

# A Logic Programming Framework for Modelling Temporal Objects

F. Nihan Kesim\*     Marek Sergot†

July 1993; December 1994; July 1995

## Abstract

We present a general approach for modelling temporal aspects of objects in a logic programming framework. Change is formulated in the context of a database which stores explicitly a record of all changes that have occurred to objects and thus (implicitly) all states of objects in the database. A snapshot of the database at any given time is an object-oriented database, in the sense that it supports an object-based data model. An object is viewed as a collection of simple atomic formulas, with support for an explicit notion of object identity, classes and inheritance.

The event calculus is a treatment of time and change in first-order classical logic augmented with negation as failure. The paper develops a variant of the event calculus for representing changes to objects, including change in internal state of objects, creation and deletion of objects, and mutation of objects over time. Implementation strategies for practical application to temporal object-oriented databases are sketched. The concluding sections present two natural and straightforward extensions, to deal with versioning of objects and schema evolution.

**Keywords:** Object-oriented databases, object versioning, deductive databases, temporal databases, temporal reasoning, event calculus, logic programming.

## 1 Introduction

The object-oriented and deductive approaches have generated considerable interest in the database and programming language fields. In databases, one of the main driving forces behind the recent interest in object-oriented languages is their support of a rich collection of data modelling and manipulation concepts. Another feature of the approach is the promise it shows in overcoming the so called impedance mismatch between programming languages used to code applications and database languages used to retrieve data. In parallel, the deductive approach has gained popularity as a candidate to solve this mismatch problem, since logic can be used as a computational formalism as well as a database specification and query language. A substantial amount of recent research has aimed at integrating

---

\*Department of International Relations, Bilkent University, Bilkent, Ankara, Turkey.  
nihan@bilkent.edu.tr

†Department of Computing, Imperial College of Science, Technology and Medicine, 180 Queen's Gate, London SW7 2BZ, UK. mjs@doc.ic.ac.uk

these two paradigms to provide a single powerful framework for future database systems. Although there is still no general agreement on how this integration should be carried out — some authors even argue that one cannot have a system that is both truly deductive and truly object-oriented because of the conceptual mismatch between *value-oriented* logic programming and the notion of *object* as imported from object-oriented programming — there have been some promising developments, especially in the emergence of logics for objects with identity and complex internal structure. Existing proposals are summarised in section 2.

Most of this work, however, has ignored dynamic aspects, that is to say, the complications that arise when objects evolve over time or mutate into objects with different internal structure. Work on the representation of temporal phenomena, on the other hand, has tended not to involve any explicit notion of ‘object’. In temporal databases, research is dominated by approaches based on the relational model, although there are some exceptions. Some references are provided at the end of this section. Outside the database field, in AI in particular, there is extensive work on temporal reasoning, but here again ‘fluents’ — the propositions whose truth value varies over time — are typically represented as ground terms of some first-order language.

In this paper, we address the representation of temporal information in object-oriented databases. We do this by developing a variant of the event calculus [46], which we call *Object-based Event Calculus* (or OEC in short), for describing and reasoning about changes to objects in a logic programming framework.

The event calculus was introduced in [46] as a general logic programming treatment of time and change. Events, which are taken as the primitive temporal notion, mark the occurrence of change, and initiate and terminate periods of time for which facts hold. Given a record of events that have happened or that may happen, the event calculus can be used to determine what facts hold at any given time, or to compute the periods — the maximal intervals of time — for which a fact holds continuously. In the standard versions these time-varying facts are represented as (ground) first-order terms. From the database perspective, they can be seen as tuples of relations: the event calculus is then a method of deriving, for every such tuple, the periods of time for which it holds (the ‘lifespan’ of the tuple in the terminology of [27]). A snapshot of what holds at any time has the form of a relational database.

The timestamping of relational tuples with intervals is a common technique in many temporal database systems. The main difference in the event calculus is that these intervals are not inserted and modified explicitly but are derived from the record of event occurrences as required: the events effectively give some semantic structure to the end-points of intervals. A similar idea, expressed in terms of an extended relational algebra, has recently been proposed in [71]. The record of event occurrences is there called a ‘journal’. Operators are proposed to derive what holds at intermediate times: a ‘history’ operator converts event data to intervals, and a ‘snapshot’ operator determines what holds at a given point in time.

The event calculus is also intended to contribute to the treatment of database updates (see in particular [44]). This is not an aspect that we shall discuss in this paper, however. Similarly, the event calculus has been extended and applied to the construction of temporal databases that support both ‘valid time’ and ‘transaction time’ [63, 64]; again this is a further development we do not undertake.

In this paper we construct a version of the event calculus — the OEC — for dealing with changes to objects. As in the original (relational) event calculus, the changing world is described in terms of a record of events (a ‘journal’), from which the OEC can reconstruct and access the states of objects at any time. We get a database in which all states of objects are stored (implicitly). A snapshot of this database at any given time is an object-oriented database — ‘object-oriented’ in the sense that it supports an object-based data model. For reasons explained later, we shall adopt the view of an object as a collection of simple atomic formulas with a standard first-order semantics.

This paper is an expansion and further development of our previous presentation [34] where we discussed the evolution of objects using the event calculus. We now develop the basic idea and explore other temporal aspects of objects as well.

The paper has three main components:

(1) Since an object, however understood, is a more complicated structure than a collection of relational tuples, several different kinds of change can be identified, each requiring its own treatment. We examine the main kinds of change in detail — in sufficient detail that the resulting formulation can be executed directly, as a Prolog program, say. The same problems arise whatever representational formalism is employed. The formulations proposed could be reconstructed, if preferred, in some other representational formalism, such as the situation calculus.

(2) In common with much current usage, the term ‘database’ is used in this paper to refer to a wide, loosely defined class of applications, not just to large-scale database systems, narrowly understood. In the first instance, the OEC is intended to be used in the construction of ‘database’ or ‘knowledge-base’ applications where the problems of scale and performance associated with large-scale database systems are not a major factor. Some examples are mentioned in the text. However, it is also our aim to develop the OEC as a basis for practical, large-scale temporal database systems. In this last respect we shall be concerned with explaining how previously proposed implementational strategies in the temporal database literature may be adapted for use with the OEC.

(3) The OEC’s mechanism for maintaining the state history of objects leads to the *versioning* of objects as a natural and straightforward development. Event descriptions can be used to keep parallel histories of objects, and these can be used to model multiple versions of the same object at a time.

The paper is organized as follows. Section 2 presents a brief survey of the existing work on reasoning with complex objects for the purpose of identifying, in section 3, the basic data model that will be supported by our object-based variant of the event calculus. The OEC itself is presented in three separate sections. In section 4 we present the basic formulation and discuss how it can be applied to describe changes in objects. In section 5 we address the mutation of objects, where objects are allowed to change their classes during their evolution. And in section 6 we extend the formulation to incorporate some other object-oriented features, specifically multi-valued attributes and methods for derived attributes. Practical considerations and implementation strategies for temporal databases are discussed in section 7. Sections 8 and 9 show how the OEC can be adapted in a natural way to deal with versioning of objects and schema evolution respectively.

The literature on temporal reasoning and temporal databases is very extensive and we do not attempt a full survey here. For various extensions and applications of the event

calculus see, e.g., [11, 15, 20, 23, 24, 51, 54, 59, 62, 65]. Comparisons of the event calculus with situation calculus are provided in [53] and [45]. For temporal databases, [43] provides a recent bibliography of work in this area together with pointers to previous bibliographies. The collection [69] gives an excellent overview of the main approaches and discusses many of the issues that are studied in this field. As already mentioned, most work in temporal databases has been undertaken in the context of the relational model. Exceptions include [21, 32, 56, 57, 68, 73]. Comparisons with other proposals and references to specific points are given as they arise in the text.

## 2 Complex Objects

The purpose of this section is to identify and motivate the choice of data model we have adopted for the OEC.

Although there has been much confusion and controversy about the meaning of object-orientation in general and object-oriented databases in particular, a number of concepts have emerged as characteristic of this approach. Several papers [6, 5, 67] have now proposed a set of base features for object-oriented databases, that is, databases which support an object-oriented data model. There is no single standard model, but there is a set of basic concepts common to all object-oriented programming and knowledge representation languages.

A great number of attempts have been made to use logic in establishing a formal semantics for object-oriented concepts. Some of the existing works take deductive databases as the basis and extend the existing systems with the concept of a structured object. Most of the work in this approach follows the research on non-1NF relations, in order to extend the data structures of logic programming with sets and complex terms [1, 14, 47, 74]. Others attempt to formalize the basic object-oriented concepts by developing a new logic to support various features of complex objects [8, 13, 37, 39, 49]. There is also another stream of work which approaches the problem from a programming language perspective. Here the aim is to extend the logic programming languages with some object-oriented features such as methods and message passing [19, 25, 50, 75]. These proposals are of less interest in the context of this paper since their primary concern is with programming constructs.

When we compare the existing work, we see that the semantics of a complex object differs widely. In the proposals which extend deductive databases with sets and complex terms a complex object is viewed as a tuple in a higher-order relation. In object logics a complex object is either an element in some partially ordered domain or a collection of simple atomic formulas. Below we summarize these different semantics of objects and assess them according to their ability to support the representation of changes to objects.

### 2.1 Higher-Order Relations

One way of incorporating complex objects is to extend predicate calculus to a higher-order logic so that the value of an argument of a predicate can also be a relation built by using tuple and set constructors. Several higher-order extensions of logic programming, such as LDL [14] and COL [1] have been proposed. They view a complex object as a tuple in a higher-order relation. For example the COL statement:

`person(john, 38, "London", {chess, tennis}, {tom, sue}).`

describes information about a person and his hobbies and children. Sets can be represented either explicitly as in this example or by data functions. In LDL, a *grouping* construct is used to construct a set by using a rule. For example, the rule

`children(X, <Y>) ← parent_of(X, Y).`

groups together all the children of each person. Here `<>` is the grouping operator.

These proposals can be characterized as attempts to incorporate some notions from the object-oriented paradigm, without compromising the goal of having the relational model as the basis of the extended model. Thus they are often said to be value based. Although higher-order logic provides a formal framework for nested relations and complex objects, it also has some disadvantages. The higher-order semantics of sets presents severe semantic problems for logic programs in these languages, and it is difficult to develop an efficient query evaluation in these approaches.

Another disadvantage is that representing a complex object by a nested tuple is practically not very convenient. Because of syntactical limitations (e.g. fixed arity) these languages do not provide access to sub-structures of complex objects in a homogeneous way. They are unsuitable for deductive retrieval at arbitrary levels. In the above example, to find the age of the person `john`'s child `tom`, one must start from the top predicate `person` and then continue down to the sub-structures where the required information can be found. In the case of updates, semantic problems also arise. For example if `john` develops a new hobby, adding the new information will yield a completely different tuple which does not have any semantic relationship with the original one. Omitting the information about address will produce a tuple of a completely different type.

## 2.2 Object Logics

The other main stream of work aims at developing a new logic to support various features of complex objects. It is argued that just as for relational databases, a logical framework can be established for object-oriented databases also. The underlying logic must be different from first-order predicate logic because most features of object-orientation require higher orderness. On the other hand it is desirable to have a logic with first-order properties: following the terminology of [12] the language (syntax) of an object model must be higher order to be able to manipulate such concepts while the semantics must be restricted enough to satisfy first orderness. A number of such object logics have been proposed.

The first work, influenced by the  $\psi$ -terms of LOGIN [4], was Maier's O-logic [49] which was later extended by a number of proposals, namely C-logic [13], extended O-logic [39], and F-logic [37]. From the object-oriented world these logics acquire the notion of object identities, complex objects, a mechanism for object classification and a structure for property inheritance. From the logic programming world they absorb the concepts of unification, answer substitution and a strategy for deductive query processing.

In these logics, an object is represented as a complex term in the language. For instance in O-logic [39], the `person` object illustrated in the previous section can be represented by a complex term as follows:

```

person : john[ age → 28,
               address → "London",
               hobbies → {chess, tennis},
               children → {person : tom, person : sue}]

```

Here **person** is the class name, **john** is the object identity and the labels denote attributes.

The syntax of these complex terms is influenced by the language LOGIN but their semantics is different. In LOGIN complex terms denote *types* and inheritance is incorporated into the unification algorithm. In the object logics, complex terms are formulas in their own right: written as a formula, a complex term asserts that an individual object with that structure exists. More complex formulas are built by combining object terms using the usual truth-functional connectives and quantifiers.

Although the syntax and the informal reading of complex terms are quite similar in all of these object logics, the precise semantics given to the complex terms varies. Extended O-logic and F-logic view an object term as an element in a partially ordered domain. Partial orderings on class names and object identities are defined and using these orderings a partial ordering over complex object terms is obtained and used to capture sub-object or sub-type relationships.

The major disadvantage of this approach is that the logic becomes more complicated as more features, such as methods and inheritance, are introduced. A natural concern is whether there might be an efficient evaluation procedure for queries. Another important concern, directly related to the main topic of this paper, is the ability of this approach to model the dynamic behaviour of objects. This question is not addressed in the current literature, and many of the difficulties seem not to have been anticipated. For instance, some common types of change would seem to require changes to the partial ordering on class names and object identities, and hence effectively to the language itself. It is difficult to see how such changes could be accommodated smoothly, and no suggestions on these points have appeared, to our knowledge.

C-logic [13] takes a different approach. Here, complex object descriptions are considered as collections or conjunctions of atomic properties. Each attribute label is viewed as a binary predicate and each class symbol as a unary predicate. An object with several labels can then be described as a conjunction of several atomic formulas. For example the term

$$\text{john}[\text{name} \Rightarrow \text{"John Smith"}, \text{age} \Rightarrow 28]$$

can be considered as

$$\text{john}[\text{name} \Rightarrow \text{"John Smith"}] \wedge \text{john}[\text{age} \Rightarrow 28]$$

or as

$$\text{name}(\text{john}, \text{"John Smith"}) \wedge \text{age}(\text{john}, 28)$$

in first-order logic.

This formula approach makes it possible to understand the semantics of complex objects within the predicate logic. Chen and Warren give a semantics to C-logic directly, and also by transformation to an equivalent first-order formula which uses unary predicates for types and binary predicates for attributes. This makes proof procedures and associated computational developments in first-order logic readily available for complex objects.

One advantage of the formula approach taken by C-logic is that it allows information about an entity to be specified and accumulated piecewise, which facilitates the update of subparts of an object independently of others. The explicit notion of object identity also makes sharing and updates easier. Adding new information about an object is just a matter of adding one or more binary predicates. The subparts of an object can be retrieved by using the identity of the object and the attribute name describing the subpart.

### 3 The Data Model

In this section we present the data model that the OEC will support.

The data model provides a basic set of features associated with structural object orientation: object identity, complex objects with both single-valued and multi-valued attributes, methods for derived attributes, classes, class hierarchies, and inheritance. This is the set of features identified as the essential ingredients of object-oriented data models in [5, 67], to which we have added derived attributes since they are useful in applications and can be supported straightforwardly. The treatment we adopt follows the formula approach exemplified by C-logic [13], as summarised in the previous section.

We view an object as a named collection of *object-attribute-value* triples. Every object is abstracted by a unique identity which distinguishes it from other objects. Following Kifer and Wu [39] we use individual terms to denote object identities. A term representing the object identity is composed of function symbols, constants and variables in the usual way. (We use the standard Prolog convention for constants and variables throughout the paper: strings beginning with an upper-case letter are variables.) For example `john`, `X`, `children(john, mary)`, `path(X,Y)` can all be terms denoting identities. The set of all ground identity terms plays a role analogous to that of the Herbrand Universe in classical logic. Function symbols are used to construct new object identities out of existing ones. The objects have attributes whose values can be other objects (or more precisely their identities).

Objects are organized into class hierarchies, defined explicitly by asserting `is_a` relationships among classes. A class denotes a set of object identities. Each class has a unique name to distinguish it from other classes. The class-subclass relation (`is_a`) is to be read as the subset relation: the set of objects represented by a class includes all the objects belonging to the subclass(es) of that class.

The relation between a class and its instances is represented by the `instance_of` relation. The set of instances of a class changes as new objects are created and cease to exist. This time-dependent behaviour of the `instance_of` relation will be discussed in section 4.5.

A class describes the internal structure of its instances by attribute names. This structure is asserted by the predicate `attribute`. A subclass inherits the structure of its superclass(es).




Figure 1: A simple class hierarchy

As an example consider the class hierarchy shown in figure 1. Classes **student** and **employee** are subclasses of **person**. The attributes common to all persons (i.e., **name**, **address**) are defined in the class **person** and are inherited by the subclasses. The subclasses also define additional (more specific) attributes. The class hierarchy is described as follows:

```

is_a(student, person).
is_a(employee, person).

attribute(person,name).
attribute(person,address).
attribute(student,section).
:
  
```

For the purposes of this example, we have assumed that all attributes are single-valued. As will be shown in the next section, the functionality constraint of such attributes is satisfied within the formulation of the OEC. The extension to allow multi-valued attributes in addition is straightforward but for explanatory purposes we leave this aside, together with methods for derived attributes, until section 6.

In order to formulate the inheritance of attribute names by the subclasses we define the predicate **attribute\_of** in the following way:

```

attribute_of(Class, X) ← attribute(Class, X).
attribute_of(Sub, X) ← is_a(Sub, Class), attribute_of(Class, X).
  
```

This formulation for objects and classes allows a very simple form of inheritance. It is limited to the subset relation between classes and monotonic inheritance of attribute names. Multiple inheritance without overriding can also be expressed by the **is\_a** predicate. This simple type of multiple inheritance causes no additional difficulty and is not mentioned again.

## 4 The Object-Based Event Calculus

In this and the following two sections we shall present the OEC, a version of the event calculus that supports the data model described in the previous section. Given a description of event occurrences (changes in the world) the OEC can reconstruct the state of any object in the database at any point in time. It can also be used to compute the period(s) of time for which an object ‘exists’ (its ‘lifespan’ [27]) and the periods of time



for which given attributes of objects have given values. For simplicity we shall assume for the time being that all attributes are single-valued and we shall ignore methods for derived attributes. These features of the data model will be re-introduced in section 6. The formulations presented in these sections can be executed as Prolog programs; however we shall comment separately on practical implementation strategies in section 7.

The OEC is based on a simplified, asymmetric case of the event calculus, where facts are assumed to persist forwards in time until they are terminated by some subsequent event. Correspondingly, the assimilation of events into the database is assumed to keep step with the occurrence of changes in the world. This is in contrast to the original formulation of the event calculus [46] which treats past and future symmetrically and can deal with the case where events are not necessarily recorded in the same order in which they actually occur.

This simplified version of the event calculus corresponds closely to updates in conventional databases [44]. It blurs the distinction between an event occurrence (a change in the world) and the recording of that event in the database. Accordingly, the database that is maintained by the OEC can be seen either as a historical or ‘valid time’ database recording the evolution of some set of objects in the world, or as a system in which all past states of an object-oriented database are accessible. (And if valid times and transaction times are distinguished but are exactly correlated, then it can be seen as a ‘degenerate bitemporal’ database [29, 30].)

It would be possible to construct a version of the OEC without these assumptions following the symmetrical treatment of past and future of the original event calculus [46]. It would also be possible to extend the treatment to support both ‘valid time’ and ‘transaction time’ as done by Sripada [63, 64] for the relational versions. We do not attempt these further developments in this paper. Similarly, although it is only the relative ordering of events that is significant in the event calculus we shall assume that the times of all event occurrences are given since this is often useful in practice.

We present the OEC in stages. We begin with the simplest kind of change, which is change to an existing object’s internal state.

## 4.1 Change of Internal State

The state of an object is determined by the values assigned to its attributes. Change in internal state corresponds to changing the value of any of the attributes. The basic idea in dealing with the evolution of an object over time is to parametrize its attributes with times at which these attributes have various values. Formulation of this idea within the spirit of the event calculus is straightforward. Events initiate and terminate periods of time for which a given attribute of a given object takes a particular value. Figure 2 shows the history of an employee object. Here **john** is the identity of the object and **rank**, **dept**, **age**, **address** are the attributes that are initiated to different values at different times. The object-based event calculus constructs such a state history of objects.

The effects of events are described by the predicates **initiates** and **terminates** by means of assertions (or more generally rules) of the form:

`initiates(EventType, Object, Attribute, Value).`




Figure 2: State History

(and similarly for **terminates**). For example an event of type ‘promote employee **X** to new rank **R**’ initiates a period of time for which employee **X** holds rank **R** and terminates whatever rank **X** held at the time of the promotion:

`initiates(promote(X, R), X, rank, R).`

Here `promote(X, R)` is a term representing the type of the event. Since we are dealing with single-valued attributes it is not necessary to specify explicitly that the old rank is terminated. The details are shown in a moment.

Given a fragment of data :

```
happens(promote(jim,assistant), 1986).
happens(promote(jim,lecturer), 1988).
happens(promote(jim,professor), 1991).
```

the OEC can be used to compute values of attributes of objects at given times, as in the following two queries:

```
?- holds_at(jim, rank, R, 1989).
?- holds_at(jim, Attr, Val, 1989).
```

The formulation of `holds_at` in terms of `initiates`, `terminates` and `happens` is shown presently.

The OEC can also compute the periods of time for which an object’s attributes have particular values. In the example, for instance, the query

```
?- holds_for(jim, rank, R, Period).
```

would generate the answers

```
R = assistant, Period = 1986–1988;
R = lecturer,   Period = 1988–1991;
R = professor,  Period = since(1991).
```

A term of the form  $T_s$ – $T_e$  denotes a time interval (closed on the left and open on the right) with start and end points  $T_s$  and  $T_e$  respectively;  $\text{since}(T_s)$  denotes an open-ended interval, the set of time points later than or equal to  $T_s$ . (Time points need not be years, as in this example.) Notice that we do *not* include in the time line any distinguished time point ‘now’ or ‘unchanged’ as seems to be common in many temporal database systems (see for example the collection of papers [69]).

The following is the basic formulation of the OEC to derive the value of an attribute of an object at a specific time:

$$\text{holds\_at}(\text{Obj}, \text{Attr}, \text{Val}, T) \leftarrow \begin{array}{l} \text{happens}(\text{Ev}, T_s), T_s \leq T, \\ \text{initiates}(\text{Ev}, \text{Obj}, \text{Attr}, \text{Val}), \\ \text{not broken}(\text{Obj}, \text{Attr}, \text{Val}, T_s, T). \end{array}$$

$$\text{broken}(\text{Obj}, \text{Attr}, \text{Val}, T_s, T) \leftarrow \begin{array}{l} \text{happens}(\text{Ev}^*, T^*), \\ T_s < T^* \leq T, \\ \text{terminates}(\text{Ev}^*, \text{Obj}, \text{Attr}, \text{Val}). \end{array}$$

Informally, these clauses may be read procedurally as follows: in order to find the value  $\text{Val}$  of an attribute  $\text{Attr}$  of an object  $\text{Obj}$  at time  $T$ , find an event  $\text{Ev}$  which happened before time  $T$  and initiated the value of that attribute; and then check that no other event which terminates that value has happened in the meantime. The interpretation of **not** as negation by failure in the last condition for **holds\_at** gives a form of default persistence: the value of an attribute is assumed to hold at all times after its initiation by event  $\text{Ev}$  unless there is information to the contrary.

The constraint that attributes are single-valued implies that the value of an attribute is terminated if an event initiates it to another value. This is represented by adding the following general rule:

$$\text{terminates}(\text{Ev}^*, \text{Obj}, \text{Attr}, \_) \leftarrow \text{initiates}(\text{Ev}^*, \text{Obj}, \text{Attr}, \_).$$

(The use of the anonymous Prolog variable ‘\_’ in this clause is just to cover the unlikely case that an event is specified to re-initialise an attribute to its existing value.)

The computation of periods of time is obtained by the following:

$$\text{holds\_for}(\text{Obj}, \text{Attr}, \text{Val}, \text{during}(T_s, T_e)) \leftarrow \begin{array}{l} \text{happens}(\text{Ev}, T_s), \\ \text{initiates}(\text{Ev}, \text{Obj}, \text{Attr}, \text{Val}), \\ \text{terminated}(\text{Obj}, \text{Attr}, \text{Val}, T_s, T_e). \end{array}$$

$$\text{terminated}(\text{Obj}, \text{Attr}, \text{Val}, T_s, T_e) \leftarrow \begin{array}{l} \text{happens}(\text{Ev}, T_e), T_s < T_e, \\ \text{terminates}(\text{Ev}, \text{Obj}, \text{Attr}, \text{Val}), \\ \text{not broken}(\text{Obj}, \text{Attr}, \text{Val}, T_s, T_e). \end{array}$$

We require another clause to deal with periods that have no end-point (i.e., for the case where the value of an attribute is initiated but there is no subsequent event which terminated the value):

```
holds_for(Obj, Attr, Val, since(Ts)) ← happens(Ev, Ts),
                                       initiates(Ev, Obj, Attr, Val),
                                       not_terminated_later(Obj, Attr, Val, Ts).
```

```
terminated_later(Obj, Attr, Val, Ts) ← happens(Ev, Te), Ts < Te,
                                       terminates(Ev, Obj, Attr, Val).
```

Given a set of events, the object-based event calculus can be used to answer queries such as finding out the value of an attribute of an object at a specific time, or the period of time for which an attribute of an object has a given value. We can determine the state of an object at any time by finding out the values of all its attributes.

Of course execution of this event calculus, in Prolog say, does not yield an object-oriented style of computation. But conceptually, in object-oriented terminology, we could consider events as messages to modify object states. The specification of how events affect the state of objects would then correspond to methods, and the predicates `initiates` and `terminates` would be the system primitives by which the methods are implemented.

So far we have discussed how event calculus can be used to describe changes to the internal states of objects, i.e., to values of attributes of objects. Apart from the events that cause changes of state of existing objects, there are other kinds of events which cause the creation of new objects or deletion of objects. Before moving on to present other kinds of changes, we wish to make a remark about the representation of events.

## 4.2 Digression: Representation of Events

For the formulation of the OEC we have adopted C-logic's formula approach for the treatment of objects in the data model. In this paper we also use C-logic syntax as a convenient shorthand for describing events. The transformation of C-logic to Prolog (see section 2) allows us to mix C-logic and standard Prolog syntax freely, and this is particularly convenient when describing events. For example, an event which is described in Prolog by the following assertions

```
event(e1).
act(e1, promote).
object(e1, jim).
newrank(e1, prof).
happens(e1,1989).
```

can be written equivalently and more concisely using C-logic syntax as follows:

```
event:e1[act⇒promote, object⇒jim, newrank⇒prof].
happens(e1,1989).
```

We could also write, for example,

```
happens(event : e1[act ⇒ promote, object ⇒ jim, newrank ⇒ prof], 1989).
```

Generally we prefer to separate the structure of the event from the record of its occurrence (**happens**), as in the first C-logic version above. Whichever formulation is chosen, the C-logic to Prolog transformation makes all of them equivalent.

It is important to note that the C-logic representation of events is not essential to the main theme of the paper. We are primarily concerned with the treatment of changes to objects, and for this we have followed C-logic’s formula approach for the semantics of complex objects. It is of minor importance that we have also chosen to use C-logic syntax for the representation of events. To put it another way, the C-logic representation of events could just as well be employed to re-formulate the original (relational) event calculus, but that would not alter the nature of the data model supported by that version of the event calculus.

### 4.3 Creation of Objects

Creation of a new object of a given class means adding new information about an entity to the database. We can think of describing object creation by events — birth of a person, manufacturing a vehicle, hiring a new employee — whose specifications provide the necessary information about the initial state of the object.

In object-oriented databases, classes provide an instantiation mechanism for creating their new instances. Instantiation is performed by calling on a class to create a new object based on the information given in the class. This object is then initialized by giving each of the attributes an appropriate initial value.

In the context of temporal databases, objects are not created and destroyed. What changes is whether an object with a given identity exists or not. But since every object in our framework belongs to a class, it is unnecessary to introduce a separate **exists** predicate: instead we make class membership, **instance\_of**, a time-varying relationship. (In other words, to determine whether an object  $x$  of class  $C$  ‘exists’ at time  $t$ , determine whether  $x$  is an instance of class  $C$  at time  $t$ .) The ‘creation’ of an object is then a matter of assigning it to a chosen class and specifying its initial state.

We handle creation of objects by specifying which events assign objects to which classes, employing for this purpose a new predicate **assigns**. We use the same event description to initialize the state of the object. As an example consider registration of a student. The description of a specific registration event might be as follows:

event : e23 [act  $\Rightarrow$  register, object  $\Rightarrow$  ali, section  $\Rightarrow$  lp, supervisor  $\Rightarrow$  bob].

The rules that specify the effects of such registration events are:

assigns(event : Ev[act  $\Rightarrow$  register, object  $\Rightarrow$  Obj], Obj, student).

initiates(Ev, Obj, section, S)  $\leftarrow$   
     event : Ev[act  $\Rightarrow$  register, object  $\Rightarrow$  Obj, section  $\Rightarrow$  S].

initiates(Ev, Obj, supervisor, S)  $\leftarrow$   
     event : Ev[act  $\Rightarrow$  register, object  $\Rightarrow$  Obj, supervisor  $\Rightarrow$  S].

The **assigns** statement is used to assign the identity of the object **Obj** to the class **student**; the **initiates** statements are used to initialize the object’s state. In initializing the state

of the object, not all attributes need to be assigned to values. Some attributes may not have any values or they may have “undefined” as a value. The occurrence of the specific registration event described above is recorded by :

`happens(e23, 1991).`

There is one further point of detail. Assimilation of new event descriptions into the database will generally require introducing one or more new object identities (`e23` and presumably `ali` in the above example). In a practical implementation, generation of unique new identities can be left to the system. But notice that generation of object identities is not the same problem as ‘creation’ of new objects.

Recall from our presentation of the basic data model that instances of a class are also instances of all the superclasses. It is therefore necessary to arrange that any new instance of a class should automatically become a new instance of the superclasses. There are several ways of arranging for this, of which the simplest is to include the following rule:

`assigns(Ev, Obj, Class) ← is_a(Sub, Class), assigns(Ev, Obj, Sub).`

For the time being we assume that once an object is assigned to a class, it remains an instance of this class throughout its lifetime. That is, objects exist (i.e., they are in the database) or cease to exist at various times. Their existence is described by assigning their identities to their class. Once an object is assigned to a class it remains as an instance of that class during its lifetime and never changes class.

#### 4.4 Deletion of Objects

Deletion of objects can also be described by events. There are two kinds of deletions that we are going to discuss in this paper. One is absolute deletion of an object where the object is removed from the database: more precisely, since we are dealing with temporal databases, the object ceases to exist, or rather, ceases to be an instance of any class. The other form of deletion deletes an object from its class but keeps it as an instance of another class, possibly one of the superclasses. The second case is related to mutation of objects over time, which will be discussed in section 5.

For the purposes of this section, we assume that when an object is ‘deleted’ not only does it cease to belong to the set of instances of its class and the superclasses, but also all of its attribute values are terminated. The reason is that attributes represent the internal structure of an object. If an object ceases to exist then it is no longer meaningful to speak of its internal structure.

We use another new predicate **destroys** to specify events that ‘delete’ objects. The rule:

`terminates(Ev, Obj, Attr, _) ← destroys(Ev, Obj).`

has the effect that all attributes defined in the class of the object and also those inherited from superclasses are automatically terminated when the object ceases to exist.

There is one point to consider when deleting objects in object-oriented databases. If we delete an object `x`, there might be other objects that have stored the identity of `x` as

a reference. The deletion therefore can lead to ‘dangling references’ [77]. We eliminate dangling references by adding another general rule for the **terminates** predicate:

$$\text{terminates}(\text{Ev}, \text{Obj}, \text{Attr}, \text{ValObj}) \leftarrow \text{destroys}(\text{Ev}, \text{ValObj}).$$

The effect is that the value **ValObj** of the attribute **Attr** is terminated by any event which destroys the object **ValObj**.

## 4.5 Class Membership

As objects are created and deleted/destroyed, the instances of a class change in time. This temporal behaviour of class membership can be handled by parametrizing the **instance\_of** relation with times. We now have events that initiate and terminate periods of time for which an object *O* is an instance of a class *C*. The **instance\_of** relation is affected when a new object is assigned to a class or when an object is destroyed. By analogy with **holds\_at**, the following finds the instances of a class at a specific time:

$$\begin{aligned} \text{instance\_of}(\text{Obj}, \text{Class}, \text{T}) \leftarrow & \text{happens}(\text{Ev}, \text{Ts}), \text{Ts} \leq \text{T}, \\ & \text{assigns}(\text{Ev}, \text{Obj}, \text{Class}), \\ & \text{not removed}(\text{Obj}, \text{Class}, \text{Ts}, \text{T}). \end{aligned}$$

$$\begin{aligned} \text{removed}(\text{Obj}, \text{Class}, \text{Ts}, \text{T}) \leftarrow & \text{happens}(\text{Ev}^*, \text{T}^*), \\ & \text{Ts} < \text{T}^* \leq \text{T}, \\ & \text{destroys}(\text{Ev}^*, \text{Obj}). \end{aligned}$$

With this time-variant class membership we can ask queries to find the instances of a class at a specific time. For example:

$$\text{?- instance\_of}(\text{Obj}, \text{employee}, 1980).$$

We can also write the analogue of **holds\_for** (i.e., **instance\_for**) to compute the periods of time for which an object belongs to a class (or ‘exists’). Note that an object can have several distinct periods of membership (or ‘existence’). We omit the details of **instance\_for** since they can be reconstructed straightforwardly by comparison with the earlier formulation of **holds\_for**.

## 4.6 Discussion and Related Work

We have introduced two separate sets of predicates, one for dealing with change in internal state of objects and one for creation/deletion of objects. The internal states of objects are derived by use of the predicates **holds\_at** and **holds\_for**. These are defined in terms of predicates **initiates** and **terminates** which specify the effects of events on objects’ internal states. The temporal class membership is derived by the predicates **instance\_of** and **instance\_for**; predicates **assigns** and **destroys** are used to specify how events affect class membership. The formulation of these two sets of predicates are direct analogues of one another. We could combine them into one set of predicates, with one general formulation, and thereby dispense with one set of predicate names altogether. We have not done so

because we want to emphasize the conceptual difference between changes in an object's state on the one hand and changes to class membership on the other.

The treatment of change formulated in the OEC is appropriate under the assumption that facts and properties of objects persist over time — that, once initiated, each fact continues to hold without interruption until it is terminated by some subsequent event. Such facts have been termed ‘stable’ or, perhaps more perspicuously, ‘stepwise constant’ [58] in the literature on temporal databases. The OEC can be extended to accommodate other kinds of time-varying behaviour by incorporating various extensions that have been developed for the original, relational event calculus. In particular continuous change can be treated using the ‘trajectories’ of [59]. We do not present the details here. The treatment can be imported from the relational versions without modification, and is actually slightly more convenient to formulate within the OEC, since it is continuous change of values of attributes that is of interest; it is difficult to imagine what continuous change of membership of a class would correspond to. Other extensions, such as allowing for different granularities of time within the same data model [24, 51], could also be adapted straightforwardly.

In the database field, the modelling of temporal information has been dominated by approaches based on the relational model. There are exceptions (see e.g. [21, 32, 56, 57, 68, 73] and [61] for a comparative survey). These proposals differ in the range of modelling features they provide. More significantly, they differ also in their general approach to the representation of temporal information, and to the notion of ‘object’ itself. We select here two examples, each of which is intended to be representative of a general class of approaches.

The first example is the extended entity-relationship model described in [21]. This has many features in common with the data model supported by the OEC. It has entities, time-invariant identities (there called ‘surrogates’ for entities), attributes, and time-varying membership of classes and sub-classes. The temporal extension records a ‘lifespan’ with each entity, with each attribute-value, and with each class membership instance. These lifespans correspond exactly to the time periods computed by the **instance-for** and **holds-for** predicates of the OEC. The model of [21] also supports relationships between entities, a feature not provided by the OEC. A corresponding extension of the OEC, discussed briefly in [34], could be obtained by combining the OEC with the relational version of the event calculus, but this is something we have not undertaken yet. The crucial difference between the OEC and the temporal entity-relationship model of [21] lies in their respective treatments of lifespans/periods and their treatments of updates. In [21] updates are treated as *operations* on the database, and the system provides a number of basic operations for this purpose: to create new entities of a given type, to insert and delete entities from classes, and to modify attribute values. In each case, the corresponding lifespan must be manipulated explicitly. The main motivation for the OEC, in contrast, as in the original event calculus, is to provide some semantic structure for updates and for the computation of ‘lifespans’. Updates, corresponding to assimilation of information about changes in the world, are treated by *adding* an appropriate *event description* to the ‘journal’ stored in the database. The effects on the objects and their attributes are not modelled as operations on the database but are *derived* from the event descriptions via the specifications given by the **initiates-terminates** and **assigns-destroys** predicates. We should note, however, that other kinds of change, such as corrections of mistakes, are not supported directly by the OEC and must be treated as modifications of the ‘journal’. A more powerful and general



treatment would require the introduction of a separate transaction time dimension. As already mentioned, this has been done for versions of the relational event calculus [63, 64] but not yet for the OEC.

The second representative is the temporal object-oriented system described in [73]. That system, however, places considerable emphasis on encapsulation and abstract data types. These are object-oriented programming features that are generally not mimicked directly in logic-based treatments of objects. (See the discussion in section 2.) The proposal in [73] also treats time points and intervals as abstract data types, which we do not attempt.


To avoid any misunderstanding, we should perhaps state explicitly that we are not claiming in this paper to have introduced or invented some new modelling concept — temporal or object-oriented — that is not already found in the literature. Conversely: our intention is to show how a base set of object-oriented features may be provided in a natural fashion in an integrated temporal/deductive framework; to demonstrate that this framework can be further extended to provide a wider range of features, some of which we present in detail; and to indicate that the resulting system can be used as it stands for the construction of practical small-medium applications and has the prospect of further development to large-scale database applications. We want to emphasise that it is the general framework that is the basis for meaningful comparisons with other work, and not the list of features currently supported by the OEC.

## 5 Mutation of Objects: Changing the Class

We have so far assumed that objects exist, cease to exist in the database, but never change class. However in the real world it is common that objects evolve over time. Consider the representation of an employee instance again. We have represented the rank of an **employee** object by including an attribute **rank** whose value can change over time:

```
employee : jim[... rank ⇒ lecturer, ... ]
employee : mary[... rank ⇒ professor, ... ]
      :
```

But suppose that instead of using the attribute **rank**, we had chosen to divide the class of employees into various distinct subclasses:



Then employees of different ranks would be considered as different classes, represented using **is\_a**:

```
is_a(lecturer, employee).
is_a(professor, employee).
```

The choice between the two representations is a data modelling issue. If employees have different additional attributes according to their ranks then it is appropriate to represent different ranks as subclasses. However, even if the structure of all ranks is identical, the choice between the two representations can become significant, if we consider the dynamics of the ‘promotion’ event. In the first representation, values of the rank attribute can be changed straightforwardly to model the effects of promotion. In the second representation, modelling a promotion from lecturer to professor requires destroying the **lecturer** object and creating a new **professor** object. But then how do we relate the new **professor** and the old **lecturer**, and how should we preserve the values of unaffected attributes common to all employees?

## 5.1 Classes and types

While researchers in knowledge representation systems have been aware of the problem of mutating objects for some time [10, 66] it has only recently begun to receive attention in the object-oriented database community [76, 72]. In most object-oriented database systems the permanent binding of an object’s identity to a single class constrains the tracking of real world entities over time. This is a critical problem, because one is forced to model entities that evolve dynamically with objects that cannot.

The ability to change the class of an object provides support for object evolution. It lets an object change its structure and behaviour, and still retain its identity. In [76] a type system which allows this kind of evolution is presented. An object  $x$  can have a set of types, and the change from one type to another is a process of selectively adding and deleting types to the set of types of  $x$ . The notion of typing is retained whilst allowing some flexibility in system evolution.

In our present framework there is no notion of type. We support the grouping of objects according to common structure and properties by means of class, which is a dynamic notion. This gives more flexibility for representing class changes. However there are other advantages to be gained from having a type system in addition. A further typing mechanism could be added as an extension to our basic data model. There is a tendency in the literature to use the terms ‘type’ and ‘class’ interchangeably. For us they are distinct notions: one (*type*) is a static, syntactic feature of the representation language; the other (*class*) is a dynamic grouping of objects according to their structure and properties.

## 5.2 Realization

We deal with the evolution of objects by allowing events that change an object’s class, and some or all of its attributes. For example, graduation causes a student to change class. The effects can be described by removing the student object from class **student** and terminating those attributes he has by virtue of being a student; attributes he has by virtue of belonging to class **person** should however be retained. The selective termination of attributes is obtained by using schema information. In order to deal with this type of class change we introduce another predicate **removes** in place of **destroys**. Again, new predicate names are used in order to emphasise conceptual differences.

Consider figure 3. A graduation event causes **ali** to move up the class hierarchy. When the student **ali** graduates, he is removed from **student** and becomes an instance of the




Figure 3: Class Changes

**person** class only. His attributes by virtue of being a student are also terminated. The effects of the graduation event are described by the following rule:

$$\text{removes}(\text{event: Ev}[\text{act} \Rightarrow \text{graduate}, \text{student} \Rightarrow S], S, \text{student}).$$

There is in addition a general rule, that removing an object from a class terminates all attributes specific to instances of that class:

$$\text{terminates}(\text{Ev}, \text{Obj}, \text{Attr}, \_) \leftarrow \text{removes}(\text{Ev}, \text{Obj}, \text{Class}), \\ \text{attribute}(\text{Class}, \text{Attr}).$$

The overall effect of a graduation event for **ali** is that, for times after the graduation, it is no longer possible to derive `instance_of(ali, student)`, nor values for any of his **student**-specific attributes since these are all terminated automatically by the graduation event.

Now consider hiring **ali** as an employee. This will cause his class to be changed from **person** to **employee**. Since the class **employee** has some additional attributes (**dept**, **rank**), the specification of this event will include values to initiate these attributes. Thus the effects of the hiring event are described by assigning him to the class **employee**, and initiating his **employee**-specific attributes. The description of such an event might be:

$$\text{event : e21}[\text{act} \Rightarrow \text{hire}, \\ \text{object} \Rightarrow \text{ali}, \\ \text{dept} \Rightarrow \text{cs}, \\ \text{rank} \Rightarrow \text{lecturer}]$$

The effects of hiring events in general can be specified as follows:

$$\text{assigns}(\text{event: Ev}[\text{act} \Rightarrow \text{hire}, \text{person} \Rightarrow P], P, \text{employee}).$$

$$\text{initiates}(\text{event: Ev}[\text{act} \Rightarrow \text{hire}, \text{person} \Rightarrow P, \text{dept} \Rightarrow D], P, \text{dept}, D).$$

$$\text{initiates}(\text{event: Ev}[\text{act} \Rightarrow \text{hire}, \text{person} \Rightarrow P, \text{rank} \Rightarrow R], P, \text{rank}, R).$$

Note that in changing **ali**'s class first from **student** to **person** then from **person** to **employee**, **ali** has not been removed from the class **person** and has retained all his **person**-specific attributes. More importantly the identity of the changing **ali** object remains the same throughout.

We have described moving an object up and down the class hierarchy by two separate event occurrences. We can also imagine a single event which would cause an object to change its class from **student** to **employee** directly ('hire-student' say). The effect of this type of event could be specified as follows:

```
removes(event: Ev[act  $\Rightarrow$  hire-student, student  $\Rightarrow$  S], S, student).
assigns(event: Ev[act  $\Rightarrow$  hire-student, student  $\Rightarrow$  S], S, employee).
```

As in the case of two separate events, we do not lose the values of the **person**-specific attributes, and we do not remove the object from class **person**.

The question naturally arises of what happens to attribute values as the object moves across the hierarchy. In our framework, the relationships between old and new values in the sibling classes, if any, can be specified explicitly using **initiates** statements, just as in the specification of the initial state of a newly 'created' object (e.g. 'hiring' above). There is nothing special about class-changing events in this respect. We do not believe that any useful general rules can be formulated, even for the case where the sibling classes contain attributes with the same name. It might be supposed that in such a case the values of the common attributes should remain unchanged. We believe this would be a mistake. If the common attribute has a different intended meaning in the two classes, then there is no reason why the two values should be the same, except by coincidence. In the case where the common attribute does have the same intended meaning in both classes (e.g. if an attribute **age** indicates the age of a student and also the age of an employee), then this suggests an inadequacy in the modelling scheme itself. If we want such attributes to retain their values during a class change, then they should belong to a common superclass of the two classes involved. (In the example, **age** should be an attribute of **person** and not of the more specialised sub-classes **student** and **employee** separately.) The whole point is that common attributes should be defined as part of the structure of a general class, with each sub-class further introducing the additional attributes specific to its instances.

We have illustrated three kinds of simple class change: changing from a class  $C$  to a direct superclass of  $C$ , changing from  $C$  to a direct subclass of  $C$  and changing from  $C$  to a sibling class of  $C$  in the hierarchy. In the general case, changing an object from class  $C1$  to class  $C2$  involves finding a path in the class hierarchy and using rules similar to the preceding ones to move along this path.

### 5.3 Remarks

In general, class changing events affect both the internal state of objects and also the class membership relation. The **holds\_at** and **holds\_for** clauses for deriving internal states of objects require no modification, since the effect is accommodated by making **terminates** dependent on **removes** as shown already. However we do need to modify the clauses for **instance\_of** (and **instance\_for**) to take **removes** into account. The modified formulation of **instance\_of** is as follows:

```
instance_of(Obj, Class, T)  $\leftarrow$  happens(Ev, Ts), Ts  $\leq$  T,
                                assigns(Ev, Obj, Class),
                                not removed(Obj, Class, Ts, T).
```

$$\begin{aligned} \text{removed}(\text{Obj}, \text{Class}, \text{Ts}, \text{T}) \leftarrow & \text{happens}(\text{Ev}^*, \text{T}^*), \\ & \text{Ts} < \text{T}^* \leq \text{T}, \\ & \text{removes}(\text{Ev}^*, \text{Obj}, \text{Class}). \end{aligned}$$

$$\text{removes}(\text{Ev}, \text{Obj}, \_) \leftarrow \text{destroys}(\text{Ev}, \text{Obj}).$$

The last clause states that if an event destroys an object, then that event also removes the object from all its classes.

There are two details left, one comparatively trivial and one more substantial. We take them in order.

In the formulation as we have presented it, the general recursive rule for **assigns**, which describes the subset relation (see section 4.3), causes the assignment of an object to a class redundantly when there are also other more specific **assigns** statements present. For instance when hiring person **x** as an employee, a new period of time for the fact **instance\_of(x, person)** is initiated even though this fact is already current (**x** is already an instance of **person** when the hiring takes place). This duplication of periods occurs with every class-changing event for which an **assigns** statement assigns an object to a subclass of its current class. The problem manifests itself when the database is queried about the **instance\_for** relation because then several different but overlapping periods of time can be generated as answers. There are various solutions to this problem. The simplest is to take into account the possibility of these different time periods in the formulation of **instance\_for** so that all these separate periods are amalgamated into one. (This requirement has been termed *coalescing* in the temporal database literature [27, 31].)

The second point is more substantial. Allowing objects to change their class presents a potential problem which is analogous to dangling references. The problem arises when an object, which is the value of an attribute **Attr** of some other object, changes class in such a way that it can no longer be regarded as a meaningful value for attribute **Attr**. For example, assume that the staff instances have an attribute **student** which takes a student as a value. Further assume that the student **ali** is a student of the staff member **john**. When **ali** graduates and changes class to employee, the student attribute of **john** is not valid any longer and should be terminated. The graduation, which is defined as a class-changing event, takes care of terminating and initiating attributes of the student, but the other objects referring to this object as a student are not changed. This problem, sometimes called the *dangling domain problem* [76], would be a type violation in a typed system. In our present framework the problem can be avoided by writing event specifications appropriately, but this obviously requires that all ‘dangling domain’ problems are identified and accounted for explicitly. This is clearly unsatisfactory. The natural solution is to refine the schema so that the type (or perhaps class) of the value of an attribute is specified as well. In principle, this additional schema information would allow for a re-formulation of the OEC so that attributes affected by the ‘dangling domain’ are terminated automatically in the same kind of way that ‘dangling references’ are eliminated. However the details turn out to be quite complicated, and we have not yet explored all the possibilities.

## 6 Further Considerations

For the purpose of focussing attention on the different kinds of changes to objects and how they can be formulated, we have presented a simplified form of the OEC so far. In this section we sketch how the OEC is modified to provide other features which are required for practical applications. The two main extensions to the basic data model are to allow *multi-valued* as well as single-valued attributes, and to support a wider range of *methods* than those introduced so far, specifically for deriving new information from the existing states of objects.

### 6.1 Multi-valued Attributes

Multi-valued attributes are supported very straightforwardly in our framework, since it is actually single-valued attributes that are the special case and that impose additional requirements. We do not attempt to support *set-valued* attributes.

A multi-valued attribute denotes a one-to-many relation which maps the identity of an object to one or more objects. Such a relation can be thought of as a set of binary predicates, as in C-logic [13], for example. Indeed in C-logic *all* attributes are multi-valued, since an attribute in that language is semantically the same as a binary predicate. Thus the C-logic term

person: john[children  $\Rightarrow$  {tom, sue, mary}].

is just a shorthand notation for the complex term

person: john[children  $\Rightarrow$  tom, children  $\Rightarrow$  sue, children  $\Rightarrow$  mary].

which is semantically equivalent to the following set of assertions in first-order logic:

children(john, tom).  
children(john, sue).  
children(john, mary).

This notion of multi-valued attribute should be contrasted with approaches where (single-valued) attributes are allowed to take *sets* of objects as values. Set-valued attributes provide increased expressive power (they are no longer just first-order) but at the cost of introducing the very severe semantical and computational problems associated with sets and set unification [47]. As argued in [13], the simpler notion of multi-valued attribute already satisfies most of the practical and expressive requirements of set-valued attributes, at enormously reduced complexity.

The re-formulation of the object-based event calculus to allow for multi-valued attributes requires just a slight generalisation of what has been presented in the previous sections. The main adjustment is a refinement of the clause which is used to implement the functionality constraint for single-valued attributes, viz:

terminates(Ev, Obj, Attr, -)  $\leftarrow$  initiates(Ev, Obj, Attr, -).

If we wanted to support only multi-valued attributes then it would be sufficient to delete the ‘functionality’ clause altogether and retain the rest without change. Since we want to support both single and multi-valued attributes in one system, the obvious solution is to make the clause above applicable to single-valued attributes only. For this it is necessary to extend the schema information so that it also specifies whether an attribute (in a given class) is single-valued or not. A third argument is therefore added to the predicate **attribute** to specify the kind of the attribute: **single** or **multi**.

The clause implementing the functionality constraint now requires an extra condition so that it does not apply to multi-valued attributes. Since the kind of an attribute (**single** or **multi**) may depend on class as well as the attribute name, it is necessary to include the class name as an additional argument in the **initiates** and **terminates** predicates. The modified clause is:

$$\begin{aligned} \text{terminates}(\text{Ev}, \text{Class}, \text{Obj}, \text{Attr}, \_) \leftarrow \\ \text{attribute\_of}(\text{Class}, \text{Attr}, \text{single}), \\ \text{initiates}(\text{Ev}, \text{Class}, \text{Obj}, \text{Attr}, \_). \end{aligned}$$

The definition of **attribute\_of** must also be altered to take inherited single-valued attributes into account, as follows:

$$\begin{aligned} \text{attribute\_of}(\text{Class}, \text{X}, \text{Type}) \leftarrow \text{attribute}(\text{Class}, \text{X}, \text{Type}). \\ \text{attribute\_of}(\text{Class}, \text{X}, \text{Type}) \leftarrow \text{is\_a}(\text{Class}, \text{Super}), \\ \text{attribute\_of}(\text{Super}, \text{X}, \text{Type}). \end{aligned}$$

Including the class name as an argument in the **initiates** and **terminates** predicates makes it necessary to specify the class of an object at query time. Moreover, the clauses of the object-based event calculus presented earlier all need to be adjusted to take the presence of this new class argument into account. This is a very simple modification that raises no additional questions and so we do not show the whole modified version of the OEC again.

## 6.2 Derived-attribute methods

In object-oriented systems, methods are operations to describe the behaviour of objects. This includes both modification and manipulation of the state of objects. We have so far considered only methods that modify the state of objects. Now we want to extend the kind of methods that are supported to include also methods which can derive new information from the existing state of objects. We call such methods *derived-attribute methods* or sometimes just *rules*. For instance the age of a person can be derived from the date of birth; the boss of an employee can be derived from the manager of his department, and so on.

Derived-attribute methods are included in the schema definition according to the classes. There have been numerous proposals for how to define and implement such methods in a logic programming framework (e.g. [19, 25]). Representing methods as deductive rules is the most common approach in the existing languages and the one we follow here.

The definition of derived-attribute methods can be given in a syntax similar to C-logic or other object-logic languages. For instance, in F-logic [37], the boss of an employee is defined by the following rule:

$E [\text{boss} \Rightarrow M] \leftarrow E: \text{employee}[\text{dept} \Rightarrow D: \text{dept}[\text{manager} \Rightarrow M]]$

stating that the boss is the manager of the department in which the employee works. This kind of syntax can be translated straightforwardly into an internal form which is manipulated by the object-based event calculus. The head of the rule contains the name of the derived attribute and the body contains the object-attribute-value information. Each rule is associated with a class. We employ the following representation in order to index methods by the class names:

`method(ClassName, Object, AttrName, Value, Body).`

The first argument, `ClassName`, refers to the name of the class for which the method is defined. `Object` is the identity of the object for which the method is invoked. The third argument, `AttrName`, is the name of the derived attribute (in the object-oriented terminology this corresponds to the message used to invoke the method). The fourth argument, `Value`, denotes the value returned as the result of the method. And finally, `Body` is the translated form of the body of the rule.

For the purposes of this paper we assume that the body of a derived-attribute method is a conjunction of complex object terms (i.e., object-attribute-value information). Other more general forms can be allowed but we present only the most basic form here. Again following C-logic [13], the relational semantics for complex object terms allows every such term to be decomposed into a conjunction of atomic object terms, and so the body of a rule can always be expressed as a conjunction of conditions of the form `o-term(Class, Obj, Attr, Val)`. Hence the example above would be translated into the form:

`method(employee, E, boss, M,  
[o-term(employee, E, dept, D), o-term(dept, D, manager, M)]).`

Note that conjunctions are here represented as *lists*. We employ the standard Prolog notation, where `[X|Y]` is the list with head element `X` and tail `Y`, the term `[a,b,...,c]` stands for the list of elements `a`, `b`, ..., `c`, and `[]` denotes the empty list, which in the present context represents the empty conjunction ('true'). This is the most convenient internal representation; it is very easy to provide an interface to translate derived-attribute methods in some chosen syntax into this internal form.

Given such a representation, we now require the ability to compute not only the conclusions that are derivable from the rules, but also the time periods for which such conclusions hold. There are two dimensions in dealing with time-varying aspects of derived information. One is that the rules themselves are subject to change with time. That is, the rules defined in a class may be applicable during different periods of time. For instance in a database dealing with legislative matters, the rules for computing taxes of employees may need to be changed as the government changes its tax policies. (See [63, 64] for a detailed treatment.) The other dimension is the actual computation of the value of the derived attribute at a particular time or for a given period of time. The first aspect of the problem is related to schema evolution. As rules are defined in the schema, any change to the definition of the rules requires modification of the schema. This problem will be discussed in section 9. The second aspect is related to query evaluation and is discussed here.



In its full generality, temporal reasoning with derived information raises a number of unresolved questions which are the subject of much current research. The usual solution, and the simplest, is to distinguish between *base* and *derived* information and treat these separately. The effects of changes (here, events) are then described by specifying only how they affect *base* information. The effects on derived information are obtained indirectly, because derived information is computed from the base information when it is required.

This is the approach that we follow also. We do not specify how events affect derived attributes directly. Values of derived attributes at any given time are determined by finding the corresponding method in the schema definition and then solving each condition of the rule body at the specified time instant. This is accomplished by adding another clause to the definition of `holds_at`:

```
holds_at(Class, Obj, Attr, Val, T) ←
    method(Class, Obj, Attr, Val, Body),
    solve_at(Body, T).
```

Here `solve_at(Body, T)` can be seen as a ‘meta-interpreter’ which executes the `Body` at the specified time `T` :

```
solve_at([], _).
solve_at([A|Rest], T) ← solve_condition_at(A, T), solve_at(Rest, T).

solve_condition_at(o-term(Class, Obj, Attr, Val), T) ←
    holds_at(Class, Obj, Attr, Val, T).
```

More general forms of rules for derived attributes are accommodated by allowing other kinds of conditions and adding further clauses for `solve_condition_at`.

The period of time for which a derived attribute takes a particular value can also be determined. Intuitively, the time period for which the conclusion of a rule holds is computed from the time period for which the rule itself holds and the time periods for which each of its conditions hold: if we think of time periods as sets, then the result is the intersection of these time periods. In order to determine the time periods for which a derived-attribute holds a particular value, we add another clause to the definition of `holds_for`:

```
holds_for(Class, Obj, Attr, Val, P) ←
    method(Class, Obj, Attr, Val, Body),
    solve_for(Body, P).
```

The predicate `solve_for` determines the intersection of time periods for the conditions in the body of a derived-attribute method. In database terms, it can be thought of as computing a ‘temporal intersection join’ [18, 48]. `solve_for` can be defined as follows:

```
solve_for([], always).
solve_for([A|Rest], P) ←
    solve_condition_for(A, P1), solve_for(B, P2),
    intersect(P1, P2, P).
```

`solve_condition_for(o-term(Class, Obj, Attr, Val), P) ←  
holds_for(Class, Obj, Attr, Val, P).`

Here **always** denotes the universal time period (the intersection of any time period **P** with **always** is **P**). There is a more efficient, tail-recursive formulation of **holds\_for** which we omit. As in the case of **solve\_condition\_at**, more general forms of rule for defining derived-attributes are treated by adding further clauses for **solve\_condition\_for**.

## 7 Implementation Issues

The OEC presented in the preceding sections can be implemented in different ways. It can be executed as a Prolog program, by transcribing the clauses given into Prolog syntax, although for reasons identified below the performance of the resulting program is likely to be poor, even in small examples. For practical applications we make use of a Prolog implementation that incorporates an additional ‘tabulation’ or ‘lemma generation’ mechanism, described presently. The OEC could also be implemented in one of the object-oriented logic programming languages now available (see section 2) in order to provide the kind of ‘clustering’ normally associated with object-oriented implementations; we have conducted some preliminary experiments using L&O [50] as an implementation language. Or the OEC as presented could be seen as an (executable) specification and suitable algorithms constructed to perform the same tasks in a procedural language.

We have employed the OEC in a number of small–medium applications, of which the most developed is a database dealing with the activities of a University department (from which the examples in this paper have been taken). To give some indication of size, object instances in these applications are numbered in hundreds and event occurrences in thousands. We have recently embarked on the reconstruction of an application in cardiovascular medicine originally constructed using the standard (relational) event calculus [62].

The purpose of this section is to summarise the implementation techniques used in these applications. As indicated in the introduction to the paper, our aim is also to employ the OEC as the basis of large-scale temporal database systems; in this regard, we wish to explain why implementation techniques under development for temporal databases can be adapted for use with the OEC.

### 7.1 Current implementations

The problems encountered in the efficient execution of the OEC can be illustrated by reference to the execution of **holds\_at** queries. Exactly similar points can be made for more general **holds\_at** queries, for **holds\_for** (the computation of time periods), and for the class membership analogues **instance\_of** and **instance\_for**.

Obviously, there is the problem of searching efficiently for relevant candidate events that could initiate or terminate the value of the attribute in question, for which some form of indexing is required. But the main factor affecting performance of the OEC (and of the standard relational versions of the event calculus) is the need to determine what the effects of each such candidate event are (what it initiates and terminates). In general, this is not

just a matter of looking up the **happens** assertions (the ‘journal’ of event occurrences), since to determine whether an event actually does initiate or terminate a value for a given attribute may require some further computation; and in a naive implementation this in turn may generate re-computation of the same facts over and over again, whenever **initiates** or **terminates** statements are context-dependent, that is, in the case where the value initiated or terminated by an event at time **t** depends on what other values are current at time **t**. In these circumstances naive execution of the event calculus clauses can lead to very severe redundancies in the computation, significantly affecting performance even in small applications.

Dramatic improvements can be obtained by adding a bottom-up component to the evaluation mechanism, by incorporating some form of ‘lemma generation’. This is a standard technique in logic programming, and in deductive databases where it is often referred to as ‘tabulation’. (See e.g. [70] for a survey and general account of tabulation techniques in deductive databases.) For the OEC implementation, all facts regarding the states of objects and the time periods initiated and terminated by the recorded events are stored as they are deduced: a very simple tabulation mechanism is sufficient. We shall refer to these tabulated facts as the Object DataBase (ODB) for convenience. The ODB can be generated bottom-up when a new event is assimilated into the database, or — as we prefer — by tabulation during top-down query evaluation. However it is produced, the ODB contains all the information about objects: their states, their classes and the necessary information to derive the time periods for which these hold. When a query is posed to the system only the contents of the ODB are accessed without searching all the events again.

The obvious set of ‘lemmas’ or ‘tabulated results’ to store in the ODB are all the time periods, analogously to the scheme proposed in [64] for (relational) temporal databases based on the event calculus. We prefer an alternative which is much more flexible and easier to maintain: for each tuple of the form (**Obj**, **Attr**, **Val**) we record the starting time(s) at which that tuple is initiated, and, separately, a record of the time(s) at which the tuple is terminated by another event. The time periods for which the tuple holds are easily derived from these start and end points as required. Similarly the class(es) to which an object belongs in time are stored as tuples of the form (**Obj**, **Class**) together with the start and end times for each. A Prolog implementation with this mechanism and a simple form of indexing gives quite acceptable performance for the applications mentioned above. The table in figure 4 gives some sample timings. The queries were executed on a database of the kind used as the source of examples in this paper, containing approximately 10,000 events (1000 objects). These timings are just intended to be indicative of performance. Further details regarding implementation techniques are provided elsewhere [36].

## 7.2 Future work

One of our longer term aims is to develop the OEC so that it can support database applications proper. More sophisticated indexing techniques will then be required. We do not want to give the impression that we underestimate the difficulty of the task, but we do want to indicate why we believe this is not an unrealistic ambition. The point is that implementation techniques being developed for temporal (relational) databases are not incompatible with use of the OEC; many can be adapted, or even applied directly.

For example, the techniques described in [28] for implementing ‘backlogs’ in transaction-

Query (all solutions)	Net time (millisec)	
	(plain OEC)	(with ODB)
holds_at(manager1,name,N,100)	219.66	1.434
holds_at(manager1,A,V,100)	1679.20	13.165
instance_of(manager1,C,100)	207.33	0.783
instance_of(O,person,100)	5755.85	46.085
holds_for(student5,address,A,P)	26415.00	8.333
holds_for(student5,A,V,P)	1.797e+05	64.500
holds_at(assistant15,A,V,100)	152.34	13.432

The database contains approx. 10,000 events (1000 objects).  
(Quintus Prolog Release 3.1.1 on a SUN Sparc workstation)

Figure 4: Some sample timings

time databases can be applied directly. The backlog of a relation is a system-generated relation that contains the complete history of updates to that relation. Each tuple is time-stamped with the times at which it was inserted, modified or deleted. A time-slice operator is provided to reconstruct the state of the relation at any given time; further techniques employing ‘differential files’ have also been developed so that query evaluation can make use of previously computed results. Although these techniques are designed for the implementation of *transaction-time* databases, the feature on which they rely is the ‘stepwise-constant’ nature of the time-varying data. Since this is also the kind of change supported by the OEC, the same techniques are applicable for implementation of the ODB: indeed, the structures we store in the ODB have (almost exactly) the form of backlogs. Indexing methods, such as the B<sup>+</sup>-tree techniques described in [22], can also be applied, with some modification, to the ODB, or more directly, to provide indexing on the record of event occurrences (the ‘journal’). Indeed, a very rudimentary form of this idea is the basis of an event calculus implementation described in [26]. Development of these ideas, and of associated query optimisation techniques, remain topics for future work, however.

In addition to temporal indexing, some form of structural indexing — some method of storing the objects and events so that search is restricted to the potentially relevant candidates — is also required. At present all information in the ‘journal’ of event occurrences and in the ODB takes the form of relational tuples. (In the Prolog-based implementation simple forms of indexing these tuples have proved adequate.) Beyond some preliminary experiments with the L&O language [50], as mentioned above, we have not addressed the question of how to provide indexing methods that reflect the object-oriented structuring at the implementational level.

## 8 Versioning of Objects

Although version control is not considered to be a mandatory feature of object-oriented databases, it has often been associated with such databases because of the importance of versioning of objects in design applications. There have been many proposals to incorporate versioning in object-oriented databases. Up to now this problem has not received

a complete and satisfactory solution mainly because several distinct models have been proposed for complex objects, but also because it is not clear what are the main concepts which concern object evolution. Usually the proposed version models are specific to the application domain and show different characterizations according to the needs of the application.

The simplest kind of versioning is to keep object state histories [2, 9, 17]. It is based on retaining information about all past states of objects. Any update to an object will generate a new version. Checkpointing is another kind of versioning which adds facilities for managing a sequence of identifiable points in time of the state history [3, 9]. At every checkpoint the state of the object is marked as being a version. Some other applications support parallel versions [7, 16]. After the initial creation of a design object, new versions of the object can be derived from it and new versions can in turn be derived from them. Therefore at any one time several valid representations of an object can coexist. This scenario necessitates the creation of a hierarchy or a graph, rather than a linear succession, to capture the evolution history of versions from the initial design.

In this section we describe how versioning models can be formulated within the OEC in a straightforward way. A feature of this treatment is that the version states are not copies of parent objects: only the differences between their state and the parent's state are recorded. Versions share the history of their parent objects for unchanged attributes. As a natural by-product we can also update and keep the history of versions, which is not possible in other proposals.

The formulation of the OEC presented so far provides us with the state history of objects, and already supports the simplest kind of versioning. At the application level we can query the state of an object at a specific time point using the `holds_at` predicate. This state can be viewed as a version of the object and by querying the object at different times we can keep track of its evolution. Checkpointing and parallel versioning can also be accommodated in the object-based event calculus with some modifications to the original. These various modifications are presented and discussed in more detail elsewhere [35]; here we summarize the modifications for parallel versioning only.

Versions are objects which are derived from existing objects as a result of some requirements and modifications. Versions are closely related to their parent versions, but they are still different objects and they must be uniquely identifiable. Also, as versions share some of their properties with the object from which they are derived, there must exist an easy way of finding the previous version.

As defined earlier, an object identity is a first-order term which is composed of constants, variables and function symbols. We now use a naming convention that uniquely identifies the versions. Suppose we have an object with the identity  $\mathfrak{o}$  and we create versions of this object. One way of distinguishing these versions is to number them. The first version of  $\mathfrak{o}$  will be  $v(\mathfrak{o},1)$ , the second  $v(\mathfrak{o},2)$ , the  $n$ th  $v(\mathfrak{o},n)$  and so on. This naming provides a total ordering on the names of versions if we satisfy the condition that the number of a new version will always be greater than its parent version. When versions of versions are created, the same naming convention can be used to identify the new versions. The first version of the object  $v(\mathfrak{o},1)$  for example, will be named as  $v(v(\mathfrak{o},1),1)$ , the second will be  $v(v(\mathfrak{o},1),2)$  and so on.

The basic idea to represent an object having several versions at a time is to keep parallel histories for the object. Each history is identified by a version identity. Each

version has its own history starting from its creation time. Once a version is created it is treated in the same way as other objects in the database. It can be updated, deleted or versioned. Meanwhile its parent object can be directly updated, even after one or more versions have been derived from it: the derived versions will not be able to see the updates in the parent object.

Creation of a version is described by events. In keeping the state history of objects, every event which causes a change in the value of the object's attribute creates a new state of the object, which can be viewed as a new version. In parallel versioning not every event is considered as a version-creating event. Only certain events can cause the creation of identifiable versions. Some attributes can be classified as version-significant attributes, whose update would force the creation of a new version bearing the modified value of that attribute. Events that are specified as having effects on these attributes can be defined to be version-creating events.

The effects of version creating events are specified by the predicate `creates_version` which is used to mark the occurrences of such events in the object's history and also to generate a unique identity for the version.

For example consider the design of a VLSI chip. Different versions of the chip may be derived, say to reduce the chip size or reduce the power consumption, etc. Every time a user performs a "reconfigure" operation, a new version is assumed to be created. This can be achieved by the following statement

$$\text{creates\_version}(\text{Ev}, \text{v}(\text{C}, \text{N})) \leftarrow \\ \text{event} : \text{Ev}[\text{act} \Rightarrow \text{reconfigure}, \text{chip} \Rightarrow \text{C}, \text{number} \Rightarrow \text{N}].$$

The functional term  $\text{v}(\text{C}, \text{N})$  is the identity of the new version where  $\text{C}$  denotes the parent object, and  $\text{N}$  is an integer denoting the version number.

Versions are also instances of the class to which their parent version belongs. Continuing with our example, the event of reconfiguring a chip assigns the new version to the class `VLSI-chip`, say:

$$\text{assigns}(\text{Ev}, \text{v}(\text{C}, \text{N}), \text{VLSI-chip}) \leftarrow \\ \text{event} : \text{Ev}[\text{act} \Rightarrow \text{reconfigure}, \text{chip} \Rightarrow \text{C}, \text{number} \Rightarrow \text{N}].$$

The version objects can be changed like any other object. Reasoning with the changing state of versions is done by similar axioms, but now the time of creation of the version is taken into account as well. The values of the attributes which are not changed at the creation time or later are derived using the parent version.

Consider figure 5 which shows a section of an object's version derivation history. Here `Vid` is a version identifier and `Oid` is the object from which it is derived. `Oid` can be another version or the initial object. `Tc` denotes the creation time of the version `Vid`, `T` denotes the time at which we query its state. The history of `Vid` starts at `Tc` and at that time `Vid` as its initial state has the same state as `Oid` has as of time `Tc`. At any time after its creation, say at `T*`, the version's state can be changed by an event. If no such event happens between `Tc` and `T`, then the state of the version at time `T` is the same as the state of the object `Oid` at time `Tc`. However if there are some events that have happened after `Tc` and have changed the values of one or more attributes of the version, then we have to consider the effects of these events as well.

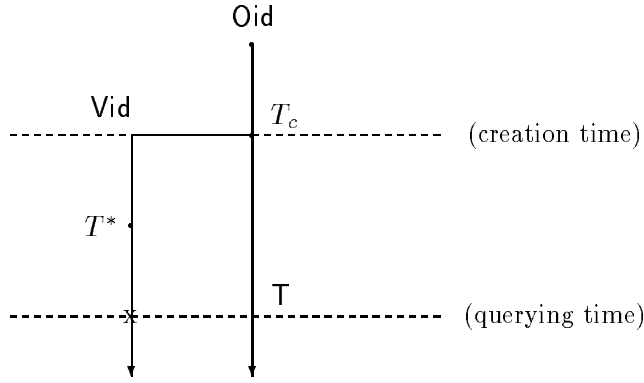


Figure 5: A section of an object derivation hierarchy

Here is our formulation of `holds_at` to reason with the state of objects and versions:

```
holds_at(Vid, Attr, Val, T) ←
  happens(Ev, Tc), Tc ≤ T,
  creates_version(Ev, Vid),
  ( happens(Ev*, T*),
    Tc ≤ T* ≤ T,
    initiates(Ev*, Vid, Attr, Val),
    not broken(Vid, Attr, Val, T*, T) )
or
  ( prev_version(Vid, Oid),
    holds_at(Oid, Attr, Val, Tc) ).
```

This formulation can be used to navigate through the versions easily, terminating when the recursion reaches the root object. This modified `holds_at` can also be used to reason with the state of objects without any versions, provided that the event creating the object is defined to be a version-creating event as well (i.e., a `creates_version` statement defines it also as an event which creates the object as the first (maybe the only) version of itself).

## 9 Schema Evolution

In this section we look at the problem of changes in class definitions (i.e., the database schema). As well as the objects in the database, class definitions may also be modified. We address the problem of maintaining consistency between a set of objects and a set of class definitions that can change.

### 9.1 Problem Definition

Schema evolution is not inherent to the object-oriented paradigm. However it is natural to assume that in the course of time, views of the world (i.e., class definitions) will change.

Relational database systems provide limited facilities to accommodate changes at the schema level. Because of the independence of relations, only a few types of schema change are provided, including adding a new relation to the database, dropping a relation from the database, adding a new column to an existing relation, and dropping a column from a relation. Further, a change to a relation does not impact on other relations or their tuples. The situation is more complex in object-oriented databases. The spectrum of possible modifications is wider due to the increased complexity of the model. Moreover, because of inheritance along the class hierarchy, a change to one class can affect other classes and their instances as well as its own instances.

The schema of an object-oriented database is the class hierarchy (or lattice). Accordingly, two kinds of changes to the schema are meaningful [42]: changes can be made to the class definitions, and changes can be made to the structure of the class hierarchy. Changes to the class definitions include adding and deleting attributes and methods. Changes to the class structure include creation and deletion of a class and alteration of the `is_a` relationship between classes (adding and deleting the superclass-subclass relationship between a pair of classes).

There have been several proposals for supporting schema evolution in object-oriented database systems. (See e.g. [55] for a bibliography.) The existing approaches can be classified according to the strategy they take to bring existing objects in line with a modified class. There are mainly two approaches [42]: *screening* and *conversion*.

Screening is to defer (possibly indefinitely) modifying instances. The values are either filtered or corrected before they are used. Several systems adopt this approach [41, 60]. Conversion is to require with each schema change that all instances be converted to the new class definition immediately after the change. Converting an instance means new attributes may be added into the object with default values and other attributes may be deleted from the object according to the new class definitions. Class change must be propagated to the subclasses. GemStone [52] is an example of an object-oriented database system using this approach as its class modification methodology.

Most of the existing works support single schema modification. That is, at every time instant, there exists only one logical schema that can be used to view the objects. Past states of the schema are not retained. There is another dimension to the schema evolution problem: *schema versioning*, where all states of the schema are accessible. Just as an object may be versioned, the schema itself may be versioned as well. In this case the system will manage more than one logical schema for one common database and present different views of the database through different versions of the schema. Although there is a considerable amount of work in single schema modification, there are only a few works on the semantics and implementation of schema versioning (e.g. [40, 60]).

There are two shortcomings to single schema modification. One is that it does not allow the history of modification of objects to be preserved. For example if an attribute of a class is dropped, the values of the attribute in existing instances are irretrievably lost; even if the attribute is added later, it will be treated as a different attribute and the old values of the attribute cannot be seen. Another shortcoming is that it does not prevent a schema change by one user from impacting all other users' views of the database. For example, once any user deletes an attribute or changes the superclass/subclass relationship between a pair of classes, all other users will see the changes.

Schema versioning removes the shortcomings of single schema modification. For ex-



ample, if a user wishes to drop an attribute from a class in one version of the schema, he may create a new version of the schema. The user will not see the values of an attribute in existing instances of the class. However if the user later chooses to access the class through the previous version of the schema, he will be able to see the values of the attribute in all instances of the class that existed before he created the new version of the schema.

We have extended the OEC to describe changes to the schema so that we can keep all states of the schema as well as the object states. This leads us to model schema versioning in a deductive framework. Our approach is different from the existing approaches since we develop the idea of having time-dependent views of the database.

## 9.2 Realization in the Event Calculus

The schema information in our framework consists of class definitions, which are asserted via **attribute** and **method** predicates, and the class hierarchy, which is defined by **is\_a** relationships. If we allow schema modifications, we have to consider a time-dependent **is\_a** relationship and time-dependent class definitions. When evaluating queries related to the states of instances, the state of the schema should also be considered. In querying the state of an object, its current class and the definition of the class will determine the valid attributes for that instance. Apart from querying instances at different times, querying the state of the database schema at different times will be possible.

We now have two levels of data that change: schema and objects. We use the object-based event calculus to describe changes at both levels. The object state history is described by a set of real world events and the schema state history is described by a set of system events. Once we are able to access all states of the schema, it is possible to derive the state of objects according to different schemas at different times. This provides us with schema versioning.

Schema changing events initiate and terminate periods of time for which a class is a subclass (**isa\_at**) of another class or a class has a certain attribute or method. We introduce time arguments to the **is\_a** and **attribute** relations to model this time-varying behaviour. The following predicates are used in reasoning with the schema states:

**attribute\_at**(Class, Attr, Type, T) : *Attr is an attribute (single or multi) of Class at time T.*

**method\_at**(Class, Obj, Mesg, Val, Body, T) : *Mesg is a method valid for Class at time T.*

**isa\_at**(Class, Super, T) : *Super is a superclass of Class at time T.*

With these predicates it is possible to keep the history of the class hierarchy and class definitions. In the following, the clauses for **is\_a** are presented. We have omitted the clauses for the other two predicates which are similar.

$$\begin{aligned} \text{isa\_at}(\text{Class}, \text{Super}, T) \leftarrow & \\ & \text{shappens}(\text{Ev}, T_s), T_s \leq T, \\ & \text{adds}(\text{Ev}, \text{Class}, \text{Super}), \\ & \text{not dropped}(\text{Class}, \text{Super}, T_s, T). \end{aligned}$$

```

dropped(Class, Super, Ts, T) ←
    shappens(Ev, T*), Ts ≤ T* < T,
    drops(Ev, Class, Super).

```

The occurrence of a schema event is recorded by the predicate **shappens**; a new predicate is introduced in order to avoid unnecessary search of real world events (recorded with **happens**) when the schema information is derived. The role of the predicates **adds** and **drops** is analogous to that of the predicates **initiates** and **terminates** respectively. They initiate and terminate periods of time for which a class is defined to be a subclass of another class. It is also possible to write the analogue of **holds\_for** to compute periods for these relations (e.g. **attribute\_for**, **isa\_for**) [33]. Note also that we could again reduce the number of different predicates, but our present aim is to emphasize the different conceptual aspects of a changing schema.

With the introduction of the time argument to the **is\_a**, **attribute** and **method** relations, inheritance becomes time dependent as well. For each state of the schema the structure of the instances may be viewed differently according to the definitions of their classes and superclasses. Thus we write:

```

attribute_at(Class, Attr, Type, T) ←
    isa_at(Class, Super, T),
    attribute_at(Super, Attr, Type, T).

method_at(Class, Obj, Mesg, Val, Body, T) ←
    isa_at(Class, Super, T),
    method_at(Super, Obj, Mesg, Val, Body, T).

```

Having two time dimensions gives the user the ability of querying the state of an object at time  $T$  according to the schema at time  $T_s$ . Thus a query:

```
?- holds_at(Class, Obj, Attr, Val, T, Ts).
```

asks for the state of an object **Obj** at time  $T$  according to the schema at time  $T_s$ . The period of time for which an object holds a particular state can be queried in a similar fashion:

```
?- holds_for(Class, Obj, Attr, Val, P, Ts).
```

In these queries the schema time  $T_s$  can be seen as a filtering mechanism for viewing the state of objects. In a model of the real world, each object has one history: by changing the schema we change our view of the real world, not the facts about objects. The facts about an object are visible only if the definition of the class of the object exists in the schema version at the specified time. For instance, in the above query the required information can be acquired only if the attribute **Attr** is valid for the class **Class** and also the object is an instance of that class according to the schema at time  $T_s$ .

The formulation of **holds\_at** and **holds\_for** are changed slightly to take the schema time into account:

$$\begin{aligned} \text{holds\_at}(\text{Class}, \text{Obj}, \text{Attr}, \text{Val}, T, T_s) \leftarrow \\ & \text{attribute\_at}(\text{Class}, \text{Attr}, T_s), \\ & \text{instance\_of}(\text{Obj}, \text{Class}, T, T_s), \\ & \text{happens}(\text{Ev}, T_1), T_1 \leq T, \\ & \text{initiates}(\text{Ev}, \text{Obj}, \text{Attr}, \text{Val}), \\ & \text{not broken}(\text{Obj}, \text{Attr}, \text{Val}, T_1, T). \end{aligned}$$

$$\begin{aligned} \text{holds\_for}(\text{Class}, \text{Obj}, \text{Attr}, \text{Val}, \text{during}(T_1, T_2), T_s) \leftarrow \\ & \text{attribute\_at}(\text{Class}, \text{Attr}, T_s), \\ & \text{initiated}(\text{Obj}, \text{Attr}, \text{Val}, T_1), \\ & \text{terminated}(\text{Obj}, \text{Attr}, \text{Val}, T_2, T_s). \end{aligned}$$

The introduction of a separate schema time requires some changes to the formulation of the `instance_of` relation:

$$\begin{aligned} \text{instance\_of}(\text{Obj}, \text{Class}, T, T_s) \leftarrow \\ & \text{happens}(\text{Ev}, T_1), T_1 \leq T, \\ & \text{assigns}(\text{Ev}, \text{Obj}, \text{Class}, T_s), \\ & \text{not removed}(\text{Obj}, \text{Class}, T_1, T, T_s). \end{aligned}$$

The schema time is passed as an argument of the `assigns` predicate. This accomplishes the establishment of the subset relation according to the schema version at that time.

Existing proposals in the literature provide different approaches to schema versioning. For example, in [60] each class, rather than the entire schema, is treated as a versionable object. Since the schema itself is not versioned, a ‘virtual’ version of the schema is constructed as a lattice of versioned class objects, having only one version of a class object included in any ‘virtual’ version of the schema. In [40], the entire schema is viewed as a versioned object. Any number of new versions of schema may be derived at any time from any existing schema version. The access scope of a schema version is the set of objects created under that version and those objects in the inherited access scopes of the ancestor schema versions. Thus it is possible to view and manipulate different sets of objects under different versions of schema.

In this section we have presented a different approach for dealing with versions of schema. We keep the history of the database schema by recording every event that causes a change in the schema. Thus it is possible to access different states of the schema at different times. The access scope of each schema version is the subset of all objects in the database whose classes are defined by the schema version. All updates to objects under one schema version will become visible to all schema versions which include the definition of the classes of the objects.

## 10 Conclusions

The integration of deductive and object-oriented approaches presents many open questions and issues to be overcome. The main problem is caused by the opposition between a value based approach and an identity based approach. After an analysis of existing proposals, we were attracted by the virtues of the simplest approaches where a framework for complex objects is built within first-order logic. The formula approach, exemplified by C-logic, gives

a semantics to an object by viewing it as a named collection of atomic formulas. Based on this approach we presented an object-based data model with support for object identities, single-valued and multi-valued attributes, class hierarchies, and derived-attribute methods.

The main contribution of this paper is the detailed development of the Object-based Event Calculus (OEC), which is intended as a general approach for representing and manipulating temporal objects in a logic programming framework. We have shown how the OEC may be used to represent and manipulate complex objects in a natural and descriptive way. We are not aware of any other work which deals with the different kinds of changes to objects in a single logical framework. We are also not aware of any other work which incorporates temporal information to the object states in a deductive framework.

From the representational point of view, there are some benefits offered by the OEC over the standard relational versions of the event calculus. Organising the specification of events by the class of object affected gives more structure to the representation, which can be of significant value in practical applications. Other benefits arise because the structure of objects, attributes, values and classes is richer than that of the relational data model, and this structure can be exploited. For example, the use of single-valued attributes and their treatment within the OEC reduces the need for general forms of integrity constraint, which otherwise are required for use with the event calculus. Similarly, much of the detail in the formulation of the OEC is concerned with the indirect effects of creating and deleting objects which themselves may be the values of attributes of other objects. As a result, the grouping of objects into classes and sub-classes gives a comparatively simple device for dealing with some of the more common types of indirect change, or ‘ramification’, a problem which in its general form is a topic of much current research in temporal reasoning.

From the point of view of temporal databases, we have presented an approach to the construction of historical (‘valid time’) databases in which all states of objects are stored (implicitly) and are accessible (by deduction). We addressed several different kinds of change that can be identified in the context of an object-based data model, and we proposed a general approach for modelling these changes in a declarative way. For practical application, we were particularly concerned to explain how the first-order semantics of objects allows implementations of the OEC to take advantage of indexing and other implementational techniques that are being developed for (relational) temporal databases.

## Acknowledgements

The bulk of this work was completed while Nihan Kesim was supported by the Scientific and Technical Research Council of Turkey (TUBITAK). The authors would like to thank an anonymous referee for suggesting a number of useful references to recent work in temporal databases and for many helpful suggestions which have improved the presentation of the paper.

## References

- [1] S. Abiteboul and S. Grumbach. COL : A logic-based language for complex objects. In *International Conference on Extending Database Technology – EDBT’88*, pages 271–293, Venice, Italy, March 1988.
- [2] M. Adiba and N.B. Quang. Historical multi-media databases. In *Proceedings of the 12th International Conference on VLDB*, pages 63–70, Kyoto, Japan, August 1986.
- [3] M.E. Adiba. Histories and versions for multimedia complex objects. *IEEE Database Engineering*, 7(4):181–188, 1988.
- [4] H. Ait-Kaci and R. Nasr. Login: A logic programming language with built-in inheritance. *The Journal of Logic Programming*, 1986.
- [5] M. Atkinson, F. Bancilhon, D. DeWitt, K. Dittrich, D. Maier, S. Zdonik. The Object-Oriented Database System Manifesto. In *Proceedings of the First International Conference on Deductive and Object-Oriented Databases*, pages 40–57, 1989.
- [6] F. Bancilhon. Object-oriented database systems. In *Proceedings of the 7th ACM-SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 152–162, Austin, Texas, March 1988.
- [7] D. Beech and B. Mahbod. Generalized version control in an object-oriented database. In *Proceedings of the 4th International Conference on Data Engineering*, pages 14–22, Los Angeles, CA, February 1988.
- [8] C. Beeri. Formal models for object oriented databases. In *Proceedings of the First International Conference on Deductive and Object-Oriented Databases*, Kyoto, Japan, December 4-6 1989.
- [9] A. Bjornerstedt and C. Hulten. Version control in an object-oriented architecture. In Won Kim and F.H. Lochovsky, editors, *Object-Oriented Concepts, Databases and Applications*, pages 451–485. ACM Press, 1989.
- [10] D. Bobrow and T. Winograd. An overview of KRL, a knowledge representation language. *Cognitive Science*, 1(1), 1977.
- [11] I. Cervesato, A. Montanari and A. Proveti. On the Non-monotonic Behaviour of Event Calculus for Deriving Maximal Time Intervals. *The International Journal of Interval Computation*, 2:83–119 (1994).
- [12] W. Chen, M. Kifer and D.S. Warren. Hilog: A first-order semantics for higher-order logic programming constructs. In *North American Conference on Logic Programming*, October 1989.
- [13] W. Chen and D. Warren. C-logic of complex objects. In *Proceedings of the 8th ACM SIGACT-SIGMOD-SIGART Symposium on the Principles of Database Systems*, 1989.
- [14] D. Chimenti, R. Gamboa, R. Krishnamurthy, S. Naqvi, S. Tsur, C. Zaniolo. The LDL system prototype. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):76–90, March 1990.

- [15] L. Chittaro, A. Montanari and A. Provetti. Skeptical and Credulous Event Calculi for Supporting Modal Queries. In *Proceedings of the European Conference on Artificial Intelligence ECAI'94*, Amsterdam, 1994.
- [16] Hong-Tai Chou and Won Kim. A unifying framework for version control in a CAD environment. In *Proceedings of the 12th International Conference on VLDB*, pages 336–344, Kyoto, Japan, August 1986.
- [17] J. Clifford and A. Croker. Objects in time. *IEEE Database Engineering*, 7(4):189–196, 1988.
- [18] J. Clifford and A. Croker. The Historical Relational Data Model (HRDM) Revisited. Ch.1 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [19] M. Dalal and D. Gangopadhyay. OOLP: A translation approach to object-oriented logic programming. In *Proceedings of the First International Conference on Deductive and Object-Oriented Databases*, pages 555–568, Kyoto, Japan, December 4-6 1989.
- [20] M. Denecker, L. Missiaen and M. Bruynooghe. Temporal reasoning with abductive event calculus. In *Proceedings of the European Conference on Artificial Intelligence ECAI'92*, Vienna, 1992.
- [21] R. Elmasri, G.T.J. Wu and V. Kouramajian. A Temporal Model and Query Language for EER Databases. Ch.9 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [22] R. Elmasri, G.T.J. Wu and V. Kouramajian. The Time Index and the Monotonic B<sup>+</sup>-tree. Ch.18 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [23] K. Eshghi. Abductive planning with event calculus. In *Proceedings of the International Conference on Logic Programming 1988*, MIT Press, 1988.
- [24] C. Evans. The Macro-Event Calculus: Representing Temporal Granularity. In *Proceedings Pacific Rim International Conference on AI*, Nagoya, Japan, 1990, pp363–368.
- [25] K. Fukunaga and S. Hirose. An experience with a Prolog-based object-oriented language. In *OOPSLA '86 Proceedings*, pages 224–231, 1986.
- [26] R.V. Indiketiya. *Event Calculus Based Temporal Database Management System*. MSc Thesis. Department of Computing, Imperial College, London, 1992.
- [27] C.S. Jensen, J. Clifford, R. Elmasri, S.K. Gadia, P. Hayes, S. Jajodia. A glossary of temporal database concepts. *SIGMOD Record*, 23(1), March 1994.
- [28] C.S. Jensen and L. Mark. Differential Query Processing in Transaction-Time Databases. Ch.19 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.

- [29] C.S. Jensen and R.T. Snodgrass. Unifying temporal data models via a conceptual model. *IEEE Information Systems* 19(7):513–547, 1994.
- [30] C.S. Jensen and R.T. Snodgrass. Temporal specialization. In *Proceedings of the International Conference on Data Engineering*, Tempe, AZ, Feb.1992, IEEE, pp 594–603.
- [31] C.S. Jensen, M.D. Soo and R.T. Snodgrass. Temporal specialization. In *Proceedings of the International Conference on Data Engineering*, Tempe, AZ, Feb.1992, IEEE, pp 594–603.
- [32] W. Kaefer, N. Ritter and H. Schoening. Support for temporal data by complex objects. In *Proceedings 16th International Conference on Very Large Data Bases*, Brisbane, 1990.
- [33] F.N. Kesim. *Temporal Objects in Deductive Databases*. PhD thesis, Department of Computing, Imperial College, London, 1993.
- [34] F.N. Kesim and M.J. Sergot. On the evolution of objects in a logic programming framework. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, volume 2, June 1992.
- [35] F.N. Kesim and M.J. Sergot. Versioning of objects in deductive databases. In *Proceedings of the Third International Conference on Deductive and Object-Oriented Databases*, December 1993.
- [36] F.N. Kesim and M.J. Sergot. Implementing an object-oriented deductive database using temporal reasoning. Department of Computing, Imperial College, London, July 1994. Submitted for publication.
- [37] M. Kifer and G. Lausen. F-logic: A higher-order language for reasoning about objects, inheritance, and scheme. In *Proceedings of the 8th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 134–146, 1989.
- [38] M. Kifer, G. Lausen and J. Wu. Logical foundations of object-oriented and frame-based languages. Technical report, Department of Computer Science, SUNY at Stony Brook, June 1990.
- [39] M. Kifer and J. Wu. A logic for object-oriented logic programming (Maier’s O-logic revisited). In *Proceedings of the 8th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, 1989.
- [40] W. Kim and H.T. Chou. Versions of schema for object-oriented databases. In *Proceedings of the 14th International Conference on VLDB*, pages 148–159, Los Angeles, CA, 1988.
- [41] W. Kim, J.F. Garza, N. Ballou, D. Woelk. Architecture of the ORION next-generation system. *IEEE Transactions on Knowledge and Data Engineering*, 2(1):109–124, 1990.
- [42] Won Kim and F.H. Lochovsky, editors. *Object-Oriented Concepts, Databases, and Applications*. ACM Press, 1989.

- [43] N. Kline. An update of the temporal database bibliography. *SIGMOD Record*, 22(4), December 1993.
- [44] R.A. Kowalski. Database updates in the event calculus. *Journal of Logic Programming*, 12:121–146, 1992.
- [45] R.A. Kowalski and F. Sadri. The Situation Calculus and Event Calculus Compared. In *Proceedings of the International Symposium on Logic Programming 1994*, M. Bruynooghe (ed.), MIT Press 1994, pp539–553.
- [46] R.A. Kowalski and M.J. Sergot. A logic-based calculus of events. *New Generation Computing*, 4:67–95, 1986.
- [47] G.M. Kuper. Logic programming with sets. In *Proceedings of the 6th ACM-SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, San Diego, CA, 1987.
- [48] T.Y.C Leung and R.R. Muntz. Stream Processing: Temporal Query Processing and Optimization. Ch.14 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [49] D. Maier. A logic for objects. In *Proceedings of the Workshop on Foundations of Deductive Databases and Logic Programming*, pages 6–26, Washington D.C., August 1986.
- [50] F.G. McCabe. *Logic and Objects*. Prentice Hall, 1992.
- [51] A. Montanari, E. Maim, E. Ciapessoni, E. Ratto. Dealing with Time Granularity in the Event Calculus. In *Proceedings of the Fifth Generation Computer Systems Conference FGCS'92*, Tokyo, 1992, pp702–712.
- [52] D.J. Penney and J. Stein. Class modification in the GemStone object-oriented DBMS. In *OOPSLA '87 Proceedings*, pages 111–117, 1987.
- [53] J. Pinto and R. Reiter. Temporal reasoning in logic programming: a case for the situation calculus. In *Proceedings of the International Conference on Logic Programming 1993*, MIT Press, 1993, pp203–221.
- [54] A. Provetti. Hypothetical reasoning about actions: From situation calculus to event calculus. *Computational Intelligence*, 12:2:1–24.
- [55] J. Roddick. Schema evolution in database systems - an annotated bibliography. In *SIGMOD Record*, volume 21, pages 35–40, December 1992.
- [56] E. Rose and A. Segev. TOODM—A temporal object-oriented data model with temporal constraints. In *Proceedings of the 10th International Conference on the Entity-Relationship Approach*, October 1991, pp205–229.
- [57] E. Rose and A. Segev. TO-Algebra—A temporal object-oriented algebra. Technical Report LBL-32013, University of California, Berkeley, January 1992.



- [58] A. Segev and A. Shoshani. A Temporal Data Model Based on Time Sequences. Ch.11 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [59] M.P. Shanahan. Representing continuous change in the event calculus. In *Proceedings of ECAI-90*, Stockholm, Sweden, 1990.
- [60] A.H. Skarra and S.B. Zdonik. Type evolution in an object-oriented database. In B. Shriver and P. Wegner, editors, *Research Directions in Object-Oriented Programming*, pages 393–413, MIT Press, 1987.
- [61] R.T. Snodgrass. Temporal Object-Oriented Databases: A Critical Comparison. Ch. 9 in W. Kim (ed.), *Modern Database Systems: The Object Model, Interoperability and Beyond*, Addison-Wesley/ACM Press, 1994.
- [62] P.J. Soper, G. Abeysinghe and C. Ranaboldo. A temporal model for clinical and resource management in vascular surgery. In *Proceedings of the International Conference on Database and Expert Systems Applications*, pages 549–552, Berlin, 1991.
- [63] S.M. Sripada. A logical framework for temporal deductive databases. In *Proceedings of the International Conference on Very Large Databases*, pages 171–182, 1988.
- [64] S.M. Sripada. *Temporal Reasoning in Deductive Databases*. PhD thesis, Department of Computing, Imperial College, London, 1991.
- [65] S.M. Sripada. A Metalogic Programming Approach to Reasoning about Time in Knowledge Bases. In *Proceedings of IJCAI'93*, Chambery, 1993, pp860–865.
- [66] M. Stefik and D.G. Bobrow. Object-oriented programming: Themes and variations. *AI Magazine*, January 1986.
- [67] M. Stonebraker, L.A. Rowe., B. Lindsay, J. Gray, M. Carey, M. Brodie, P. Bernstein, D. Beech. Third-Generation Database System Manifesto. In *SIGMOD Record*:19(3):31–44, September 1990.
- [68] J. Su. Dynamic constraints and object migration. In *Proceedings of the Conference on Very Large Data Bases*, Barcelona, 1991.
- [69] A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [70] D.S. Warren. Memoing for Logic Programs. *Communications of the ACM* 35:3:94–111, 1992.
- [71] G. Wiederhold, S. Jajodia and W. Litwin. Integrating Temporal Data in a Heterogeneous Environment. Ch. 22 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [72] K. Wilkinson, P. Lyngbaek and W. Hasan. The Iris Architecture and Implementation. *IEEE Transactions on Knowledge and Data Engineering* 2(1): 63–75, March 1990.

- [73] G.T.J. Wu and U. Dayal. A Uniform Model for Temporal and Versioned Object-oriented Databases. Ch.10 in A. Tansel, J. Clifford, S. Gadia, S. Jajodia, A. Segev, R. Snodgrass (eds.). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings, 1993.
- [74] C. Zaniolo. The representation and deductive retrieval of complex objects. In *Proceedings of Very Large Databases*, page 458, Stockholm, 1985.
- [75] C. Zaniolo. Object-oriented programming in Prolog. In *Proceedings of the 1984 International Symposium on Logic Programming*, Atlantic City, New Jersey, February 1984.
- [76] S.B. Zdonik. Object-oriented type evolution. In F. Bancilhon and P. Buneman, editors, *Advances in Database Programming Languages*, pages 277–288. ACM Press, 1990.
- [77] S.B. Zdonik and D. Maier, editors. *Readings in Object-Oriented Database Systems*, chapter 4, page 239. Morgan Kaufmann, 1990.