

A Collaborative Environment for Authoring Large Knowledge Bases

Peter D. Karp, Vinay K. Chaudhri, Suzanne M. Paley
Artificial Intelligence Center
SRI International
333 Ravenswood Avenue, Menlo Park, CA, 94025, USA
Phone: (415)859-6375 Fax: (415)859-3735

April 30, 1997

Abstract

Collaborative knowledge base (KB) authoring environments are critical for the construction of high-performance KBs. Such environments must support rapid construction of KBs by a collaborative effort of teams of knowledge engineers through reuse of existing knowledge and software components. They should support the manipulation of knowledge by diverse problem-solving engines even if that knowledge is encoded in different languages and by different researchers. They should support large KBs and provide a scalable and interoperable development infrastructure. In this paper, we present an environment that satisfies many of these goals.

We present an architecture for scalable frame representation systems (FRSs). Multi-user access is controlled through a collaboration subsystem that uses a novel optimistic concurrency control algorithm. The Generic Frame Protocol (GFP) provides infrastructure for reuse of software components. It is a procedural interface to frame representation systems that provides a common means of accessing and modifying frame KBs. The Generic KB Editor (GKB-Editor) provides graphical KB browsing, editing, and comprehension services for large KBs. All the results have been implemented and tested in the development of several real KBs.

Keywords: Knowledge bases, knowledge base management systems, knowledge representation, storage management, concurrency control

1 Introduction

Collaborative knowledge base (KB) authoring environments allow multiple, geographically distributed users to collaborate in the development of large KBs. The first generation of FRSs (Karp 1992) provided only single-user KB authoring environments whose engineering limitations constrained the size of the resulting KBs, and did not permit distributed KB access. This report describes a next-generation, reusable environment for collaborative KB authoring that consists of the following components:

- The Generic Frame Protocol, which provides infrastructure for software and knowledge reuse. It is a procedural interface to FRSs that provides a common means of accessing and modifying frame KBs.
- The Storage Subsystem, which provides scalable storage of frame KBs in a commercial DBMS.
- The Collaboration Subsystem, which provides optimistic concurrency control over the KB updates made by multiple distributed KB authors.
- The GKB-EDITOR, which provides KB browsing and editing services for large KBs.

This work makes a number of contributions to the field of knowledge base management systems. We identify design requirements for collaborative KB authoring environments, and present an architecture that satisfies those requirements and an implementation of that architecture. The architecture includes both the Storage Subsystem, for expanding the storage capabilities of FRSs, and novel “optimistic” concurrency control techniques for coordinating KB updates made by multiple simultaneous users. The GKB-EDITOR completes this authoring environment by providing four different viewer and editor utilities for browsing and modifying large KBs. All of these tools have been used in the development of several real-world KBs, including a planning ontology, and the EcoCyc biochemical-pathway KB, which contains 12,000 frames.

These results are made more significant because the implementations have been reused across several FRSs. The Storage Subsystem has been used with LOOM (MacGregor 1991), THEO (Mitchell et al. 1989), and OCELOT; the GKB-EDITOR has been used with LOOM, THEO, SIPE-2 (Wilkins 1990), OCELOT, and partially with ONTOLINGUA (Gruber 1993) and CLASSIC (Borgida et al. 1989). This reusability supports the generality of our approach, and provides the most

substantial example to date of the type of software reuse envisioned by the Knowledge Sharing Initiative (Patil et al. 1992; Neches et al. 1991). We discuss various obstacles to software reuse that we encountered, and the solutions that we devised.

2 Design Requirements

We derived the design requirements for KB authoring environments from our experiences with KBs in several domains. We present two scenarios of the use of these environments to illustrate the motivations behind these requirements.

Scenario 1 involves the distributed development of a planning ontology by multiple DARPA contractors who are distributed throughout the United States, as part of the DARPA /Rome Laboratory Planning Initiative. Contractors at roughly a dozen sites might want to access the LOOM KB that implements a planning ontology, to browse the current state of the ontology, to add new class definitions interactively, and to edit existing definitions. The LOOM classifier runs during the editing process to infer relationships among new and existing classes. Users might also execute ontology translators to convert the planning ontology to other forms, such as a relational database.

Scenario 2 involves a biological KB that describes the biochemical reaction network within the *E. coli* cell. The KB models biological objects such as enzymes and biochemical pathways. Expert biologists from around the world interact with the KB in different ways. Some biologists want to update the KB by editing that region of the KB that falls within their expertise. Other biologists require only read-only access to the KB. They will browse and query the KB using a graphical interface. Others will execute qualitative and quantitative simulation programs that model the metabolism of the cell. Other scientists will execute programs that redesign the biochemical network of *E. coli* for commercial purposes in bio-technology. Still other users will evaluate complex queries over the KB to answer scientific questions.

We have extracted the following design requirements from these scenarios.

- User sessions will involve several distinct patterns of KB operations:
 - Interactive browsing and editing sessions that access a small portion of a KB. Editing sessions may last hours or even days, and include changes to many different frames. The patterns of most KB modifications differ from database modifications that update a small number of isolated values.

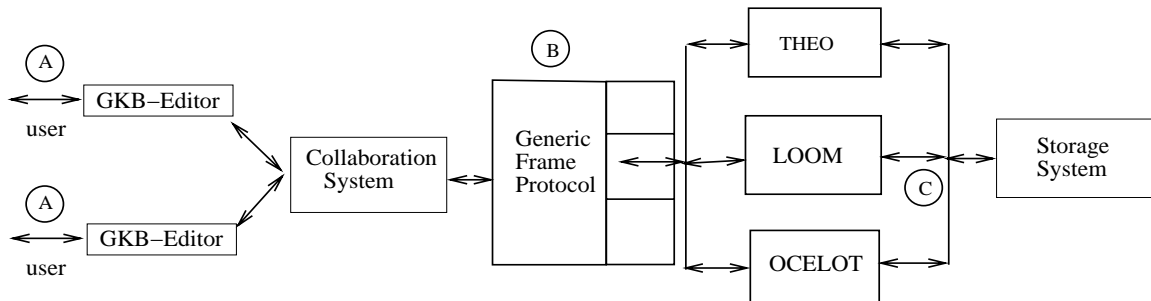


Figure 1: System architecture

- Complex computations such as simulation, design, and expert systems that repeatedly access significant subsets of the KB.
 - Traditional database-like complex queries that return small amounts of data to the application.
- The next generation of KB sizes will range from 10^4 to 10^7 frames. Users may be geographically distributed, and may access the KB by using long-distance Internet connections.
 - The schema-evolution capabilities of most FRSs should be retained. The KB schemas are usually larger, more complex and more dynamic than the one for databases. As KBs grow, we expect their schemas to grow also, as well as the need to modify class and relation definitions dynamically.

3 Architecture

Our system architecture is shown in Figure 1. The GKB-EDITOR is a graphical tool for interactive KB editing and browsing. The GKB-EDITOR is reusable across multiple FRSs because all of its KB-access operations are implemented using the Generic Frame Protocol (GFP). GFP provides a uniform procedural interface to FRS applications. Our architecture includes a storage subsystem that allows KBs to be stored within an ORACLE DBMS. Frames are retrieved on demand from the DBMS into an FRS that acts as a client. The storage subsystem records what frames have been modified, and can incrementally save modified frames to the DBMS. Updates from multiple users are synchronized using a concurrency control method that detects, at commit time, the conflicts between the updates made by the committing user and other recent updates performed by other

concurrent users. Our technique not only detects conflicts but also provides user assistance in resolving those conflicts. Unlike traditional database techniques that keep an item locked for the whole duration of a transaction, we use locks only while the conflict-free updates are being deposited in the database.

The architecture of Figure 1 allows distributed operation because network links can be inserted between components at several places. We can insert a network link at (C) by transmitting Structured Query Language (SQL) calls to the DBMS server over a network. We can insert a network link at (B) by creating a remote-procedure call version of GFP, in which every GFP call is sent over a network. We can insert a network link at (A) by allowing the X-window graphics of the GKB-EDITOR to flow between an X client and X server.

In most functional implementations of our system, the network link has been inserted at C. The GKB-EDITOR, collaboration system, GFP, and the FRS act as a client program and the DBMS acts as a server. This approach allows computationally-intensive applications such as planners, design programs, and simulators to operate in parallel on multiple client machines, on regions of the KB that are cached in the local memory of those clients. AI applications tend to make a very large number of KB accesses, sometimes accessing the same frames repeatedly.

4 Generic Frame Protocol

The goal of GFP is to allow reuse of knowledge and software by mixing and matching knowledge and software components in a large AI system. By reuse of knowledge, we mean transforming a knowledge base or ontology developed in one FRS in another FRS, for example importing an ontology developed in LOOM into ONTOLINGUA. By reuse of software, we mean an easy porting of the software across FRSs, for example, reusing a system such as the GKB-EDITOR across multiple FRSs. GFP accomplishes its goals by defining a set of Common LISP generic functions that constitute a common application-program interface (API) to FRSs. These functions constitute a wrapping layer for FRSs; that wrapping layer is implemented by creating FRS-specific methods that implement GFP operations. The rest of the section describes GFP in more detail.

4.1 Design Principles of GFP

- **Simplicity:** The protocol should be simple and reasonably quick to implement for a particular FRS, even if this means sacrificing theoretical considerations or support for idiosyncrasies

of that FRS.

- **Generality:** The protocol should apply to many FRSs, and support the most common FRS features.
- **No legislation:** The protocol should not require substantial changes to an FRS for which the protocol is implemented. That is, the protocol should not mandate the method of operation of an underlying FRS.
- **Performance:** Inserting the protocol between an application and an FRS should not introduce a significant performance cost.
- **Consistency:** The protocol should exhibit consistent behavior across implementations for different FRSs, that is, a given sequence of operations within the protocol should yield the same result over a range of FRSs.
- **Precision:** The specification of the protocol should be as precise and unambiguous as possible.
- **Language independence:** Ideally, GFP should be independent of programming language.

Satisfying these objectives simultaneously is impossible because many of them conflict. Different FRSs behave differently, and unless we mandate a minutely detailed behavioral model for FRSs (which no developers will subscribe to anyway), we cannot force these systems to exhibit uniform behavior. The GFP uses a knowledge model that encompasses many FRSs and is detailed enough to be useful in practice, but is not so detailed as to exclude every FRS from its model. Another example of conflicts among our objectives is that to precisely specify the semantics of the GFP function that retrieves the values for a frame slot, we must specify the inheritance semantics to be used. However, different FRSs use different inheritance mechanisms (Karp 1992). Conformance with a specific semantics for inheritance would require either altering the inheritance mechanism of a given FRS (violating the no-legislation goal), or emulating the desired inheritance mechanism within the implementation of the protocol (violating performance and generality, since the inheritance method used by that FRS is inaccessible through the protocol).

Currently, GFP has many LISP dependencies. Connecting GFP to an FRS implemented in some other language should be straightforward using LISP foreign-function calls. Also, precision and generality are conflicting goals. If we precisely model various characteristics of the FRSs in

the GFP, the protocol will not be general, because the detailed features of one FRS may not be shared by another.

4.2 Proposed Design of the GFP

The most important decision in the design of the GFP was whether the GFP should be the “least-common-denominator” or a “superset” of a class of FRSs. In the least-common-denominator approach, the GFP would contain only those features that are supported by most of the systems. In the superset approach, it would aim to support the union of the features of all FRSs. Our design is a hybrid of the two approaches described. The core design of the GFP is based on the least-common-denominator approach. It can be parameterized in various ways to capture the peculiarities of an FRS that are not covered in the core model. We, however, did not design the GFP to be a superset of all FRSs, because that would make the protocol cumbersome, complex, and difficult to use.

The design of the GFP is based on a comprehensive survey of FRSs. The survey revealed a large variety of system designs (Karp 1992). Some of the differences among these systems were significant, whereas others were superficial. We identified those commonalities that are useful for a broad range of applications. The knowledge model of GFP, as presented here, is based on an axiomatic formalization of classes, relations, and functions for the frame ontology in ONTOLINGUA (Gruber 1993) but extends the ontology to include aspects relevant to operational applications (e.g., kinds of inheritance, facets). Let us consider the design of GFP in more detail.

4.2.1 Representational Primitives

A *frame* is an object with which facts are associated. Each frame has a unique name. Frames are of two kinds: classes and instances. A *class* frame represents a semantically related collection of entities in the world. Each individual entity is represented by an *instance* frame. A frame can be an instance of many classes, which are called its *types*. A class can also be an instance, that is, an instance of a class of classes (a *meta-class*).

Information is associated with a frame via *slots*. A slot is a mapping from a frame and a slot name to a set of values. A slot value can be any LISP object (e.g., symbol, list, number, string). Slots can be viewed as binary relations; GFP does not support the explicit representation of relations of higher arity.

Facets encode information about slots. Some facets pertain to the values of a slot; for example,

a facet can be used to specify a constraint on slot values or a method for computing the value of a slot. Other facets may describe properties of the slot, such as documentation. Facets are identified by a facet name, a slot name, and a frame.

A *knowledge base*, or KB, is a collection of frames and their associated slots and values. Multiple KBs may be in use simultaneously within an application, possibly serviced by different FRSs. Frames in a given KB can reference frames in another KB, provided that both are serviced by the same FRS.

4.2.2 Inference Mechanisms

Our survey of FRSs revealed that three forms of inferences are most prevalent: *subsumption reasoning*, limited forms of *constraint checking*, and *slot value inheritance*. Consequently, GFP supports only these three types of inference.

In GFP, it is possible to specify and query subsumption (or class-subclass) relationships. Some FRSs require subsumption relationships to be specified when frames are created. In contrast, FRSs that perform automatic classification infer the subsumption relationships by comparing class definitions. The GFP operations allow the user to interrogate any class-subclass and class-instance relationships, no matter how the relationships were derived.

GFP recognizes type and number restriction constraints on slot values that are specified as facets. The GFP relies on the underlying FRS to check those constraints.

Inheritance in GFP is based on the use of *template* and *own* slots. A template slot is associated with a class frame and may be inherited by all the instances of that class. An own slot can be associated with a class or instance frame and cannot be inherited. Inheritance in GFP can be characterized as follows. The value of a slot in a frame F is computed by combining the values of own slots and template values for that slot in F and the superclasses of F , provided those values do not conflict. GFP allows the use of different semantics for “combining” and “conflict,” to support a range of inheritance methods.

4.2.3 Parameterization Using Behaviors

GFP supports extensions to the core knowledge model to capture the features of FRSs that are not covered by the core model. This diversity is supported through behaviors, which provide explicit models of the FRS properties that may vary. An application program can query the value of behaviors and use the result in executing the code specific to the value of that behavior.

For example, some FRSs may allow a user to define new facets. In such cases, in addition to the standard GFP facets, an application program must take into account the user-defined facets. When an FRS allows user-defined facets, we can set the `:user-defined-facets` to true. An application program can query the value of the behavior `:user-defined-facets`, and if returned true, can execute the code to handle user-defined facets.

4.2.4 Set of Operations

The GFP defines a programmatic interface of common operations that span the different object types in the knowledge model, namely, *knowledge bases*, *frames*, *classes*, *instances*, *slots*, and *facets*. Three categories of operation are supported for each object type: *retrieval operations*, *manipulator operations*, and *iterators*. Retrieval operations extract information about objects and slot values. Manipulator operations create, destroy, and modify objects. Iterator operations loop over a set of objects.

As an example, consider operations on slots. The most commonly used retrieval operation on slots is `get-slot-values`, which retrieves the values of a slot given a frame name and a slot name. The operation `member-slot-value-p` tests whether a given value is one of the values of the slot of a given frame. An example of manipulator operation on slots is `create-slot`, which is used to define a new slot of a class. To operate on all the values of a slot, one may use the iterator function `do-slot-values`. It can have a body consisting of a valid LISP code that is evaluated by iterating over the slot values.

In addition, GFP supports operations on behaviors. Retrieval operations obtain information about the behaviors supported by GFP in general, the behaviors that a given FRS supports, and the behaviors that are enabled for a particular KB.

4.3 Implementation

Each GFP operation is implemented as a Common LISP Object System (CLOS) method. We identified a subset of GFP operations, called the *kernel* operations, which can be used to implement every other operation not in the kernel. We have created a default method for all the nonkernel operations that defines them in terms of kernel operations. For example, the default method for `member-slot-value-p` calls the kernel operation `get-slot-values`. The kernel consists of roughly 30 operations, whereas the total GFP operations are over 200. The default methods can be overridden to improve efficiency or for better integration with development environments. We

can create GFP implementations for new FRSs quickly because only the kernel methods need to be implemented.

GFP back ends exist for LOOM (MacGregor 1991), OCELOT, SIPE-2 (Wilkins 1990), THEO (Mitchell et al. 1989), and ONTOLINGUA (Gruber 1993). The LOOM and ONTOLINGUA back ends have been used in conjunction with the GKB-EDITOR. A read-only back end exists for CLASSIC (Borgida et al. 1989). Additional back ends are planned for Algernon (Crawford 1990) and object-oriented databases. GFP is being used in DARPA's high-performance KB initiative (see <http://www.teknowledge.com:80/HPKB/>). It is now sufficiently mature that the implementors of FRSs can consider it as a serious candidate for an API for their system.

A network version of the GFP (NGFP) is also available. The NGFP allows us to invoke the GFP operations on KBs that are not necessarily resident on a local machine. In Figure 1, the NGFP supports the network link at B. The NGFP implementation consists of a client and server that may be running on different machines on the internet and may be written using different programming languages. The clients open a connection to the server by specifying its address, port number, user name, and password. Once a connection has been created, the user can create a KB on the remote server. All subsequent operations that use the remote KB as the KB argument are transparently executed on the remote server. The NGFP clients are available for C, LISP, and JAVA. The NGFP implementation supports two performance improvement techniques — client side cacheing and remote procedures. Client cacheing facility records all the remote GFP query operations and their results in a table. As a result, if a query is repeated, it can be answered locally instead of executing it on the remote server. Whenever there is an update, the cache is flushed. Since updates are infrequent, such a simple scheme for client consistency has proven to be adequate. The NGFP allows a user to define procedures that consist of several GFP operations. The procedures are executed on the server in one network call instead of one network call for each GFP operation thus improving the performance.

4.4 Logging Facilities

The GFP implementation provides a facility to capture, in the form of a log, all KB update operations executed in a user session. The log can be used in a variety of ways. For example, in Section 7, we describe how we use the logging capabilities of GFP to support multiuser access to knowledge bases. The log can be used for propagating updates between replicated copies of a knowledge base at multiple sites: the updates are transmitted as a log and applied to the

local copies of the KB. The log can be used by the GKB-EDITOR to support “undo” for user operations: undoing an operation involves applying, to the current state of the KB, the inverse of every operation in the log.

4.5 Evaluation of the Protocol

In this section, we discuss our experiences in using the GFP with LOOM and CLASSIC. The knowledge model of LOOM fits the GFP model fairly closely. For most GFP operations, we were able to identify equivalent LOOM functions. The key difficulties we faced were in capturing the concept-definition language of LOOM and representing context information. Attributes of LOOM classes are specified through complex definition expressions. As GFP does not support concept definition languages per se, we developed a mapping between the LOOM language and GFP facets. For example, the `:at-most` concept construct in the LOOM concept definition language maps to the `:numeric-maximum` facet in GFP. But there is no way to represent the `:not-filled-by` construct of LOOM in GFP (`:not-filled-by` construct specifies the values that are not valid for a slot). In FRSS such as LOOM, CYC, and ONTOLINGUA, one can associate a context with each assertion in a knowledge base. For example, one may state that in the context of theoretical computer science, “Karp” is a professor at UC Berkeley, but in the context of knowledge based systems, “Karp” is a scientist at SRI. Contexts have a variety of uses such as for storing the alternative world states that are generated by a planner. At present, there is no notion of contexts in GFP, and each GFP KB maps to a single LOOM context. This representation does not capture the feature that contexts in LOOM can be organized into a hierarchy, and can inherit assertions from the parent contexts.

CLASSIC differs from GFP to a greater extent than LOOM does. Here, we consider only three examples to illustrate the differences. First, in CLASSIC, a slot must be defined before the concept that utilizes it, whereas the GFP assumes the opposite. This makes it difficult to implement the GFP operation `create-slot`, which has two required arguments: slot name and the classes it is attached to. We faced a similar problem while working with LOOM, but in LOOM, either slots can be created before classes or *vice versa*. Therefore, our LOOM/GFP back end assumed that a class has been defined before its slots, and thus provides a subset of the functionality supported by LOOM. A possible solution to solve the problem is to revise the `create-slot` operation to accept only the slot name as the required argument. Second, CLASSIC does not have strong typing, that is, given a slot, we can determine if it is valid in a KB, but we cannot determine its domain. To

capture this in the CLASSIC/GFP back end, we indicate the applicability of a slot to a class by giving a restriction on the slot, specified by means of a facet, that specifies the classes to which it is applicable. Finally, CLASSIC supports open world reasoning by offering an operation for “closing” slots on an instance to indicate that there can be no more values of this slot. Such a behavior can be represented in the GFP by introducing a new facet called `:closed` that can have values `T` or `nil`.

In summary, GFP was able to cover a substantial part of the FRS features that we considered. It was not complete in many cases, but we were able to deal with them by minor extensions to our initial design. Therefore, we believe that GFP effectively meets the goal of a generic API for knowledge bases.

Our experimental evaluation of GFP back end for LOOM indicates that the performance penalty for using GFP is acceptable. The overhead incurred by the use of the GFP is largest for fast operations and smallest for slow operations. Using a LOOM implementation of GFP, we compared the running times of key GFP kernel operations with the corresponding LOOM operations. The results showed the GFP operations to be 1% to 50% slower (depending on the operation). The high overhead costs resulted for operations without direct counterparts in LOOM. For example, GFP provides an operation for retrieving a frame when given a frame name, but LOOM has no such operation. Instead, it provides separate operations for instances and classes. The GFP operation must consider whether the name corresponds to a class or an instance in order to invoke the appropriate underlying LOOM operation. We note that on an absolute scale, the overhead is very small in this case (approximately 0.02 millisecond).

For directly comparable operations, the upperbound on overhead was 35%. Much of the increased execution time results from activities common to all GFP operations. Thus, the overhead is high on a percentage basis for fast operations such as slot value retrievals (35% for a 0.3-millisecond operation), but low for more expensive operations such as retrieving all instances of a class (1% for a 16-millisecond operation).

4.6 Related Work

Both GFP and Knowledge Interchange Format (KIF) (Genesereth and Fikes 1992) seek to provide a domain-independent medium that supports the portability of knowledge across applications. KIF is more expressive than the GFP, as KIF is a comprehensive first-order representation formalism whereas the GFP captures a subset of first-order logic that represents class hierarchies.

ONTOLINGUA (Gruber 1993) is a set of tools for writing and analyzing KIF knowledge bases along with translators for mapping KIF KBs to specific FRSs. KIF and ONTOLINGUA are declarative representation languages; GFP is a procedural interface for accessing representation structures. KIF and ONTOLINGUA are designed for use in sharing a large corpus of knowledge at specification time, through the use of translators. GFP is designed for runtime access and modification of existing KBs. GFP is similar to Knowledge and Query Manipulation Language (KQML) (Patil et al. 1992) in that it provides a set of operations defining a functional interface for use by application programs. For example, an agent may query an FRS using a KQML “performative.” The KQML allows an agent to express the action of querying, but provides no language to express the query itself. (The query could be expressed using GFP.) Thus, GFP is complementary to KQML. GFP plays the same role for FRSs as ODBC (Open Database Connectivity) does for relational DBMSs (Geiger 1995).

5 Generic Knowledge Base Editor

The knowledge representation community has long recognized the need for graphical knowledge-base browsing and editing tools to facilitate the development of complex knowledge bases. However, the past approach of developing KB editors that were tightly wedded to a single FRS is impractical (Kehler and Clemenson 1984; Lenat and Guha 1989). The substantial efforts required to create such tools are lost if the associated FRS falls into disuse. Since most FRSs share a common core functionality, a more cost-effective approach is to amortize the cost of developing a single FRS interface tool across a number of FRSs. Another benefit of this approach is that it allows a user to access KBs created using a variety of FRSs through a single graphical user interface (GUI), thus simplifying the task of interacting with a new FRS.

The GKB-EDITOR is a generic editor and browser of KBs and ontologies — generic in the sense that it is easily portable across several FRSs. This generality is possible because the GKB-EDITOR performs all KB access operations through GFP. To adapt the GKB-EDITOR to a new FRS, we need only create a GFP implementation for that FRS — a task that is considerably simpler than implementing a complete KB editor. The GKB-EDITOR and the Stanford Ontology Editor (Farquhar et al. 1995) have been the most significant applications driving the development of the GFP. They have driven the addition of new operations to GFP, and they have challenged the portability of GFP because the GKB-EDITOR has been used in conjunction with several FRSs.

The GKB-EDITOR contains a number of relatively advanced features, such as incremental browsing of large graphs, KB analysis tools, operation over multiple selections, cut-and-paste operations, and both user and KB-specific profiles. The GKB-EDITOR is in active use in the development of military-application planning KBs and ontologies for LOOM (MacGregor 1991) at several sites, for SIPE-2 KBs, and for THEO (Mitchell et al. 1989) KBs. It is used daily in the development of EcoCyc, a biological KB, based on the OCELOT FRS that contains more than 12,000 frames and is accessed daily via the World Wide Web (WWW) by scientists from around the world (Karp et al. 1996).

5.1 Viewing and Editing Knowledge Bases

The GKB-EDITOR offers four different ways to view and edit parts of a KB. The user can view the KB as a class-instance hierarchy graph, as a set of inter-frame relationships (this view is roughly analogous to a conceptual graph representation, a semantic network, or an entity-relationship diagram), by examining the slot values and facets of an individual frame, or by viewing a collection of slot values within a spreadsheet. Each viewer is optimal for different viewing and editing tasks.

5.1.1 Class-Instance Hierarchy Viewer

In the class-instance hierarchy viewer, each node in the graph represents a single class or instance frame, and directed edges are drawn from a class to its subclasses and from a class to its instances. Multiple parentage is handled properly.

The hierarchy is normally browsed incrementally. The roots of the hierarchy graph are either computed or specified by the user, and the graph is expanded to a specified depth. If a particular node has more than a designated number of children, the remaining children are condensed and represented by a single node. Unexpanded nodes are visually distinguished from expanded nodes. The user can browse the hierarchy by clicking on nodes that are to be expanded or compacted.

The hierarchy viewer can also be used to modify the class-instance hierarchy. Operations such as creating, deleting, and renaming frames, and altering superclass-to-subclass links and class-instance links can all be accomplished with a few mouse clicks.

5.1.2 Relationships Viewer

It is often useful to visualize other relationships in a KB than the parent-child relationship, such as the part-of relationship. Such relationships are encoded in the KB as slot values. For example,

if frame B is a value of slot X in frame A , then an edge can be drawn from node A to node B , labeled X . If we recognize that slot X represents a relationship between frames A and B , then this view is analogous to the view of a KB as a conceptual graph (although our displays do not use all the visual conventions of the conceptual graph community) or to a semantic network. Like the hierarchy view, a relationships view is browsed incrementally. This viewer supports a variety of relationship editing operations.

5.1.3 Frame Editing Viewer

The frame-editing viewer allows the user to view and edit the contents of an individual frame. The user may select a frame from the hierarchy viewer or the relationships viewer and display it in a frame-editing viewer, which presents the contents of a frame as a graph. Each slot name forms the root of a small tree; its children are individual slot values and facet values. Inherited items are distinguished visually from local items, and cannot be edited (although they can be overridden where appropriate).

In addition to duplicating, renaming, or deleting the viewed frame, the editing operations available in this viewer permit adding, deleting, replacing, and editing of slot and facet values. Slots themselves may be added, removed, or renamed, when classes are edited.

5.2 Related Work

Our work builds on ideas from previous graphical browsers and editors that were built for individual FRSs, including KnEd (Eilerts 1994), CODE4 (Skuce and Lethbridge 1995), HITS (Lenat and Guha 1989), and Protege-II (Eriksson et al. 1994).

All of the tools are restricted to use with a single FRS. Stanford's Ontology Editor (Farquhar et al. 1995) is a browser and editor for ontologies (Gruber 1993). Currently, the Ontology Editor operates only on ONTOLINGUA ontologies, but because it is implemented using GFP, it could be used to browse and edit KBs for a variety of KR systems. The WWW implementation of the Ontology Editor is both its biggest advantage and its biggest drawback. The advantage is the easy accessibility of the server; its drawbacks result from the many limitations of the HTTP protocol: most information is presented in textual form, rather than graphically; displays cannot be updated incrementally, as they can in the GKB-EDITOR.

6 Storage System

Most existing FRSs such as LOOM and ONTOLINGUA fully load a KB into memory before accessing any part of it. To provide persistence, KBs are written in their entirety to flat files on secondary storage. This approach is not scalable because loading and saving time are proportional to the size of the KB rather than to the volume of information accessed in a given session, or to the number of frames modified in a given session.

A simple way to achieve our objective is to use an indexed file that lets us selectively retrieve frames by some key (for example, frame name). An indexed file, however, does not provide fault tolerance: a long update operation, if terminated abnormally, may corrupt the file. An indexed file also does not control the access by multiple clients. Therefore, our storage system design submerges a commercial DBMS within an FRS. The FRS and associated application code act as a client that accesses a DBMS server using a DBMS application-programming interface. Our storage system, PERK (for **PER**sistent **K**nowledge), retrieves frames incrementally, on demand, from the DBMS. Because fetching of frames on demand from the DBMS, or *demand faulting*, is analogous to page faulting in operating systems, we call this process *frame faulting*. PERK tracks which frames have been modified and transmits those frames back to the DBMS when the KB is saved.

Given the basic architecture, other choices must be made: How should the FRS information be organized in the DBMS? In what order should the frames be loaded into the memory of an FRS client? How can we utilize idle periods to speed up the loading time?

6.1 Design of a DBMS Schema

To store a KB in a DBMS, we must define a mapping from the knowledge model of an FRS to the data model of a DBMS. The mapping can be either KB-specific or generic. In a KB-specific mapping, the database schema is different for different KBs (Chen and Markowitz 1995; Copeland and Khoshafian 1985; Valduriez, Khoshafian and Copeland 1986; McKay, Finin and O'Hare 1990). For example, when mapping from a KB to an object-oriented DBMS, the classes may correspond to classes in the KB. A similar mapping can be designed to tables in a relational DBMS. In a generic mapping, we define one schema that can hold any frame-based KB. One way to construct such a schema is to store an entire frame as a single uninterpreted string in the database.

The main disadvantage of the KB-specific mapping is that neither the relational nor the object-

Frames			
KB ID	Frame Name	Frame Bodies	Type
1	Physical Entity	...	class
1	Nation	...	class
1	Diesel	...	class
1	Norway	...	instance

Relations		
KB ID	Relation Name	Body
1	Maker	...
1	Propulsion	...

KB Mapping	
KB Name	KB ID
Naval Theory	1
Aircraft	2

Supers		
KB ID	Class Name	Super
1	Ship	Vehicle
1	Vehicle	Physical-Entity

Instance Classes		
KB ID	Instance Name	Class Name
1	Norway	Nation
1	Canada	Nation

Figure 2: Relational schema used to store LOOM KBs in an DBMS, with sample data

oriented (OO) data models can capture the rich semantics of the knowledge model used by FRSs (Borgida and Brachman 1993). For example, the relational model does not allow inheritance of properties from one class to another. The OO data model solves the problem of inheritance, but it does not support the full power of FRSs such as runtime inheritance of default values or production rules attached to slot values or (in many cases) dynamic changes to the schema. In the generic-mapping approach, we perform all the processing of frames within the FRS so that no frame semantics are lost.

The advantage of a KB-specific mapping over generic mapping is that the DBMS query-processing engine can operate on KB data on the server side. The server can achieve high efficiency because it can make use of indices, and because it can process the query locally, without transmitting data across a network. Since the DBMS cannot represent the full semantics of data, it will not be able to process the full set of queries.

Indexed Slots	
KB ID	Slot Name
1	Maker
1	Propels
1	...
1	...

Character Slot Values			
KB ID	Instance Name	Slot Name	Slot Value
1	Akula-Class	Maker	Russia
1	Akula-Class	Propulsion	Nuclear

Figure 3: Tables to support indexed queries on the DBMS server

We chose to design a generic relational schema that can be extended in limited ways to different KBs. That schema consists of five core tables and four tables to support extensibility. An example of the core tables is shown in Figure 2. The *Frames* table contains frame bodies. A *body* of a frame is a string that provides an FRS with all the information necessary to create the frame. There are occasions when a frame is referenced without its type being known. Therefore, we defined a type field in the Frames table that identifies the frame as a concept or instance. A KB identifier is included in each table, to enable multiple KBs to be stored in the same Frames table. The *KB Mapping* table associates a KB name with a unique identifier.

One way to store the bodies in the DBMS is to use the LISP printed representation of the s-expressions corresponding to a frame body. In practice, we found that the standard LISP mechanism for reading and writing an s-expression is extremely slow. Therefore, we developed an encoding of s-expressions into strings that is not human-readable but is more efficient to read and write.

The tables *Supers* and *Instances* enable reconstruction of the concept and instance hierarchy. The former lists the superclass-subclass relationships between concepts; the latter stores the relationship between instances and their parent concepts. The tables are indexed on Supers, Classes, and Instances. This information is necessary because for a concept or instance to be defined in an FRS (for example LOOM), all parent concepts must already be loaded, or LOOM will not be able to classify the new frame. Thus, we must be able to quickly determine the parents of a given frame so that those parents can be retrieved if necessary. These tables can also be used to answer the transitive closure queries on class-subclass relationships without actually faulting in all the frames in the FRS client.

The core tables are adequate if all the queries are to be processed on the client side. In some cases, especially when the answer of a query is a small subset of all possible answers, it may be desirable to process some queries on the server. To do that, we restrict the DBMS query processing

to only those frame slots whose semantics can be captured by the DBMS schema — such as slots whose values are not inferred using defaults or production rules. We extended the generic schema in Figure 2 in a manner that indexes specific KB slots within the DBMS. Thus, the slots to be indexed can be KB-specific. We defined three additional tables that store slot values, and one table that stores the slot names that should be indexed for a given KB. In Figure 3, we show representative samples for two of the four tables. Each of the three slot value tables contains two columns: one stores a slot name, the second stores a slot value (the three tables are for slot values that are strings, numbers, and long strings, respectively). We separate the types in different tables because retrieval from an index on numbers is much faster than from the one on strings. The long values cannot be indexed, but they should be available so that we can perform a sequential search on them whenever there is a query on a slot with character values.

This approach increases the storage space requirements because slot values are stored as part of the body as well as in the slot value tables. Our scheme, therefore, trades the generality and storage space for the speed of retrieval. PERK uses the frames tables during demand faulting and uses the slot tables for processing queries on the slot values.

6.2 Frame Faulting

When a user begins a session, she first accesses a KB by opening it. During the process of opening a KB, PERK establishes a connection with the DBMS server, and all the root classes and slot definitions (relations) in the KB are faulted into the FRS client. We fault in root classes, because many of them will be faulted in anyway when a frame is requested by an application. We fault in all the relations, because usually a KB contains a small number of them, and faulting them for once simplifies the faulting of class and instance frames.

A frame fault occurs when an application (or the FRS itself) references a frame F that has not been fetched into the FRS client from the DBMS. When faulting a frame into an FRS client, PERK retrieves its body from the DBMS by issuing one or more SQL queries. PERK then calls standard FRS functions to add the frame to the KB.

Among the difficulties we faced in implementing demand faulting for LOOM was that if a frame F refers to some other frame G , LOOM requires that G must exist in the KB. F might refer to G because G is a parent or a child of F , or because G is referenced in a slot of F . If G is a parent of F , and G is not currently in the memory of the FRS client, PERK generates a frame fault for G to allow proper operation of LOOM. But if G is an instance of F , or if G is referenced in a slot

of F , we create a place holder (stub frame) for G . G itself will be faulted at a later time if there is some reference to it by the application.

The classifier code in LOOM is not re-entrant, that is, while we are classifying a concept A into the subsumption hierarchy, we cannot always start classifying another concept B until classification of A is complete. As a result, frame faults generated while classification is in progress cannot be serviced, because if we try to do so, the classifier may find some of the internal data structures in an inconsistent state. Therefore, before we start classifying a concept, we must make sure that all the concepts that will be referenced in the process are already faulted in so that no frame faults are triggered while the classification is in progress.

We found it difficult to predetermine all the frames that will be used while defining a given concept because there is no document describing the classification algorithm of LOOM. We addressed the problem by testing our system with several real KBs and identifying common problem cases.

For example, when a concept A is referenced in the definition of concept B , we can notice it while parsing the definition of A , and demand fault A before processing B . Another simple case arises while dealing with the super concepts: before a concept is loaded, we make sure that its direct supers have been loaded. It is easy to determine supers because the supers of a concept can be determined using the Supers table.

A somewhat more involved case arises while classification is in progress. Suppose concept A has two subconcepts B and C . When we load A , we will create stubs for B and C . When there is a demand fault for B , we fetch it from the DBMS. While B is being classified, the classifier considers classifying B under C , causing a demand fault for C . No matter whether we load B or C first, while classifying one of them, the other one will have to be faulted in. In this particular instance, we know that B cannot be classified under C , because otherwise this fact would have been stored in the Supers table in the database, and we would have faulted C before faulting B . Therefore, we modified the classifier to not consider C during the classification process.

Inverse relationships need special treatment because of the circular dependencies that they may generate. For example, consider a slot A that is the inverse of another slot B . Suppose we retrieve A when B is not loaded into LOOM. If, while loading A , there is a frame fault for a concept C (perhaps C appears in the definition of A) that has B as one of its slots, we cannot define C because B is not yet defined. We could not go ahead and load B , because B is defined in terms of A , and any attempt to load B will trigger a cyclical frame fault for A . We addressed the problem

by delaying the assertion of the slot values for C until both A and B have been faulted in.

Even though we have empirically tested frame faulting for several KBs, we have no way to guarantee that incremental loading has not changed the behavior of classification or that we have taken into account all possible situations under which a frame fault may be generated during classification. A principled study of classification in conjunction with demand faulting remains a problem open for future research.

6.3 Prefetching

We can decrease the overall latency of PERK by decreasing the number of demand faulting operations it performs. The number of demand faults will decrease if frames that would have been demand faulted are already in memory at the time the demand fault would have occurred. We can achieve that state of affairs by *prefetching* frames from the DBMS before they are demanded by the application, assuming that the cost of prefetching is less than the cost of demand fetching or that the prefetching can be performed during user idle time.

In our current system, once a frame has been brought into the client memory, it is never discarded. Under this assumption, if prefetching occurs only during client idle time, it is likely to improve performance since it will eliminate some future demand-fetch operations. Prefetching, however, does place an additional load on the DBMS server; a large number of clients prefetching from the same server will ultimately decrease overall system performance. Prefetching can also lead to a greater number of page faults on the client.

A prefetching system must decide which frames should be prefetched. It is better to first prefetch those frames that are likely to be demand-faulted in the near future. The principle of locality suggests that the frames most likely to be referenced in the future will be related to those referenced most recently. We consider three types of frame relationships: a frame that fills a slot in a recently fetched frame X , a frame that is a subconcept of X , and a frame that is an instance of X . Since concepts are more likely to be accessed than instances, the subconcepts of X are the first candidates for prefetching. Next, we prefetch the frames that fill some slot of X . We do not prefetch instances of X , however, because the number of instances of a class can be large and the probability of access to a prefetched instance is low. More details on the implementation of prefetching may be found elsewhere (Karp and Paley 1995).

6.4 Implementation

We have implemented PERK in conjunction with LOOM, THEO, and OCELOT FRSs, which are all implemented in LISP. Most of the PERK code is also written in LISP. We have used the ORACLE DBMS in our system. In our initial work we had experimented with the EXODUS extensible object-oriented DBMS developed at the University of Wisconsin, Madison (Franklin et al. 1992). We implemented separate prototype storage systems based on EXODUS and on ORACLE, and evaluated them empirically (Karp, Paley and Greenberg 1994; Karp and Paley 1995). We found that ORACLE is much easier to work with from a practical point of view, because SQL provides a much higher level of interaction than does the extensive C++ programming necessary to interact with EXODUS.

PERK interacts with ORACLE by using Oracle Call Interface (OCI). The OCI allows a user to embed SQL calls in an application program. The OCI calls appear as foreign function calls in our storage system. We were able to reuse some of the code for OCI calls from the Intelligent Database Interface (IDI) (McKay, Finin and O'Hare 1990). We had to optimize IDI to reuse the parsed representation of queries and the buffers used to store the query results. The implementation has been thoroughly tested with both Lucid Common LISP and Allegro Common LISP environments.

6.5 Empirical Results

The experimental setup for evaluating the storage system consists of test KBs, and software used to conduct the experiments. We considered two sets of KBs. The first set consisted of two real KBs that we obtained from the Information Sciences Institute at USC, which covered the naval and aircraft domains. The second set was randomly generated so that the characteristics of KBs approximated real KBs. All of the KBs in the second set had 100 concepts, all primitive, with just one super each. The concept hierarchy in each KB was the same, regardless of the number of instances. Different test KBs were generated with 1000, 2000, 3000, and 5000 instances. We refer to these KBs as RKB-1000, RKB-2000, RKB-3000, and RKB-5000 respectively. Instances averaged five slots apiece, with an average of two fillers per slot. Half the slots were filled by integers and the other half filled by symbols.

Experiments were run using LOOM 2.1, running on Allegro Common LISP 4.3. Similar experiments were done for the THEO FRS, and the results are available elsewhere (Karp and Paley 1995). Both the FRS and the ORACLE server were running on the same workstation, a SPARC-

KB	Time to Open a KB		Time to Initiate a Browse	
	Flat file	PERK	Flat file	PERK
Naval Theory KB	5.75	2.78	2.20	2.8
Aircraft KB	3.36	0.86	0.77	2.336
RKB-1000	100.96	1.45	0.876	9.99
RKB-3000	218.4	1.43	1.08	11.345
RKB-5000	355.0	1.44	1.32	22.58

Table 1: Interactive browsing time (seconds)

station, model Ultra 1 with 64 MB of physical memory running Solaris 5.2. LISP was restarted before every trial, to avoid caching effects, and a garbage collection was executed immediately before timing. The DBMS server was configured with large enough buffer so that the KB could be memory resident. All the experiments assume a warm start for the DBMS server. Overall elapsed times were measured using the LISP `time` function. Each trial was repeated enough times so that the confidence intervals on the mean elapsed time were less than 5% of its length. These experiments measured demand fetching times only; the prefetcher was disabled.

6.5.1 Experiment 1: Interactive Browsing

The first experiment considers a scenario where users are interested in interactively browsing and editing a KB or an ontology by using a tool such as the GKB-EDITOR. The users in this scenario perform the following operations frequently: (1) opening an existing KB (2) initiating a browse of the class hierarchy, (3) searching for a node and displaying its contents. Therefore, we measured the times to perform these operations for our test KBs.

In the second and third columns of Table 6.5.1, we show the time to open a KB by using a flat file system and PERK. The flat file numbers are obviously higher, because the whole KB must be loaded. While opening a KB, PERK loads only the root classes, the slot definitions, and any classes that may be needed in the process. The benefits of PERK can, however, depend on how interconnected the KB is. The Naval theory KB is highly interconnected, and while faulting in the root classes and relations, several other concepts are referenced and must be faulted in. As a result, for the Naval Theory KB, the flat file system and PERK differ by only a factor of 2, whereas for RKB-5000, the difference is two orders of magnitude. The time to open a random KB does not change much with its size, because in each case, the number of root classes is the same.

The experiment suggests that with PERK, a user does not have to face a long startup time before beginning to work with the KB.

In Table 6.5.1, we show the time to start a typical browse that starts from all the root classes in the KB and expands the subclasses and instances up to a maximum depth of 3. Depth is defined as the length of the shortest path from the root to a leaf node. The actual number of frames that are displayed on the screen is different for each KB. We can see that the time required by PERK is significantly higher. That is natural because PERK faults in frames that are necessary to start the browse, whereas for the flat file system all the frames are in memory. The time to initiate a browse while using PERK is still considerably smaller than the time to open a KB while using the flat file system, and the waiting periods for the user are better spread out over a period of time.

6.5.2 Experiment 2: Batch Queries

In the second experiment, we consider a scenario where users are interested in running complex programs over a KB. The access pattern in such a case can be modeled by a query that randomly references a certain fraction of the KB. We did not consider a more detailed access pattern (such as depth-first search) because for the type of result we want to show here, it would not have made any difference. Each reference faults in at least one frame from the DBMS server. In our experiment, we assume that there is no intermediate time between successive frame references. In practice, there will be some time between the requests, and during the intermediate time the prefetcher could have already brought in the frame that is needed next. Therefore, our results present a worst case scenario.

Figure 4 lets us evaluate the relative merits of loading frames from the DBMS versus loading from flat files. We consider only random KBs here because they let us study the effect of varying the KB size on the loading time. Clearly, when a small fraction of the KB is to be referenced in a given session (25%), the DBMS outperforms the flat file. When a larger fraction of a KB is to be referenced (75%), for smaller KBs, it is better to use a flat file, but for larger KBs, loading from the DBMS is still faster. For example, for a KB of size 5000, it is faster to reference 75% of the frames from the DBMS than from a flat file.

6.5.3 Experiment 3: Batch Updates

We measured the time required to save updates to some number of randomly chosen instances from KBs of various sizes. To be consistent with traditional LOOM behavior, updates are not




Figure 4: Frame loading time. The solid line shows the time required to load entire KBs of varying sizes from a flat file. The dashed lines show times required to fault in frames from the DBMS due to references to instances by an application. All times refer to total elapsed times. The vertical ordering of dashed lines in each graph and in its legend are the same.

written as they occur. Rather, we wait until the user issues a command to save updates, and then all are written in a single transaction. We varied the percentage of frames updated between 0 and 100. Selected results are shown in Figure 5. For comparison, we have included the time to save KBs of varying sizes to LOOM flat files (the time is constant for a given KB regardless of the number of frames updated in that KB).

Figure 5 demonstrates that, as expected, our architecture achieves the goal of saving KB changes in time linearly with the number of updates. In most cases, the DBMS outperforms the flat file system for saving the updated instances.




Figure 5: Update times. Each data point represents the time to save a given fraction of instances. The solid line shows the time required to save entire KBs of various sizes to LOOM flat files. The dashed lines show, for KBs of various sizes, the times required to store a given fraction of updated instances to the DBMS. In most cases, saves to the DBMS are faster than saves to the flat file.

6.6 Related Work

The Knowledge Engineering Environment Connection couples the KEE FRS with a relational DBMS (Abarbanel and Williams 1986), and the IDI couples LOOM with a relational DBMS (McKay, Finin and O’Hare 1990). Both systems employ a KB-specific mapping. The advantage of this architecture is that it allows existing information from a database to be imported into an AI environment. Its drawback is that the storage capabilities of the FRSs are not enhanced transparently, as in our approach. Users of KEEconnection (and of the IDI) must define mappings between class frames and tables in the DBMS; KEEconnection creates frame instances from analogously structured tuples stored in the DBMS, and can store instance frames out to the DBMS. However, only slot values in instance frames can be transferred to the database — class frames

cannot be persistently stored using database techniques and cannot be accessed by multiple users. Our approach allows *all* information in a KB to be permanently stored in the DBMS.

Groups at IBM and at MCC have coupled FRSs to object-oriented DBMSs (Mays et al. 1991; Ballou et al. 1988). The IBM effort differs from our approach in that a KB is read from the DBMS in its entirety when it is opened by a K-REP user, which we believe will be unacceptably slow for large KBs.

Markowitz and Chen have developed a system called *OPM* that implements an object-oriented data model on top of a relational DBMS (Chen and Markowitz 1995). They have developed automated methods for mapping an OPM object-oriented schema into a relational schema, and for mapping queries in the OPM query language into queries to the underlying relational schema.

Borgida and Brachman have connected the CLASSIC FRS to a database by using a two-phase approach. In the first phase, they load the individual objects and associated facts into an FRS client through a back-door where the usual CLASSIC inference machinery is turned off. Instead, many inferences are performed by the DBMS by asking the appropriate queries. In the second phase, for the inferences that cannot be computed by the database, they use queries to eliminate from consideration those instances that will not add any new information, so that the inference needs to be performed on only a small subset of all the instances. Their approach is geared toward a system that has numerous instances already stored in the database, a complex conceptual schema available in the FRS, with a goal of rapidly loading those individuals into the FRS. Their approach cannot take a complex CLASSIC schema and store it in the database.

7 Collaboration Subsystem

The development of large KBs and shared ontologies requires collaboration of multiple people who make simultaneous contributions. Most existing FRSs, however, are single-user systems and allow only one person at a time to update a KB, and therefore are inadequate to support collaborative work. It is neither viable to maintain multiple copies of the KB for each of its users, nor feasible to restrict access to the KB to one user at a time. Using an existing commercial DBMS is an unacceptable alternative for the following reasons. The commercial DBMSs control the operations of multiple users by locking the data items that are accessed. The net effect is to isolate the database in a way that each user gets the illusion that she is the sole user of the database. KBs are highly interconnected, and locking one object may require locking several other objects.

Transactions in a KB usually access a large number of entities, for example, while performing an inference using backward chaining. If we use locking to control concurrent access, large portions of a KB need to be locked for long periods of time, thus limiting the potential concurrency. Locking model of controlling concurrent access may work well for short online transactions, but in a collaborative environment, it prevents users from working together. Therefore, we need a facility that makes users aware of each other and *helps* them to work together instead of *preventing* them from creating conflicting updates.

7.1 Proposed Solution

To deal with the problem described above, we allow users to make independent changes to the KB; when their changes are complete, we merge their changes with the KB. The user changes merged with the public copy of the KB are those that do not *conflict* with any updates that have been applied to the KB since the begin time of the user's session.

Our solution is similar to an optimistic concurrency control technique in which the users make independent changes to the database, but one (or more) of the users making a conflicting change must abort the changes and restart (Bernstein, Hadzilacos and Goodman 1987). In our framework, instead of issuing an abort, we go a step further and assist the users in identifying and resolving their conflicting updates. We use locking only while the conflict-free updates are being deposited into the KB.

Our work so far has considered only those operations that change slot values and do not involve any schema changes. We plan to consider schema changes in our future work.

7.1.1 Model of Collaboration

We maintain a *public* copy of the KB that is readable by multiple KB developers, but cannot be modified directly. All updates to a KB occur in *private workspaces* (or simply *workspaces*). A developer who wants to modify an existing public KB creates a private workspace as a child of that public KB; by default the workspace contains all of the information present in its parent, but is not necessarily obtained by copying the parent KB. The developer then modifies the workspace by creating new frames, and updating or deleting existing instance frames. Those updates do not affect the public KB, and a private workspace can be accessed by only a single user at a time. A developer who has brought the workspace to a satisfactory state, then *merges* the workspace with the newest state of the public KB. The merge operation detects and resolves conflicts (incon-

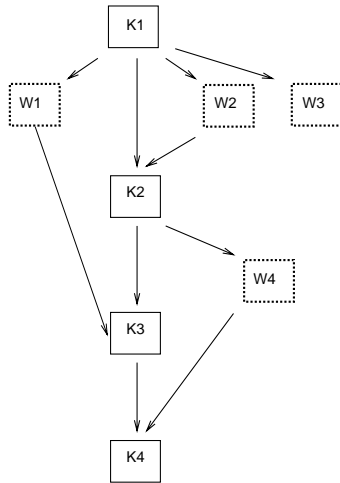


Figure 6: A sample set of relationships among public KBs (K1–K4), and private workspaces (W1–W4)

sistencies) between the updates made in the private workspace, and updates by other users from which the public KB has been derived.

Figure 6 illustrates these concepts more clearly. The boxes labeled K1–K4 represent successive states of a public KB. The boxes labeled W1–W4 represent private workspaces. Each workspace has a single parent that is a public KB. Each state of a public KB either is the initial state, or is derived from a merge of the previous state of that public KB with a private workspace. For example, K2 is derived from a merge of K1 with W2.

Because K1 is the parent of W2, and is also the public KB with which W2 is merged, no conflicts can occur during that merge. However, when W1 is merged with K2 to create K3, conflicts may occur between the modifications made in W2 and the modifications made in W1. Developers were updating these two private workspaces simultaneously, and no attempt was made by the KBMS to enforce any consistency with respect to those updates. Consistency is enforced at merge time. Merging will require detections and resolution of conflicts (which we term *arbitration*), and creation of a new state of the public KB.

The users always work with the state of the public KB with which they started. For example, in Figure 6, even after KB is in state K2, the user of workspace W1 does not see any changes and continues to work with K1.

7.1.2 Recording User Updates

We assume a KB that uses the knowledge model of GFP (Section 4.2.1); that is, it consists of classes, instances, slots, facets, and values. The user interacts with the KB by means of GFP operations. A transaction is a sequence of update GFP operations.

With the public copy of the KB, we also record a log, called the *net-log*, of all the changes that have occurred since a given time in the past. When the updates of a transaction are merged with the public copy, conflicts are detected by comparing the operations in the transaction with the operations in the net-log. For detecting conflicts, only those net-log operations need to be considered that occurred since the begin time of the transaction being merged.

A log consists of a sequence of log records. Each log record is a list with two elements. The first element is a list containing the name of the GFP operation and the arguments with which it was invoked. The second element is the list of any values that are being overwritten by the current operation; it has a non-null value only when some value is being overwritten and the old value is not available as part of the GFP operation itself. For a **put-slot-value** operation, we must record the old values of the slot. For a **replace-slot-value** operation, the old value of the slot is one of the arguments of the GFP operation, and we do not need to record it separately. The old values are useful for inverting an update operation when necessary.

7.1.3 Implementation of KB States

When the user requests a value from the KB, the FRS must return a value that corresponds to the state in which the user started his or her work. For example, consider a KB state KB_1 that contains a frame **person** with a slot **friend** and an instance **John** whose slot **friend** initially contains the value **Peter**. Suppose a user updates the KB to add a new value **Adam** as a friend of **John**. Let the resulting KB states be K_2 . If another concurrent user who started when the KB was in state K_1 , and queries for the friends of **John**, she should get **Peter** as an answer even if the query is executed when the public KB is in state K_2 . There are two approaches to implement such behavior: *delta* and *interval*.

In the *delta* approach, we physically store only one state of the KB, and use the log to compute other states of the KB. The delta approach can be implemented in two ways: positive and negative. In the *positive delta* approach, we store the oldest state and apply the necessary log records to the KB to compute a later state with respect to which the query is to be answered. In the *negative delta* approach, we store the newest state of the KB and negatively apply the updates to answer

the queries with respect to the older states of the KB. If we apply the negative delta approach to the example considered in the previous paragraph, we store K_2 , negatively apply the entries in the log that are between the KB states K_2 and K_1 (which in this case is an operation adding **Adam** as a friend of **John**) to obtain the desired answer (**Peter** is a friend of **John**). Saving the updates from a workspace is more expensive with the negative delta approach, because in addition to saving the log of changes, we must save the updated frames. On the other hand, purging a log is more expensive for the positive delta approach, because in addition to deleting the log entries, we must save the new state of the updated frames. In general, more users are interested in the recent states of the KB, and therefore the negative delta approach is likely to be more efficient for answering the queries than the positive delta approach.

In the *interval* approach, we associate an interval with each entity in the KB indicating the duration for which it exists in the KB. For example, if the value of a slot salary is 20K in interval t_1 and 25 K during interval t_2 , the KB contains both of these values. We associate the interval t_1 with the first value and t_2 with the second value. To evaluate any queries on the slot salary, we can use a temporal join index (Shrufi and Topaloglou 1995). While the query is evaluated, only the answers that are valid for the duration specified in the query are returned. In the example considered in the previous paragraph, we get only **Peter** as an answer, because that is the only value valid for the KB state K_1 .

Over a period of time, the delta approach accumulates a log and the interval approach accumulates past values of entities. Therefore, both approaches require a purging process to remove the information that is no longer necessary. In the negative delta approach, one needs to periodically purge the net-log, and in the interval approach, one needs to periodically purge the old versions of entities. Clearly, the interval approach requires more storage space than the delta approach, because even the entities that have been deleted must be stored. The delta approach has higher runtime cost because the desired answer must be computed by evaluating the relevant log records. If we assume that the purging overhead is comparable for the two approaches, we are faced with the classical computing versus storage tradeoff. A more detailed evaluation of the two approaches is planned for future work. In our current system, we have implemented the negative delta approach.

7.1.4 Definition of Conflict

We say that two operations o_1 and o_2 are *conflict-free* if, for any KB state, executing o_1 followed by o_2 leaves the KB in the same final state and returns the same values for each as executing o_2 followed by o_1 . For example, inserting two frames with different names is conflict-free. But if a user wants to change a slot value from 2 to 5 and another user wants to change that same value from 2 to 10, then their operations are not conflict-free.

It is useful to consider alternative (and perhaps weaker) notions of freedom from conflict. We say that operations o_1 and o_2 are *partially conflict-free* if, for any KB state, executing o_1 followed by o_2 leaves the KB in the same final state but may not necessarily return the same values for each one of them as executing o_2 followed by o_1 . For example, if two users were trying to create a frame with the same name, then their operations leave the KB in the same state but the user who executes the operation first is successful and the operation of the second user is a no-op. We call such operations *partially conflict-free*. Under many situations, it is acceptable to have partially conflict-free operations.

7.1.5 Conflict Detection

For conflict detection with respect to a transaction T , we examine the portion of the net-log that contains operations executed after the begin time of T . To check the conflicts between a user transaction and the net-log, we must check for each operation in the transaction, if the net-log contains an operation that performs a conflicting update. Each update GFP operation can, in general, involve multiple updates. For example, the `put-slot-values` operation deletes the old slot value(s) and inserts several new slot values. Furthermore, the number of possible GFP operations is large (over 200), so that analyzing conflicts between the GFP operations can be cumbersome. Therefore, before analyzing the conflicts, we translate the operations into a canonical set of three operations — `INSERT`, `DELETE`, and `REPLACE` — on the nodes and edges of the underlying KB graph (defined below). A `REPLACE` operation modifies an existing value. Since the number of operations in the canonical set is considerably smaller than the number of GFP operations, conflict analysis is substantially simplified.

We view the KB as an undirected graph. The nodes represent classes, instances, slots, and values. The edges represent class-subclass, class-instance, class-slot, instance-slot, and slot value relationships. Thus, there are four types of a node and five types of an edge. We use the generic term *entity* to denote either a node or an edge.

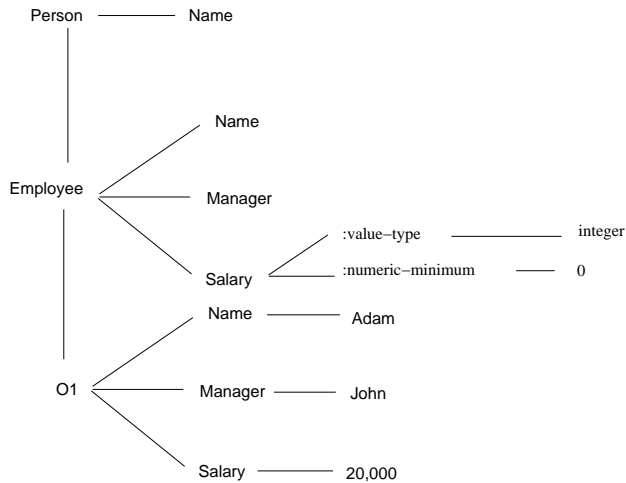


Figure 7: The KB KB_1

In Figure 7, we show a KB KB_1 in a graph form. **Employee** and **Person** are classes and represented as nodes. The subclass relationship between them is represented by an edge. O1 is an instance of **Employee**. **Person** has one slot called **Name**. It is inherited by **Employee**, which has two local slots — **Manager** and **Salary**. The slot values for the instance O1 are shown as nodes with edges from the slot nodes for O1. Since a slot can appear several times in the graph, to identify a slot node uniquely, we must associate it with the frame node it is attached to. For example, the slot **Salary** of the frame O1 could be identified as (O1, **Salary**).

A replace operation for edges is not defined. An insert operation on a node is always accompanied by an edge insertion. For example, in the KB KB_1 , if for frame O1, we add the value of 20000 for the slot **Salary**, we are inserting a value node 20000 and a slot value edge from the node (O1, **Salary**). Similarly, a node deletion is always accompanied by an edge deletion. To delete Adam's **Salary**, we delete the node associated with the value and at the same time delete the edge between the (O1, **Salary**) slot and the value node. We are interested only in sensible operations. Under an assumption of no duplicates, sensible operations can be defined to mean that one can delete (insert) an entity only if it exists (does not exist) in the KB. The set of all operations is the cross product of the set of entity types and the set of operations.

With the above model, the conflict analysis between GFP operations reduces to conflict analysis between graph operations. When the operations involve distinct nodes or edges, they are trivially conflict-free. Therefore, we consider only the situations in which the operations involve the same

entity.

Table 2 shows the conflict matrix for operations such that both of them operate on either a node or an edge. To explain the table, and to keep the discussion concrete, consider the case of slot operations. For two slot operations to conflict, they must refer to the same frame and involve the same slot; that is, if the operation is on a slot value X, its frame name and the slot name must be the same in the two operations under consideration. Consequently, for showing the operations in Table 2, we omit the frame name and the slot name.

When both operations attempt to insert the same value in a slot, they are partially conflict-free, because only one of them succeeds. If one operation inserts a slot value and another deletes it, the operations cannot both be defined, because for a value to be inserted, it must not exist in the KB, and for it to be deleted it must exist in the KB, which is a contradiction. Therefore, we do not analyze conflicts for such situations.

The conflicts resulting from an operation replacing the slot value X of an entity can lead to three situations: two operations trying to replace X with another but same value, two operations replacing X with a different value, or an operation replacing slot value Y to X. When two operations try to replace a value of slot from X with Y, only one of them will succeed, but they will leave the KB in the same state, and therefore they are partially conflict-free. When one of the operations wants to replace X with Y and another X with Z, they conflict, because they have different effects on the KB. Finally, the situation when one operation replaces X with Y and another operation replaces Z with X cannot occur in practice, because for the former operation to happen, X must exist in the KB but for the latter operation, X must not exist, which is a contradiction. However, the operation renaming Y to X conflicts with the operation that inserts X, because executing them in a different order will have different effects on the KB. An operation that deletes X and another operation that replaces X with Y conflict because one of them intends to retain the item in the KB and the other intends to remove it.

Let us now briefly see how Table 2 is applicable for class operations. Two operations inserting a class into the KB are partially conflict-free. Replacing a class name has implications similar to modifying a slot value. For example, two operations replacing a class with different names conflict, but replacing with the same name are partially conflict-free. Table 2 also generalizes to edge operations. To specify an edge operation, however, we do need to specify both the end points, and two edge operations conflict only if they involve the same end points. Thus, two $\text{Insert}(X,Y)$ operations are partially conflict-free. Since the replace operation is not defined for edges, only the

	Insert(X)	Delete(X)	Replace(X,Y)	Replace(X,Z)	Replace (Y,X) Replace (Z,X)
Insert(X)	P	*	*	*	N
Delete(X)	*	P	N	N	*
Replace(X,Y)	*	N	P	N	*

P – Partially conflict-free, N – Not conflict-free, “*” – operations not defined

Table 2: Conflict matrix for node operations or edge operations. For example, Delete(X) and Replace(X,Y) conflict, because the effect of executing them in different order leaves the KB in different states, because only the operation that is executed first succeeds.

first two columns and first two rows of Table 2 are relevant to edge operations.

7.1.6 Merging User Logs

The merge process proceeds in the following steps: log translation, log simplification, log modification, conflict detection, conflict resolution, and log concatenation. Log translation transforms the user log, which is originally represented as a sequence of GFP operations, into a sequence of operations on the underlying KB graph. Log simplification takes a translated log, represented as a sequence of graph operations, and obtains the smallest possible log that has the same effect on the KB as the original log. The simplified log captures the net effect of the changes made by a user in a session. Log simplification is necessary, because while detecting conflicts in a later stage of merging, we do not want to consider any redundant operations. Log modification takes the simplified user log, and modifies it to take into account any changes that have taken place since the time the user started. Conflict detection compares the user log with the net-log and identifies conflicting operations. Conflict resolution resolves the conflicts identified in the conflict-detection phase. Once the user log and net-log are free from any conflicting operations, the user log can be simply concatenated to the net-log.

Log Translation. Typically, each GFP operation translates into a series of operations on the underlying KB graph. As an example of this translation, we consider a few GFP operations. For example, the `put-slot-value` operation translates to the deletion of all the nodes corresponding to old values and insertion of nodes corresponding to new values. As a more involved example, consider the `create-class` operation. The parameters of this operation also include the names of

superclasses and slots. If there is only one superclass and no slots, then the operation is equivalent to an insert-node operation, which inserts a node corresponding to the class and an edge between the class and its superclass. In addition, the class inherits all the slots (except for the local slots) from the superclass, which translate into inserting a node and an edge corresponding to each slot. If some of the inherited slots have default values, the translation will contain an insert node operation corresponding to each default value.

Log Simplification. Log simplification uses a collection of simplification rules to obtain a log that is smaller than the original log but has the same effect on the KB as the original log. For example, the insertion of an entity followed by its deletion can be simplified to a null operation. Similarly, if an operation inserts the slot value A, and a following operation replaces A with B, the two operations can be simplified to an operation that inserts the slot value as B. In general, two successive operations on the same slot value can always be simplified to one operation. Therefore, in the simplified log, there is only one operation on each entity.

Log Modification. Log modification changes the user-log to take into account that the KB is not in the same state as when the user started. For example, suppose the initial value of a slot when the user starts is 20. Suppose the net-log contains an operation from another user that adds a second value 30 to this slot. Let the user-log contain a **put-slot-value** operation asserting the value of the slot to be 40, which when executed by the first user, removes the old value of 20 and adds the new value of 40. Since by the time the log is merged, the slot has another value of 30, which should be removed by the **put-slot-value** operation in the user-log, the value 30 must be deleted, and the user should be informed of this additional effect of her operation. We use log modification to achieve this, but we must encode in the log some information that can be used during modification. We solved this problem by introducing wild cards in the user log. In the above example, in addition to representing in the log the deletion of slot value 20, we introduce a delete operation in which the slot value is represented as a wild card, meaning that all values of that slot should be deleted. The operation containing the wild card serves as a pattern that is matched against all the insert operations in the net-log. If there is an insert operation on a slot for which all values should be deleted, we generate delete operations for the slot value appearing in the insert operation.

Log Concatenation. Once conflicts have been detected and the user-log modified, it does not contain any operation that conflicts with some operation in the net-log and is ready for concatenation. Log concatenation involves simply appending the user-log to the net-log. At the

successful completion of log concatenation, all the user updates are saved to the KB and a new KB state is created.

Over a period of time, the net-log grows in size. To keep its size manageable, we purge the logs that are no longer necessary. The operations of a user-log L_1 can be purged from the net-log if all the users who started their session before L_1 have committed their changes. Alternatively, older log entries can be retained to save the history of KB modifications.

Since the net-log can be large, we store it in the ORACLE DBMS. Our storage system supports the storage and retrieval of the log from the DBMS server. Because the log is stored in a central server, different clients' sessions can retrieve log entries and reconstruct an older state of a KB if necessary.

7.1.7 Correctness

The conflict matrix defined in Table 2 should have the property that merging conflict-free updates to the public copy of the KB always lead to *correct* KB states. Correctness is commonly captured by the notion of *serializability* (Papadimitriou 1986).

Serializability can be informally defined as follows. Let a transaction be the set of operations executed by a user. An interleaved execution of a set of transactions is serializable if it is equivalent to some serial execution of the same set of transactions in the sense that it leaves the database in the same state and returns the same answers to each one of its users. The serializability graph of an execution is a graph with a node T_i for each transaction T_i and an edge (T_i, T_j) if a step of T_i precedes in the execution a conflicting step of T_j . It is well known that an execution is serializable if and only if its serializability graph is acyclic (Eswaran et al. 1976).

By induction on the number of transactions in the net-log, we can prove that the merge process described in the previous section produces serializable executions. Let L_i denote the net-log with i transactions. The base case easily follows because L_1 is trivially serializable. For contradiction, suppose L_{m+1} is not serializable. There must be a cycle in the serializability graph of L_{m+1} that includes T_{m+1} and contains an edge from some $T_i, i \leq m$ to T_{m+1} . But that is not possible because the updates of T_{m+1} are conflict-free from the updates of all $T_i, i < m$.

We must, however, make sure that the conflict matrix covers all possible cases of conflict between user operations. There are two main classes of conflicts — conflicts between update operations and conflicts between a read and an update operation. The conflict matrix defined in Table 2 covers all cases of conflicts between update operations and some of the important cases of

conflicts between a read and an update operation.

Traditionally, a read operation on an entity A by a transaction T_1 is assumed to conflict with an update operation on A by T_2 . But the two operations conflict only if T_1 's read operation influences any values written by it. In practice, it is hard to determine which read operations influence the values written by a transaction. In the KB environment, we have found that transactions read a large number of entities, but most of them do not influence the values written by it. For example, a user may explore a KB using the GKB-EDITOR, but may change the value of only one slot of a frame. In our current work, we assume that a read operation on an entity A does not conflict with an update operation on A unless we have reasons to believe otherwise.

We have considered two important situations when a read operation on an entity A conflicts with an update operation on A . The first situation is considered in the conflict matrix defined in Table 2, and arises when a transaction replaces A to B and another transaction replaces A to C . The second situation is considered in log modification when a transaction deletes all the values of a slot, and by the time its updates are merged, a new value of the slot has been added. In our future work, we plan to consider other situations where a read operation on an entity A influences the values written by the transaction.

7.2 Implementation

We have implemented the solution described above. One of the difficulties we had to address in the implementation was that while merging the updates of a transaction T , we need to consider only those operations in the net-log that were applied to it after the begin time of T . It is straightforward to identify such operations if we keep track of the begin timestamp of each transaction. Let us name the subset of the net-log that is of interest to be *sub-net-log*. The *sub-net-log* must be simplified, because we should check conflicts with respect to the net effect of the operations. We used the same procedure for simplifying the sub-net-log as we used for simplifying the user-logs. But while simplifying the sub-net-log, if two (or more) operations are simplified, we must associate the simplified operation with the original operations, because in the event that any conflict is detected, we need to determine who made the conflicting updates.

We performed extensive testing of the implementation for merging, using some sample logs collected for a KB that is under development at SRI. All the logs were conflict-free, which suggests that most of the time the users will work on disjoint portions of the KB. Log simplification resulted in simplified logs in many cases, confirming its utility in practice. A more detailed evaluation is

planned as future work.

7.3 Related Work

Many approaches have been proposed for supporting collaborative design (Nodine and Zdonik 1990; Kaiser 1990; Chou and Kim 1986; Korth and Speegle 1990). Hall has developed a model for collaborative work on design databases that is broadly quite similar to what we propose, but differs in a number of specifics. The central difference between Hall's model and our model is that Hall's model relies on the process of change notification to maintain the consistency of each object, whereas our model relies on the merge operation. In Hall's model, inconsistencies can be detected very quickly since change notifications can be sent as users perform updates. In our model, inconsistencies are detected only when an individual user merges his changes into the public KB. A limitation of Hall's model is that two users cannot simultaneously check out the same object for update in different workspaces, which retains some of the disadvantages of locking. Generally speaking, our approaches are complementary, and further knowledge of their relative merits must be gained empirically.

CYC-L (Lenat and Guha 1989) provides the most sophisticated knowledge-sharing capabilities of any FRS (but which are nonetheless insufficient for many needs): A user can copy the virtual memory image of a KB stored on a master knowledge server to the user's workstation (a process that is extremely time consuming). When the user changes the KB on his workstation, CYC-L transmits KB updates to the master server. The server maintains its KB in virtual memory, not in persistent storage, and it has no mechanism for controlling concurrent KB updates.

Mays and his colleagues added persistence and sharing capabilities to the K-REP FRS (Mays et al. 1991). They developed the notion of relaxing the strict consistency requirements of traditional databases in favor of merging privately developed workspaces (which they call *versions*) that may contain conflicts. Their merge operation is not published, but our private communications with Mays have convinced us that their models of workspaces and of merging are significantly different from ours, and that several approaches to this new paradigm of KB development should be explored.

Another proposal, similar to ours, uses "history merging" as a concurrency control mechanism for collaborative applications (Wäsch and Klas 1996). The merging algorithm, called *import algorithm*, computes the minimal subset of a history in the same way we do log simplification. It then looks for conflict-free subsets of the history and tries to merge them. The authors use

the conventional notions of conflict and assume a relational model. In contrast, we have used an object model of the database that is more current.

An extensive body of work on concurrency control in partitioned networks is relevant in our context (Davidson, Garcia-Molina and Skeen 1985). A partitioned network arises in a replicated database scenario when a subset of the sites get disconnected from the rest because of communication failure. Such a situation arises more naturally in a mobile computing environment where users are disconnected from a central server for extended periods of time (Terry et al. 1995). Since databases naturally contain logs to support recovery, our conflict detection and log merging techniques can be easily adapted for both mobile computing and replicated databases.

8 Summary and Conclusions

Our research results advance the state of the art of knowledge base management systems in several ways. The GFP provides a generic API for KBs that constitutes a solid foundation for knowledge sharing and software reuse. It unifies both description-logic and frame KBs. The strength of GFP lies in its simplicity and extensibility. Several GFP wrappers are now in active use, and GFP is being used in DARPA's high-performance KB initiative (see <http://www.teknowledge.com:80/HPKB/>). The protocol is undergoing some improvements, for example, we are enhancing it to deal with highly expressive KBs encoded in first-order logic. It is now sufficiently mature that the implementors of FRSs can consider it as a serious candidate for an API for their system.

Our storage system transparently supports efficient storage and retrieval of KBs in a DBMS to provide improved scalability for KBs. Because the storage system utilizes a KB-independent DBMS schema, it can be used with different KBs in a fashion that is transparent to the user. A frame-based schema provides fast demand faulting of frames, and a slot-based schema allows the DBMS to support fast indexed queries. Our prefetching architecture allows additional performance improvements by exploiting idle times. The storage system is portable to other FRSs, because only a small part of it is dependent on an FRS. We have demonstrated its portability by using it in conjunction with the LOOM, THEO, and OCELOT FRSs.

The collaboration subsystem permits multiple users to cooperatively update a shared KB. In our model, users perform updates in independent workspaces that allow them to maintain long-duration sessions without interference from other concurrent users. The user updates are recorded as a log of GFP operations. At the end of a session, the log is simplified to remove any redundant

operations. To enable efficient conflict detection, the operations in the user log are translated to operations on a graph that encodes the structure of the KB. Rules for checking conflicts between graph operations are defined as a matrix. Whenever possible, conflicts are resolved or support is provided for resolving the conflicts. The collaboration system has been tested in the context of an ongoing scientific KB project at SRI.

The GKB-EDITOR completes the collaborative KB-authoring environment by providing four different viewer/editor utilities for browsing and modifying large KBs: the taxonomy viewer, relationships viewer, frame viewer, and spreadsheet viewer. Each viewer is optimal for different browsing and editing tasks. The incremental browsing capabilities of the GKB-EDITOR help a user to explore large KBs as a starting point for KB reuse. The GKB-EDITOR itself has been reused across multiple domain KBs and across multiple FRSs because it performs all KB access and update operations via GFP. The GKB-EDITOR provides extensive facilities for customizing the user interface, such as changing the color and shape of nodes and edges, displaying inherited and local attributes in different colors, etc.

The next generation of KBs will be collaboratively manufactured by teams of knowledge engineers through assembly of prefabricated knowledge components and through semiautomatic population of the instances in a domain of interest. Knowledge will be manipulated by diverse problem-solving engines even though they were developed in different languages and by different researchers. The KBs will be of unprecedented size and will use a scalable and interoperable development infrastructure. Tremendous speedups in the creation of KBs will significantly enhance our ability to quickly solve complex problems. The results described in this paper provide important technologies for developing the knowledge bases of the twenty-first century.

Acknowledgements

This work was supported by DARPA contract F30602-92-C-0115, by Rome Laboratory contract F30602-96-C-0332, and by grant R29-LM-05413-01A1 from the National Library of Medicine. The contents of this article are solely the responsibility of the authors.

References

Abarbanel, R. and Williams, M. 1986. A Relational Representation for Knowledge Bases. Technical report, IntelliCorp, Mountain View, CA.

- Ballou, N., Chou, H.-T., Garza, J. F., Kim, W., Petrie, C., Russinoff, D., Steiner, D., and Woelk, D. 1988. Coupling an Expert System Shell With an Object-Oriented Database System. *Journal of Object Oriented Programming*, 1(2):12–21.
- Bernstein, P. A., Hadzilacos, V., and Goodman, N. 1987. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley Publishing Company.
- Borgida, A. and Brachman, R. J. 1993. Loading Data into Description Reasoners. In *Proceedings of the 1993 ACM SIGMOD International Conference on Management of Data*, pages 217–226, Washington DC.
- Borgida, A., Brachman, R. J., McGuinness, D. L., and Resnick, L. A. 1989. CLASSIC: A Structural Data Model for Objects. In *Proceedings of the 1989 ACM SIGMOD International Conference on Management of Data*, pages 58–67, Portland, OR.
- Chen, I. A. and Markowitz, V. M. 1995. An Overview of the Object-Protocol Model (OPM) and the OPM Data Management Tools. *Information Systems*, 20(5):393–418.
- Chou, H. and Kim, W. 1986. A Unifying Framework for Version Control in a CAD Environment. In *Proceedings of the Twelfth International Conference on Very Large Data Bases*, pages 336–344.
- Copeland, G. and Khoshafian, S. 1985. A Decomposition Storage Model. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*, pages 268–279.
- Crawford, J. 1990. *Access-Limited Logic — A Language for Knowledge Representation*. PhD thesis, University of Texas at Austin. Technical report AI90-141.
- Davidson, S. B., Garcia-Molina, H., and Skeen, D. 1985. Consistency in Partitioned Networks. *Computing Surveys*, 17(3):341–370.
- Eilerts, E. 1994. KnEd: An Interface for a Frame-Based Knowledge Representation System. Master’s thesis, University of Texas at Austin.
- Eriksson, H., Puerta, A. R., Gennari, J. H., Rothenfluh, T. E., Tu, S. W., and Musen., M. A. 1994. Custom-Tailored Development Tools for Knowledge-Based Systems. Technical Report KSL-94-67, Stanford University Knowledge Systems Laboratory, Stanford, CA.

- Eswaran, K., Gray, J., Lorie, R., and Traiger, I. 1976. The Notions of Consistency and Predicate Locks in Database Systems. *Communications of the ACM*, 19(9):624–633.
- Farquhar, A., Fikes, R., Pratt, W., and Rice, J. 1995. Collaborative Ontology Construction for Information Integration. Technical Report KSL-95-63, Stanford University, Knowledge Systems Laboratory, Stanford, CA.
- Franklin, M., Zwilling, M., Tan, C. K., Carey, M., and DeWitt, D. J. 1992. Crash Recovery in Client-Server EXODUS. In *Proceedings of the ACM SIGMOD 1992 Annual Conference*, pages 165–174, San Diego, CA.
- Geiger, K. 1995. *Inside ODBC*. Microsoft Press.
- Genesereth, M. R. and Fikes, R. E. 1992. Knowledge Interchange Format, Version 3.0 Reference Manual. Technical Report Logic-92-1, Computer Science Department, Stanford University, Stanford, CA.
- Gruber, T. 1993. A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2):199–220. URL for Ontolingua is <http://www-ksl.stanford.edu/knowledge-sharing/ontolingua/README.html>.
- Kaiser, G. E. 1990. A Flexible Transaction Model for Software Engineering. In *Proceedings of the Sixth International Conference on Data Engineering*, pages 560–567.
- Karp, P. 1992. The design space of frame knowledge representation systems. Technical Report 520, SRI International, Artificial Intelligence Center. URL <ftp://www.ai.sri.com/pub/papers/karp-freview.ps.Z>.
- Karp, P. and Paley, S. 1995. Knowledge Representation in the Large. In *Proceedings of the 1995 International Joint Conference on Artificial Intelligence*, pages 751–758.
- Karp, P., Paley, S., and Greenberg, I. 1994. A Storage System for Scalable Knowledge Representation. In Adam, N., editor, *Proceedings of the Third International Conference on Information and Knowledge Management*, pages 97–104, New York. Association for Computing Machinery. Also available as SRI International Artificial Intelligence Center Technical Report 547.
- Karp, P., Riley, M., Paley, S., and Pellegrini-Toole, A. 1996. EcoCyc: Electronic Encyclopedia of *E. coli* Genes and Metabolism. *Nucleic Acids Research*, 24(1):32–40.

- Kehler, T. and Clemenson, G. 1984. KEE: The Knowledge Engineering Environment for Industry. *Systems And Software*, 3(1):212-224.
- Korth, H. F. and Speegle, G. D. 1990. Long-Duration Transactions in Software Design Projects. In *Proceedings of the Sixth International Conference on Data Engineering*, pages 568-574, Los Angeles, CA.
- Lenat, D. B. and Guha, R. 1989. *Building Large Knowledge-based Systems: Representation and inference in the Cyc Project*. Reading, Mass.:Addison-Wesley Publishing Co.
- MacGregor, R. 1991. The Evolving Technology of Classification-based Knowledge Representation Systems. In Sowa, J., editor, *Principles of Semantic Networks*, pages 385-400. Morgan Kaufmann Publishers, Los Altos, CA.
- Mays, E., Lanka, S., Dionne, B., and Weida, R. 1991. A Persistent Store for Large Shared Knowledge Bases. *IEEE Transactions on Knowledge and Data Engineering*, 3(1):33-41.
- McKay, D., Finin, T., and O'Hare, A. 1990. The Intelligent Database Interface: Integrating AI and Database Systems. In *Proceedings of the 1990 National Conference on Artificial Intelligence*, pages 677-684. Morgan Kaufmann Publishers.
- Mitchell, T., Allen, J., Chalasani, P., Cheng, J., Etzioni, E., Ringuette, M., and Schlimmer, J. 1989. THEO: A Framework for Self-Improving Systems. In *Architectures for Intelligence*. Erlbaum.
- Neches, R., Fikes, R., Finin, T., Gruber, T., Patil, R., Senator, T., and Swartout, W. 1991. Enabling Technology for Knowledge Sharing. *AI Magazine*, 12(3):36-56.
- Nodine, M. H. and Zdonik, S. B. 1990. Cooperative Transaction Hierarchies: A Transaction Model to Support Design Applications. In *Proceedings of the Sixteenth International Conference on Very Large Data Bases*, pages 83-94, Brisbane, Australia.
- Papadimitriou, C. 1986. *The Theory of Database Concurrency Control*. Computer Science Press, Rockville, MD.
- Patil, R., Fikes, R. E., Patel-Schneider, P. F., Mackay, D., Finin, T., Gruber, T., and Neches, R. 1992. The DARPA Knowledge Sharing Effort: Progress Report. In *The Third International Conference on Principles of Knowledge Representation and Reasoning*, pages 777-788, Boston, MA.

- Shrufi, A. and Topaloglou, T. 1995. Query Processing for Knowledge Bases Using Join Indices. In *Proceedings of the 4th International Conference on Information and Knowledge Management*, Baltimore.
- Skuce, D. and Lethbridge, T. 1995. CODE4: A Unified System for Managing Conceptual Knowledge. *International Journal of Human-Computer Studies*.
- Terry, D., Theimer, M., Petersen, K., Demers, A. J., Spreitzer, M., and Hauser, C. H. 1995. Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System. In *Proceedings of the 15th ACM Symposium on Operating Systems Principles (SOSP-15)*, Copper Mountain Resort, Colorado.
- Valduriez, P., Khoshafian, S., and Copeland, G. 1986. Implementation Techniques of Complex Objects. In *Proceedings of the 12th International Conference on Very Large Data Bases*, pages 101–109, Kyoto, Japan.
- Wäsch, J. and Klas, W. 1996. History Merging as a Mechanism for Concurrency Control in Cooperative Environments. In *Proceedings of RIDE'96: Interoperability of Non-Traditional Database Systems*, New Orleans, Louisiana.
- Wilkins, D. 1990. Can AI planners solve practical problems? *Computational Intelligence*, 6(4):232–246.

Biographical sketch of the authors

Dr. Peter D. Karp, Senior Computer Scientist: Dr. Karp received his M.S. and Ph.D. in Computer Science from Stanford University. His research at SRI over the past five years has yielded new techniques for building large, shared knowledge bases. He is the principal architect of three SRI technologies: the PERK system for extending the storage capabilities of FRSs and for providing FRSs with multiuser update capabilities using optimistic concurrency control techniques, the Generic Frame Protocol (GFP) for providing standardized and portable programmatic access to FRSs, and the GKB-Editor for providing graphical browsing and editing services for multiple FRSs. Dr. Karp has applied these technologies to the development of scientific knowledge bases for the human genome project.

Dr. Vinay K. Chaudhri, Computer Scientist. Dr. Chaudhri has been a member of SRI's Artificial Intelligence Center since April 1995. At SRI, his research focuses on reusable tools for building large shared knowledge bases. He has contributed to the design and development of a collaboration subsystem, persistent store for knowledge bases, the Generic Frame Protocol, and the GKB-Editor. Dr. Chaudhri received his Ph.D. degree from the University of Toronto, Master's degree from Indian Institute of Technology, Kanpur, and Bachelor's degree from Kurukshetra University. His doctoral dissertation *Transaction Synchronization in Knowledge Bases: Concepts, Realization and Quantitative Evaluation* was one of the first attempts to address the issues of concurrent access for knowledge bases. He devised algorithms that can effectively exploit the semantic structure of a knowledge base to offer more efficient concurrency control than would be otherwise possible. His master's thesis was on interactive relational database design and was published as *Lecture Notes in Computer Science* by Springer Verlag. He worked for one and one-half years for Tata Consultancy Services, New Delhi, where he was involved in the development of large-scale information systems.

Suzanne M. Paley, Computer Scientist: Ms. Paley received her B.S. and M.S. degrees in chemistry from Stanford University, and her M.S. degree in computer science from the University of California, Berkeley. Her Master's research, conducted at SRI, explored means of extending the storage capabilities of frame representation systems by using a relational database as backing storage, including the design and implementation of a novel prefetching mechanism. Also at SRI, she has developed a graphical user interface for EcoCyc (an electronic encyclopedia of E.

coli metabolism), concentrating on automated displays of biochemical pathways. Other projects include the GKB-Editor, an application for graphically browsing and editing knowledge bases from arbitrary frame representation systems, and CWEST, a toolkit for making CLIM applications available over the World Wide Web.