

# Thémis: A Database Programming Language Handling Integrity Constraints

Véronique Benzaken and Anne Doucet

*Received June, 1993; revised version received, June, 1994; accepted March, 1995.*

**Abstract.** This article presents a database programming language, Thémis, which supports subtyping and class hierarchies, and allows for the definition of integrity constraints in a global and declarative way. We first describe the salient features of the language: types, names, classes, integrity constraints (including methods), and transactions. The inclusion of methods into integrity constraints allows an increase of the declarative power of these constraints. Indeed, the information needed to define a constraint is not always stored in the database through attributes, but is sometimes computed or derived data. Then, we address the problem of efficiently checking constraints. More specifically, we consider two different problems: (1) statically reducing the number of constraints to be checked, and (2) generating an efficient run-time checker. Using simple strategies, one can significantly improve the efficiency of the verification. We show how to reduce the number of constraints to be checked by characterizing the portions of the database that are involved in both the constraints and in a transaction. We also show how to generate efficient algorithms for checking a large class of constraints. We show how all the techniques presented exploit the underlying type system, which provides significant help in solving (1) and (2). Last, the current status of the Thémis prototype is presented.

**Key Words.** Database programming languages, integrity constraints, program analysis.

## 1. Introduction

Research in database programming languages has been devoted mainly to the definition of elaborated type systems and persistence mechanisms for those languages. The problems of polymorphism, static typing and inference, and object identity have been the main topics (Cardelli 1984, 1987, 1988; Cardelli and Wegner, 1985; Atkinson

---

Véronique Benzaken, Ph.D., is Associate Professor, Université de Paris I-Sorbonne, 12 place du Panthéon, 75005 Paris, France, and Researcher at LRI, bat 490, Université de Paris XI, 91405, Orsay, France, [benzaken@lri.lri.fr](mailto:benzaken@lri.lri.fr); and Anne Doucet, Ph.D., is Professor, Univeristé de Paris VI, Place Jussieu, LAFORIA, 75005 Paris, France, [doucet@laforia.ibp.fr](mailto:doucet@laforia.ibp.fr).

and Buneman, 1987; Hull et al., 1989; Castagna, 1995a, 1995b; Castagna et al., 1995).

In general, database programming languages are not able to express integrity constraints in a global and declarative way, although some interesting work has been done in the context of object-oriented databases (Martin, 1991).

A first specification of a language able to express integrity constraints has been proposed (Benzaken et al., 1992; Benzaken and Doucet, 1993). However, the class of constraints expressible by this language is restricted to first-order logic, well-typed formulas. Some derived data (computed attributes such as the age of a person, given the birth date, or the computation of the incoming and outgoing degrees of a directed graph) cannot be defined by first-order logic. In object-oriented languages, the only way to express derived data is by using a set of operations called methods (Section 3.1) To precisely capture the semantics of an application, some integrity constraints must consider derived data. Thus, it is necessary to introduce method calls in the language used to express constraints.

Relational and extended relational systems take integrity constraints and views into consideration (Stonebraker, 1975; Gardarin and Melkanoff, 1979; Weber et al., 1983; Sheard and Stemple, 1989). These systems provide models in which relation attribute domains are not necessarily atomic, but can be constructed using abstract types. The associated query language also can be extended to manipulate these user defined types instances. However, extended relational systems are not integrated in the sense of database programming languages. In these systems, relations are a very special kind of data type that cannot be used orthogonally to the others. In most systems, sets cannot be constructed independently of relations, and the query languages are not integrated within the language used to define the new attributes domains.

A second approach assumes transactions to be provided with the *atomicity* property, and consists of restricting the constraints to be enforced and of avoiding a retest of the portion of the database that is known to be consistent after the execution of the transaction (Nicolas, 1979; Hsu and Imielinski, 1985).

In the deductive database field, the problem of integrity constraint checking has been fully investigated (Bry and Manthey, 1986; Kowalski et al., 1987; Bry et al., 1988). Most of the techniques proposed are based on the Linear resolution with function Selection on Definite clauses with Negation as Failure (SLD/SLDNF resolutions) and theorem proving.

The work described by Sheard and Stemple (1989) consists of proving at compile time that database transactions respect integrity constraints, to reduce the overhead of unnecessary runtime tests. Their framework is the relational model. Transactions are complex updates of multiple relations, and constraints can be functional dependencies, inclusion dependencies, aggregate constraints, intersection dependencies, and inter-relational redundancies. Sheard and Stemple (1989) used the axiomatic semantics method, in which properties about language constructs are defined. These properties are found by using axioms and inference rules. Inference rules are

re-write rules on functional expressions of theorems, which allow the reduction of these expressions to true by using axioms, function definitions, and previously proven theorems. This leads to a formal proof of the property. The system uses a mechanical theorem prover in higher order computational logic to build a formal theory about database systems. That theory is extended to a specific database by generating specific knowledge from the structures and constraints contained in the schema. It is finally used by a transaction safety verifier.

We first describe the salient features of the Thémis language: types, names, classes, integrity constraints (including methods), and transactions. Then, we consider two different problems: (1) statically reducing the number of constraints to be checked, and (2) generating an efficient run time checker. Of course, in the general case, the problem is very complicated, and finding an optimal solution to (1), for instance, is undecidable. What we want to show is that, using simple strategies, we can significantly improve the efficiency of the verification. In this article, we suppose that transactions can be neither nested nor call other transactions. The general case will be the topic of a forthcoming study.

Our main goal is to fully exploit the type information to simplify constraint violation detection, and to speed up constraint checking. Not only are classes partially ordered according to an inheritance hierarchy, but we also have to face the problem of constraint checking in an environment that allows updates to be propagated among several distinct paths among objects. A first part of the article consists of using simple compilation techniques to statically determine which constraints might be violated by a transaction. The originality of this static analysis is that it captures the notion of inheritance and subtyping, and of late binding.

A second contribution consists of generating a checking algorithm from a transaction and a (restricted) constraint, which will operate on the smallest portion of the database involved by the transaction. We show how to significantly reduce the number of checking operations to be performed, relying on the underlying typing information.

This article is organized as follows. In Section 2, we summarize the main techniques that have been developed in the domain. In Section 3, we describe the salient features of the Thémis language: types, names, classes, integrity constraints (including methods), and transactions. We also present a detailed example to motivate and illustrate our language and checking techniques. Then, we consider the two steps of the verification process. In Section 4, we use simple compilation techniques to statically determine which constraints might be violated by a transaction, thus reducing the number of constraints to be checked. In Section 5, we propose the generation of constraint checking algorithms for a special class of constraints (universally quantified formulas). These algorithms are shown to significantly improve naive checking methods. In Section 6, we describe the current implementation of the Thémis prototype, which allows us to validate our work. Section 7 contains some concluding remarks.

## 2. Related Work

Relational and extended relational systems generally handle integrity by means of triggers. Triggers allow a user-defined procedure to be executed when a predicate is satisfied. Integrity constraints can be seen as rules, but they do not perform database updates. They simply return an error condition when an attribute is incorrectly modified. Although rules (or triggers) allow integrity constraints to be specified in a declarative way, it is the responsibility of the application programmer to code the procedures that will guarantee database safety. In our approach, we propose that these checking procedures be automatically generated, thus relieving the programmer of such a task, and therefore enhancing his/her productivity. To ensure database integrity, the user only describes the constraints.

Hsu and Imielinski (1985) proposed another solution, which extends Blaustein's work (1981). The constraints they considered are closed formulas of tuple calculus in prenex normal form. Here, the simplification method consists of transforming a constraint into an AND-OR tree of constraints, which is simpler to evaluate (simpler means that the checking space is reduced). Indeed, instead of testing the constraint on all the data, they only consider the data that might affect the database consistency with respect to both constraints and transactions. Interesting data are either inserted tuples or deleted tuples (updates consisting of deletions followed insertions).

Constraint simplification is performed in three steps. The constraint is first transformed into an updated form, involving the updated data. The second step consists of applying decomposition rules to the prefix of the updated constraint. These rules, which take into account only some prefix patterns, are recursively applied to the constraint, and produce either a conjunction or a disjunction of new formulas. The third step consists of eliminating the subformulas that are known to be true.

Our work adapts and extends these techniques to the object-oriented framework. More precisely, we use a similar technique in the second phase of our checking process, namely the generation of checking algorithms.

The deductive framework is well suited to integrity constraint management (Bry and Manthey, 1986; Kowalski et al., 1987; Bry et al., 1988). Deductive databases are a set of facts associated with a set of rules, which represent derived data. Integrity constraints can be expressed in this formalism as rules. Such rules, of course, do not perform updates or generate new facts. In this context, two problems are addressed: satisfaction and satisfiability. Two kinds of updates are considered for the problem of satisfaction, which are the addition of a new fact and the deletion of an existing one. According to the update, a first step consists of detecting which integrity constraint might be affected. Then, the checking process operates on the facts contained in the database. Both steps are achieved using SLD/SLDNF resolution.

For the second problem, namely satisfiability, the update considered is the addition of a new constraint. To detect whether the constraint is consistent with respect

to the existing ones, the method aims at generating a finite model, independent of the existing instance. The method has been shown to be semi-decidable. In both cases, only a restricted set of integrity constraints is handled, all of which are either universally quantified or existentially quantified constraints.

### 3. Basic Concepts of the Thémis Language

In this section, we present the basic concepts of the Thémis language. These concepts are illustrated by detailed examples (Section 3.5). Thémis is a strongly and statically typed object-oriented database language. In Thémis, a schema is defined using abstract and concrete types, classes, integrity constraints, and transactions.

#### 3.1 Types

We consider a framework in which all database manipulations are strongly and statically typed. Let us suppose the existence of the set  $\mathcal{D}$  of atomic types containing *integer*, *string*, and *boolean*. Types can either be concrete types or abstract types.

**3.1.1 Concrete Types.** The set of expressions of concrete types, denoted  $\mathcal{T}_C$ , is built by induction in the following way:

- Basic types:  $\mathcal{D}$
- If  $t_1, \dots, t_n \in \mathcal{T}_C$  and  $a_1, \dots, a_n \in \mathcal{A}$  ( $a_i \neq a_j$  for  $i, j \in \dots, n, i \neq j$  and  $n \geq 1$ ) then  $[a_1 : t_1, \dots, a_n : t_n] \in \mathcal{T}_C$ ,  $\{t_1\} \in \mathcal{T}_C$ , and  $(t_1) \in \mathcal{T}_C$   
 where  $[ ]$ ,  $\{ \}$ , and  $( )$  denote the constructors *tuple*, *set*, and *list*, respectively, and  $\mathcal{A}$  denotes the set of attribute names.

Concrete type equivalence is structural. Subtyping of concrete types is structural and inferred, following the classical rules of Cardelli (1984). For instance, we have:

```
[num: integer, label: string] <- [num: integer]
```

Concrete type instances are *non shared*, *non mutable* values.

**3.1.2 Abstract Types.** Abstract types have names. An abstract type is composed of a structural part and a behavioral part. The structural part is similar to concrete types. The behavioral part is described by a set of operations, called *methods*. Methods are defined in the following section. Instances of abstract types are objects, and have an *identity*, which is independent of their value. These instances are mutable and may be shared values. Equality of instances of abstract types is identity. Equivalence of abstract types is name equivalence. Subtyping of abstract types is explicit. The subtyping relation is declared in the definition of the abstract type.

**3.1.3 Methods.** Methods describe the behavior of the objects. They are composed of a signature and an implementation (the body of the method). Methods are not considered here as first class objects, and thus cannot be passed as parameters of

other methods. Let  $m$  be a method defined for the abstract type  $T$ . We denote its signature by  $m@T(\tau_1, \dots, \tau_{n-1}) : \tau_n$ , where  $\tau_1, \dots, \tau_{n-1}$  represents its parameter types, and where  $\tau_n$  represents the result type.

Passing a message is denoted by  $o \leftarrow m(x_1, \dots, x_{n-1})$ . This means that the method  $m$  is sent to the object  $o$ , called the receiver of the message. The  $x_i$ 's denote the actual parameters of method  $m$ . We denote  $o \leftarrow m()$  when the method  $m$  has no parameters.

Message passing can be more complex and may consist of the passing of several messages. This is denoted by:  $o \leftarrow m_1(x_1^1, \dots, x_{n_1}^1) \leftarrow \dots \leftarrow m_k(x_1^k, \dots, x_{n_k}^k)$ . When a method is redefined in a subtype hierarchy, the corresponding signatures are constrained to be *covariant*.

### 3.2 Classes

Types are used to describe the components of a database. The database can be seen as a graph of interconnected objects and values. The persistent roots of this graph are classes. Persistence is achieved through reachability.

A class gathers the set of objects having the same characteristics and the same behavior. The notion of class is an extensional notion. It represents a collection of objects of one type (abstract or not) and is characterized by a name and the type of its elements.

Classes are organized in a subclass hierarchy. The semantics of the inheritance relation is inclusion.

### 3.3 Integrity Constraints

In our framework, integrity constraints are well-typed boolean expressions, built using the names and classes of the schema and general operators. More formally, terms are defined as follows:

- Constants (e.g., true, false, nil) are terms.
- Each variable  $x$  is a term.
- Let  $t$  be a term, let  $a$  be an attribute (and not an operation),  $t.a$  is a term (provided that  $t$  is a tuple-structured term with attribute  $a$ ).
- Let  $t$  be a term,  $x_1, \dots, x_n$  be variables; let  $m$  be a method,  $t \leftarrow m(x_1, \dots, x_n)$  is a term.
- Let  $t_1$  and  $t_2$  be two terms; let  $\theta$  be an arithmetical operator ( $+$ ,  $-$ ,  $*$ ,  $\div$ ),  $t_1 \theta t_2$  is a term.

An integrity constraint,  $A$ , is an expression of the form:

$$A = Qx_1 \in S_1, \dots, Qx_k \in S_k M(x_1, \dots, x_k)$$

where  $Q \in \{\forall, \exists\}$ ,  $S_j$  is a set-structured expression, and  $M(x_1, \dots, x_k)$  is a quantifier-free formula. Expression  $Qx_1 \in S_1, \dots, Qx_k \in S_k$  is usually referred to as the constraint *prefix*, while  $M$  denotes the *matrix* of the constraint. More precisely, formulas  $M$  are defined as follows:

- Let  $\theta$  be a comparator ( $=, \neq, <, >, \leq, \geq$ ), let  $x$  and  $y$  be two terms,  $x \theta y$  is an atomic formula,
- Each atomic formula is a formula,
- Let  $F$  and  $F'$  be two formulas,  $F \wedge F', F \vee F', \neg F$  and  $(F)$  are also then formulas.

The equality operator can be applied to values of any type. The other comparators can be applied to numbers and sets.

**3.3.1 Remarks and Restrictions.** The introduction of methods into integrity constraints allows us to increase the declarative power of these constraints. Indeed, the information needed in the definition of a constraint is not always stored in the database through attributes, but is sometimes computed or derived data. This happens for information requiring important computations, or for derived data structures that cannot be defined with the first-order logic (e.g., transitive closure).

To keep the declarative aspect of a constraint, a method cannot modify the data stored in the database, but it must be allowed to define *virtual* data (methods allowed in the definition of constraints are overloaded queries). This virtual data represents the *intensional structures* of the database.

A method can appear in both the prefix and the matrix of a constraint. The signature of a method appearing in the prefix of the constraint must return a set structured result. However, in a constraint, all quantified variables denote persistent data. Therefore, to keep this property, a method appearing in the prefix of a constraint must return a set of persistent data. Hence, the body of this method can only contain a set of selections over the classes.

### 3.4 Transactions

Transactions are provided with the atomicity property: a transaction is either completely executed, or not executed at all. This mechanism allows us to overcome some errors, and to provide consistent executions. For the sake of simplicity, in this article, we consider only simple “flat” transactions (a transaction that does not call other transactions). A transaction is syntactically defined as follows:  $T = \text{trans}(\tau_1, \dots, \tau_n)\Gamma$  where  $\tau_i \in \{\mathcal{T}_C \cup \mathcal{T}_A\}^1$  and  $\Gamma$  represents the set of all elementary statements of a transaction. This set is recursively defined as follows:

- *assignment*  
 $e_1 := e_2 \in \Gamma$   
 $e_1.a := e_2 \in \Gamma$  if  $(a \in \mathcal{A})$   
 where  $e_2$  represents any expression
- *method call*  
 $o \leftarrow m(x_1, \dots, x_n) \in \Gamma$

---

1.  $\mathcal{T}_A$  denotes the set of abstract types.

- *sequencement*  
 $\forall s_1, s_2 \in \Gamma$   
 $s_1; s_2 \in \Gamma$
- *conditional test*  
 $\forall s_1, s_2 \in \Gamma$   
 if ( $b$ ) then  $s_1$  else  $s_2 \in \Gamma$   
 where  $b$  denotes a boolean expression
- *iteration loop*  
 $\forall s \in \Gamma$   
 for ( $o$  in  $x$ )  $s \in \Gamma$   
 where  $x$  denotes a set expression, and  $o$  an element of  $x$
- *set operations*  
 insert  $o$  into  $x \in \Gamma$   
 drop  $o$  from  $x \in \Gamma$   
 where  $x$  denotes a set expression, and  $o$  an element of  $x$  for the *drop* instruction

### 3.5 Example

To illustrate the concepts of Thémis, we give an example, which will be used in the remainder of this article. Let us consider the types given in Figure 1. The type *Person* has five attributes (name, age, birthday, spouse, and children), and three methods. The method *descendants()* computes the graph representing all the descendants of a given person. The method *ancestor()* computes the graph representing the ancestors of a given person. The method *genealogy()* computes both the ancestors and the descendants of a person.

A graph is represented by a pair  $\langle V; E \rangle$ , where  $V$  is a finite set of vertices, and  $E$  is a finite set of edges, each edge being a pair of vertices.

The type *Matrix* is used as an alternative representation of a graph which simplifies the implementation of some algorithms. The type *Matrix* is a list of lists of booleans. The *closure()* operation returns a matrix representing the set of all possible paths in the graph. The *connected()* operation indicates if the graph is strongly connected or not. Finally, the *non\_circuit()* operation determines if the graph has a circuit or not.

For this schema, we define the classes and constraints described in Figure 2. Constraint  $A_1$  expresses that the descendants of a given person are represented by a directed acyclic graph, while constraint  $A_2$  expresses that the genealogy of a given person is a strongly connected graph. Constraint  $A_3$  states that the age of a person ranges between 0 and 130. Constraint  $A_4$  expresses that every person is either the spouse of his/her spouse or is not married, and constraint  $A_5$  expresses that every child must be younger than his (her) parents. Finally, constraint  $A_6$  expresses that every Ferrari is owned by an instance of *Persons* older than 40. In our schema, we define the transactions given in Figure 3.



**Figure 1. A Thémis schema**

```

type Person is abstract [
  name: string,
  age: integer,
  birthday: integer,
  spouse: Person,
  children: set {Person}]
  descendants(): Graph,
  ancestors(): Graph,
  genealogy(): Graph
end
type Graph is abstract [edges: Edge, vertices: Vertex]
  add_edge( $v_1, v_2$ : Vertex),
  delete_edge( $e$ : Edge),
  matrix(): Matrix
end
type Matrix is abstract ((boolean))
  closure(): Matrix,
  connected(): boolean,
  non_circuit(): boolean
end
type Vertex is abstract [num: integer, id: Person]
  incoming_degree( $g$ : Graph): integer,
  outgoing_degree( $g$ : Graph): integer
end
type Edge is abstract [vertex1: Vertex, vertex2: Vertex, weight: integer]
end
type Vehicle is abstract [name: string, owner: Person]
end

```

#### 4. Static Analysis of a Thémis Schema

To avoid checking unnecessary constraints, we want to be able to statically characterize the integrity constraints that *may be* violated by a given transaction. Because the problem of determining if a transaction definitely will violate a constraint is undecidable, we are only looking for the set of constraints that might be violated. To characterize this superset of constraints, for a given transaction, we consider the parts of the database that are dealt with in a given constraint, and/or involved in a given transaction. A syntactic analysis of the constraints and the transactions has been defined (Benzaken et al., 1992; Benzaken and Doucet, 1993). Such an analysis consists, informally, of a set of *paths* in the database, gathering the set of classes and attributes used in the constraints and the transactions.

## Figure 2. Classes and Integrity Constraints

**class** Persons **of type** Person

**class** Vehicles **of type** Vehicle

(A<sub>1</sub>)  $\forall p \in \text{Persons}, p \leftarrow \text{descendants}() \leftarrow \text{matrix}() \leftarrow \text{closure}() \leftarrow \text{non\_circuit}();$   
 (A<sub>2</sub>)  $\forall p \in \text{Persons}, p \leftarrow \text{genealogy}() \leftarrow \text{matrix}() \leftarrow \text{closure}() \leftarrow \text{connected}();$   
 (A<sub>3</sub>)  $\forall p \in \text{Persons}, p.\text{age} \leq 130 \wedge p.\text{age} \geq 0;$   
 (A<sub>4</sub>)  $\forall p \in \text{Persons}, p.\text{spouse}.\text{spouse} = p \vee p.\text{spouse} = \text{nil};$   
 (A<sub>5</sub>)  $\forall p \in \text{Persons}, \forall c \in p.\text{children}, p.\text{age} > c.\text{age}$   
 (A<sub>6</sub>)  $\forall p \in \text{Persons}, \forall v \in \text{Vehicles}, (v.\text{name} \neq \text{"Ferrari"} \vee v.\text{owner} \neq p) \vee p.\text{age} \geq 40$

## Figure 3. Transactions

```
T1 = trans(p1. p2: Person) {
    insert p2 in p1.children }
/* this transaction adds a new child to a person */
T2 = trans(p1, p2: Person) {
    p1.spouse := p2;
    p2.spouse := p1 }
/* this transaction performs a marriage between two persons */
T3 = trans() {
    for p in Persons when (today = p.birthday) {
        print('Happy Birthday', p.name);
        p.age := p.age + 1 } }
/* this transaction updates the age of all Persons born on the current day */
```

The analysis proposed by Benzaken et al. (1992) and Benzaken and Doucet (1993) only considers the structure of the database, but does not take methods into consideration. The introduction of methods makes the situation much more complex. Indeed, the data structures they manipulate are not always explicitly present in the database, but can be defined only for computing purposes. Thus, a syntactic analysis of the methods will retrieve the set of “paths” that create these “temporary structures.”

In this section, we propose a structural and behavioral syntactic analysis of the constraints and transactions.

### 4.1 Syntactic Analysis of the Constraints

**4.1.1 Structure.** The structural analysis of the constraints is recursively defined as follows:

$\Upsilon(\text{exp1 } \theta \text{ exp2}) = \Upsilon(\text{exp1}) \cup \Upsilon(\text{exp2})$ , where  $\theta$  denotes any comparator;

$\Upsilon(C) = \bigcup_{C_i \leq C} \{C_i\}$  (the set of all subclasses of  $C$ , including itself).

$\Upsilon(\text{exp} \cdot \text{a}) = \bigcup_{t \leq \text{type}(\text{exp})} \{t \cdot \text{a}\} \cup \Upsilon(\text{exp})$ , if  $\text{type}(\text{exp})$  is an abstract type  
( $\text{type}()$  is a function which, given an expression, returns its corresponding

type);

$= \Upsilon(\text{exp})$  otherwise.

$\Upsilon(x) = \Upsilon(S)$ ,

where  $x$  represents a quantified variable in the constraint ranging over the set expression  $S$ ;

$\Upsilon(F \wedge F') = \Upsilon(F) \cup \Upsilon(F')$ , where  $F$  and  $F'$  are two quantifier-free formulas;

$\Upsilon(F \vee F') = \Upsilon(F) \cup \Upsilon(F')$ , where  $F$  and  $F'$  are two quantifier-free formulas;

$\Upsilon(\neg F) = \Upsilon(F)$ , where  $F$  is a quantifier-free formula;

$\Upsilon((F)) = \Upsilon(F)$ , where  $F$  is a quantifier-free formula.

**4.1.2 Behavior.** The syntactic analysis of a method requires more details. A method  $m$  of signature  $m@T(\tau_1, \dots, \tau_{n-1}):\tau_n$  is applied to an object  $o$  of abstract type  $T$  in the following way:  $o \leftarrow m(x_1, \dots, x_{n-1})$ . The syntactic analysis  $\Phi$  of the operations  $\Pi^m$  performed in  $m$  is defined as follows:

- **Assignment:** The syntactic analysis of an assignment “ $e_1 := e_2$ ” is the union of the syntactic analysis of the two expressions  $e_1$  and  $e_2$ .

$$\Phi(e_1 := e_2) = \Phi(e_1) \cup \Phi(e_2)$$

Indeed, in an assignment “ $e_1 := e_2$ ,” the left part ( $e_1$ ) cannot represent a persistent variable (a variable attached to a persistent root). However, the expression  $e_1$  can be built from a constructor using persistent parameters; therefore, we have to analyse the expression  $e_2$ .

- **Field extraction:**  $\Phi(\text{exp} \cdot \text{a}) = \Upsilon(\text{exp} \cdot \text{a})$  if  $\text{exp}$  is a variable attached to a persistent root.

While analyzing  $\text{exp} \cdot \text{a}$ , we do not take into account temporary variables (non-persistent variables), which appear in the bodies of the methods for computing purposes, nor instances of concrete types, because they are non-mutable, non-shared. In the case where  $\text{exp}$  is represented by the key word `self`, which refers to the receiver of the method,  $\text{type}(\text{self})$  is equal to the abstract type of the receiver.

- **Sequencement:** The syntactic analysis of a sequence of elementary statements is the union of the syntactic analysis of all these statements.

$$\Phi(s_1; s_2) = \Phi(s_1) \cup \Phi(s_2)$$

- **Iteration loop:** The syntactic analysis of an iteration loop is the union of the syntactic analysis of all the statements performed in the loop, and of the set on which the iteration holds.

$$\Phi(\text{for}(o \text{ in } x) s) = \Phi(x) \cup \Phi(s)$$

- **Conditional test:** The syntactic analysis of a conditional test is the union of the syntactic analysis of the boolean expression  $b$ , and of the statements  $s_1$  and  $s_2$ , which might be performed.

$$\Phi(\text{if}(b) s_1 \text{ else } s_2) = \Phi(b) \cup \Phi(s_1) \cup \Phi(s_2)$$

It is necessary to consider the syntactic analysis of the boolean expression (b) for the same reasons as those given above, concerning the iteration set in the case of the syntactic analysis of a conditional test.

- *Set operations:*

$$\Phi(\text{insert } o \text{ in } x) = \Phi(o) \cup \Phi(x).$$

$$\Phi(\text{drop } o \text{ from } x) = \Phi(o) \cup \Phi(x).$$

$\Phi$  allows the body of the methods invoked in the constraints to be analyzed. The methods to execute are dynamically linked to the selector (the name of the method), according to the abstract type of the receiver. For a given selector, it happens that there are several methods belonging to different abstract types, corresponding to different implementations. It is not always possible to determine at compile time which method to link. Thus, the same message passing may have different results, depending on the abstract type of the receiver.

To analyze a message passing  $o \leftarrow m$ , it is necessary to take into account the set of all methods that might be executed, according to the abstract type of the receiver. It is sufficient to consider, for each message passing  $o \leftarrow m$ , the set of all methods  $m$  declared in the subtypes  $\tau_i$  of the receiver  $\tau_o$ .

The syntactic analysis of a method call in a constraint corresponds to the syntactic analysis ( $\Phi$ ) of the operations invoked in this method:

$$\Upsilon(o \leftarrow m(x_1, \dots, x_n)) = \bigcup_{\tau_i \leq \tau_o} \Phi(\Pi^{m@ \tau_i})$$

The syntactic analysis of multiple calls is defined by:

$$\begin{aligned} \Upsilon(o \leftarrow m_1(x_1^1, \dots, x_{n_1-1}^1) \leftarrow m_2(x_1^2, \dots, x_{n_2-1}^2) \leftarrow \dots \leftarrow m_k(x_1^k, \dots, x_{n_k-1}^k)) = \\ \bigcup_{\tau_i \leq \tau_o} \Phi(\Pi^{m_1@ \tau_i}) \cup \bigcup_{\tau_i \leq \tau_{n_1}^1} \Phi(\Pi^{m_2@ \tau_i}) \cup \dots \cup \bigcup_{\tau_i \leq \tau_{n_k-1}^k} \Phi(\Pi^{m_k@ \tau_i}) \end{aligned}$$

$\bigcup_{\tau_i \leq \tau_{n_k-1}^k} \Phi(\Pi^{m_k@ \tau_i})$  represents the union of the syntactic analysis of the operations appearing in the methods  $m_k$  declared in the type  $\tau_{n_k-1}^k$  (abstract type of the result of the previous method  $m_{k-1}$ ) and in the subtypes of  $\tau_{n_k-1}^k$ .

**4.1.3 Example.** Table 1 represents the results yielded by  $\Upsilon()$  for the constraints defined in Figure 2. We detail the syntactic analysis of  $\mathcal{A}_1$ .

( $\mathcal{A}_1$ )  $\forall p \in \text{Persons}, p \leftarrow \text{descendants}() \leftarrow \text{matrix}() \leftarrow \text{closure}() \leftarrow \text{non-circuit}()$ . We have to analyze the following:

$$\Upsilon(p \leftarrow \text{descendants}() \leftarrow \text{matrix}() \leftarrow \text{closure}() \leftarrow \text{non-circuit}())$$

which is:

$$\bigcup_{\tau_i \leq \text{Person}} \Phi(\Pi^{\text{descendants}()@ \tau_i}) \cup \bigcup_{\tau_i \leq \text{Graph}} \Phi(\Pi^{\text{matrix}()@ \tau_i})$$

**Table 1. Syntactic analysis of constraints**

Constraint	Syntactic Analysis $\Upsilon(A)$
$A_1$	Persons, Person, Person.children
$A_2$	Persons, Person, Person.children
$A_3$	Persons, Person.age
$A_4$	Persons, Person.spouse
$A_5$	Persons, Person.children, Person.age
$A_6$	Persons, Vehicles, Person.age, Vehicle.owner, Vehicle.name

**Figure 4. Implementation of method descendants()**

```

{ Graph g;
Person x;
g = new (Graph);
for(x in self.children) {
    insert x ← descendants() in g;
    insert [vertex1: self, vertex2: x] in g.edges }
insert self in g.vertices;
return(g) }

```

$$\bigcup_{\tau_i \leq \text{Matrix}} \Phi(\Pi^{\text{closure}()} @ \tau_i) \quad \bigcup_{\tau_i \leq \text{Matrix}} \Phi(\Pi^{\text{non-circuit}()} @ \tau_i)$$

which yields:

$$\Phi(\Pi^{\text{descendants}()} @ \text{Person}) \cup \Phi(\Pi^{\text{matrix}()} @ \text{Graph}) \cup \\ \Phi(\Pi^{\text{closure}()} @ \text{Matrix}) \cup \Phi(\Pi^{\text{non-circuit}()} @ \text{Matrix})$$

Let us assume that method `descendants()` is implemented as shown in Figure 4. The analysis of the loop leads to  $\Upsilon(\text{self.children})$ , which is `Person.children`. The analysis of the code in the loop reduces to the analysis of `g.edges`, which is empty since `g` is a temporary variable, and of `[vertex1: self, vertex2: x]` which is `Person`. The analysis of the last assignment is again `Person`.

Furthermore,

$$\Phi(\Pi^{\text{matrix}()} @ \text{Graph}) \cup \Phi(\Pi^{\text{closure}()} @ \text{Matrix}) \cup \Phi(\Pi^{\text{non-circuit}()} @ \text{Matrix}) = \emptyset$$

because no persistent variable (no variable that might be attached to a persistent

root) appears in these methods.

The syntactic analysis  $\Upsilon(A_1)$  of the constraint  $A_1$  is thus:

{Persons, Person, Person.children}.

## 4.2 Syntactic Analysis of Transactions

The syntactic analysis of a transaction  $T$  yields the set  $\Psi(T)$  of all paths involved in  $T$ . This set again represents “paths” in the database, that is, those paths corresponding to data that might be modified by a transaction.

The syntactic analysis of transactions containing no methods is recursively built as follows :

$$\begin{aligned}\Psi(s_1; s_2) &= \Phi(s_1) \cup \Phi(s_2) \\ \Psi(\text{if}(b) s_1 \text{ else } s_2) &= \Phi(\text{if}(b) s_1 \text{ else } s_2) \\ \Psi(\text{for}(o \text{ in } x) s) &= \Phi(s) \\ \Psi(\text{exp. a}) &= \Upsilon(\text{exp. a})\end{aligned}$$

In the analysis of  $\Psi(\text{exp. a})$ , we do not take into account non-persistent temporary variables that appear in the body of the methods for computing purposes.

$$\Psi(e_1 := e_2) = \Upsilon(e_1)$$

In an assignment “ $e_1 := e_2$ ,” the left part ( $e_1$ ) can represent a variable attached to a persistent root.

$$\begin{aligned}\Psi(\text{insert } o \text{ in } x) &= \Phi(x) \\ \Psi(\text{drop } o \text{ from } x) &= \Phi(x) \\ \Psi(o \leftarrow m_1(x_1^1, \dots, x_{n_1-1}^1) \leftarrow m_2(x_1^2, \dots, x_{n_2-1}^2) \leftarrow \dots \leftarrow m_k(x_1^k, \dots, x_{n_k-1}^k)) \\ &= \Phi(o \leftarrow m_1(x_1^1, \dots, x_{n_1-1}^1) \leftarrow m_2(x_1^2, \dots, x_{n_2-1}^2) \\ &\quad \leftarrow \dots \leftarrow m_k(x_1^k, \dots, x_{n_k-1}^k))\end{aligned}$$

**4.2.1 Example.** The analysis of transactions  $T_1$ ,  $T_2$ , and  $T_3$  yields the results given in Table 2.

## 4.3 Safety Detection

While analyzing transactions and constraints, we detected the set of “paths” used in a constraint or invoked by a transaction. This set of paths gathers the various structures that are manipulated by the transactions and the constraints. Intuitively, a transaction  $T$  might violate a constraint  $A$ , if  $T$  and  $A$  manipulate the same “data structures.” This property (Benzaken et al., 1992; Benzaken and Doucet, 1993) is still valid in this context. It can be expressed the following way:

*Property 1.* Given a transaction  $T$  and a constraint  $A$ , transaction  $T$  might violate constraint  $A$  if and only if  $\Upsilon(A) \cup \Psi(T) \neq \emptyset$ .

**Table 2. Syntactic analysis of transactions**

Transaction	Syntactic Analysis ( $\Psi(T)$ )
$T_1$	Person, Person.children
$T_2$	Person, Person.spouse
$T_3$	Persons, Person, Person.age

**Table 3. Constraints hit by transactions**

Transaction	Constraints
$T_1$	$A_1, A_2, A_5$
$T_2$	$A_4$
$T_3$	$A_3, A_4, A_5, A_6$

Table 3 gathers the set of the constraints that might be violated by a given transaction. Note that  $T_3$  is detected as a transaction that might (potentially) violate the constraint  $A_6$ , even if this will never happen, since this transaction only increments the attribute age of a Person instance. This is, indeed, one limitation of this approach. Such a problem will be solved using more powerful techniques, namely, abstract interpretation of programming languages (Cousot and Cousot, 1976).

## 5. Generation of Enforcement Tests

### 5.1 Restrictions

The problem addressed in this section is to generate, given a constraint, a checking algorithm that guarantees that either the constraint is satisfied at the end of the transaction or the transaction is aborted. In a first approach, we restrict our constraints to constraints without methods. We also impose some restrictions on the constraint prefixes allowed.

Indeed, incremental checking is not possible for every constraint. Existential quantifiers, for example, cannot be simply incrementally checked as illustrated by the following example:

$$\exists x \text{ in aSet}, M(x)$$

If a transaction removes an element  $\alpha$  from aSet, and  $M(\alpha)$  is true, then there is no simple and cheap way to ensure that there is still another element that satisfies the predicate. In the remainder of this section, we consider only universally quantified constraints.

### Figure 5. Integrity constraints

- (A<sub>3</sub>)  $\forall p \in \text{Persons}, p.\text{age} \leq 130 \wedge p.\text{age} \geq 0$ ;  
 (A<sub>4</sub>)  $\forall p \in \text{Persons}, p.\text{spouse}.\text{spouse} = p \vee p.\text{spouse} = \text{nil}$ ;  
 (A<sub>5</sub>)  $\forall p \in \text{Persons}, \forall c \in p.\text{children } p.\text{age} > c.\text{age}$ ;  
 (A<sub>6</sub>)  $\forall p \in \text{Persons}, \forall v \in \text{Vehicles}, (v.\text{name} \neq \text{“Ferrari”} \vee v.\text{owner} \neq p) \vee p.\text{age} \geq 40$

### 5.2 Constraint checking

The problem of efficiently checking a constraint at the end of a transaction consists of finding the minimal set of objects involved in the process of checking. Then, the constraint will be checked only on this set, which guarantees that data consistency is ensured at the end of checking. However, this set, unfortunately, is not always reachable at run time. To illustrate this, we use the following four constraints  $A_3$ ,  $A_4$ ,  $A_5$ , and  $A_6$ , together with transactions  $T_1$ ,  $T_2$ , and  $T_3$  as shown in Figure 5.

If we consider  $T_3$  for the first constraint, we just have to collect the identifiers of every person whose age is modified. The objects collected by this process correspond to the ideal relevant set of objects on which  $A_3$  has to be checked.

For the second constraint, when executing transaction  $T_2$ , the ideal relevant set is not so easy to obtain. This set consists of the identifiers of  $p_1$  and  $p_2$ , as well as the identifiers of  $p_1.\text{spouse}$  and  $p_2.\text{spouse}$  *before* the assignment. Indeed, we need to know the former spouses of  $p_1$  and  $p_2$ , because the constraint  $A_4$  will certainly be violated for them. Of course, collecting those identifiers requires that the constraint checking manager be provided with some kind of “intelligence.” This problem is addressed in Benzaken et al. (1995), and relies on abstract interpretation techniques.

For the third constraint, when executing  $T_3$ , we have no means to collect the parents of a child whose age has been modified, because we don’t have backward pointers or indexes.

As a consequence, we do not attempt to obtain the ideal set of relevant objects. At the same time, we do not assume the existence of special access structures like indexes or backward pointers. Instead, we address the problem of finding an efficient checking algorithm that can be applied to all constraints. For constraints such as  $A_3$ , the algorithm will operate on the ideal relevant set of objects; for other constraints,



we show that the checking algorithm improves the trivial approach, which consists of performing a whole scan on the populations involved in the constraints.

Let  $T$  be a transaction and let  $A$  be a constraint. We are looking for an algorithm that satisfies the following properties:

- The evaluation of this algorithm at the end of the transaction ensures that the constraint  $A$  is still satisfied.
- The evaluation of this algorithm is more efficient than the direct evaluation of  $A$ .

At execution time, the only objects that can be collected are those instances of abstract data types whose attributes, relevant with respect to the constraints, have been modified. It may be the case that such a set of objects exactly matches the ideal relevant set, as for constraint  $A_3$ . But, in general, the set obtained at execution time only intersects the relevant set, as for constraint  $A_4$ . Therefore, we propose that checking algorithms be generated, which allows us to test the constraint on the whole set of relevant objects, thus ensuring database consistency. As a consequence, we have to perform some additional work to get these objects.

To define these algorithms, let us introduce the following definitions.

*Definition 1:*  $\Delta^C$

Given a class  $C$ , we posit  $\Delta^C$  the set of instances of class  $C$  that have been created (and inserted in  $C$ ) by a transaction.

*Definition 2:*  $\Gamma_a^{\tau(x)}$

Given an iteration variable  $x$  of type  $\tau(x)$ , and an attribute  $a$  of  $x$ , we posit  $\Gamma_a^{\tau(x)}$  the set of instances of the abstract type  $\tau(x)$ , in which attribute  $a$  has been modified by a given transaction.

This set represents information on the updates that a transaction has made on the database. The constraints considered here have the following generic form:

$$\forall x_1 \in C_1, \forall x_{1,1} \in x_1.p_{1,1}, \dots, \forall x_{1,n_1} \in x_1.p_{1,n_1}, \\ \forall x_2 \in C_2, \forall x_{2,1} \in x_2.p_{2,1}, \dots, \forall x_{2,n_2} \in x_2.p_{2,n_2}, \dots, \\ \forall x_k \in C_k, \dots, \forall x_{k,n_k} \in x_k.p_{k,n_k}, M(x_1, x_{1,1}, \dots, x_k, \dots, x_{k,n_k})$$

where  $p_{i,j}$  denotes prefix paths.

### Figure 6. Generic checking algorithm

For each class  $C_i$ , we generate the following enforcement test:

$$\forall x \in \Delta^{C_i}$$

$$\text{check } [ \forall x_1 \in C_1, \dots, \forall x_{i+1} \in C_{i+1}, \dots,$$

$$M(x_1, \dots, x, \dots, x_{i+1}, \dots) ]$$

For each path  $x.a_1 \dots a_k$  (either in the prefix or in the matrix), we generate the following enforcement test:

$$\forall x \in C_i$$

$$\forall y \in \Gamma_{a_1}^\tau(x), \text{ if } y = x,$$

$$\text{check } [ \forall x_1 \in C_1, \dots, \forall x_{i+1} \in C_{i+1}, \dots,$$

$$M(x_1, \dots, x, \dots, x_{i+1}, \dots) ]$$

...

$$\forall y \in \Gamma_{a_k}^\tau(x.a_1 \dots a_{k-1}), \text{ if } y = x \dots a_{k-1},$$

$$\text{check } [ \forall x_1 \in C_1, \dots, \forall x_{i+1} \in C_{i+1}, \dots,$$

$$M(x_1, \dots, x, \dots, x_{i+1}, \dots) ]$$

For each path  $y_i.b_1 \dots b_l$  in the matrix (where  $y_i$  ranges in  $x.p_i$ ),

we generate  $\forall x \in C_i$

$$\forall y \in x.p_i$$

$$\forall z \in \Gamma_{b_1}^\tau(y), \text{ if } z = y,$$

$$\text{check } [ \forall x_1 \in C_1, \dots, \forall x_{i+1} \in C_{i+1}, \dots,$$

$$M(x_1, \dots, x, \dots, y, \dots, x_{i+1}, \dots) ]$$

...

$$\forall z \in \Gamma_{b_l}^\tau(y \dots b_{l-1}), \text{ if } z = y \dots b_{l-1},$$

$$\text{check } [ \forall x_1 \in C_1, \dots, \forall x_{i+1} \in C_{i+1}, \dots,$$

$$M(x_1, \dots, x, \dots, y, \dots, x_{i+1}, \dots) ]$$

Let  $x$  be a variable ranging over class  $C_i$ , and let  $y_1, \dots, y_n$  be variables ranging, respectively, over  $x.p_1, \dots, x.p_n$ , where  $p_i$  denotes a prefix path leading to a set structured component of  $x$ . In Figure 6, we show how to generate generic checking algorithms.

For a given constraint  $A$ , the enforcement test generation consists of generating the above tests for each class  $C_i$  involved in the constraint prefix. Let us illustrate this on the constraints,  $A_3, A_4, A_5$ , and  $A_6$ . For the constraint  $A_3$ ,

$$(A_3) \forall p \in \text{Persons}, p.\text{age} \leq 130 \wedge p.\text{age} \geq 0;$$

and the checking algorithm is shown in Figure 7.

This can be rewritten as shown in Figure 8. In this case, the set  $\Gamma_{\text{age}}^{\text{Person}}$  actually represents the relevant set of objects on which the constraint has to be checked. Indeed, this algorithm leads to a check of the constraint on the set  $\Gamma_{\text{age}}^{\text{Person}}$ , testing if each element belongs to the class Persons. Thus, we perform as many check operations as the minimal algorithm does. Note that the trivial algorithm would have performed as many checks as the number of elements in the class Persons.

**Figure 7. Algorithm for  $A_3$** 

$\forall x \in \Delta^{\text{Persons}}$ , check ( $A_3(x)$ )

$\forall x \in \text{Persons}$

$\forall y \in \Gamma_{\text{age}}^{\text{Person}}$ ,

if  $y = x$ , check ( $A_3(x)$ )

For the constraint  $A_4$

( $A_4$ )  $\forall p \in \text{Persons}$ ,  $p.\text{spouse}.\text{spouse} = p \vee p.\text{spouse} = \text{nil}$ ;

the checking algorithm is described in Figure 9. For this algorithm, we have to scan the whole class Persons and test whether an element of  $\Gamma_{\text{spouse}}^{\text{Person}}$  corresponds with either an instance of class Persons, or to the spouse attribute of a given instance of Persons.

For the constraint  $A_5$

( $A_5$ )  $\forall p \in \text{Persons}$ ,  $\forall c \in p.\text{children}$ ,  $p.\text{age} > c.\text{age}$

the checking algorithm is shown in Figure 10.

This algorithm iterates over three sets: Persons,  $\Gamma_{\text{age}}^{\text{Person}}$ , and  $\Gamma_{\text{children}}^{\text{Person}}$ . For each element  $x$  of Persons whose age has been modified, we have to check the constraint. For each element  $x$  of Persons, if the age of one of his/her children has been modified, we have to check whether the constraint is still valid. Last, for each element of Persons whose set of children has been modified, we also have to check the constraint.

Finally, for the constraint  $A_6$ ,

( $A_6$ )  $\forall p \in \text{Persons}$ ,  $\forall v \in \text{Vehicles}$ ,  $(v.\text{name} \neq \text{“Ferrari”} \vee v.\text{owner} \neq p) \vee p.\text{age} \geq 40$

The checking algorithm is illustrated by Figure 11. This algorithm can be rewritten as shown in Figure 12.

This last example deserves some comments: checking  $A_6$  means that we check  $A_6$  with respect to all the elements in either Vehicles or Persons. Therefore, some tests are redundant. When checking the set of Persons whose age has been modified, we consider all Vehicles, particularly those Vehicles whose name or owner attribute has been updated. In the second phase of the algorithm, we test the constraint for all updated Vehicles with respect to all Persons, including those whose age has been modified. To avoid such redundant tests, we refine this algorithm in the following way. In the previous examples (for the constraints  $A_3$  and  $A_6$ ), the checking algorithms could be rewritten in an optimized form. Such an optimization can take place only for the algorithms containing no navigation in the type structures. For example, it is not possible to optimize in the way the algorithm was generated for constraint  $A_5$ , because  $y$  has to range over  $x.\text{children}$ .

**Figure 8. Optimized algorithm for  $A_3$** 

$$\forall x \in \Delta^{\text{Persons}}, \text{ check } (A_3(x))$$

$$\forall x \in \text{Persons} \cap \Gamma_{\text{age}}^{\text{Person}}, \text{ check } (A_3(x))$$
**Figure 9. Checking algorithm for  $A_4$** 

$$\forall x \in \Delta^{\text{Persons}}, \text{ check } (A_4(x))$$

$$\forall x \in \text{Persons}$$

$$\quad \forall y \in \Gamma_{\text{spouse}}^{\text{Person}},$$

$$\quad \quad \text{if } y = x, \text{ check } (A_4(y))$$

$$\quad \forall y \in \Gamma_{\text{spouse}}^{\text{Person}},$$

$$\quad \quad \text{if } y = x.\text{spouse}, \text{ check } (A_4(y))$$
**Figure 10. Checking algorithm for  $A_5$** 

$$\forall x \in \Delta^{\text{Persons}}, \text{ check } (A_5(x))$$

$$\forall x \in \text{Persons}$$

$$\quad \forall y \in \Gamma_{\text{age}}^{\text{Person}},$$

$$\quad \quad \text{if } y = x, \text{ check } (A_5(y))$$

$$\quad \forall y \in \Gamma_{\text{children}}^{\text{Person}},$$

$$\quad \quad \text{if } y = x, \text{ check } (A_5(x,y))$$

$$\forall y \in x.\text{children}$$

$$\quad \forall z \in \Gamma_{\text{age}}^{\text{Person}},$$

$$\quad \quad \text{if } z = y, \text{ check } (A_5(x,y))$$

We now give a general optimized version of this class of algorithms (Figure 13). The union of  $\Gamma_i$  denotes the set of all instances of an abstract type whose attributes relevant to a given constraint have been updated. Such an optimized version prevents us from testing the same constraint on the same objects more than one time.

**6. Implementation**

Thémis is implemented on top of the  $O_2$  system, using a preprocessing approach. The  $O_2$  integrity preprocessor takes a schema written in Thémis, and produces an  $O_2$  schema and a set of  $O_2$  executable programs, which allows us to instantiate the constraints while preserving the inclusion semantics.  $O_2$  integrity is written in  $C^{++}$ , and uses *lex* and *yacc*.

**Figure 11. Checking algorithm for  $A_6$** 

$$\begin{aligned} &\forall x \in \Delta^{\text{Persons}}, \text{ check } (A_6(x)) \\ &\forall x \in \text{Persons} \\ &\quad \forall y \in \Gamma_{\text{age}}^{\text{Person}}, \\ &\quad \quad \text{if } y = x, \text{ check } (A_6(y)) \\ &\forall x \in \Delta^{\text{Vehicles}}, \text{ check } (A_6(x)) \\ &\forall x \in \text{Vehicles} \\ &\quad \forall y \in \Gamma_{\text{name}}^{\text{Vehicle}}, \\ &\quad \quad \text{if } y = x, \text{ check } (A_6(y)) \\ &\quad \forall y \in \Gamma_{\text{owner}}^{\text{Vehicle}}, \\ &\quad \quad \text{if } y = x, \text{ check } (A_6(y)) \end{aligned}$$
**Figure 12. Optimized checking algorithm for  $A_6$** 

$$\begin{aligned} &\forall x \in \Delta^{\text{Persons}}, \text{ check } (A_6(x)) \\ &\forall x \in \text{Persons} \cap \Gamma_{\text{age}}^{\text{Person}}, \\ &\quad \text{check } (A_6(x)) \\ &\forall x \in \Delta^{\text{Vehicles}}, \text{ check } (A_6(x)) \\ &\forall x \in \text{Vehicles} \cap (\Gamma_{\text{name}}^{\text{Vehicle}} \cup \Gamma_{\text{owner}}^{\text{Vehicle}}) \\ &\quad \text{check } (A_6(x)) \end{aligned}$$
**Figure 13. Optimized algorithms**

$$\begin{aligned} &\forall x \in C_1 \cap (\cup \Gamma_1) \\ &\quad \text{check } [ \forall x_2 \in C_2, \dots, \forall x_k \in C_k, \dots, \\ &\quad \quad M(x, \dots, x_2, \dots, x_k, \dots) ] \\ &\quad \dots \\ &\forall x \in C_i \cap (\cup \Gamma_i) \\ &\quad \text{check } [ \forall x_1 \in C_1 - (C_1 \cap (\cup \Gamma_1)), \dots, \\ &\quad \quad : \forall x_{i-1} \in C_{i-1} - (C_{i-1} \cap (\cup \Gamma_{i-1})), \dots, \\ &\quad \quad : \forall x_{i+1} \in C_{i+1}, \dots, \\ &\quad \quad M(x_1, \dots, x_2, \dots, x, \dots, x_k, \dots) ] \\ &\quad \dots \\ &\forall x \in C_k \cap (\cup \Gamma_k) \\ &\quad \text{check } [ \forall x_1 \in C_1 - (C_1 \cap (\cup \Gamma_1)), \dots, \\ &\quad \quad : \forall x_i \in C_i - (C_i \cap (\cup \Gamma_i)), \dots, \\ &\quad \quad : \forall x_{k-1} \in C_{k-1} - (C_{k-1} \cap (\cup \Gamma_{k-1})), \\ &\quad \quad M(x_1, \dots, x_2, \dots, x, \dots) ] \end{aligned}$$

## 6.1 Mapping Between Thémis and O<sub>2</sub>

In this section, we describe the mapping between the Thémis language and O<sub>2</sub>.

### 6.1.1 Atomic Types

Thémis	O <sup>2</sup>
int	integer
string	string
boolean	boolean
—	real
—	bits

**6.1.2 Type Constructors.** In the O<sub>2</sub> language, it is possible to define complex objects and values by using various constructors, as in Thémis.

Thémis	O <sub>2</sub>
$[a_1 : t_1, \dots, a_n : t_n]$	$\text{tuple}(a_1 : t_1, \dots, a_n : t_n)$
$\{t_1\}$	$\text{set}(t_1)$
$(t_1)$	$\text{list}(t_1)$

**6.1.3 Types and Classes.** In the O<sub>2</sub> language, the instances of a type are values, and the instances of a class are objects. These properties are offered in Thémis through concrete and abstract types.

Thémis	O <sub>2</sub>
Concrete type	type
Abstract type	class
Classes	named values

**6.1.4 Subtyping and Inheritance.** O<sub>2</sub> and Thémis follow the same subtyping rules:

- An explicit subtyping for abstract types (Thémis) and the classes (O<sub>2</sub>).
- An implicit subtyping for concrete types (Thémis) and types (O<sub>2</sub>).

## 6.2 Constraints and Transactions

The constraints defined in Thémis are instances of a predefined class “Constraint” in O<sub>2</sub>. The transactions are translated into O<sub>2</sub> transactions, and compiled by the O<sub>2</sub>C compiler. Each time a transaction is compiled, the O<sub>2</sub> Integrity preprocessor updates a global table describing which constraints might be violated by the transaction.

Meanwhile, the corresponding checking algorithms are generated at the end of the transaction.

The user can visualize the set of constraints defined on the schema, and the global table showing the constraints that will be checked for a given transaction. Each time a constraint is actually violated by the execution of a transaction, the user is warned and the transaction is aborted.

## 7. Conclusion

This work proposes a specification of a database programming language allowing for the definition of integrity constraints in a global and declarative way. The characteristics of the object-oriented data model, in particular, inheritance and subtyping, are taken into account. The language used to express the integrity constraints is not limited to first-order logic formulas, but also includes method calls. This allows an increased declarative power of the constraints.

To detect which constraints may be violated by a given transaction, we define a syntactic analysis of both the constraints and the transactions. This analysis takes into consideration the specificities of the object-oriented model, such as inheritance, subtyping, late binding, and the persistent nature of the data. It allows us to obtain a necessary and sufficient condition to determine at compile time if a transaction might violate a constraint. A second part of this work concerns the automatic generation of constraint checking algorithms at the end of transactions. Those algorithms are generated for a sub-class of formulas: universally quantified formulas.

A first prototype of the Thémis language has been implemented. This prototype allows the proposed analysis to be validated. We propose that our work be extended in the following directions:

The analysis proposed detects transactions as being (potentially) unsafe when they are actually safe. More generally, we would like to refine our static analysis by using abstract interpretation techniques.

To be able to generate an efficient constraints checker, we extend our checking algorithms to constraints including methods and existential quantifiers.

Finally, our last aim is to build a complete compiler for the Thémis language. Such a compiler should be implemented on a persistent object manager (e.g., O<sub>2</sub> Engine, Napier88 Store).

## Acknowledgments

We would like to thank P.Y. Policella and P. Tronowski for implementing the first Thémis prototype. We also greatly acknowledge the referees for their enlightening comments and helpful suggestions.

## References

- Atkinson, M. and Buneman, P. Types and persistence in database programming languages. *ACM Computing Surveys*, 0(0):00-00, 1987.
- Benzaken, V. and Doucet, A. Thémis: A database programming language with integrity constraints. *Proceedings of the Fourth International Workshop on Database Programming Languages, Workshop in Computing*, New York, 1993.
- Benzaken, V., Doucet, A., and Schaefer, X. Integrity constraint checking optimization based on abstract databases generation and program analysis. *Journal de l'Ingénierie des Systèmes d'Information*, 1(3):9-29, 1995.
- Benzaken, V., Lécluse, C., and Richard, P. Enforcing integrity constraints in database programming languages. *Proceedings of the Fifth International Workshop on Persistent Object Systems, Workshop in Computing*, Pisa, Italy, 1992.
- Blaustein, B.T. Enforcing database assertions. Ph.D. thesis, Harvard University, Computer Science Department, Cambridge, MA, 1981.
- Bry, F., Decker, H., and Manthey, R. A uniform approach to constraint satisfaction and constraint satisfiability in deductive databases. *Proceedings of the EDBT International Conference, LNCS 303*, Venice, Italy, 1988.
- Bry, F. and Manthey, R. Checking consistency of database constraints: A logical basis. *Proceedings of the VLDB International Conference*, Kyoto, Japan, 1986.
- Cardelli, L. A semantics of multiple inheritance. In: *Semantics of Data Types, LNCS 173*, Springer-Verlag, 1984, pp. 51-67.
- Cardelli, L. Basic polymorphic type checking. *Science of Computer Programming*, 8(2):147-172, 1987.
- Cardelli, L. Structural subtyping and the notion of power type. *ACM POPL International Conference*, San Diego, CA, 1988.
- Cardelli, L. and Wegner, P. On understanding types, data abstraction and polymorphism. *ACM Computing Surveys*, 17(4):310-440, 1985.
- Castagna, G. Covariance and contravariance: Conflict without a cause. *ACM Transactions on Programming Languages and Systems*, 17(3):220-237, 1995a.
- Castagna, G. A proposal for making O<sub>2</sub> more type safe. Rapport de Recherche liens-95-4, LIENS, March 1995b.
- Castagna, G., Ghelli, G., and Longo, G. A calculus for overloaded functions with subtyping. *Information and Computation*, 117(1):115-135, 1995.
- Cousot, P. and Cousot, R. Static determination of dynamic properties of programs. *Proceedings of the Second International Symposium on Programming, Location?*, 1976.
- Gardarin, G. and Melkanoff, M. Proving the consistency of database transactions. *VLDB International Conference*, Rio, Brazil, 1979.
- Hsu, A. and Imielinski, T. Integrity checking for multiple updates. *Proceedings of the ACM SIGMOD International Conference*, Austin, TX, 1985.
- Hull, R., Morrison, R., and Stemple, D., eds. *International Workshop on Database Programming Languages*. Salishan Lodge, OR, 1989.



- Kowalski, R., Sadri, F., and Soper, P. Integrity checking in deductive databases. *Proceedings of the VLDB International Conference*, Brighton, UK, 1987.
- Martin, H. Contrôle de la cohérence dans les bases objets: Une approche par le comportement. Ph.D. thesis, Université Joseph-Fourier—Grenoble I, 1991.
- Nicolas, J.M. Logic for improving integrity checking in relational databases. Technical report, ONERA-CERT, 1979.
- Sheard, T. and Stemple, D. Automatic verification of database transaction safety. *ACM Transactions on Database Systems*, 14(3):322-368, 1989.
- Stonebraker, M. Implementation of integrity constraints and views by query modification. *ACM SIGMOD International Conference*, San Jose, CA, 1975.
- Weber, W., Stugky, W., and Karzt, J. Integrity checking in database systems. *Information Systems*, 8(2):125-136, 1983.