in the next states; so  $SB(\zeta)$  plays the role of the mentioned "memo". Recall that a substitution just represents a certain combination of data and objects to be used as actual arguments.

Thus, as many predicates are added to the schema as constraints are given in the original specification. In special cases, such predicates can be expressed in terms of existing functions and predicates so that a new predicate can be avoided. This will be illustrated in a later example.

Step II:

Let an arbitrary transaction  $T$  be given with an original specification or with a specification resulting from transformation of other constraints. Each assertion  $\{P\} \mid \{Q\}$  consisting of a pair of pre/postconditions  $P/Q$  is rewritten to:

$$\begin{array}{ll} (1) & (2) \\ \{P \wedge \neg SB\} T \{Q \wedge (\alpha \neq \neg \tau) \wedge (\alpha \wedge \neg \beta \Rightarrow SB)\} & \\ \{P \wedge SB\} T \{Q \wedge \neg \tau \wedge (\beta \Rightarrow \neg SB)\} & \end{array}$$

Here, the laws (1) and (11) from above are utilized to restrict the set of post-states. The new specification expresses the requirement that  $T$  has to monitor the given constraint as far as this can be done by inspecting the new actual database state. A distinction is made, whether or not a memo  $SB$  "exists" previously. The  $\neq$ -direction of condition (2) controls

"insertion" ( $SB$ ) or "deletion" ( $\neg SB$ ) of the memo. The  $\neq$ -direction of (2) determines if  $\neg SB$  or  $SB$  is retained. It should be noticed that the free variables in the added parts must be named differently from parameter and other variables, because those conditions shall hold for arbitrary objects.

Since postconditions affect only states following transactions, the addition in the first line above has to hold in the initial state of the database, too:

$$(\alpha \neq \neg \tau) \wedge (\alpha \wedge \neg \beta \Rightarrow SB)$$

Example. By transforming the integrity constraint

$$(I1) \text{ from produced}(c) \wedge \text{this-year} = y \\ \text{sometime registered}(c) \\ \text{before this-year} \geq y + 2$$

the following new specification is obtained for the transaction (T1) REGISTER( $c, m, \dots$ ) of section 2.2:

$$\begin{array}{l} \text{pre}_1 \text{ approved}(m) \wedge \neg SB(c, y) \\ \text{post}_1 \text{ registered}(c) \wedge \dots \\ \wedge ( (\text{produced}(c) \wedge \text{this-year} = y) \\ \Rightarrow \neg \text{this-year} \geq y + 2 ) \\ \wedge ( (\text{produced}(c) \wedge \text{this-year} = y \\ \wedge \neg \text{registered}(c) ) \\ \Rightarrow SB(c, y) ) \end{array} \} (-)$$

Obviously  $(-)$  is a tautology that may be dropped

$$\begin{array}{l} \text{pre}_2 \text{ approved}(m) \wedge SB(c, y) \\ \text{post}_2 \text{ registered}(c) \wedge \dots \\ \wedge \text{this-year} < y + 2 \\ \wedge ( \text{registered}(c) \Rightarrow \neg SB(c, y) ) \end{array}$$

For design purposes, it is recommended to introduce a more application oriented name for the predicate  $SB$ . A conversion of its parameters may be included, as it can be done for (I1):

$$\begin{array}{l} \text{to-be-registered-before: CAR year} \rightarrow \text{bool} \\ \text{where to-be-registered-before}(c, y + 2) \\ = SB(c, y) \end{array} \quad ***$$

Step III:

Now, the transaction specifications of step II can be essentially simplified.

Due to the frame rule with respect to the original specification many transactions will not change the added conditions at all, others will not do so for many substitutions. Therefore, the transformation above needs only be done for transactions which possibly change the truth value of  $\alpha$ ,  $\beta$  or  $\tau$ , i.e., some of these formulas or their negations are not in the frame of the transaction. Moreover, the new variables of step II may be replaced by those parameter variables, original variables, constants or terms which are involved in changes of the corresponding formula parts. If not all three constraint components change, whole factors of the above conjunctions may be omitted. Finally, the frame rule may be utilized for the new specification, too

To discuss typical situations, let us consider transactions  $T$  which change at most one of the constraint components  $\alpha$ ,  $\beta$  or  $\tau$  from false to true. Equivalently, either  $\neg\alpha$ ,  $\neg\beta$  or  $\neg\tau$  is not in the frame of  $T$ . The following specifications result from step III:

- (I) If  $T$  can change only  $\alpha$  to true:
  - { $P \wedge \neg SB$ }  $T$  { $Q \wedge \neg\tau \wedge (\neg\beta \Leftarrow SB)$ }
  - { $F \wedge SB$ }  $T$  { $Q$ }
- (II) If  $T$  can change only  $\beta$  to true:
  - { $F \wedge \neg SB$ }  $T$  { $Q \wedge \neg SB$ }
  - { $F \wedge SB$ }  $T$  { $Q \wedge \neg SB$ }
  - or combined: { $P$ }  $T$  { $Q \wedge \neg SB$ }
- (III) If  $T$  can change only  $\tau$  to true:
  - { $F \wedge \neg SB$ }  $T$  { $Q \wedge \neg\alpha \wedge \neg SB$ }
  - { $P \wedge SB$ }  $T$  { $Q \wedge \text{false}$ } (-)

Since the second postcondition (-) cannot hold in any state, the whole line may be dropped so that the precondition of  $T$  is restricted to { $F \wedge \neg SB$ } only.

Examples: The above refinement of the specification for REGISTER can be simplified according to case (II), because only  $\neg\beta \equiv \neg\text{registered}(c)$  is not invariant:

```
{approved(m)} REGISTER(c,m,..)
  {registered(c) ..  $\wedge \neg SB(c,y)$ }
```

After execution of REGISTER(c,m,...) no memo may exist for the car c. An imple-

mentation of the transaction has to take care of deleting a memo which possibly existed before.

Let a further transaction NEXT-YEAR be given which updates the current year this-year :

```
{this-year=y} NEXT-YEAR {this-year=x+1}
```

Referring to constraint (II), only the termination condition  $\tau \equiv \text{this-year} \leq y+2$  is changed for  $y = x-1$ , so that transformation according to case (III) and obvious simplification of the resulting postcondition yield:

```
{this-year=y  $\wedge \neg SB(c,y-1)$ }
  NEXT-YEAR {this-year=x+1}
```

Thus, right before execution of NEXT-YEAR any memo from the previous year must not exist any longer. In a way, "the clock must be stopped" in order to settle the required registrations for all cars  $c$  with memo SB(c,y-1). \*\*\*

All simplifications in step III presuppose that any partial condition or substitution which is excluded here has been checked after its last change (and in the initial state). This is the case if all transaction specifications are treated according to the rules above.

As can be seen by induction on state sequences, the transformation rules do guarantee that executable state sequences satisfy the given constraint in the sense of provisional admissibility. Thus, single verifications are rendered superfluous. Additionally, our examples demonstrate how conditions can typically be simplified further in concrete situations, so that a "readable" specification results in spite of formal transformation.

The transformation of universal constraints will only be given by analogy to existential ones. For each constraint

from  $\alpha$  always  $\beta$  before  $\tau$

a predicate always  $\beta$  before  $\tau \equiv AB$  is needed (step I). Then, the basic transformation of { $F$ }  $T$  { $Q$ } reads as follows:



$\{P \wedge \neg AB\} T \{Q \wedge (\alpha \wedge \neg T \Rightarrow \beta) \wedge (\alpha \wedge \neg T \Leftarrow \neg AB)\}$   
 $\{P \wedge AB\} T \{Q \wedge (\neg T \Rightarrow \beta) \wedge (T \Leftarrow \neg AB)\}$

This specification is constructed similarly to step II for existential constraints. The subsequent simplifications of step III apply here, too.

Example. The following constraint is to be transformed:

(I3) from owner(c)=co and co is GARAGE  
always owner(c)=co  
before owner(c) is PERSON

Among the given transactions only SELL (T2) is affected by a transformation, since no other transaction changes the function owner. After some equivalence conversions of the basic transformation a refined specification is obtained for SELL(c,newco) :

pre<sub>1</sub>  $\neg AB(c,co)$   
post<sub>1</sub> owner(c)=newco  
 $\wedge ((owner(c)=co \wedge co \text{ is } GARAGE) \Leftarrow \neg AB(c,co)) \quad \} (*)$

or:

post<sub>1</sub> owner(c)=newco  
 $\wedge (newco \text{ is } GARAGE \Leftarrow AB(c,newco))$   
 ({pre<sub>1</sub>} T {post<sub>1</sub>}) is equivalent to  
 ({pre<sub>1</sub>} T {post<sub>1</sub>}), since all omitted assertions are in the frame of T wrt post<sub>1</sub>.)

pre<sub>2</sub> AB(c,co)  
post<sub>2</sub> owner(c)=newco  
 $\wedge (\neg newco \text{ is } PERSON \Rightarrow co=newco) \quad (+)$   
 $\wedge (newco \text{ is } PERSON \Leftarrow \neg AB(c,co))$

Thus, restrictions on succession of ownerships have been incorporated into the transaction specification. Subsequent transformation of the other integrity constraint (I2) would lead to a further differentiation. \*\*\*

The second postcondition of the example (post<sub>2</sub>) strongly restricts the possibility of transferring states by this transaction: For some states satisfying the precondition, there is no state satisfying the postcondition. This specifies correctly that certain state sequences are not executable, but the situation seems to be modelled in a somehow complicated

way. To get a clearer specification it is recommended that the precondition should exclude cases when no poststate exists.

Example, cont d.: In spite of the precondition, part (+) of the postcondition cannot be satisfied for arbitrary parameters newco. In this example, it is enough to put this restriction into the precondition in its original form, since no object function being subject to change is involved in this formula. Then, the second specification part of SELL(c,newco) reads:

pre<sub>2</sub> AB(c,co)  
 $\wedge (\neg newco \text{ is } PERSON \Rightarrow co=newco) \quad (+)$   
post<sub>2</sub> owner(c)=newco  
 $\wedge (newco \text{ is } PERSON \Leftarrow \neg AB(c,co))$

Moreover, here the new predicate AB can be expressed by means of existing predicates provided that the function owner is affected by SELL only. It can be proved by induction on executable state sequences that the following equivalence holds always:

$(owner(c)=co \wedge co \text{ is } GARAGE) \Rightarrow AB(c,co)$   
 (This was mentioned in the first postcondition post<sub>1</sub> as (\*).) \*\*\*

Static integrity constraints  $\beta$  which correspond to special universal constraints from true always  $\beta$  before false

do not require introducing any additional predicate, either. For them, transformation of {P} T {Q} yields:

$\{F \wedge \beta\} T \{Q \wedge \beta\}$

In particular, this refinement might give rise to checking the original postcondition Q for consistency with the static constraint  $\beta$ .

Of course, one can imagine situations where certain integrity constraints can be concluded directly from the original transaction specifications so that no transformation is necessary. In any case, general transformations are possible and reliable; thus, they can serve the designer at least as a guideline to refining the schema specification. Besides, the exam-

ples show that systematic simplifications can lead to formulations easy to survey and to grasp, where verification has already been done. So transaction specifications become a sufficient and useful interface to the implementor.

#### 4. RELATED WORK AND CONCLUSIONS

The literature based upon similar calculi (temporal logic and pre/postconditions) contains first approaches to comparable multi-level database specifications [CCF82/ CVF84/ VF85, ku84/ ku85b, SFNC84]. Requirements and methods of verifying temporal constraints are analyzed presuming that complete descriptions of transactions are given beforehand. Already [GM79] have studied verification of static constraints by checking invariant pre/postconditions of transactions. Besides, [ku84/ku85a] is concerned with consistency of schema specifications.

Only [VF85] give some basic methodical hints that the specification - and also the selection - of transactions should be refined in a stepwise way. The above approaches, however, cannot consider "some-time" constraints in refinements, although they may be formulated partially; for this, our notion of "provisional admissibility" proves to be appropriate. Transformation methods for dynamic integrity constraints have not yet been discussed by other authors.

The specification method presented here depends on the principle to fix a module of basic transactions for all database updates. This module becomes responsible for the correctness of dynamic database behaviour. Its design is based on the transformation of the original schema specification: the designer is guided to recognize the impacts of the global specification on the transactions.

The design procedure itself might be supported by an interactive system that controls and manages stepwise transformations and that offers simplifications of

logical formulas. Moreover, those transformations may be used to generate a prototype of a database application system from its original specification. As shown by [VF85], preliminary implementations of transactions can be derived from certain restricted pre/postconditions automatically.

Concerning theoretical foundations, the simplification step (III) could be treated more formally than in this paper, but our emphasis was on introducing typical elements of a specification methodology. For simplifications, inference rules are needed to derive frame formulas and to detect irrelevant variable substitutions. The approach of [Ni82] on static constraint simplification should be examined if it can be extended to general transaction specifications. In addition, new monitoring principles for dynamic constraints [LSE86] might be utilized to transform a larger class of temporal formulas.

#### REFERENCES

##### Abbreviations:

- FODS Proc. ACM Symp. on Principles of Database Systems
- SIGMOD Proc. Int. ACM-SIGMOD Conf. on Management of Data
- TFAIS Proc. IFIP Work. Conf. on Theoretical and Formal Aspects of Information Systems (A. Sernadas et al., eds.), North-Holland, Amsterdam 1985
- VLDB Proc. Int. Conf. on Very Large Data Bases

- [CCF82] Castilho, J.M.V.de/ Casanova, M.A./ Furtado, A.L.: A Temporal Framework for Database Specifications. VLDB 1982, 280-291
- [CVF84] Casanova, M.A. / Veloso, P.A.S. / Furtado, A.L.: Formal Database Specification - An Eclectic Perspective. PODS 1984, 110-118
- [Ce83] Ceri, S. (ed.): Methodology and Tools for Data Base Design. North-Holland, Amsterdam 1983
- [ELG84] Ehrich, H.-D. / Lipeck, U.W. / Gogolla, M.: Specification, Semantics and Enforcement of Dynamic Database Constraints. VLDB 1984, 301-308
- [Fu84] Furtado, A.L.: An Informal Approach to Formal Specifications. SIGMOD Record 14,1 (1984), 45-54

- [GM79] Gardarin,G./Melkanoff,M.: Proving Consistency of Database Transactions. VLDB 1979, 291-298
- [ISO82] ISO/TC97/SC5/WG3 : Concepts and Terminology for the Conceptual Schema and the Information Base. (J.J. van Griethuysen, ed.) 1982
- [ku84] Kung,C.H. : A Temporal Framework for Database Specification and Verification. VLDB 1984, 91-99
- [ku85a] Kung,C.H. : A Tableaux Approach for Consistency Checking. TFAIS 1985, 191-210
- [ku85b] Kung,C.H. : On Verification of Temporal Database Constraints. SIGMOD 1985, 169-179
- [LE85] Lipeck, U.W. / Ehrich, H.-D. / Gogolla,M. : Specifying Admissibility of Dynamic Database Behaviour Using Temporal Logic. TFAIS 1985, 145-157
- [LSE86] Lipeck,U.W./Saake,G./Ehrich,H.D.: Monitoring Dynamic Database Integrity by Transition Graphs. 1986 (submitted for publication)
- [MP81] Manna,Z./Pnueli,A.: Verification of Concurrent Programs: The Temporal Framework. in: The Correctness Problem in Computer Science (R.S. Boyer/ J.S. Moore, eds.), Academic Press, 1981, 215-272
- [Ni82] Nicolas,J.-M.: Logic for Improving Integrity Checking in Relational Data Bases. Acta Informatica 18 (1982), 227-253
- [SFNC84] Schiel,U./ Furtado,A.L./ Neuhold, E.J./Casanova,M.A.: Towards Multi-Level and Modular Conceptual Schema Specifications. Information Systems 9 (1984), 43-57
- [TF82] Teorey,T.J./ Fry,J.P.: Design of Database Structures. Prentice-Hall, Englewood Cliffs 1982
- [VCF81] Veloso,P.A.S./ Castilho,J.M.V.de/ Furtado,A.L.: Systematic Derivation of Complementary Specifications. VLDB 1981, 409-421
- [VF85] Veloso,P.A.S./ Furtado,A.L.: Towards Simpler and Yet Complete Formal Specifications. TFAIS 1985, 175-189

#### ACKNOWLEDGEMENT

I wish to thank Klaus Drostén for helping me to improve the English of this paper.