Deriving Transaction Specifications from Deductive Conceptual Models of Information Systems

Maria Ribera Sancho
Antoni Olivé

Facultat d'Informàtica, Universitat Politècnica de Catalunya
Pau Gargallo 5, 08028 Barcelona - Catalonia
e-mail:{riberalolivé}@lsi.upc.es

Abstract. We review the main components of a Deductive Conceptual Model (DCM) of an IS, and introduce a logic-based language for its specification.

We then present a new, formal method for the derivation of transaction specifications implied by a given DCM. The method is based on the SLDNF proof procedure, and can be implemented easily in Prolog environments.

The method requires the development of an Internal Events Model (IEM). We present such a model, point out how can it be automatically obtained and discuss its use in transaction derivation.

1 Introduction

A general trend in information systems engineering is the adoption of knowledge engineering techniques to enhance existing methodologies as well as to introduce new and more productive development paradigms, methods and supporting tools [Bub86]. This trend is based on the fact that information systems development is a knowledge intensive task [MBJ90] and, thus, it is not surprising that a lot of research has been directed toward providing knowledge-based tools to support the entire information systems development life cycle, from high-level design to low-level code generation [Fre85].

This paper describes our work on deriving transaction specifications from a Deductive Conceptual Model. Our approach uses a logic-based language for the specification of conceptual models, and applies logic-based techniques for the generation of a system design from a conceptual model.

The main originality of this work is that our language is based on the "deductive" approach to conceptual modelling, instead of the traditional, "operational" approach. Both approaches can provide a complete specification of the static and dynamic aspects of an information system, but they differ in the way the dynamic aspect is modelled.

In the operational approach, changes to the Information Base (IB), corresponding to changes in the Universe of Discourse (UoD), are defined by means of operations. The occurrence of a real-world, external event triggers the execution of an operation (transaction), which reflects the effect of the event on the IB. These effects consist usually of insertions, updates or deletions to the IB. On the other hand, operations, as well as queries and integrity constraints, usually have only access to the current state of the IB.

In the deductive approach, the IB is defined only in terms of the external events, by means of deductive rules, and queries and integrity constraints are defined as if the complete history of the IB were available. A deductive conceptual model (DCM) is a

specification of an IS in the deductive approach. Examples of conceptual modelling languages using the deductive approach are DADES [Oli82] and CIAM [GKB82].

A detailed comparison of the operational and deductive approaches can be found in [BuO86, Oli86]. The main conclusions are that DCMs provide more local definitions, are easier to change and to accommodate new requirements, and provide more design freedom. However, DCMs are much more difficult to implement than operational models. The reason is that an operational model already embeds some architectural design decisions, which are not made in DCMs [Oli89].

The main design decisions required to implement an information system from a DCM are data base design and transaction design. Usually, many valid alternatives in data base design exist, and the designer must choose the alternative that he/she considers most appropriate. Complete automation of this decision is not possible, but there is a place for CASE tools that aid the designer by confirming the consistency of his/her decisions and by evaluating their impact in terms of performance.

In transaction design, the designer must decide which transactions will exist and, for each of them, when will be executed, its pre-conditions and the actions to be performed, including data base updates and output production.

Contrary to data base design, transaction design from a DCM can be completely automated. We can build a transaction for each external event, to be executed when that event occurs in the real world. Transaction pre-conditions can be determined from the integrity constraints defined in the DCM. Data base updates can be determined from an analysis of deductive rules, and output production can be determined from queries definition.

In this paper we present a new, formal method for generating transaction specifications from a DCM. We extend here the work in the ODISSEA project reported in [San92], where more details on the general framework can be found. The method is based on the use of the SLDNF proof procedure and, thus, it can be implemented easily in Prolog environments.

To our knowledge, there is no similar work to ours in the deductive approach, although there exists some similar research in the context of the operational approach. We can mention here the recent work from the DAIDA project, reported in [CKM91], where additional references to previous research can be found. DAIDA proposes a dependency-based framework for the mapping from a requirements specification into a system design. The framework is dependency-based in the sense that the mapping of parts of the requirements specification is guided by predefined allowable dependencies. At the same time, the framework is goal-oriented, in the sense that non-functional requirements are treated as possibly conflicting goals to be satisfied by the generated design.

In DAIDA, requirements specification prescribe not only the behaviour of the system, but also the environment within which it will function. Instead of this, we focus only on the system to be developed, and our specifications are executable. This allows us to automate part of the mapping from specification to design. On the other hand, we have not considered yet the role of non-functional requirements in the generation of designs.

The paper is organised as follows. Section 2 briefly introduces the main components of a DCM, including an example that is used throughout the paper. Section 3 presents the Internal Events Model (IEM), a key concept in our approach to design generation. Section 4 discusses the use of the IEM in transaction generation, and gives a formal method for deriving transaction specifications from a DCM. Finally, Section 5 summarises the results of our work and points out future research.

2 Deductive Conceptual Models

We characterise the deductive conceptual modelling approach in a first order logic framework.

Time plays a major role in this approach. Every possible information i is associated with a time point $T(i)$, which states when the information holds. We will assume that times are always expressed in a unique time unit (such as second, day, etc.) small enough to avoid ambiguities. By life span T of an information system we mean the time interval in which the system operates. It is defined as an ordered set of consecutive time points $T = \{t_0 \dots t_f\}$, where t_0 and t_f are the initial and final times, respectively, and where each $t \in T$ is expressed in the given time unit. We can then say that, for any information i , $T(i) \in T$.

A deductive conceptual model (DCM) of an IS consists of five sets: A set B of base predicates, a set D of derived predicates, a set IC of integrity constraints, a set Q of predefined queries and a set A of alerts. In the following, we briefly describe each of these sets. Figure 1 shows an example that will be used throughout the paper.

```

base predicates
new_person(Person,Time)
new_subject(Subject,Time)
offer(Course,Title,Subject,Time)
enroll(Person,Course,Time)
cancel(Course,Time)

derived predicates
subject(S,T)  $\leftarrow$  new_subject(S,T1),  $T1 \leq T$ 
person(P,T)  $\leftarrow$  new_person(P,T1),  $T1 \leq T$ 
course(C,Ti,T)  $\leftarrow$  offer(C,Ti,S,T1),  $T1 \leq T$ ,
                    not  $\exists T2(\text{cancel}(C,T2), T2 > T1, T2 \leq T)$ 
subject_of_course(S,C,T)  $\leftarrow$  subject(S,T), course(C,Ti,T),
                    offer(C,Ti,S,T1),  $T1 \leq T$ 
enrolled(P,C,T)  $\leftarrow$  course(C,Ti,T), enroll(P,C,T1),  $T1 \leq T$ 
interested(P,S,T)  $\leftarrow$  enroll(P,C,T1),  $T1 \leq T$ , subject_of_course(S,C,T1)

integrity constraints
ic1(T)  $\leftarrow$  new_subject(S,T), subject(S,T-1)
ic2(T)  $\leftarrow$  new_person(P,T), person(P,T-1)
ic3(C,T)  $\leftarrow$  offer(C,Ti,S,T), course(C,Ti1,T1),  $T1 < T$ 
ic4(C,S,T)  $\leftarrow$  offer(C,Ti,S,T), not subject(S,T-1)
ic5(C,T)  $\leftarrow$  offer(C,Ti,S,T), offer(C,Ti1,S1,T),  $Ti \neq Ti1$ 
ic6(C,T)  $\leftarrow$  offer(C,Ti,S,T), offer(C,Ti1,S1,T),  $S1 \neq S1$ 
ic7(T)  $\leftarrow$  cancel(C,T), not  $\exists Ti(\text{course}(C,Ti,T-1))$ 
ic8(T)  $\leftarrow$  enroll(P,C,T), not person(P,T-1)
ic9(T)  $\leftarrow$  enroll(P,C,T), not  $\exists Ti(\text{course}(C,Ti,T-1))$ 
ic10(T)  $\leftarrow$  enroll(P,C,T), enrolled(P,C,T-1)

```

Figure 1. Example of Deductive Conceptual Model

Base predicates correspond to the external event types. They are the inputs to the IS. Each fact of a base predicate, called base fact, is an occurrence of an external event. We assume, by convention, that the last term of a base fact gives the time when the event occurred and was communicated to the IS. If $p(a_1, \dots, a_n, t_i)$ is a base fact we say that $p(a_1, \dots, a_n)$ is true or holds at t_i .

In the example of figure 1 we have five base predicates: *new_person*, *new_subject*, *offer*, *enroll* and *cancel*. A base fact *new_person*(p,t) means that at time t, p becomes a person. A base fact *new_subject*(s,t) reports that at time t, s becomes a new subject. A base fact *offer*(c,t,i,s,t) means that at time t, course c is offered with title ti and subject s. A base fact *enroll*(p,c,t) means that person p is enrolled to course c at time t. Finally, a base fact *cancel*(c,t) reports that course c is cancelled at time t. We take as time unit a second.

Derived predicates model the relevant types of knowledge about the Universe of Discourse. Each fact of a derived predicate, called derived fact, represents an information about the state of the UoD, at a particular time point. We also assume that the last term gives the time when the information holds. Thus, for example, a derived fact *person*(p,t) might mean that p is a person at time t.

Each derived predicate is defined by means of one or more deduction rules. A deduction rule of predicate p has the form $p(X_1, \dots, X_n, T) \leftarrow L_1, \dots, L_m$, where $p(X_1, \dots, X_n, T)$ is an atom denoting the conclusion and $L_1, \dots, L_m$ are literals representing conditions. Each L_j is either an atom or a negated atom. Variables in the conclusion or in the conditions are assumed to be universally quantified over the whole formula. The terms in the conclusion must be distinct variables, and the terms in the conditions must be variables or constants.

Condition predicates may be ordinary or evaluable ("built-in"). The former are base or derived predicates, while the latter are predicates, such as the comparison or arithmetic predicates, that can be evaluated without accessing a database.

We assume every rule to be allowed [GMN84], i.e. any variable that occurs in the rule has an occurrence in positive condition of an ordinary predicate. We also require every rule to be time-restricted. This means that for every positive literal $q(\dots, T1)$ of a base or derived predicate q occurring in the body, the condition $L_1, \dots, L_m \rightarrow T1 \leq T$ must hold. This ensures that $p(X_1, \dots, X_n, T)$ is defined in terms of q-facts holding at time T or before.

In the example, there are six derived predicates, with their corresponding (and hopefully self-explanatory) rules. For the sake of clarity, our examples do not follow strictly the above format, but can be transformed into it using the procedure given in [LIT84].

Integrity constraints are closed first-order formulas that base and/or derived facts are required to satisfy. We deal with constraints that have the form of a denial $\leftarrow L_1, \dots, L_m$, with $m \geq 1$, where the L_j are literals, and variables are assumed to be universally quantified over the whole formula. More general constraints can be transformed into this form as described in [LIT84]. For the sake of uniformity, we associate to each integrity constraint an inconsistency predicate *icn*, with at least a time term, and thus it has the same form as the deductive rules. We call them integrity rules.

In the example of figure 1 we show ten inconsistency predicates, with their rules. To see how an inconsistency may arise, assume that base facts *new_person*(john,10), *new_subject*(maths,10), *offer*(c1,ti,maths,15) and *enroll*(john,c1,18) were received at times 10, 15 and 18, respectively. Thus, the facts that hold at time 19 are:

person(john,19), subject(maths,19), course(c1,ti,19), enrolled(john,c1,19) and interested(john,maths,19). Now, if at time 20, the IS receives offer(c1,ti,electronics,20) the inconsistency facts ic3(c1,20) and ic4(c1,electronics,20) will hold, because course(c1,ti,19) does hold and subject(electronics,19) does not hold, respectively. Therefore, the base fact offer(c1,ti,electronics,20) would be rejected.

Outputs from an IS may be requested by the users (queries) or triggered internally by the system when some condition holds (alerts). Each query is defined by a name, a number of parameters, to which the user will give values when he/she makes the query, and a body. The answer to a query is the set of values that satisfies the conditions given in the body. (In order to focus on our objective, we omit in this paper output definition and handling).

As can be observed, there is a strong similarity in form between a DCM and a deductive database. However, there are some fundamental differences between both. An explanation can be found in [Oli89]

3 The Internal Events Model

We have seen, in the previous Section, the main components of a DCM. Now, we start to describe our approach to the design and implementation of an IS from its DCM. The key concept of our approach is the Internal Events Model (IEM)[Oli89,San90]. In the following, we briefly describe the main concepts of an IEM, and show its application to the example.

3.1 Classification of predicates

Predicates defined in a DCM can be classified according to their temporal behaviour. For our purposes, the most important classification is the following. Let p be a predicate and k a vector of constants. Assume that fact $p(k)$ holds at time $T-1$. What can we say about the truth of $p(k)$ at time T ? Three cases are possible:

- a) $p(k)$ will be true at time T . Then we classify p as *P-steady*.
- b) $p(k)$ will be false at time T . Then we classify p as *P-momentary*.
- c) $p(k)$ can be true or false at time T . In this case, assume that no external events happen at time T . We have three subcases:
 - c1) $p(k)$ will be true at time T . Then we classify p as *P-state*.
 - c2) $p(k)$ will be false at time T . Then we classify p as *P-transient*.
 - c3) $p(k)$ can be true or false at time T , depending on the truth value of some condition that must be evaluated at time T . Then we classify p as *P-spontaneous*.

Base predicates are assumed to be *P-transient*. In our DCM example of figure 1, predicates "subject", "person" and "interested" are *P-steady*. The other derived predicates are *P-state*.

3.2 Internal events

The concept of internal event tries to capture, in a natural way, the notion of change in the extension of a predicate. We associate to each predicate p an insertion internal event ip , and to each *P-state* or *P-spontaneous* derived predicate q a deletion internal event δq .

Insertion internal events are defined as follows. Let p be a *P-steady*, *P-state* or *P-spontaneous* predicate, then:

$$(1) \quad \forall X, T \quad (\text{ip}(X, T) \leftrightarrow p(X, T) \wedge \neg p(X, T-1))$$

If p is P-momentary or P-transient, then:

$$(2) \forall X, T \ (tp(X, T) \leftrightarrow p(X, T))$$

Observe that if p is a base predicate, tp facts represent external events (given by the environment) corresponding to insertions of base facts. If p is a derived predicate, then tp facts represent induced insertions of derived facts. Finally, if p is an inconsistency predicate, tp facts correspond to violations of its integrity constraint. In this case, since we assume that the IB is consistent at time $T-1$, $p(X, T-1)$ is always false, and the literal $\neg p(X, T-1)$ can be removed from (1).

We similarly define *deletion internal events*. Let p be a P-state or P-spontaneous derived predicate, then:

$$(3) \forall X, T \ (\delta p(X, T) \leftrightarrow p(X, T-1) \wedge \neg p(X, T))$$

If p is a derived predicate, then δp facts represent induced deletions of derived facts. Note that, for inconsistency predicates, δp facts cannot happen in any transition, since $p(X, T-1)$ is always false.

Rules (1), (2) and (3) are called *internal events rules*. An internal event rule is a rule that defines the conditions upon which an internal event happens in a given transition.

For example, the rules:

$$icourse(C, Ti, T) \leftarrow toffer(C, Ti, S, T)$$

$$\delta course(C, Ti, T) \leftarrow course(C, Ti, T-1), \cancel{cancel}(C, T)$$

are an insertion and a deletion internal event rule, respectively. The first rule states that the occurrence of an $toffer$ fact (in this case, the insertion of an offer base fact) induces a corresponding $icourse$ fact. The second rule states that the occurrence of a fact $\cancel{cancel}(c)$ at time t induces a fact $\delta course(c, ti)$ at time t if $course(c, ti)$ was true at previous time.

Internal events rules can be automatically obtained using a procedure based on a transformation of DCM rules (see [Oli89, San90] for details). These rules can be further simplified by taking into account the constraints [San93]. Part of the resulting IEM for the DCM example of figure 1 is shown in figure 2.

IDR.1	$isubject(S, T) \leftarrow unew_subject(S, T)$
IDR.2	$iperson(P, T) \leftarrow unew_person(P, T)$
IDR.3	$icourse(C, Ti, T) \leftarrow toffer(C, Ti, S, T)$
IDR.4	$\delta course(C, Ti, T) \leftarrow course(C, Ti, T-1), \cancel{cancel}(C, T)$
IDR.5	$isubject_of_course(S, C, T) \leftarrow toffer(C, Ti, S, T)$
IDR.6	$\delta subject_of_course(S, C, T) \leftarrow subject_of_course(S, C, T-1), \delta course(C, Ti, T)$
IDR.7	$ienrolled(P, C, T) \leftarrow course(C, Ti, T-1), not \delta course(C, Ti, T), ienroll(P, C, T)$
IDR.8	$\delta enrolled(P, C, T) \leftarrow enrolled(P, C, T-1), \delta course(C, Ti, T)$
IDR.9	$iinterested(P, S, T) \leftarrow ienroll(P, C, T), subject_of_course(S, C, T-1),$ $not \delta subject_of_course(S, C, T), not interested(P, S, T-1)$
IDR.10	$ic6(C, T) \leftarrow toffer(C, Ti, S, T), toffer(C, Ti1, S1, T), S1 \neq S1$
IDR.11	$ic8(T) \leftarrow ienroll(P, C, T), not person(P, T-1)$
IDR.12	$ic9(T) \leftarrow ienroll(P, C, T), not \exists Ti(course(C, Ti, T-1))$
IDR.13	$ic10(T) \leftarrow ienroll(P, C, T), enrolled(P, C, T-1)$

Figure 2. Part of the internal events model of the DCM example.

4. Deriving Base Transactions

4.1 Implementation of a DCM

Implementation of a DCM comprises two main design decisions: data base and transactions. The simplest implementation consists in storing in the data base (Extensional Data Base, EDB) only the base facts. There is a transaction (that we call Base Transaction, BT) for each base predicate. The role of each BT is to read the corresponding base fact, check the relevant integrity constraints and store the fact in the EDB. Derived facts could be computed using the DCM rules when requested, either during constraints checking or query answering. This implementation, however, would not be acceptable in practical situations from an efficiency point of view.

In general, the EDB comprises all facts explicitly stored in the data base. These facts do not need to be the base facts defined in the DCM. Indeed, it is quite usual to store in the EDB some derived facts, and not to store all base facts. Usually, there are many valid alternatives and the designer must choose in each case the most appropriate. The EDB schema is characterized by: (1) A set of base or derived predicates of the DCM to be included in the EDB (we call them *stored predicates*), and (2) for each of them, a time interval for which its facts will be stored. The Intensional Data Base, IDB, comprises all facts that can be derived from the EDB using the DCM rules.

Then, the role of each BT is to read the corresponding base fact, check the relevant integrity constraints and update the EDB (inserting and/or deleting one or more EDB-facts). In this section, we present a method for deriving the specifications of each BT. These specifications include two parts:

- (1) The transaction pre-conditions, which is the set of conditions the base fact must satisfy in order to guarantee EDB integrity.
- (2) The transaction updates, which are the insertions and/or deletions that must be done to the EDB.

The design and implementation of each BT will then be dependent on the particular execution environment. We will not deal with this aspect here, but we mention that it is quite straightforward to build such transactions on top of a deductive DBMS (see [San92, MSS92] for more details).

4.2 The approach

As we mentioned before, there is a BT for each base predicate defined in the DCM. We can also derive composite transactions, corresponding to the simultaneous occurrence of two or more base facts, but this extension will not be considered here.

Let $TR = \{ib(k, t)\}$ be the transaction corresponding to the insertion of a base fact $b(k, t)$, where k is a vector of constants denoting the transaction parameters, and t the occurrence time. We derive its effects in the following way: For each integrity constraint ic_j specified in the DCM, we check its violation by evaluating the internal event predicate tic_j . In the same way, for each stored predicate $p(X, T)$, we derive the induced insertion (resp. deletion) of a fact $p(x, t)$ by evaluating the internal event predicate $tp(X, t)$ (resp. $\delta p(X, t)$). As IEM rules are defined in terms of data base predicates extensions at time $t-1$ or before, to evaluate internal event predicates we also need the IDB rules, which relate DCM predicates to the facts stored in the EDB.

To obtain these evaluations we use an SLDNF-based proof procedure. More precisely, let $I = [i/\delta]p(X, t)$ be the internal event predicate to be evaluated. We say that

TR induces I if goal $\{\leftarrow I\}$ succeeds from input set $IEM \cup IDB \cup EDB \cup TR$. If every branch of the SLDNF-search space for $IEM \cup IDB \cup EDB \cup TR \cup \{\leftarrow I\}$ is a failure branch, then TR does not induce I.

Obviously, the evaluation can be completely done at execution time, when the concrete values of EDB and transaction parameters are known. However, we can do some preparatory work at compilation time, by partially evaluating $IEM \cup TR$ wrt I [LIS91]. The result of the partial evaluation is a set E of n rules ($n \geq 0$):

$$I \leftarrow C_i \quad i = 0..n$$

where C_i is a conjunction of literals such that the evaluation of $\{\leftarrow I\}$ at execution time from $E \cup IDB \cup EDB$ gives the same result as the evaluation of $\{\leftarrow I\}$ from $IEM \cup IDB \cup EDB \cup TR$. If $n=0$ then TR does not induce I. If C_i is empty then TR unconditionally induces a fact I. In this last case, if I corresponds to the insertion of an inconsistency predicate, the transaction can never happen because it violates the corresponding integrity constraint.

Applying this procedure to our example, at compilation time, assuming that the current state of all derived predicates are stored in the EDB, we have to partially evaluate, for each possible transaction, internal event predicates $uc1(t)$, $uc2(t)$, $uc3(C,t)$, $uc4(C,S,t)$, $uc5(C,t)$, $uc6(C,t)$, $uc7(t)$, $uc8(t)$, $uc9(t)$, $uc10(t)$, corresponding to the integrity constraints, and $isubject(S,t)$, $iperson(P,t)$, $icourse(C,Ti,t)$, $\delta course(C,Ti,t)$, $isubject_of_course(S,C,t)$, $\delta subject_of_course(S,C,t)$, $tenrolled(P,C,t)$, $\delta enrolled(P,C,t)$, and $tinterestedf(P,S,t)$ to derive the update conditions of stored predicates.

Take, as an example, the transaction: $TR = \{tenroll(p,c,t)\}$, where p and c are parameters and t the occurrence time of TR. After application of our method, the set of rules obtained at compilation time will be:

$$\begin{aligned} uc8(t) &\leftarrow \text{not person}(p,t-1) \\ uc9(t) &\leftarrow \text{not } \exists Ti(course(c,Ti,t-1)) \\ uc10(t) &\leftarrow \text{enrolled}(p,c,t-1) \\ tenrolled(p,c,t) &\leftarrow course(c,Ti,t-1) \\ tinterested(p,S,t) &\leftarrow subject_of_course(S,c,t-1), \text{not interested}(p,S,t-1) \end{aligned}$$

meaning that constraint ic8 will be violated if p was not a person at previous time, constraint ic9 will be violated if c was not a course at previous time and constraint ic10 will be violated if p is already enrolled to course c. If the constraints are satisfied, the fact $enrolled(p,c)$ must be inserted into the EDB if c was a course with title Ti at previous time, and the fact $interested(p,S)$ must be inserted into the EDB if it was not already true, being S the previous value of $subject_of_course(S,c)$. Note, however, that $tenrolled(p,c,t)$ can be unconditionally induced, instead of depending on the condition $course(c,Ti,t-1)$, because ic9 guarantees that there exists some $course(c,Ti,t-1)$.

4.3 Example

Before giving the formal definition of the method, we illustrate our approach with an example.

Assume that we want to derive the transaction $TR = \{tenroll(p,c,t)\}$ and that the current state of all derived predicates is stored in the EDB. We have first to partially evaluate $IEM \cup TR$ wrt ucj , corresponding to the integrity constraints. Consider, for example, the partial evaluation wrt $uc8(t)$. Possible sets of conditions C_i that would lead to the violation of $ic8(t)$ are obtained by having some failed derivation of $IEM \cup$

$TR \cup \{\leftarrow \text{tic8}(t)\}$ succeed. Figure 3 shows this derivation tree, where the circled labels are references to the rules of the method, defined in section 4.4.

Figure 3: derivation tree for tic8(t)

Steps 1 and 2 are SLDNF resolution steps. At step 1 rule IDR11 of the IEM acts as input clause, while at step 2 the input clause is the base fact from TR. At step 3 the selected literal $\text{not person}(p,t-1)$ corresponds to an stored predicate and, therefore, it can not be evaluated at compilation time because it depends on the concrete value of the database at execution time. As a consequence, this literal is included in the condition set C and its evaluation is delayed until execution time.

The set of conditions obtained for tic8(t) adds the following rule to the transaction definition:

$\text{tic8}(t) \leftarrow \text{not person}(p,t-1)$

Rules for $\text{tic1}(t), \dots, \text{tic10}(t)$ can be similarly obtained.

We also have to derive, for each stored predicate $p(X,T)$, the set of conditions to be checked at execution time, in order to induce an insertion or a deletion of a fact corresponding to this predicate, by partially evaluating $\text{IEM} \cup \text{TR}$ wrt $\text{tp}(X,t)$ and $\delta p(X,t)$. Consider, for example, the partial evaluation for $\text{tenrolled}(P,C,t)$. Possible sets of conditions C that would lead to the insertion of a derived fact $\text{enrolled}(p,c,t)$ are obtained by having some failed derivation of $\text{IEM} \cup \text{TR} \cup \{\leftarrow \text{tenrolled}(P,C,t)\}$ succeed. Figure 4 shows this derivation tree.

Step 1 is an SLDNF resolution step where rule IDR7 of the IEM acts as input clause. At step 2 the selected literal is an external event literal $\text{tenroll}(P,C,t)$ that can be resolved with the base fact in TR. At step 3 the selected literal is $\text{not } \delta \text{course}(c,T_i,t)$. In order to get a successful derivation, the SLDNF-search space must fail finitely for $\{\leftarrow \delta \text{course}(c,T_i,t)\} \cup \text{IEM} \cup \text{TR}$. The failure for $\delta \text{course}(c,T_i,t)$ is shown in figure 5. Finally, at step 4 the selected literal $\text{course}(c,T_i,t-1)$ corresponds to an stored predicate. Therefore, it is included in the condition set C and its evaluation delayed until execution time.

Figure 4: derivation tree for $\text{tenrolled}(P, C, t)$

Figure 5 shows the failure tree for $\delta\text{course}(c, Ti, t)$. Steps 1 and 2 are SLDNF resolution steps. At step 1 rule IDR4 of the IEM acts as input clause. At step 2, the selected literal is the external event $\text{tcancel}(c, t)$ which has to be resolved with facts in TR. Given that there is no fact $\text{tcancel}(c, t)$ in TR, the tree fails unconditionally.

Figure 5: failure tree for $\delta\text{course}(c, Ti, t)$

As a consequence, the following rule will be added to the BT definition:
 $\text{tenrolled}(p, c, t) \leftarrow \text{course}(c, Ti, t-1)$

Rules corresponding to the insertions and deletions of the other stored predicates can be similarly obtained.

4.4 The method

This section describes, in a formal way, the method explained in sections 4.2 and 4.3. Our proposal is an adaptation of the partial evaluation procedure given in [LIS91]. The main difference consists on the treatment given to the case in which a negative and non-ground literal is selected during the SLDNF refutation. In some cases, for our particular

application, the failure of such kind of literals can be guaranteed by the subsidiary derivation and, therefore, they can be evaluated at compilation time.

As we know, the transaction will be derived by partially evaluating $IEM \cup TR$ wrt internal events predicates corresponding to integrity constraints and stored predicates.

If I is the internal event predicate to be evaluated, C will be a set of conditions of I if there exists a **constructive derivation** from $(I \ \{ \})$ to $(\{ \} \ C)$, having as input set $IEM \cup TR$, where $TR = \{ \text{tb}(k, t) \}$, being b a base predicate and k a vector of constants.

We will call:

- **Non-event literal** to a literal corresponding to a base, derived or inconsistency predicate.

- **Accessible literal** to a non-event literal corresponding to an stored predicate or to a predicate that can be deduced from stored predicates.

Constructive derivation. A constructive derivation from $(I_1 \ C_1)$ to $(I_n \ C_n)$ via a selection rule R , that selects literals not corresponding to evaluable predicates with priority, is a sequence:

$(I_1 \ C_1), (I_2 \ C_2), \dots, (I_n \ C_n)$

such that for each $i \geq 1$, I_i has the form $\leftarrow L_1, \dots, L_k$, $R(I_i) = L_j$ and $(I_{i+1} \ C_{i+1})$ is obtained according to one of the following rules:

A1) If L_j is a positive internal event or transition literal, and S is the resolvent of some clause in IEM with I_i on the selected literal L_j , then $I_{i+1} = S$, and $C_{i+1} = C_i$.

A2) If L_j is a positive, external event literal $\text{tp}(X, t)$, and S is the resolvent of fact in TR with I_i on the literal L_j using substitution σ , then $I_{i+1} = S$, and $C_{i+1} = (C_i)\sigma$.

A3) If L_j is a negative, external event literal "not $\text{tp}(X, t)$ ", and TR can not be unified with $\text{tp}(X, t)$ then $I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i$.

A4) If L_j is a negative, external event literal "not $\text{tp}(X, t)$ ", and TR can be unified with $\text{tp}(X, t)$ using substitution $\sigma = \{ X_1/k_1 \}$, then $I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup \{ X_1 \neq k_1 \}$.

A5) If L_j is an accessible literal then $I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup \{ L_j \}$.

A6) If L_j is an evaluable literal then $I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup \{ L_j \}$.

A7) If L_j is a ground internal event or transition literal "not Q " and there exists a **consistency derivation** from $(\leftarrow Q \ \{ \})$ to $(\{ \} \ C')$ then

$I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup C'$.

A8) If L_j is a ground internal event or transition literal "not Q " and it does not exist a **consistency derivation** from $(\leftarrow Q \ \{ \})$ to $(\{ \} \ C')$ then

$I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup \{ L_j \}$.

A9) If L_j is a non-ground internal event or transition literal "not Q " and there exists a **consistency derivation** from $(\leftarrow Q \ \{ \})$ to $(\{ \} \ \{ \})$ then

$I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i$.

A10) If L_j is a non-ground internal event or transition literal "not Q " and it does not exist a **consistency derivation** from $(\leftarrow Q \ \{ \})$ to $(\{ \} \ \{ \})$ then

$I_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k$, and $C_{i+1} = C_i \cup \{ L_j \}$.

The step corresponding to rule A1) is an SLDNF resolution step. Note that we also handle transition literals, which have not been explained in this paper (see [San90]). The step of rule A2) is also an SLDNF step, but now L_j is resolved with TR . In rule A3) not $\text{tp}(X, t)$ is true because $\text{tp}(X, t)$ can not be unified with TR .

In A4) $\text{not } \text{up}(X,t)$ will be true only if $X_1 \neq k_1$. Assume, for example, that the current goal is $\leftarrow \text{not } \text{tenroll}(P,c,t), \alpha$ and that $\text{TR} = \{\text{tenroll}(p,c,t)\}$. By applying A4) we get the new goal $\leftarrow P \neq p, \alpha$.

At steps A5) and A6) accessible and evaluable literals are added to the condition set C because their evaluation have to be delayed until execution time. Steps corresponding to rules A7) and A8) deal with the evaluation of ground and negative internal event or transition literals "not Q". If the failure of Q can be ensured via a consistency derivation (A7), the set of conditions is added to the condition set C. On the contrary (A8), the literal itself has to be added to the condition set C because it can not be evaluated at compilation time. Finally, steps corresponding to rules A9) and A10) perform the evaluation of non-ground and negative internal event or transition literals "not Q". If a consistency derivation ensures an unconditional failure of Q (A9), the selected literal can not fail, and thus, it can be eliminated from the goal I in order to have a failed derivation of $\text{IEM} \cup \text{TR} \cup \{\leftarrow I\}$ succeed. On the contrary (A10), the literal has to be added to the condition set C because its evaluation has to be delayed until execution time.

There are different ways in which a constructive derivation can succeed. Each one may lead to different insertion or deletion rule.

Consistency derivation. A consistency derivation from $(F_1 \ C_1)$ to $(F_n \ C_n)$ via a safe selection rule R, that selects literals not corresponding to evaluable predicates with priority, is a sequence:

$(F_1 \ C_1), (F_2 \ C_2), \dots, (F_n \ C_n)$

such that for each $i \geq 1$, F_i has the form $\{\leftarrow L_1, \dots, L_k\} \cup F'_i$ and for some $j=1 \dots k$, $(F_{i+1} \ C_{i+1})$ is obtained according to one of the following rules:

B1) If L_j is a positive internal event or transition literal, and S' is the set of all resolvents of clauses in IEM with $\leftarrow L_1, \dots, L_k$ on the literal L_j , then $F_{i+1} = S' \cup F'_i$ and $C_{i+1} = C_i$.

B2) If L_j is a positive external event literal, and TR can not be unified with L_j then $F_{i+1} = F'_i$, and $C_{i+1} = C_i$.

B3) If L_j is a ground positive external event literal $\text{up}(x,t)$, and S' is the resolvent of TR with $\leftarrow L_1, \dots, L_k$ on the literal L_j , and $[\] \notin S'$, then $F_{i+1} = S' \cup F'_i$ and $C_{i+1} = C_i$.

B4) If L_j is a ground negative external event literal "not $\text{up}(x,t)$ " and TR can be unified with $\text{up}(x,t)$ then $F_{i+1} = F'_i$, and $C_{i+1} = C_i$.

B5) If L_j is a negative external event literal "not $\text{up}(X,t)$ ", and TR can not be unified with $\text{up}(X,t)$ then $F_{i+1} = \leftarrow L_1, \dots, L_{j-1}, L_{j+1}, \dots, L_k \cup F'_i$, and $C_{i+1} = C_i$.

B6) If L_j is a ground negative internal event or transition literal "not Q" and there exists a constructive derivation from $(\leftarrow Q \ \{ \})$ to $([\] \ C')$ then $F_{i+1} = F'_i$, and $C_{i+1} = C_i \cup C'$.

Steps corresponding to rules B1) and B3) are SLDNF resolution steps. In case B2) and B4) the current branch fails and thus, it can be eliminated. In case B5) the selected literal can not fail, and thus, it is eliminated from the subgoal set F_i in order to make a successful SLDNF branch fail. Finally, in case B6) the current branch can be dropped if there exists a constructive derivation for the negation of the selected literal. This ensures failure for it.

5 Conclusions

We have presented the main components of a Deductive Conceptual Model (DCM). We have then discussed the main decisions involved in the design and implementation process: data base design and transaction design. Data base design can not be completely automated, although some tools to aid the designer in alternatives' generation and analysis can be built.

Once the data base has been designed, transactions can be derived automatically. Preconditions of the transactions can be determined from the integrity rules, and database updates from the deductive rules of the DCM. We have presented a formal method to derive the transactions from the DCM. The method is based on the use of the SLDNF proof procedure and can be implemented easily in Prolog.

We plan to extend our work in several directions. First, we would like to be able to simplify even more the transaction specifications. Second, we plan to consider the case of transactions consisting of multiple base facts, instead of just a single one. Finally, it might be convenient to consider the design and implementation of transactions in conventional architectures. In this sense, our current transaction must be seen as a transaction specification.

Finally, we would like to build a library of rule schemes, such that a particular DCM, and its corresponding IEM, could be developed by reusing the components of the library.

Acknowledgements

The authors wish to thank D. Costal, C. Martin, E. Mayol, J.A. Pastor, C. Quer, J. Sistac, E. Teniente and T. Urpí for their comments and suggestions on earlier drafts. This work has been supported by the CICYT PRONTIC program, project TIC 680.

References

- [Bub86] Bubenko,J.A. "Information system methodologies - A research view". In [OSV86], pp. 289-318.
- [BuO86] Bubenko,J.A.;Olivé,A. "Dynamic or temporal modelling?. An illustrative comparison". SYSLAB Working Paper No. 117, University of Stockholm, 1986.
- [CKM91] Chung,L.;Katalagarianos,P.;Marakakis,M.;Mertikas,M.;Mylopoulos,J.; Vassiliou,Y. "From information system requirements to design: A mapping framework", Information Systems, Vol.16, No.4, 1991, pp. 429-461.
- [Fre85] Frenkel,K.A. "Toward automating the software development cycle". Comm. of the ACM, Vol.28,No.6,June 1985, pp. 578-589.
- [GKB82] Gustafsson,M.;Karlsson,T.;Bubenko,J.A. "A declarative approach to conceptual information modelling". In [OSV82], pp. 93-142.
- [GMN84] Gallaire,H.;Minker,J.;Nicolas,J-M. "Logic and databases: A deductive approach". ACM Computing Surveys, Vol. 16, No. 2, June 1984, pp. 153-185.
- [LIS91] Lloyd,J.;Shepherdson,J.C. "Partial evaluation in logic programming". J. Logic Programming, 1991, No.11, pp. 217-242.
- [LIT84] Lloyd,J.W.;Topor,R.W. "Making Prolog more expressive". J. Logic Programming, 1984, No.3, pp. 225-240.

- [MBJ90] Mylopoulos,J.;Borgida,A.;Jarke,M.;Koubarakis,M. "Telos: Representing knowledge about information systems", ACM Trans. on Information Systems, Vol.8,No.4, Oct.1990, pp 325-362.
- [MSS92] Mayol,E.; Sancho,M.R.; Sistac,J. "The ODISSEA project: An environment for the development of Information Systems from Deductive Conceptual Models". Proc. Intl. Workshop on the Deductive approach to IS and DB, Roses (Catalonia), 1992, pp. 17-47.
- [Oli82] Olivé, A. "DADES: A methodology for specification and design of information systems". In [OSV82], pp. 285-334.
- [Oli86] Olivé,A. "A comparison of the operational and deductive approaches to conceptual information systems modelling". Proc. IFIP 86, North-Holland, Dublin, 1986, pp. 91-96.
- [Oli89] Olivé,A. "On the design and implementation of information systems from deductive conceptual models". Proc. of the 15th. VLDB, Amsterdam, 1989, pp. 3-11.
- [OSV82] Olle,T.W.;Sol,H.G.;Verrijn-Stuart,A.A. (Eds.) "Information systems design methodologies: A comparative review". North-Holland, 1982.
- [OSV86] Olle,T.W.;Sol,H.G.;Verrijn-Stuart,A.A. (Eds.) "Information systems design methodologies: Improving the practice". North-Holland, 1986.
- [San90] Sancho,M.R. "Deriving an internal events model from a deductive conceptual model". Proc. Intl. Workshop on the Deductive approach to IS and DB, S'Agaró (Catalonia), 1990, pp. 73-92.
- [San92] Sancho,M.R. "The ODISSEA approach to the design of Information Systems from deductive conceptual models", Proc. of the IFIP World Congress, Madrid, September 1992, vol. 3, pp. 182-188.
- [San93] Sancho,M.R. "Deriving transactions from deductive conceptual models", PhD. thesis. In preparation.