

Nelson Mendonca Mattos

University Kaiserslautern, Department of Computer Science,  
P.O. Box 3049, D-6750 Kaiserslautern, West Germany  
e-mail: mattos@uklirb.uucp

## Abstract

This paper discusses architectural issues of Knowledge Base Management Systems and describes the architecture of KRISYS, a system whose goal is the effective and efficient management of large, shared knowledge bases. Focal points are primarily the design decisions and the system's features: knowledge independence, object-orientation, mechanisms for knowledge organization, data-driven mutation, inheritance mechanisms, reasoning facilities, etc. Generally, some of these issues which we have been combining to form a KBMS context, are similar to approaches developed in different isolated projects. Instead of giving a detailed comparison of these approaches and those of other projects, we show in this paper how they can be architecturally combined to build realistic KBMS. Keywords: AI architectures, Knowledge Base Management Systems, Database support for Knowledge-Based Systems

## 1. Introduction

Technology has produced a variety of knowledge-based systems ranging from simple expert systems to complex natural language understanding systems. When used for large-scale applications, KS are faced with problems of managing very large volumes of knowledge: virtual memory sizes are not large enough to store the knowledge to be handled, and operations on knowledge (i.e. inference) are computationally intolerable when knowledge bases are maintained on secondary storage devices.

These problems show that the applicability of KS is limited, since appropriate systems for the efficient knowledge management do not exist. Approaches combining KS with traditional Database Management Systems (DBMS) for this purpose have failed for several reasons (see [1] for a description of the deficiencies of DB support for KS).

A solution to KS problem is to develop a new generation of systems aimed at the efficient management of large, shared KB. Analogous to DBMS, these systems are called Knowledge Base Management Systems (KBMS) [2].

This observation has motivated our research efforts to identify important design issues regarding an architecture for KBMS. Our investigations have shown that these issues are strongly influenced by three classes of requirements. Firstly, KBMS should satisfy the requirements of their applications (i.e. KS or end-users). Secondly, they must support the needs of the KB-designer, who plays a very important supporting role in this context. And thirdly, some implementation aspects in terms of data structures and algorithms should be taken into account, in order to manage the knowledge efficiently. Thus, KBMS should provide features obtained from three different points of view. So, we believe that KBMS should be architecturally divided in three different layers: implementation layer, engineering layer and application layer, which respectively support the above classes of requirements.

We feel that the above points are involved with a number of quite new ideas, which we incorporated in a multi-layered prototypical KBMS. Clearly, some of these ideas, which we combined to apply to a KBMS context, are related to approaches developed in different projects [3,4]. However, these projects handle one or more of these ideas in an isolated manner, not taking into account the practical use of them in KBMS. In this paper, we describe our prototypical architecture for KBMS, called KRISYS, showing above all how these and other approaches can be combined in order to build realistic KBMS. To motivate our ideas, we first review the results of our investigation: the KBMS architectural issues that arise when addressing the above mentioned three classes of requirements.

## 2. Architectural Issues

Traditionally, knowledge representation systems or KS components responsible for the knowledge management have been designed just to support what we called the engineering layer. Since the manner of knowledge organization and access is visible at the external interface of those systems, KS possess information about the knowledge organization and retrieval possibilities built in their logic (i.e. embedded in their programs). Any modification of the knowledge structures therefore requires program modifications. Changing the kind of representational framework used, for example from frames to semantic networks, would be even impossible, since this would mean throwing the entire KS away and implementing a new one.

However, with complex KS and very large KB on the horizon, the dependence on the framework supported by a knowledge representation system promises to be very problematic. *Knowledge independence* as an analogy to data independence seems to be the key answer to this problem. Knowledge representation systems must be characterized not in terms of the representational framework they use, but functionally, in terms of what they know about the KS domain.

This idea of abstraction aimed at the independence of knowledge motivated us to introduce another layer (the application layer) on top of the engineering layer in the architecture of KBMS. At the object-orientation interface, KS can, therefore, work independently from the specification of the representational framework supported by the engineering layer. In fact, KS are not interested in things like the complexity of frame structures, the variety of links in a semantic network, the properties of inheritance mechanisms or the power of reasoning facilities. They are really interested in what they can ask or tell the KBMS about the knowledge of their domain, which is stored in the KB managed by the KBMS. This motivated us also to believe that the KS interface is the one supported by the application layer (object-abstraction interface). The interface for the KB-designer is, however, a different one. He (in opposition to the KS) is concerned with the aspects of the representational framework. He decides, whether a specific information is to be represented a

## 2.1 The Application Layer

At knowledge independence, the application layer should be knowledge functionally, in terms of only two basic types of operations supported at this level: one to enable the KS to ask the KBMS questions to be answered on the basis of the knowledge kept in the KB and another one to permit the KS to tell the KBMS new knowledge to be maintained in the KB. Therefore, at the object-abstraction interface the KB can be compared with an abstract data type that interacts with the end users or KS through a set of "ask" and "tell" operations. Thus, the way the knowledge is captured or organized is hidden from the KS. Because of this, no distinction can be made between knowledge that is extensionally stored and that which has to be implied (intensionally). Whether just simple retrieval capabilities or inference of some kind is required to answer a question, is not to be decided by the KS, but in the layer.

The functional view of the KB has also been introduced in KRYPTON [3], where ask and tell operations are split in two different interfaces: a terminological Tbox, and an assertional Abox. However, the issues supported by KRYPTON do not meet the idea of knowledge independence. Before making changes in a KB a KRYPTON designer has to decide for example, whether the KB's theory of the domain should imply these changes (tell operation at the Abox) or the KB vocabulary should include them (tell operation in the Tbox). In other words, he has to decide whether these changes should be extensionally represented or intensionally. We argue that such a decision should not be made by the user but in the application layer, so that the way knowledge is represented can be hidden from the KBMS users. We believe that this is the only possible way to view the KB as an abstract data type. Furthermore, KRYPTON's interface has not been developed for a KBMS context and consequently is neither flexible nor powerful enough to meet the requirements of real world applications. (For example, ask statements are restricted to simple questions).

The query language provided at the object-abstraction interface naturally permit adequate and flexible ways to select the application's domain knowledge and to cause changes in the KB corresponding to the many ways knowledge can be accessed and organized. Furthermore, it must be set-oriented. This permits the KS to obtain several "pieces-of-knowledge" with just one operation, reducing communications overhead between KS and KBMS and offering an enormous optimization potential for the lower layers. Additionally,

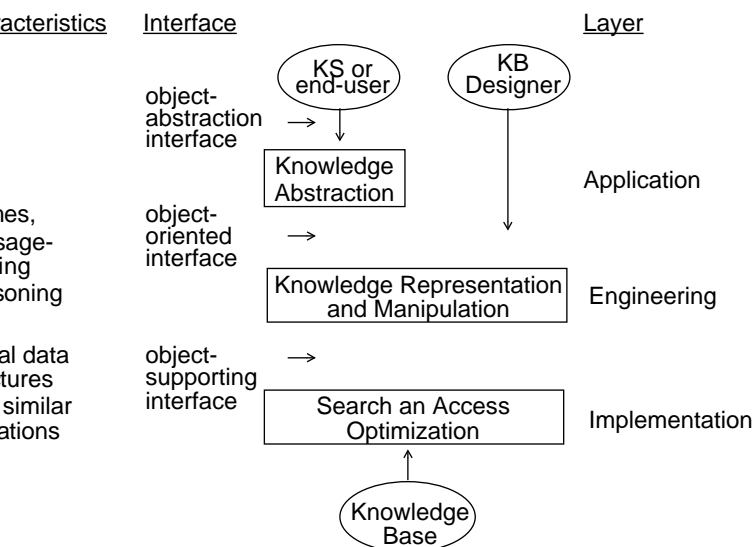


Figure 1: Overall architecture of a KBMS

## 2.2 The Engineering Layer

The engineering layer focusses on a KBMS from the point of view of the KB-designer. Here the KBMS is seen in terms of how the knowledge of an application can be represented, organized, and manipulated. In other words, this layer deals with descriptive, organizational and operational aspects of a knowledge representation model supported for the KB-designer at the engineering layer's interface (object-oriented interface). Concerning the descriptive principles:

- Entities in the application's world should correspond to exactly one object in the model.
- Each object should be composed of attributes expressing the properties of the entities.
- Attributes are multiple-valued and should be further described (e.g. data type of the attribute values, etc.).
- Objects are described by object types (i.e. classes), which should also be objects of the model.
- Classes characterize prototypical instances, i.e., properties described in the classes might be contradicted by particular instances.
- Objects can be instances of several classes.
- Relationships between entities should be expressed as properties of objects [5].

Concerning the organizational principles, the abstraction concepts [6,7,8], extensively used in AI knowledge representation systems, are the best candidates:

- Objects are to be grouped into classes (classification), which, in turn, should be organized into class hierarchies providing the means for the overall organization of the KB (generalization).
- Inheritance must be supported, in order to eliminate unnecessary definitions. Multiple inheritance should also be allowed and the notion of inheritance should be that of default [9,6].
- Similar element objects should be organized into set-element hierarchies (association).
- Aggregation of component objects should be supported, and might be applied recursively, so that another organizational dimension of the KB can be defined.

Operational issues should primarily support object orientation. This will allow the attributes of the objects to store either extensional or intensional data, which are accessed in a uniform way (i.e. sending object messages). Moreover, object orientation makes it possible to extend the functionality of the layer interface very easily, since it permits the KB-designer to define his own functions, which as attributes of objects can also be directly addressed at the interface by message passing. Reasoning facilities should also be provided. The knowledge representation model should offer already implemented general purpose reasoning functions, not prohibiting however, that these functions may be substituted by special ones developed by the designer. Mechanisms for checking integrity constraints are to be provided too. To define constraints this layer should enable the KB-designer to define procedures that are linked to attributes of objects and are automatically activated before or after such attributes are accessed or updated. These procedures attached to data will then perform the necessary checks, and the corresponding actions to keep the KB in a consistent state.

supply to the engineering and application layers. For this reason, many of the issues here are related to traditional DB prob-

requisite to identify the specific requirements to be supported is the investigation of the behavior of KS running on second-storage environments. In order to pursue this investigation, we adapted some available KS to work on such environments [10]. Observations may be summarized as follows:

The accesses made at the implementation layer interface (object-supporting interface) are mainly to tiny granules referring to individual attributes of objects rather than to objects as a whole.

The attributes' access frequencies differ very much. Dynamic attributes (i.e., the specific knowledge of a consultation) are accessed with very high frequencies; static attributes (i.e., the expert knowledge of the KS) on the other hand with very low frequencies.

The dynamic knowledge can be kept temporarily, since it only expresses information of a particular consultation.

In multi-user environments, the KBMS must maintain as many versions as the number of users working with the KB for the dynamic knowledge. These versions are then accessed individually, and that synchronization should only be controlled for the static knowledge.

In each phase of the problem solving process, accesses concentrate on just some objects of the KB. These objects build a context which contains the knowledge needed to infer the specific goal of that phase. The context needed in each phase can be established from some information that the KS deduces during the preceding phases (i.e., dynamic knowledge), so that it is possible to determine the next context needed at the end of a phase. This enables the dynamic preplanning of the accesses to the KB that should be made at this layer for optimization purposes.

The storage structures for the KB should give more priority to the optimization of the retrieval operations, since modifications in the structures of the KB (e.g. changing object types) are very seldom.

### 3. The KRISYS-Prototype

#### 3.1 Overview of the System Architecture

KRISYS is architecturally divided in three different modules, which correspond to the aspects of the previously discussed layers (Figure 1). The application layer corresponds to the Query-Processing Manager (Figure 2). Here, knowledge independence is supported by the KRISYS Object-Abstraction Language (KOALA), which characterizes the object-abstraction interface. To implement the engineering layer, we choose a Frame-System (KRISYS Frame System KFS). This, in turn, supports the interface to the KB-Designer and to the Query-Processing Manager, i.e., the object-oriented interface. The implementation layer requirements are fulfilled by the Working-Memory and Context Manager (WMCM) and by a DBMS Kernel. The WMCM embodies the nearby application locality concept, i.e., the reservation of the locality of the application's object processing

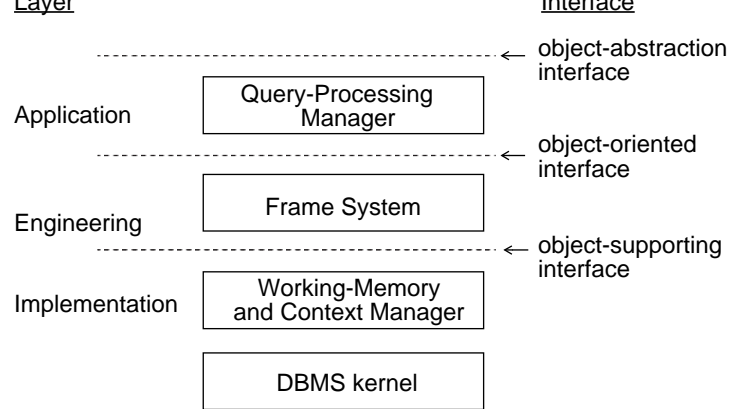


Figure 2: The KRISYS Architecture

#### 3.2 Query-Processing Manager

The application layer of KRISYS can be better characterized by a description of its object-abstraction interface, the language KOALA, which defines all interactions between a user or application and the system, which occur with two types of operations:

- ASK** - to ask KRISYS, whether an "expression" is true, retrieving the qualified knowledge, a
- TELL** - to tell KRISYS, that an "expression" is true, changing the KB appropriately

##### The Ask-Statement

An Ask-statement contains two parts; a projection and a selection (i.e. qualification) part.

ASK[<projection>][<selection>]

The last one specifies the expression to be proved, which is structured in accordance to the first-order predicate calculus. The selection is evaluated following the model-theoretic approach [11]. In other words, ask-statements are evaluated as true or false with respect to the KB. During evaluation, the Closed World Assumption (CWA) [12], the Unique Name Assumption (UNA), and the Domain Closure Assumption (DCA) are taken as granted. For example, the query (ASK (IS\_INSTANCE tweety penguin)) would be true only if there exists both tweety and penguin as objects of the KB related to the abstraction concept of classification.

The logical formulas are built with predicates (e.g. IS\_INSTANCE in the above example). These use either constants, variables or functions as their terms. The last one can, in turn, use variables or constants as terms again, in order to address the objects of the KB. KOALA supports functions and predicates to be applied to objects (schemas), attributes (slots), attribute values (slot values) and attribute descriptions (aspects). Some important object predicates are those used to express the abstraction concepts of classification: instantiation (IS\_INSTANCE) generalization/specialization (IS\_SUPERCLASS, IS\_SUBCLASS) and association (IS\_ELEMENT, IS\_SUBSET). The next example shows a query which asks whether or not elephant is a specialization of mammal and an element of the animals\_from\_africa:

```
(ASK (AND (IS_SUBCLASS elephant mammal)
          (IS_ELEMENT elephant animals_from_africa)))
```

To manipulate attributes, KOALA offers functions to select all slots of an object, only the inherited ones, the inherited ones from a particular object, the slots defined for an object (i.e. not inherited), instance\_slots, its own\_slots, its standard\_slots or its user\_defined\_slots. These functions have as their result a set of slots which are to be interpreted as the components of the specified object (abstraction concept of aggregation). If someone would like to a

(IS\_IN nutriment (INHERITED\_SLOTS camel  
(SUPERS mammal))))

itions about slot values can be expressed by a lot of predicates,  
n use functions to access the slot value itself (e.g. SLOTVALUE  
<slot\_name> <schema\_name>). The equality predicate can be used  
for example, to express that a slot value is equal to a specific  
:

(EQUAL milk (SLOTVALUE nutriment camel))

erical comparison predicates are also provided. Furthermore,  
other functions, it is possible to evaluate arithmetic expres-  
such as, sum, average, minimum and maximum value of a set  
members or even to count the number of elements of a set.

pute description predicates allow the application to formulate  
itions about slot value restrictions, cardinality specifications,  
for example, the query

(IS\_POSSIBLE\_VALUE fly (ASPECTVAL  
POSSIBLE\_VALUE  
form\_of\_locomotion mammal))

s, whether the value fly is allowed in the slot  
of\_locomotion of the object mammal.

orientation is achieved by using variables in the logical formu-  
variables express either objects, attributes, attribute values or  
pute descriptions of the KB. During evaluation they are instan-  
l with all values, which satisfy the logical formulas. Because of  
CA, formulas, which contain variables that cannot be instan-  
l, are assigned 'false'. To ask whether there exists any bird  
cannot fly, one can use variables to create the following ques-

(EXIST ?X (AND (IS\_INSTANCE ?X bird)  
(NOT(IS\_IN fly (SLOTVALUE  
form\_of\_locomotion ?X)))))

bles can also be used to construct questions involving more  
one object (roughly analogous to a database join). The user can  
build queries combining variables that represent different  
s (i.e. objects, attributes, attribute values or attribute descrip-  
). For example, to know whether there is any object in the KB  
has stored a specific value in any of its attributes, someone  
d pose the question

(EXIST ?X ?Y (AND (IS\_SCHEMA ?X)  
(IS\_IN ?Y (ALL\_SLOTS ?X))  
(EQUAL specific\_value  
(SLOTVALUE ?Y ?X)))))

ently, the user is also interested to see the instances, which  
y his logical formulas (e.g. to identify the dogs descending from  
ame parents), in addition to the boolean value of these formu-  
le can specify this in the projection part of the ask-statement.  
example, if someone would like to know which mammals can  
, he should ask

(?X) ( AND (IS\_INSTANCE ?X mammal)  
(IS\_IN swim (SLOTVALUE  
form\_of\_locomotion ?X)))))

projection also enables the user to express exactly which parts  
accessed information (objects, attributes etc.) he would like to  
This can be achieved by using specific projection clauses or by  
ining different variables (e.g. variables that represent objects  
those that represent attributes). For example, if someone  
s to see the dogs stored in a KB together with the value of the  
they have in common with cats, one would ask KRISYS the fol-  
g query:

Recursive queries may be expressed using the PATH predic-  
with which a special class of recursion equations, called generali-  
transitive closure (GTC) [13], can be specified. Unfortunately,  
to space limitations it is not possible to discuss these aspects he-  
(For a detailed description see [1]).

### The Tell-Statement

The tell-statement takes a sentence and asserts that it is true. T-  
effect is to change the KB into one whose contents imply that s-  
entence. Naturally, it can also occur that the contents of the KB  
already imply the sentence asserted. In this case, no changes will  
made, since they are not required. For example, if the KS asse-  
that penguin is a subclass of bird, expressing "(TE  
(IS\_DIRECT\_SUBCLASS penguin bird))" two situations can occ-  
Either the KB already contains this information, requiring  
changes, or the KB does not contain it and consequently chan-  
must be made. Here, many things can happen. If neither bird  
penguin exist as objects in the KB, both will be created and rela-  
to each other by the abstraction concept of generalization. If o-  
one of them exists, the other one will then be created and related  
the first one as specified above. And if both of them exist only  
generalization relationship will be built.

TELL<assertion>[WITH<selection>]

The sentences to be asserted are specified in the assertion par-  
the tell-statement. An assertion is syntactically similar to the se-  
tion part (see ask-statement), however, much more restrictive. It  
for example, not possible to specify formulas combined with logi-  
connectors, in order to avoid ambiguities (note, that if someone  
serts p v q, it is impossible to know whether p, q or both are tru-  
Nevertheless, several assertions can be specified, which will then  
interpreted independently, i.e., KRISYS will make each one tru-

(TELL (IS\_DIRECT\_SUBCLASS penguin bird)  
(IS\_DIRECT\_INSTANCE tweety penguin)  
(EQUAL frankfurt\_zoo (SLOTVALUE address tweety)))

If the same assertions should be made valid for several objects, a  
can specify this requirement in just one tell-statement by us-  
variables to represent these objects. The variables would then be  
stantiated during the evaluation of the selection part, which has  
same syntax and works just like the selection of the ask-statement.  
After this evaluation, the assertions will then be applied for each  
stance of the variables that satisfy the selection. The following  
ample shows a query expressing the assertion that every anim-  
that lives in water (here represented as elements of the set wa-  
animal) can swim:

(TELL (IS\_IN swim (SLOTVALUE locomotion ?X))  
WHERE (IS\_ELEMENT ?X water\_animal)).

Another example for the use of variables shows the following qu-  
which expresses that every owner of a dog of the KB should also  
an object of the KB:

(TELL (IS\_SCHEMA ?X)  
WHERE (FOR ALL ?Y WITH (IS\_INSTANCE ?Y dog)  
(IS\_IN ?X (SLOTVALUE owner ?Y)))).

Summarizing this chapter, we can say that the object-abstract  
interface is completely specified by the above two operations. "T-  
takes a sentence and asserts that it is true, changing the KB i-  
one whose contents imply the assertion. "Ask" takes a sentence a-  
checks on the basis of the current contents of the KB whether i-  
true or not. Schematically, they can be described as

**TELL : KB x sentence ==> KB Sentence is true !**  
**ASK : KB x sentence ==> T/F Is sentence true ?**

hidden from the user. The KB is characterized as an abstract type, specified only by these two operations rather than by a n implementation structure.

### 3.3 The Frame-System

Choose a frame system [4] to implement our engineering layer, use we believe that it offers the necessary framework to represent the descriptive, organizational and operational aspects of the main knowledge of any KS.

#### Object-Oriented Representation

In the KRISYS Frame-System (KFS) the three aspects mentioned are incorporated in its basic concept: the schema. A schema is a symbolic representation of a real world entity. It is composed of a schema name and of a set of attributes. Attributes can be of two types:

- **Instance-Attributes**, used to describe the descriptive and organizational aspects of the schema, and
- **Operational-Attributes**, used to describe its operational aspects.

Therefore, KFS could be named object oriented. It allows both attribute-like characteristics and procedural properties of the real world objects to be integrated into a schema.

Operations in KFS occur by sending messages to the objects, which can, in turn, communicate with other objects by message passing. This approach supports the important principle of data abstraction. That is, in order to request the methods of some objects to be performed, no assumptions have to be made about the implementation and internal representation of the objects (i.e. slots and methods).

Attributes and methods have the same structure. Both possess name, type, and a schema-name, that specifies where the attribute is defined (origin). The type indicates whether an attribute represents characteristics of the schema itself, or of its instances. Characteristics of a schema are described by the types OWNSLOT and OWNMETHOD, whereas the instance ones by INSTANCESLOT and INSTANCEMETHOD, respectively.

Associated with the attributes there can be aspects, which are used to describe the attribute more exactly. Five aspects are predefined for slots: possible-values, cardinality, comment, default and comment. Figure 3 shows the structure of the schema elephant. Since the attributes shown in the example describe characteristics of the schema itself, they have type OWNSLOT or OWNMETHOD (the meaning of the slots INSTANCE-OF and ELEMENT-OF will be defined later).

elephant			
INSTANCE-OF	mammal	OWNSLOT	GLOBAL
ELEMENT-OF	savanna-animal	OWNSLOT	GLOBAL
form-of-locomotion	(walk swim)	OWNSLOT	animal
POSSIBLE_VALUES	(walk fly swim)		
CARDINALITY	[1,3]		
DEFAULT	(walk)		
habitat	(savanna)	OWNSLOT	animal
POSSIBLE_VALUES	INSTANCE-OF	habitat-class)	
CARDINALITY	[0, ∞]		
COMMENT	(This slot contains the places, where this animal is found in the nature)		
eat	(<code>)	OWNMETHOD	animal
PARAMETER	(food)		

Figure 3: Example Definition of the Schema Elephant

slots, that are automatically activated when these slots are accessed. This concept of data-oriented computation is very useful to represent intensional data or to check complex integrity constraints. For example, if someone wants to represent some information whose value changes with respect to other data (e.g. the exact age of a person, which changes every day), he needs a mechanism that generates the extensional value of this information automatically each time this information is accessed. This can be realized in KFS by using demons.

Demons are stored in a KFS schema (like any other), which has however, the predefined schema DEMON as its most superclass. As an instance of DEMON this schema inherits the method attributes GET, PUT, ADD and RETRACT, where the code for the respective demons are stored (represented by schema D1 in Figure 4).

The linkage between slot and corresponding demons is done by specifying in the demons aspect of the slot (S1) the schema name of the respective demons (D1). Whenever an access to the slot is made KFS checks if there is a demons aspect defined in this slot, activating by sending a message to the schema specified there, demanding evaluation of the corresponding demon. That is, the method specified in GET will be invoked if a get-access has been issued to KFS, the method PUT is invoked by issuing a put-access etc. By not storing the demon-code directly in the aspect of the slot, as is done in many systems, KFS allows many slots to use the same demon without having to introduce redundancy in the representation. KFS additionally allows the KB-designer to specify when the demon should be activated (i.e. before or after the access to the slot). This flexibility allows the use of demons for many different purposes. Demons applied to check integrity constraints should be activated before updating the value of a slot, whereas those used to trigger actions when particular slots are accessed should be activated after the access itself.

```

schema-x
:
S1 <S1-value> OWNSLOT <S1-origin>
POSSIBLE-VALUES (...)
:
DEMONS D1
:
:
D1
INSTANCE_OF (DEMON)
GET <get-code> OWNMETHOD DEMON
PUT <put-code> OWNMETHOD DEMON
ADD <add-code> OWNMETHOD DEMON
RETRACT <retract-code> OWNMETHOD DEMON

```




Figure 4: Activation of Demons

#### Knowledge Organization

In order to enable the KB-designer to structure the KB adequately, KFS supports various abstraction concepts. These are represented as relationships between objects specified by predefined slots, which are contained within every schema. The abstraction concept of generalization is represented by the slots SUBCLASS\_OF and HAS\_SUBCLASSES. INSTANCE\_OF and HAS\_INSTANCES are used to specify the classification relationship. The association concept

any special slot for its representation.

ating objects to another, hierarchies are built, where a particular schema can play the role of a class (i.e. object type), an instance, a set and/or an element. Since a schema is able to reflect an instance and a class as well, we need two different slot and method to know whether a slot represents characteristics of itself (INSTANCESLOT and OWNMETHOD) or of its instances (INSTANCEMETHOD and INSTANCEMETHOD).

### heritance

of the advantages of organizing the schemas into abstraction hierarchies are the built-in reasoning facilities they provide [6]. By coupling objects with abstraction relationships, special "automatic" reasoning facilities are supported as part of each modification and evaluation operation. The most usual and important of these reasoning facilities is the one built-in the is-a hierarchy i.e. inheritance. In this approach, only instanceslots and instancemethods are inherited, since ownslots and ownmethods represent characteristics of their own objects. Between schemas linked with association relationships no inheritance occurs. Here, as well as over aggregation relationships, several kinds of built-in reasoning facilities are provided. For a detailed description of these facilities see [6].

### Reasoning

KFS also supports general reasoning facilities. Rules in KFS are represented as schemas, which are also organized in special hierarchies having two predefined schemas as superclasses. One of these schemas specifies the structure of the rules, i.e. this schema (RULE-SETS) contains some instanceslots (condition, action, etc.) which are inherited by each rule. In order to construct a rule base, the KB designer has only to store the particular contents of the condition, action, etc. in the respective slot of each rule. The other top level schema (RULE-SETS) contains methods corresponding to the standard inference strategies, which are then inherited by each set of

particular rule of KFS can, therefore, be used in both reasoning directions (i.e. forward-reasoning and backward-reasoning) without having to be stored redundantly. The reasoning direction of the rules are dynamically defined when the KS requests the activation of one of the inference strategies. Since these are implemented as methods, the activation of reasoning mechanisms is made by sending a message to a rule set. This will then activate its corresponding method (forward-reasoning or backward-reasoning) evaluating one of its rule elements.

Summarizing this chapter, we can argue that KFS supports all is needed at the engineering layer. Descriptive aspects are satisfied by the rich declarative part of our frame model, i.e. the structure of schemas, slots, slot values and aspects. Organizational aspects are supported by the abstraction hierarchies and inheritance. Operational requirements are met by allowing the KB-designer to specify his own methods (object orientation), by supporting standard procedures (data-oriented computation) and by offering reasoning mechanisms.

## 3.4 Working-Memory and Context Manager

As already mentioned, the goal of this component of KRISYS is to provide a framework for the exploitation of the application's locality. It is, therefore, desirable to have a mechanism that enables reduction of DBMS calls and the reduction of the path length when accessing the objects of the KB, allowing the application to reference objects almost directly. This is reached by storing needed objects temporarily in a special main memory structure called working-

memory to the stored objects.

The approach used to support locality, is to dynamically extract the context (see 2.3) needed in each phase of the application's problem solving process from the DBS. Therefore, during changes of processing phases the old context would be discarded from the working memory and the new one loaded into it. Since the application's accesses will then be concentrated on the objects of the loaded context, only a few DBS calls will be made during the next processing phase. Due to the set-oriented specification of the required objects, the DBS can use its optimization potential, reducing I/O and transaction overheads. Furthermore, the tiny granules of the application's accesses will not bring any inefficiency, since the path length when accessing the objects is now very short. (Details about the internal structure and implementation aspects of the working-memory manager are provided by [1]).

	duration of consultation (CPU-sec)	ratio to main memory approach
main mamory approach	146	1
direct coupling	28105	~ 170
coupling with application buffer	2946	~ 18
WMCM	643	~ 4

Figure 5: Performance Comparison of Different Coupling Approaches between KFS and DBS

The performance of WMCM has been compared with the performance of other coupling approaches between KFS and DBS (Figure 5). In this case we use a main memory-based form (i.e. when the whole KB is stored in main memory) as basis for this comparison. Coupling KFS and DBS directly seems to be the most inefficient alternative, due to the long path length (from application to DB-buffer) by the application's very tiny access granules. This problem is eliminated by using an application buffer, which uses LRU as a cache allocation/deallocation strategy, to keep the most recently used objects. This approach needs, however, to perform very many DBS calls since individual objects are frequently being extracted from the DBS and stored into it. This is particularly critical after change processing phases, when most of the needed objects are not found in the buffer. We solved this problem by a set-oriented fetch, as described above, which accomplishes an efficiency very close to the main memory approach. Indicative performance figures of a particular example as illustrated in Figure 5 are derived in [10].

## 3.5 DBMS Kernel

Due to space limitations it is not possible to discuss the aspects of the DBMS kernel and the mapping of frames with it in sufficient depth.

The DBMS kernel architecture chosen for KRISYS [14, 15,16] was developed for the support of the so called non-standard applications. For this reason, our kernel, named PRIMA, offers neutral yet powerful mechanisms for managing KB of all types of KS efficiently: storage techniques for a variety of object sizes, flexible representation and access techniques, basic integrity features, locking and recovery mechanisms, etc. [17]. PRIMA is a *PRO*totypic Implementation of the *MAD* model. MAD provides dynamic definition and handling of objects, based on direct and symmetric mana-

### 3.6 Summarizing Example

In this chapter, we illustrate the control flow through the KRISYS system by giving a simplified example of the operations needed to evaluate a KS query at each of the system interfaces (we consider the Working-Memory as been empty at the beginning of an evaluation). Since a complete description is too space-consuming, we restrict ourselves to essential properties expressing the flavor of the control operations.

The structure of Figure 6 indicates the calling hierarchy needed for the evaluation of an ASK-statement (where control operations etc. are dropped for simplicity reasons). Exploitation of information encountered in the Working-Memory would save (expensive) calls to the DBMS kernel interface (with accesses to the DBMS buffer or even to the disks). Perfect locality preservation would avoid any kernel calls for repeated object references.

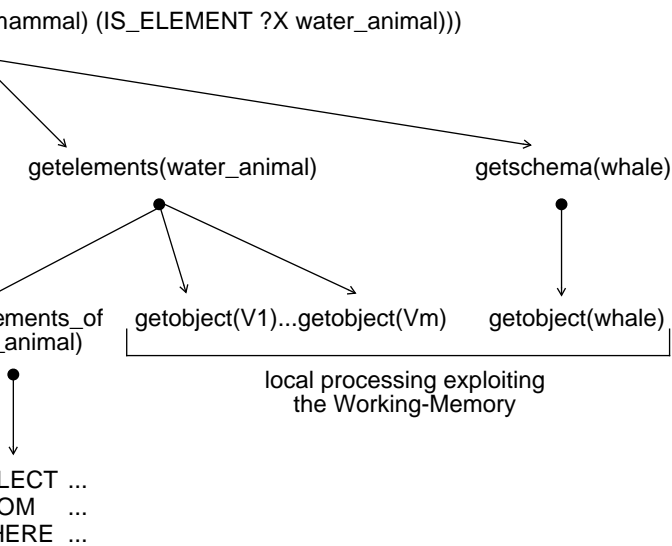


Figure 6: Control flow of an ASK-statement

presented the design of our prototype, named KRISYS. We have advocated the division of KBMS in three different layers, which respectively support the requirements of the application, the needs of the KB-designer and the requirements of an efficient management of the KB. The main philosophy of the system is the idea of abstraction aimed at the independence of knowledge. At the system interface, knowledge is seen not in terms of the flexible representational framework supported by the engineering layer but, functionally, in terms of the application's beliefs about its domain. In order to achieve this functional view of the KB, it is necessary to restrict the communication between application and KBMS. ASK and TELL operations, however, providing:

- multiple ways to select the application's domain knowledge,
- user defined projections,
- flexible forms to cause changes in the KB,
- set-orientation, and
- recursion.

A second very important philosophy is the effective support of the needs of the KB-designer. Clearly, this philosophy has also been the goal of most existing knowledge representation systems which try to support these needs by focusing on the improvement of the expressiveness (i.e. descriptive aspects) of their representational model. However, in a KBMS context expressiveness is just one of the three aspects to be considered at the engineering layer. Hence, organizational and operational aspects are equally as important as descriptive ones. Because of this, an effective support of the needs of the KB-designer is to be achieved by a mixed knowledge representation framework offering:

- object-orientation,
- data-driven computation,
- mechanisms for knowledge organization (i.e. all abstract concepts),
- inheritance and other built-in facilities, and
- reasoning mechanisms,

all of them uniformly integrated with a very rich descriptive power.

A number of concepts used in the KRISYS implementation pay attention to performance requirements. Most important is the framework provided by the implementation layer for the exploitation of the application locality. A single-user version of KRISYS is now complete. KRISYS is to be considered a research vehicle for a variety of KS. It is intended, therefore, to run as a tool for building KS and as a generic system, aimed at the effective and efficient management of large knowledge bases.

### Acknowledgement

I should like to thank T. Härder and B. Mitschang for the helpful comments on the revision of this paper, and H. Neu and I. Litt for preparing the manuscript.

### References

- [1] Mattos, N.M.: KRISYS - A Multi-layered Prototype KBMS Supporting Knowledge Independence, Research Report ZRI No. 2/88, University of Kaiserslautern, 1988.
- [2] Mylopoulos, J.: On Knowledge Base Management Systems, in: [21], pp. 3-8.
- [3] Brachman, R.J., Fikes, R.E., Levesque, H.J.: Krypton: A Functional Approach to Knowledge Representation, *Computer*, Vol. 16, No. 10, 1983, pp. 67-73.
- [4] Fikes, R., Kehler, T.: The Role of Frame-based Representation in Reasoning, in: *Communications of ACM*, Vol. 28, No. 9, September 1985, pp. 904-920.
- [5] Borgida, A.: Survey of Conceptual Modeling of Information Systems, in: [21], pp. 461-470.

Schrefl, M., Tjoa, A.M., Wagner, R.R.: Comparison-Criteria for Semantic Data Models, in: Proc. IEEE 1st Int. Conf. on Data Engineering, Los Angeles, 1984, pp. 120-125.

Brodie, M.L., Mylopoulos, J., Schmidt, J.W. (eds.): On Conceptual Modelling (Perspectives from Artificial Intelligence, Databases, and Programming Languages), Topics in Information Systems, Springer, New York, 1984.

Borgida, A., Mylopoulos, J., Wong, H.K.T.: Generalization/Specialization as a Basis for Software Specification, in: [8], pp. 87-114.

Härder, T., Mattos, N., Puppe, P.: On Coupling Database and Expert Systems (in German), in: State of the Art, Vol. 1, No. 3, 1987, pp. 23-34.

Gallaire, H., Minker, J., Nicolas, J.-M.: Logic and Databases: A Deductive Approach, in: ACM Comp. Surveys, Vol. 16, No. 2, June 1984, pp. 153-185.

Reiter, R.: On Closed World Databases, in: Logic and Databases, (eds. Gallaire, H. and Minker, J.), Plenum, New York, pp. 56-76.

Dayal, U., Smith, J.M.: PROBE: A Knowledge-Oriented Database Management System, in: [21], pp. 227-257.

Härder, T., Reuter, A.: Architecture of Database Systems for Non-Standard Applications (in German), in: Proc. of the GI-Conf. on Database Systems for Office, Engineering and Science Environments, 1985, pp. 253-286.

Dadam, P., et al.: A DBMS Prototype to Support Extended NF2-Relations: An Integrated View on Flat Tables and Hierarchies, in: Proc. ACM SIGMOD Conf., Washington, D.C., 1986, pp. 356-367.

Paul, H.-B., Schek, H.-J., Scholl, M.H., Weikum, G., Depisch, U.: Architecture and Implementation of the Darmstadt Database Kernel System, in: ACM SIGMOD Conf., San Francisco, 1987, pp. 196-207.

Härder, T., Meyer-Wegener, K., Mitschang, B., Sikeler, A.: PRIMA - A DBMS Prototype Supporting Engineering Applications, in: Proc. 13th VLDB Conf., Brighton, 1987, pp. 433-442.

Mitschang, B.: A Molecule-Atom Data Model for Non-Standard Applications - requirements, data-model design, and implementation concepts (in german), PhD Thesis, University of Kaiserslautern, 1988.

Mattos, N.M.: Mapping Frames with the MAD model (in German), Research Report No. 164/86, Univ. Kaiserslautern, 1986.

Härder, T., Mattos, N., Mitschang, B.: Mapping Frames with New Data Models (in German), in: Proc. GWAI'87, Springer, 1987, pp. 396-405.

Brodie, M.L., Mylopoulos, J. (eds.): On Knowledge Base Management Systems, Springer, New York, 1986.



