The ORDBMS-based SFB 501 Experience Base -Exemplified by the SDL-Pattern Approach-

**Raimund L. Feldmann, Birgit Geppert,
Wolfgang Mahnke, Norbert Ritter, Frank Rößler**

SFB 501 Bericht 08/99

The ORDBMS-based SFB 501 Experience Base -Exemplified by the SDL-Pattern Approach-

Raimund L. Feldmann[*], Birgit Geppert[+],
Wolfgang Mahnke[#], Norbert Ritter[#], Frank Rößler[+]
{r.feldmann, b.geppert, f.roessler}@computer.org
{mahnke, ritter}@informatik.uni-kl.de

Sonderforschungbereich 501
Technical Report 08/1999

[+]*Computer Networks Group*
[#]*Databases and Information Systems Group*
[*]*Software Engineering Group*

Fachbereich Informatik
Universität Kaiserslautern
Postfach 3049
67653 Kaiserslautern
Germany

ABSTRACT

Comprehensive reuse and systematic evolution of reuse artifacts as proposed by the Quality Improvement Paradigm (QIP) do not only require tool support for mere storage and retrieval. Rather, an integrated management of (potentially reusable) experience data as well as project-related data is needed.

This paper presents an approach exploiting object-relational database technology to implement the QIP-driven reuse repository of the SFB 501. Requirements, concepts, and implementational aspects are discussed and illustrated through a running example, namely the reuse and continuous improvement of SDL patterns for developing distributed systems. Based on this discussion, we argue that object-relational database management systems (ORDBMS) are best suited to implement such a comprehensive reuse repository. It is demonstrated how this technology can be used to support all phases of a reuse process and the accompanying improvement cycle.

Although the discussions of this paper are strongly related to the requirements of the SFB 501 experience base, the basic realization concepts, and, thereby, the applicability of ORDBMS, can easily be extended to similar applications, i. e., reuse repositories in general.

Keywords

Quality Improvement Paradigm (QIP), reuse repositories, comprehensive reuse, Object-Relational DataBase Management Systems (ORDBMS), SDL-pattern approach

Table of Contents

1 Introduction . . . 1
2 Continuous Improvement of SDL Patterns . . . 3
2.1 The Basic SDL-Pattern Approach . . . 3
2.2 SDL-Pattern Based Improvement Cycle . . . 5
2.3 Requirements on Data Management . . . 7
3 An ORDBMS-based Approach . . . 11
3.1 Assessment of Database Technology . . . 11
3.2 The ORDBMS-based SFB 501 Experience Base . . . 12
4 Related Work . . . 23
4.1 Repository Structures . . . 23
4.2 Technical Realizations of Repositories . . . 23
5 Conclusions . . . 25

List of Figures

Fig. 1.1: Exploiting the SFB-EB within a QIP-based improvement cycle . . . 2
Fig. 2.2: Interaction of reuse process and improvement cycle . . . 5
Fig. 2.1: Measuring the reuse process . . . 4
Fig. 3.1: Structure of an experience element . . . 12
Fig. 3.2: UML schema of the SFB-EB (excerpt) . . . 13
Fig. 3.3: Architectural overview . . . 17
Fig. 3.4: Step function for similarity-based search . . . 19

List of Tables

Tab. 2.1: Requirements vs. QIP steps . . . 9
Tab. 3.1: Requirements vs. ORDBMS technology . . . 11

1 Introduction

Learning from experience gained in past projects is seen as a promising way to improve software quality in upcoming projects. As a result, (anti-)patterns, frameworks, and code fragments are being developed to capture the gained experience of software already developed. Recently, SDL patterns [16] have been introduced to increase reusability for developing distributed systems. They can be considered a representative example of such reuse artifacts. But experience is not only represented in the form of (directly) reusable software artifacts. To allow comprehensive reuse, as, for instance, proposed in [5], a large variety of different reusable elements exists. Process descriptions, for example, are used to precisely describe how to develop software in a certain domain, or lessons learned [6] can be used to store qualitative experience. Consequently, every kind of (software engineering) experience, independent of its type of documentation, is regarded as *experience element* [12] in this paper.

However, the benefits that can be achieved by reusing such experience elements strongly depend on their quality. Thus, they always have to represent the latest state of the art. Hence, checking and continuously improving their quality becomes a crucial issue. The Quality Improvement Paradigm (QIP) [3] as suggested by Basili et. al. fosters this problem by integrating systematic evolution and comprehensive reuse of experience elements into an improvement cycle. Within this context, each project can be regarded as an experiment, focusing not only on the developed software products, but also on learning about and improving the applied experience elements. A QIP cycle consists of six steps, with steps 1–4 dealing with the planning and executing of an experiment, step 5 with analyzing its results, and step 6 with packaging the gained experience for later reuse. To store experience elements and offer them, on demand, to the (re-)user, a reuse repository called the Experience Base (EB) [3] is introduced. Within the EB, experience is only stored after it is carefully analyzed (QIP step 5) and packaged (QIP step 6) for reuse. However, information concerning the experiments in which these experience were gained is not directly integrated into the EB. Therefore, we developed a conceptual extension of the EB [12, 17] that allows for storing complete experiment documentations, structured in accordance with QIP steps 1–4. Our resulting SFB 501 Experience Base structure consists of two logically disjunct sections called Organization-Wide Section (OWS) and Experiment-Specific Section (ESS), where the OWS is an instantiation of Basili's EB and the ESS holds the experiment documentations.

To evaluate the usefulness of the conceptual extension, we implemented a web-based prototype of the SFB 501 Experience Base (SFB-EB) and applied it to the SDL-pattern approach [17]. Based on the positive results we now concentrate on the technological advancements for such QIP-driven reuse repositories. With this paper we demonstrate that Object-Relational DataBase Management Systems (ORDBMSs), the most recent trend in the field of database technology [32], are best suited to meet the requirements for implementing such comprehensive reuse repositories. We will show that by using an ORDBMS, some shortcomings of current repository implementations can be avoided. It is, for instance, possible to provide data to tools as files stored within an experiment-specific

file-system area called Working Area (WA) and link these files to database entries in a way allowing the ORDBMS to control access and maintain consistency. Figure 1.1 illustrates our proposed integration of the ORDBMS-supported SFB-EB sections into the QIP-based improvement cycle.

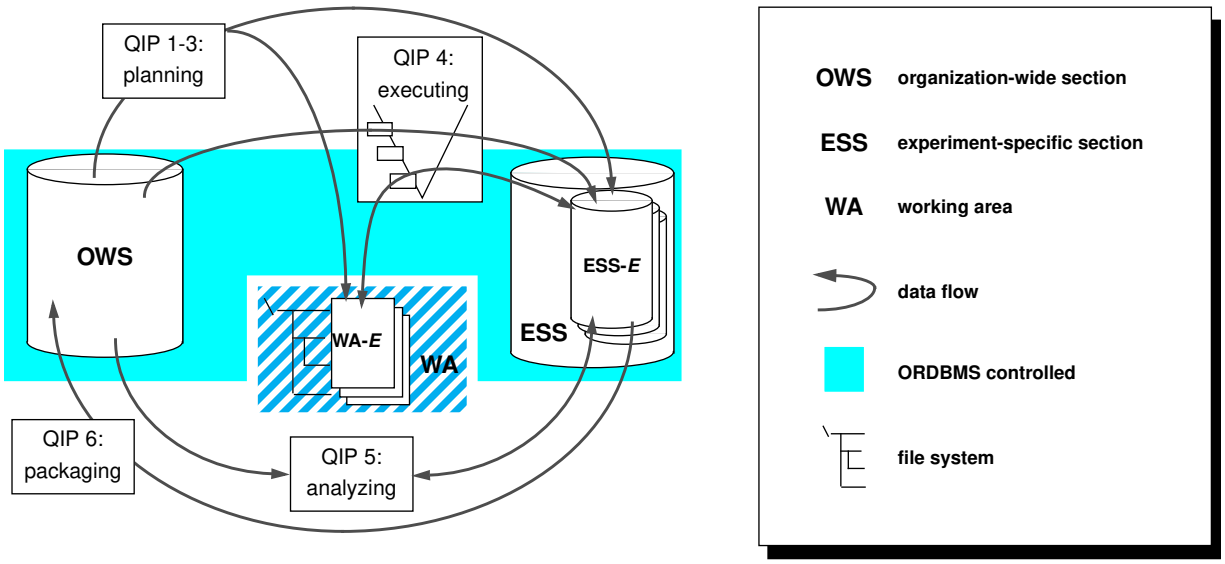

Fig. 1.1: Exploiting the SFB-EB within a QIP-based improvement cycle

This paper is organized as follows: In Section 2, the SDL-pattern approach is outlined. It serves as a representative, running example for discussing and illustrating requirements for an advanced reuse repository supporting a QIP-based improvement cycle. Section 3 justifies our opinion that ORDBMSs are the best choice of database technology for providing and realizing the SFB-EB. Related work is discussed in Section 4. Finally, we conclude with a summary in Section 5.

2 Continuous Improvement of SDL Patterns

The formal description technique SDL [22] is in widespread use in industry as a design language for reactive, distributed systems. Recently, SDL patterns have been introduced as a new concept for increasing reusability in SDL-based system design [16]. In the following, we sketch the basic SDL-pattern approach and discuss its integration with continuous quality improvement. Resulting requirements on data management are discussed subsequently.

2.1 The Basic SDL-Pattern Approach

The basic SDL-pattern approach extends pattern-based reuse for engineering distributed SDL systems. It comprises a product model of reuse artifacts and an incremental, use-case driven design process.

SDL Patterns. Scattered parts of a given SDL design may together provide a certain functionality. By analysis, abstraction, and documentation, such a design solution can be reused whenever the design problem arises again. Roughly speaking, an SDL pattern is defined as a reusable software artifact representing a generic solution for a recurring design problem with SDL as applied design language.

Design reuse as in the case of patterns is more flexible than code reuse, but the learning curve required before a pattern can be reused could be very high. It is therefore important to keep the specification of SDL patterns precise but intelligible. For this purpose, a standard description template is defined [14]. Physically, SDL patterns can be stored in different formats, e. g., as FrameMaker documents, in HTML, or in Postscript.

In order to support a stepwise and systematic evolution of SDL patterns, we defined different maturity levels [10] that allow developers to participate in their improvement, while reusing premature artifacts. This allows an incremental investment in reuse and contributes to quality, because good artifacts depend on practical and long-term experience. It can generally be stated that the effort in reusing an artifact decreases with higher maturity levels. In other words, the higher the previous investment in reuse, the higher the benefits.

Reuse Process. Generally, an SDL system development process encompasses the following development phases: object-oriented analysis, SDL-based design, formal validation, and, finally, code generation. In [15] we present an incremental, use-case driven design process tailored to SDL patterns (upper part of Figure 2.1). First, system requirements are decomposed into single functionalities (the use cases), which are further analyzed and then implemented one after the other in separate development steps. Analysis includes developing (and improving) an outline system architecture, finding and exploring collaborations that refine the use cases, and breaking them down into smaller patterns (pattern selection). Detailed SDL design is done incrementally by incorporating the identified collaborations step by step. Each development cycle applies predefined SDL patterns for implementing the current

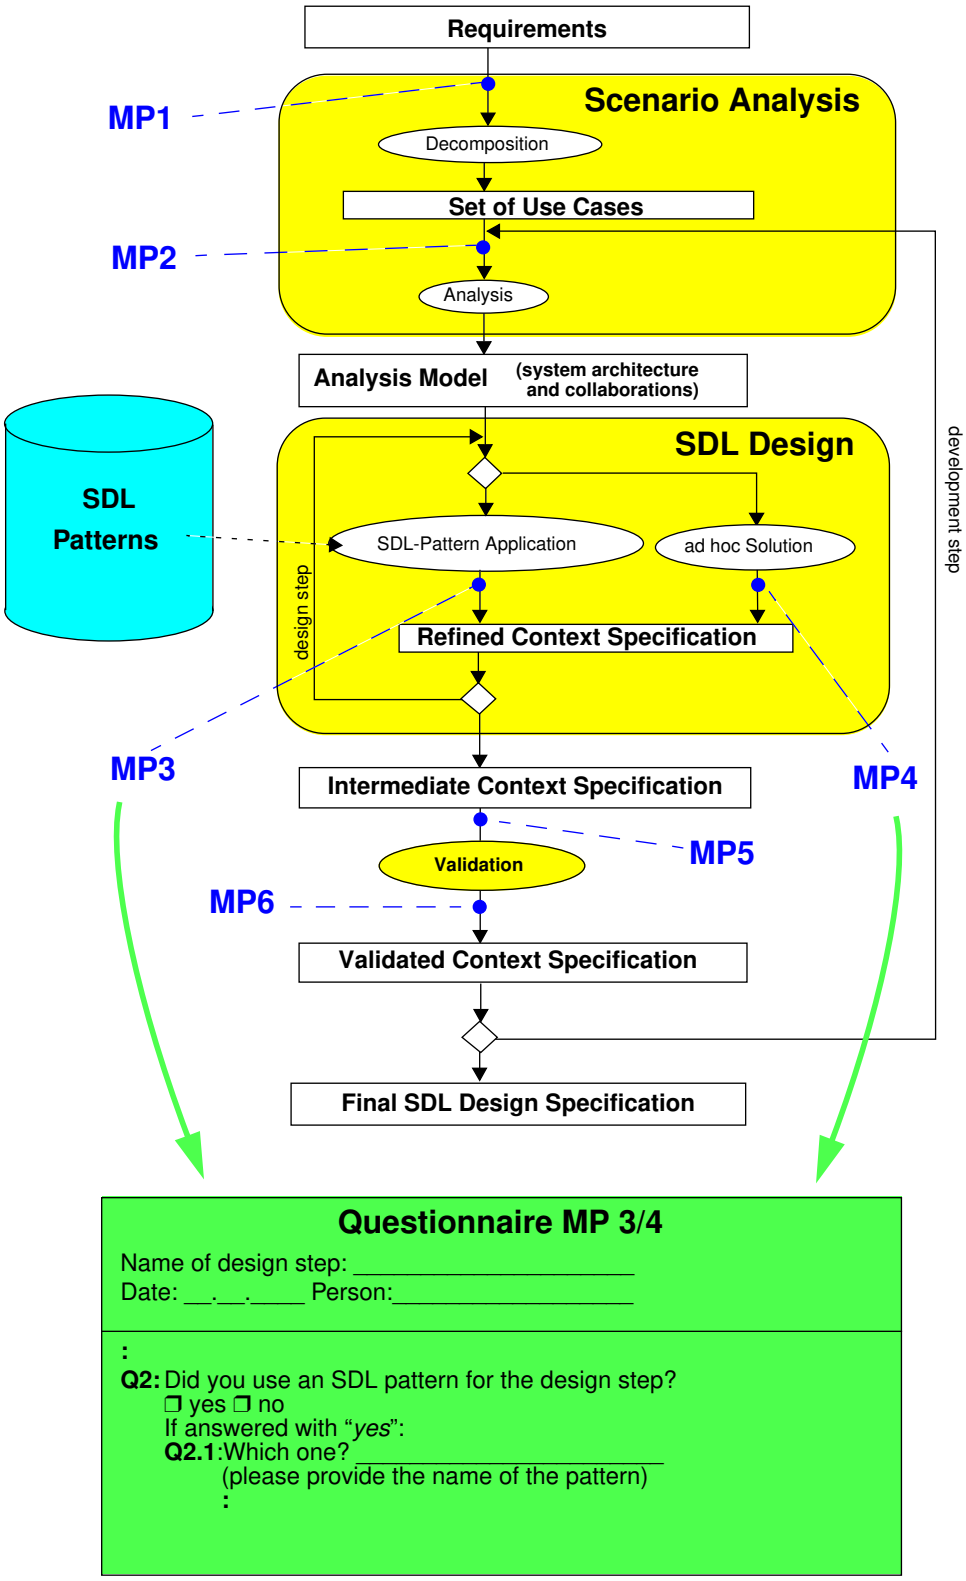

Fig. 2.1: Measuring the reuse process

collaboration or develops ad hoc solutions, if necessary. Note that application of an SDL pattern means adaptation and subsequent composition with the embedding context specification according to its application rules. The result of each development step is an executable SDL design specification, which is validated before entering the next step. For implementing the final SDL specification several SDL development environments provide automatic code generators.

2.2 SDL-Pattern Based Improvement Cycle

Continuous quality improvement as applied to the SDL-pattern approach supports an incremental evolution of reuse artifacts that is essentially driven by practical experience and triggered upon user demands. Figure 2.2 shows a

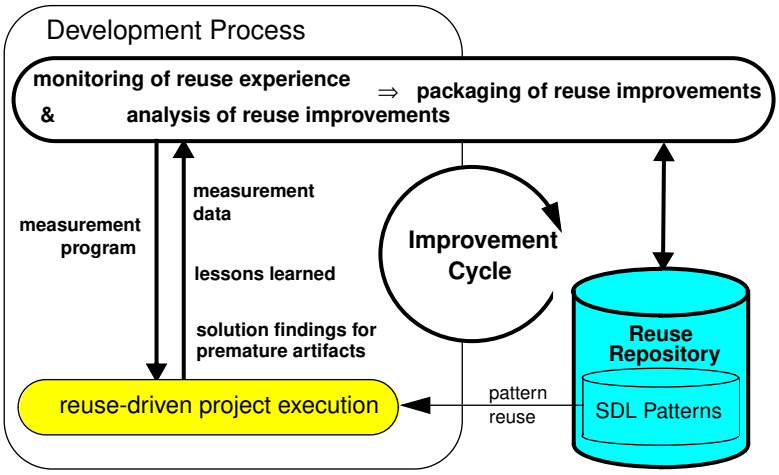

Fig. 2.2: Interaction of reuse process and improvement cycle

graphical representation of the improvement cycle (bold lines) and its interaction with the reuse process (thin lines). We illustrate this interaction by means of usage scenarios that follow the QIP steps.

Planning a new project (QIP steps 1–3). When setting up a new SDL-pattern based project, a new project database and working area is instantiated where all kinds of documents coming up during the project can be stored. The first entry is the characterization of the project, including a description of its goals and possible time restrictions. With this information at hand, we can browse other SDL-pattern projects stored in the repository. This guides the project manager in defining a concrete project plan and determining necessary resources. Furthermore, the specific goals of the project (including learning-specific goals) must be fixed quantitatively according to the GQM paradigm [4]. This results in one or more GQM plans, which are integrated by a measurement plan. The measurement plan defines measurement points when different metrics are to be collected. For data collection, questionnaires are developed that must be filled out by the development team during project execution. Figure 2.1 shows example measurement points MP1 – MP6 together with an excerpt of a questionnaire used in one of our case studies [11]. For some standard learning goals such as monitoring the quality of applied patterns, predefined measurement points and ques-

tionnaires are supplied in the repository. All documents prepared in this step are deposited in the project database of the repository for later access. This serves as documentation of the project history, which is decisive for analyzing and interpreting the measurement results in QIP step 6. Finally, it must be checked if needed tools (e. g., FrameMaker) and training material (e. g., an SDL tutorial) are ready to use. SDL tools that may be applied are, for instance, SDT (SDL design tool) [37] and the pattern-specific tool components SPEAR (SDL-pattern pool administration tool) and SPEEDI (SDL-pattern editor) [9].

Executing the project (QIP step 4). The project is executed according to the project plan, which is an instantiation of the incremental SDL-pattern based reuse process. Tools that can be applied for scenario analysis are, for instance, the UML [31] and MSC [23] editors of the SDT tool set. During SDL design, the SDL-pattern pool is searched for suitable patterns. This is supported by the SPEAR tool offering a search facility and by linking patterns with other related ones. Patterns are normally represented in HTML. However, SPEAR can also generate the internal pattern representation needed by the SDL-pattern editor SPEEDI. For SDT, pattern-based design is not supported. That means, when using the SDL editor of the SDT tool set –instead of SPEEDI–, the patterns must be inserted manually. However, SDT offers other advantages, such as tool-supported validation of an SDL design. During project execution, the questionnaires are filled out according to the predefined measurement program and, in addition, project-specific lessons learned are determined by the developers. The products produced, such as the SDL design specifications, are added to the project database of the repository.

Analyzing the project data (QIP step 5). The data that has been collected during project execution is now processed and analyzed. This comprises the quantitative measurement data as well as textually fixed lessons learned. In order to get the right interpretation, special events or exceptions that occurred during project execution must be considered. This is made possible by the project trace stored in the project database of the repository.

Due to the detailed process and product models of the SDL-pattern approach, project analysis according to the pattern-specific learning goals proceeds straightforward. As illustrated in [11], the outputs from the measurement program capture many areas of possible improvements, such as increasing the maturity level of an existing pattern, so that “ad hoc” analysis can often be avoided. Another example is the detection of a new pattern event that may result in a proposal for a new SDL pattern.

Before the analysis results can be deposited in the reuse repository, they must be compared and related to previous experience. Only if the suggested improvements turn out to be general enough, can they be stored in the repository as common experience. Otherwise, they keep their preliminary status, but may serve as recommendation for planning similar future projects.

Packaging of analysis results (QIP step 6). To allow effective storage, search, and retrieval of experience data in the repository, a unique representation of similar experience is necessary. Therefore, the analyzed project data is stored according to predefined templates and/or style guides. General lessons learned, for example, are best described according to the template defined in [6]. SDL patterns are represented according to their standard description template. Support for modifying the description of an existing SDL pattern or extending the pattern pool is another functionality of the tool SPEAR. Modifying an existing pattern may result in upgrading its maturity level. When adding a new pattern, it has to be decided with which maturity level to start with. Additionally, new patterns must be linked with related patterns of the pool.

2.3 Requirements on Data Management

The discussion in Section 2.2 implicitly addressed various requirements pertaining to data modeling as well as data manipulation. For the sake of clarity, we discuss only the most important requirements.

2.3.1. Data Modeling Needs

R-I: Integrated management of experience data and project-related data. For the purpose of maintaining consistency and avoiding redundancy, it is reasonable to manage experience elements and experiment-specific data (project data) in an integrated manner (within the same database). For example, after having identified an SDL pattern as a reuse candidate, it is generally helpful to have hints to projects in which this particular experience element has been applied before. Furthermore, integrated management is advantageous for the QIP steps analyzing and packaging, because it makes it easier to realize new and unique experience and, thereby, helps avoiding redundancy as well as inserting gained experience at the right place.

R-II: Semantic classification of experience elements. The discussion in Section 2.2 clearly identified that the database area storing experiment elements (in Section 1 identified as the organization-wide section, OWS) must be able to manage semantically different data structures. For example, structures for storing SDL patterns, measurement program elements, lessons learned, and so forth. Similarly, the experiment-specific section (ESS) must also be able to manage such data structures in order to handle project-related data.

R-III: Relations between experience elements. The experience elements managed in the above-mentioned classification must be connected by various semantic relations carrying crucial information for the analyzing and packaging steps of the QIP. For raising an SDL pattern's maturity level, for example, the pattern engineer needs to retrieve all information concerning occurrences of this pattern in former projects. Therefore, a `uses/used_in` relation between the SDL pattern and former projects is needed.

R-IV: Multiple representations of experience elements. We learned that different tools need SDL artifacts in different formats, e. g., HTML, text, proprietary tool formats, etc.. This concerns both OWS and ESS. Since this is a very typical scenario in design applications, there must be data structures allowing to store design artifacts in different/alternative representations.

R-V: Management units for project-related data. The SFB-EB has to support several projects simultaneously. Since project work usually results in different release levels of the created products (private data, data released for the project team, data released for public access), private as well as project-related workspaces must be managed. This requirement has already been addressed in Figure 1.1 showing the project-specific areas ESS-E and WA-E.

R-VI: Maintaining a project history. Within the context of a single project, different objects have to be associated with a (project) history, which is a crucial prerequisite for the step of analyzing design data. For example, the possibility of tracking the project history often helps to validate and generalized the lessons learned of the project (QIP step 5 in Section 2.2).

2.3.2. Data Manipulation Needs

R-VII: Role and authorization concept. A prerequisite for data manipulations is an adequate role concept and corresponding authorization mechanisms. The previous discussion already indicated that different roles for managing the OWS (e. g., the SDL pattern engineer), leading a project, working within a project team as a designer or as a quality manager are needed. The different roles are to be associated with different areas of visibility w. r. t. the data stored. Preliminary data of a project, for example, should not be visible to members of a different project's team.

R-VIII: User interface. The different tasks associated with different roles require corresponding functions to be provided at the SFB-EB interface. Generally, functions are needed for browsing the SFB-EB data, traversing object structures, searching for artifacts, and manipulating the data stored within the SFB-EB. For example, the person responsible for managing experience data has to be supported by functions allowing to access the complete OWS. This includes retrieval and manipulation functions. The latter must not be provided to unauthorized personnel.

R-IX: Distributed access. Regarding the often occurring situation that teams are geographically dispersed, and the cooperative working style needed in software development projects, it must be possible to contact the SFB-EB from different nodes of a distributed environment.

R-X: Tool integration. Besides a (specially designed) interface providing different roles with corresponding access functions, an infrastructure allowing tools to access data stored within the SFB-EB must be provided. For

example, the SDL-pattern based reuse cycle employs applications of the tools SDT, SPEEDI, and SPEAR. Where own implementations like SPEEDI and SPEAR can use the API (Application Programming Interface) of the ORDBMS, commercial tools like SDT require flexible coupling mechanisms for their integration.

Finally, Table 2.1 lists all requirements and there relevance in accordance to the different QIP steps.

Requirement	QIP step 1-3	QIP step 4	QIP step 5	QIP step 6
R-I	X		X	X
R-II	X	X		
R-III	X		X	X
R-IV	X	X	X	X
R-V		X		
R-VI			X	X
R-VII	X	X	X	X
R-VIII	X	X	X	X
R-IX	X	X		
R-X	X	X		

Tab. 2.1: Requirements vs. QIP steps

3 An ORDBMS-based Approach

In this section, we discuss how object-relational database technology can serve as an appropriate foundation for the realization of the SFB-EB fulfilling the needs mentioned above. But before detailing our ORDBMS-based approach, we justify our decision for object-relational technology.

3.1 Assessment of Database Technology

Current activities in developing object-relational database systems [32] aim at integrating object-oriented concepts into the relational model. Although the evolutionary process of designing the object-relational data model and identifying the components of a corresponding data management system is by far not completed, the *extensibility* property of ORDBMS is extremely beneficial for realizing the SFB-EB. Extensibility allows for:

- adding user-defined data types consisting of specially designed data structures and corresponding operations to the database schema; the possibility of exploiting large (binary) objects allows for capturing arbitrary formats within user-defined data types;
- linking externally stored data with database entries and controlling access to these (externally stored) data by database mechanisms;
- exploiting pre-defined extensions for accessing the database via the web, for managing data of various formats (text, HTML, video, audio, image, etc.), and for dynamically transforming externally stored data into relations, which may be incorporated into SQL processing.

Since neither object-oriented [24] nor well-known relational DBMS provide comparable capabilities, we decided on ORDBMS[30]. Table 3.1 summarizes which requirements can specifically be covered using ORDBMS technology and in which of the following sections they are further discussed. A '-' in the second column of Table 3.1 denotes that the corresponding requirement can be supported by exploiting ORDBMS features, but it could also be meet by using other kinds of database systems. For example, the traversal operation demanded by **R-VIII** can best be supported by the concept of references supported by object-oriented database systems [24]. Nevertheless, considering the entirety of requirements, ORDBMS are best suited.

Requirement	covered by ORDBMS technology	discussed in
R-I	yes (advanced modeling concepts)	Section 3.2.1.
R-II	yes (advanced modeling concepts)	Section 3.2.1.
R-III	yes (advanced modeling concepts)	Section 3.2.1.
R-IV	yes (extensions)	Section 3.2.1.
R-V	yes (extensions)	Section 3.2.1.
R-VI	-	Section 3.2.1.

Tab. 3.1: Requirements vs. ORDBMS technology

R-VII	-	Section 3.2.2.
R-VIII	-	Section 3.2.3.
R-IX	yes (extensions)	Section 3.2.3.
R-X	yes (extensions)	Section 3.2.4.

Tab. 3.1: Requirements vs. ORDBMS technology

3.2 The ORDBMS-based SFB 501 Experience Base

In the remainder of this section, we introduce our approach of the ORDBMS-based SFB-EB. First, we describe the data structures resulting from the data modeling needs (**R-I** to **R-VI** in Section 2.3.1.). Then, data manipulation needs (**R-VII** to **R-X** in Section 2.3.2.) are elaborated one after the other in further subsections.

3.2.1. Data Structures

Many different kinds of design artifacts have to be stored in the SFB-EB, like source code, SDL patterns, process models, reports and so on (**R-IV**). Those *representations* of *experience elements* (EEs) are given in many different data formats. To easily handle the different formats (and to be open for new formats), EE representations are saved as plain data type in the SFB-EB without respect to special data formats. ORDBMS offer a special kind of data type for this purpose, called BLOB (binary, large object). We have learned from Section 2 that our EE structure must be capable of capturing several alternative representations. An SDL pattern, for example, may occur as a FrameMaker, Postscript, or HTML document. Thus, several alternative representations of the same EE can occur. Furthermore, each of these alternative representations can itself be composed of several parts. A framed HTML document, for instance, consist of many files. As a consequence, an EE in the SFB-EB has an internal structure as illustrated in Figure 3.1.

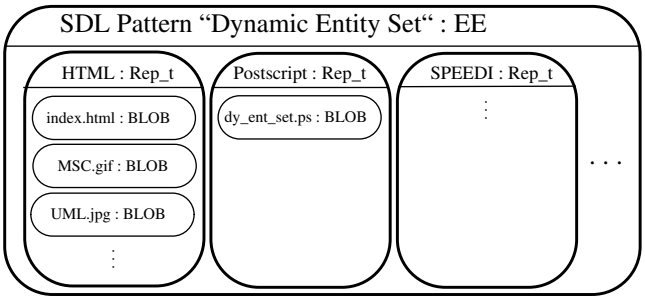

Fig. 3.1: Structure of an experience element

For retrieval purposes, each EE is associated with a so-called *characterization vector* (CV) containing further describing data, e. g., relevant for (similarity-based) search of potentially reusable design artifacts. For performance reasons, EEs (containing large objects) and corresponding CVs (comparably small amount of data) are stored separately. In Figure 3.2 the schema of the SFB-EB is shown in an (intentionally incomplete) UML notation.

It illustrates the separation of EEs and CVs and shows that each EE is associated with exactly one CV via a (1:1)-relation.

CVs can be viewed as the representatives of EEs in our database schema. Hence, they also implement the required partitioning of EEs (**R-I**, **R-II**). Therefore, the CV (attribute) structure depends on the section (OWS, ESS), the corresponding EE belongs to. More precisely, the sections are divided into logical areas, and the areas determine the CV attributes. Figure 3.2 illustrates the inheritance hierarchy. Subclasses of the abstract class `Organization-Wide Section` have special attributes of the different kinds of EEs. For example, `Qualitative Experience` contains the inspection techniques used, whereas `Process Modeling` contains valid process models. Note that some CV attribute values can be determined automatically (e. g., author, creation date, etc.) while others have to be specified manually.

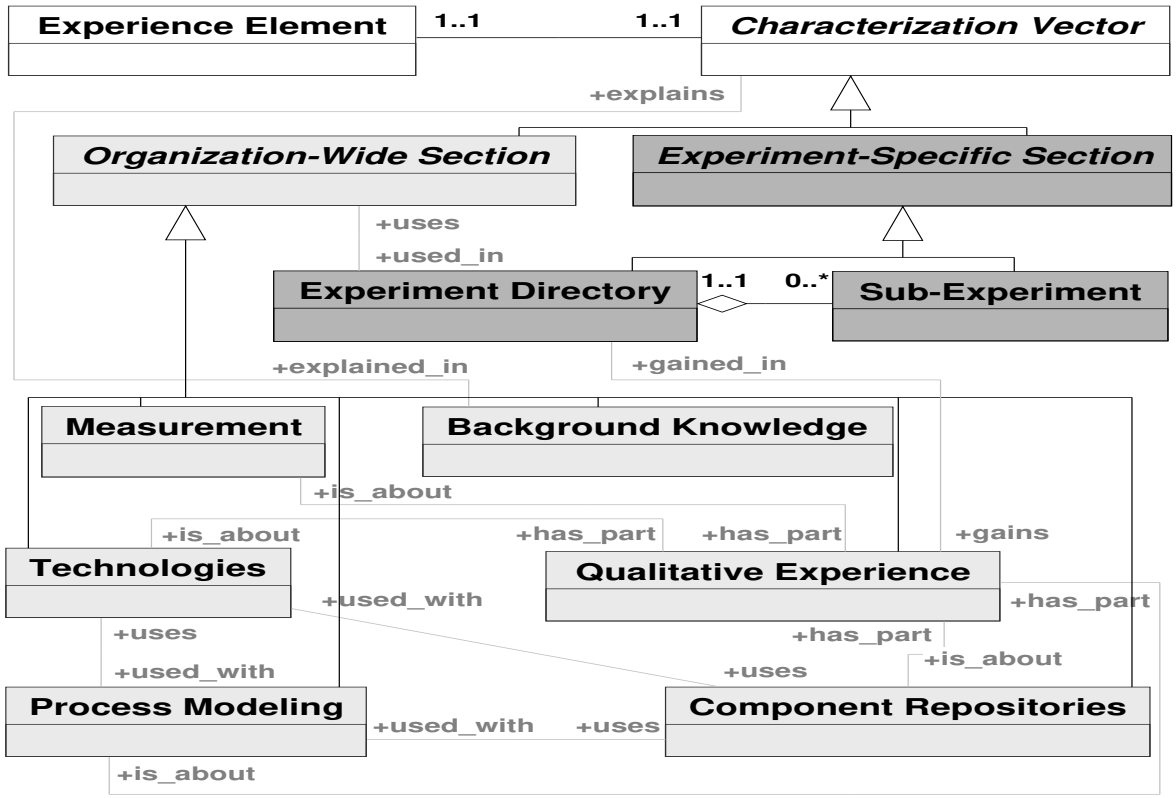

Fig. 3.2: UML schema of the SFB-EB (excerpt)

The (abstract) class `Experiment-Specific Section` has the subclasses `Experiment Directory` and `Sub-Experiment`. Remember that the ESS stores the data of running and completed experiments (**R-I**). Each experiment consists of one `Experiment Directory` object and many `Sub-Experiment` objects (**R-V**). The general data of an experiment like project-manager notes, project type etc. are stored in the `Experiment Directory` while the data of modules, measurements, reports etc. are stored in the `Sub-Experiment` objects. Figure 3.2 shows the corresponding aggregation

between these two classes. Besides other attributes, these classes contain attributes allowing to capture the project history, e. g., version identificators and timestamps (**R-VI**). Figure 3.2 also shows many associations between the different `Characterization Vector` subclasses denoting semantic dependencies (**R-III**).

Obviously, the different CV classes are arranged within an inheritance hierarchy. While RDBMS do not support table hierarchies, ORDBMS do. Therefore, the UML classes depicted in Figure 3.2 can easily be mapped to a table hierarchy in the ORDBMS schema. Some CV attributes can be mapped to the common data types of an RDBMS (like String or Integer), but we also use complex data types (like multisets) and user-defined types (UDT) as supported by ORDBMS. In order to provide similarity-based search on CVs, we also use user-defined functions (UDF) allowing to integrate application functionality into the DB Server.

3.2.2. Users, Roles, and Authorization

Usually a large number of persons are involved in different processes of the QIP. Consequently a role management and authorization concept is installed to grant access to the supporting SFB-EB.

Role Management. Since our approach aims at supporting the overall improvement cycle, all persons involved in this process have to be provided with adequate functionality. Whether or not an interface function (see below) is available to a user, is determined by the role s/he uses to access the system. To fulfill **R-VII**, the following roles are defined in our repository:[1]

- *EB Manager* (managing OWS data);
- *EB Assistant* (supporting the EB Manager);
- *Project Manager* (managing experiment-specific parts of the ESS);
- *Quality Manager* (assuring quality of the development process and corresponding products);
- *Developer* (attending an experiment).

Furthermore, the role contributes to determine which data is visible, may be modified or even deleted within a user session. Regarding running projects, however, the role is not sufficient to determine corresponding rights (and duties). Thus, a Project Manager's possibilities also depend on the project's name. In order to control the data access of a (project) team member, besides the role, the system needs to know the project's and the Developer's name.

1 For the sake of clarity, we do not consider more roles than those listed here. Taking additional roles into account would not require additional concepts. In either case, besides the mentioned roles, a *DB Administrator* is needed.

Authorization. The user and role management component manages users (names) and allows to associate users with roles and projects. At the beginning of a session, the user is required to deliver name, password, and, possibly, the name of the experiment intended to join. The system checks security. This also contributes to fulfilling **R-VII**.

3.2.3. A User Interface supporting the QIP steps

In this section we outline the functions provided at the SFB-EB interface in order to support the QIP steps identified in Section 2.2.

Planning a new project (QIP steps 1-3). A new project is set up by *creating a project-specific data area* ESS-E, initially taking up project describing data, e. g., description of the project goals. A corresponding function is provided at the SFB-EB interface.

Now, facing the project goals, the Project Manager starts investigating the OWS in order to identify reusable EEs. This step is supported by *browsing and navigation* functions, and, additionally, by a specially designed *search* function, which, as we will see in the following, is not only beneficial for planning but also during the steps executing, analyzing, and packaging. In the planning step, descriptions of similar, past projects and corresponding training material are usually looked for (Section 2.2).

When running earlier (not DBMS-based) SFB-EB prototypes [17], it already turned out that this retrieval operation should be *similarity-based*. Thus, the system provides datatype-specific measures and similarity functions, which may be adapted to the semantics of attributes (associated with CVs). A user, e. g., a Project Manager, may specify comparison values for an arbitrary selection of attributes (associated with CVs) as parameters for his search. Additionally, s/he may specify a weight for each of these attributes expressing its importance for his interests. The system answers such a query by, first, identifying those CV types containing the attributes addressed in the user's query. Second, the comparison instance specified in the query is compared to corresponding database entries by computing a weighted sum of attribute comparison values, respectively. Those experience elements that are considered to be similar to the comparison instance specified in the query are provided to the user in a structured form. An example will be given later on in this section.

After having identified reusable EE, (e. g., the SDL reuse process description) the Project Manager can now use a further interface function designed to support the planning step. It allows for *transferring* (copying) EEs from the OWS to ESS-E. This function is designed to automatically add a `uses/used_in` relation between the original EE (in the OWS and its copy (in the ESS-E).

Executing the project (QIP step 4). As a default, ESS-E data are not visible to anyone but the project team consisting of the Project Manager, a Quality Manager and a number of Developers. The data initially transferred from the OWS to the ESS-E (results of the planning step) may be read by all team members. Developers have private working areas (WA-E) in which preliminary results are to be managed. Developers may grant team members or the overall team (read and/or write) access to private data. Write access does not comprise deletion; thus, the Developer who initially created the data remains the owner of these data and, consequently, may withdraw the rights s/he previously granted.

The ESS-E contains data entries which may exclusively be accessed by the Project Manager or the Quality Manager, respectively. In cooperation with Project Manager and affected Developer, the Quality Manager may release ESS-E data to public. This means, data which has reached a certain stability (w. r. t. product quality) and, therefore, may be beneficial for others, can be made visible. As soon as the EB Manager acknowledges this step, the corresponding data acquire the (default) status of OWS data (regarding visibility, see below). If necessary, the Quality Manager can withdraw this release, as long as the corresponding experiment is active.

This short description of the executing step clarifies that functions for *browsing, navigating, searching, manipulating, and releasing* ESS-E data must be provided at the SFB-EB. The challenge in developing these functions is to appropriately manage and ensure the respective persons' rights during function evaluations.

Naturally, during the executing step OWS data can also be further investigated in order to react to the current project state. Actually, all functions described in the previous subsection can also be applied in the executing step. It might, for example, be reasonable to query the OWS for certain SDL patterns during the executing step, because corresponding needs may evolve during project work.

Analyzing the project data (QIP step 5). After the experiment has been completed, the Quality Manager uses *browsing, navigation, and search* functionality to consider ESS-E data and can perform *manipulations* in order to fix or increase the reuse potential. For that purpose, the overall set of functions provided to support the executing step are also at his disposal in this step. Additionally, functions for *browsing, navigating, and searching* OWS data are available, helping to identify new, non-redundant experience data (Section 2.2). An intuitive example for actions to be performed during this step is determining or increasing the maturity level of an SDL pattern, as described in Section 2.2. The Quality Manager uses another interface function to *mark* those parts of the experiment-specific data which should be provided to the public as gained experience.

Packaging of analysis results (QIP step 6). In this final QIP step, the EB Manager transfers the marked data into appropriate OWS structures in order to commit the final release. Furthermore, s/he (semantically) integrates these

new experience data by establishing relations, e. g., `gains/gained_in` relations, among new entries as well as among new and previously integrated data (Section 2.2). To do so, s/he applies a couple of interface functions that are at his disposal not only to specifically support the packaging step, but also to support him in generally managing OWS data as described in the following.

As a default, all users may read OWS data, but the EB Manager has the possibility of restricting read access by explicitly detracting certain EEs (to be determined by a predicate on the corresponding CV type) from the visibility of certain users or roles. Furthermore, also as a default, only the EB Manager is allowed to insert, modify, or delete OWS data; but, again, there are possibilities for delegating work, e. g., packaging work. Thus, the EB Manager may explicitly grant modification (including deletion) rights on certain EEs as well as the right to insert data into the OWS to EB Assistants (e. g., the SDL pattern engineer). Naturally, s/he may withdraw these rights as soon as reasonable from his point of view.

Obviously, again, *retrieval and manipulation* functions are needed that observe user rights. The function for *inserting* or *manipulating* OWS data, for example, must not be provided to anyone but the EB Manager or, if specifically authorized, to EB Assistants.

Realization aspects and example. Regretfully, due to space restrictions, we could only outline the functionality provided at the interface of our system. Nevertheless, it should have become clear that the functions addressed in this section establish the SFB-EB interface fulfilling **R-VIII**. The presentation will be amended in the following subsection by discussing the coupling of tools (used by Developers to fulfill their tasks) with the ORDBMS. The cooperation control facilities realized in our approach are similar to those introduced in [29]. In order to further demonstrate the ORDBMS-based approach, we now give an example illustrating several of the aspects mentioned so far and emphasizing the benefits of ORDBMS for that purpose.

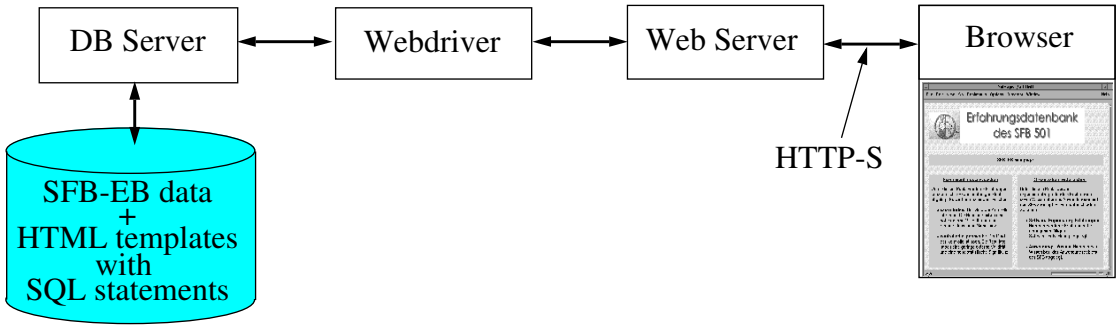

Fig. 3.3: Architectural overview

In order to fulfill **R-IX,** we decided to exploit the web infrastructure provided by ORDBMS to realize the interface functions. As realization platform, we exploited the ORDBMS Informix IDS/UDO [21] including Informix WebBlade as web infrastructure. The architecture of our system is shown in Figure 3.3.

The database stores EEs, CVs, pre-defined HTML pages as well as extensions (special tables, user-defined functions (UDFs)) needed to dynamically generate HTML pages and answer user requests. In order to answer special user requests, HTML templates including SQL statements can be stored. A user request which has been specified at the browser and passed to the DB Server may address such an HTML template. The UDF *webexplode* (offered by the WebBlade) evaluates the SQL statements contained in the template, incorporates the results into the template, and sends it back to the browser. As we will see in the example below, there is also the possibility of adapting or even generating the SQL statement(s) to be evaluated in order to create the resulting HTML page dynamically.

In order to demonstrate the approach, we have chosen the similarity-based search as an example. After the user has started a user session by authenticating himself, s/he may start working at the SFB-EB interface by invoking the functions outlined in the previous section. For security purposes, we use the net protocol HTTP-S. Since all HTTP protocols are context-free, the current session context, represented by an authentication number, is communicated each time a user request or a resulting HTML document is transferred. An authentication number is generated each time a message is sent for safety reasons.

The authorization component of the SFB-EB consists of a couple of database relations and corresponding UDFs checking authentication and generating authentication numbers. Whereas this is sufficient to prohibit unauthorized access to the SFB-EB, more fine-granular access control mechanisms are needed to control database queries, like similarity-based search. It is necessary to control access to tuples (e. g., a single experience element) as well as attributes (e. g., notes of the EB Manager). For these purposes, again, the DBMS has been extended by some management tables and corresponding UDFs.

In order to generate a mask for the similarity-based search, a UDF is used expecting the current role as input parameter, and delivering a list of all accessible CV attributes as output. After the user has specified the comparison instance, the request is sent to the server and there transformed into one or more SQL queries, making heavy use of predefined UDFs basically observing rights and computing similarities. Such a generated SQL statement may look as follows.

```
SELECT name,(SIM(validity,5,1,1,1)*10+
      SIM(impl_technique,"SDL/MSC",...)*3+...)
      /TW AS Total_Sim
FROM     Component_Repositories
WHERE    Total_Sim > 0.5
AND      RightsOK(:RID, ID)
```

The UDF `RightsOK()` ensures that only tuples are taken into account that are enabled for the user/role. The similarity of an EE is computed from the similarities of the regarded CV attributes (validity, impl_technique, etc.) and depends on the weight of these attributes. The weight of each attribute can be specified in the mask for the similarity-based search. In order to facilitate the use of the SFB-EB, default values are given, but each user may store his/her own default values. Since similarity values are between 0 and 1.0, the sum of attribute similarities has to be divided by the total weight of all attributes (TW). In the query above, only experience elements with a similarity of at least 0.5 are taken into account.

The UDF `SIM()` represents the similarity function associated with a given attribute. For each data type a special similarity function is needed, so the `SIM()` function is overloaded. Usually, it is a step function, as shown in Figure 3.4. The signature of such a UDF is `SIM(attribute, search value, delta similarity, delta value, direction)`. The direction can restrict the validity of the step function (Figure 3.4 illustrates a function that only takes into account values higher than the comparison value; to other values, the similarity zero is assigned). The similarity function is equipped with the parameters direction, delta similarity, and delta value, in order to allow users to bring in individual notions. If the user does not want to specify values for these parameters and did not specify individual default values previously, the system uses default values.

Obviously, the generated SQL query (and the resulting similarity values) to some extent depends on the current user, resp. the user-specific values. The function as illustrated in Figure 3.4 only works for scalar data types, but similar principles hold for other types of similarity functions. Thus, users may influence the computation of similarities.

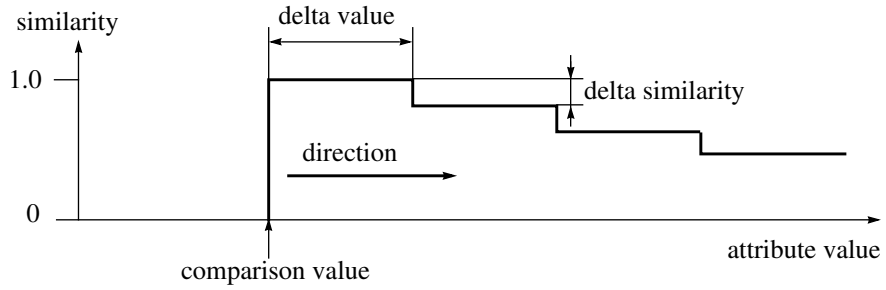

Fig. 3.4: Step function for similarity-based search

3.2.4. Tool Integration

In order to support the overall improvement cycle, the specially designed interface functions outlined in the previous section are not sufficient. Rather, we also have to provide an infrastructure for coupling tools with the SFB-EB. Especially Developers who are in charge of certain design tasks depend on the possibilities of accessing the

SFB-EB via tools. In our approach we aim at supporting four different tool integration modes as outlined in the following.

White-Box Integration. Implementing tools from scratch allows to exploit the ORDBMS API for data access purposes. Thus, transaction control and data manipulation statements can be incorporated into the tool program. A tool that has been integrated in this way can directly work on data stored within the ESS-E. In the context of our project (SFB 501), white-box integration is used for tools that are implemented by project partners.

Note, the primary characteristic of the white-box integration is that the tool directly calls DBMS API functions. Possibilities of incorporating externally (in the file system) stored data into database processing may additionally be applied. For example, files that might have been created by auxiliary tools may be incorporated by using ORDBMS mechanisms as IBM's Table Functions [7] or Informix' Virtual Table Interface [33]. These possibilities are completely orthogonal to the chosen tool integration mode and may also be applied with the other modes below.

White-box integration is the easiest way of tool integration, but often not feasible, if commercial tools are to be integrated.

Integration via communication infrastructures. Some commercially available tools are designed to be connected to data sources via standardized communication infrastructures, as ODBC/JDBC [13] or CORBA [28]. Other infrastructures [34] also belonging to this category are DCOM/OLE/ActiveX (Microsoft) and Java/JavaBeans (Sun). We are currently examining these possibilities; realization of a CORBA integration, in particular, is planned.

Wrapper Integration. One way to integrate a commercial tool with the SFB-EB is to put a so-called mediator [39] or wrapper in between tool and ORDBMS. This way, for example, file access operations of a tool developed to work on a file-system API may be 'wrapped' onto database operations. Obviously, this approach requires the file system to allow exchanging or overloading some of its functions, e. g., *open file* or *read file*.

A representative of this integration mode is the DataLinks concept [26, 27]. DataLinks technology provides (OR)DBMS management of file data which is stored external to the (OR)DBMS, while still providing referential integrity, access control, and coordinated backup and recovery as if the files were managed by the (OR)DBMS. This idea is realized by extending the file system by functionality, causing it to collaborate with the (OR)DBMS in managing data access. In the original DataLinks approach [26, 27], file access operations issued by an application are 'wrapped' as follows. First, a URL (Uniform Resource Locator, as in the WWW) is requested via the

(OR)DBMS API, so that the DBMS may apply its access control mechanisms. If the DBMS delivers the URL, then the application may access the content of the corresponding file via the file system API.

One possibility of exploiting similar concepts in our approach for tool integration purposes is as follows. Instead of BLOBs, EEs are represented by URLs. These URLs 'point at' files stored within the experiment's working area WA, which is located in the file system (see Figure 1.1). Thus, URLs may be considered to be a further kind of EE representation (Section 3.2.1.). File system access of tools is wrapped as outlined above. Newly created files have to be registered manually by the responsible user at the SFB-EB interface.

We are currently following this approach in connecting the tools mentioned in Section 2.2 (SDT, SPEEDI, SPEAR) to the SFB-EB.

Manual Control. Since not all commercially available ORDBMS offer mechanisms comparable to the DataLinks concept, and since some operating system platforms do not allow exchanging/overloading functionality, this last possibility for tool integration is meant to be an alternative to the previous one.

This possibility again is based on storing all the data within database BLOBs. Thus, files created by tools are 'sucked' into the database. Users are instructed to manually control manipulations as follows. After having identified the data entry to be manipulated via a retrieval function, another interface function allows to copy the data into a file which is expected by the tool to be applied. Now the tool may be applied to the file. After the tool has finished its work, another interface function may be used to propagate results of the tool application from the file system into the database.

This approach, on one hand, gives rather weak guarantees (in comparison to the previous integration mode) and requires more user discipline, but, on the other hand, it is easier to implement.

Further Remarks. The different integration modes mentioned in this section imply different possibilities of transaction control. For example, the white-box integration suggests exploiting the transaction concept supported by the ORDBMS. This means that a tool application corresponds with a database transaction. Manual control basically means that only checkout/checkin are mapped to short database transactions and that the tool's application is not protected at all, unless the user is made responsible for taking care. Integration via communication infrastructures and wrapper integration allows several, possibly tailored modes of transaction (especially concurrency) control. Currently, we follow the approach that only white-box integrated tools do completely isolate database resources. Other integration modes, in the same way as user sessions exploiting the functions mentioned in Section 3.2.2., are mapped to several short database transactions, respectively. Although interface functions are and all tool integra-

tion modules, e. g., wrapper, can especially be designed to preserve database consistency, we intend to realize an application server on top of the ORDBMS, allowing to handle versions of EEs appropriately and to control concurrent access to these versions as well.

Supporting the various integration modes discussed in this sections leads to a very flexible infrastructure. Thus, **R-X** can also be considered to be fulfilled by our approach.

4 Related Work

The description of related work is twofold: First, we take a look at the structure of existing reuse repositories in Section 4.1. Then, we will compare the technical realization of other repositories with our ORDBMS-based implementation. (Section 4.2). Due to space restrictions, a comparison with AI systems (like Case-based reasoning systems [1]), which could be used to support the intelligent search and retrieval of EEs, is beyond the scope of this paper.

4.1 Repository Structures

In [19,20] Henninger discusses a ***repository for reusable software components***. The paper focuses on the indexing structure of such a repository. A method is proposed supporting index implementation "*with a minimal up-front structuring effort*". Therefore, a rudimental set of reusable components is stored in the repository with a simple, basic index structure to get the reuse activities started. Then, while the components are being reused, the index structure is incrementally improved. The idea of starting the reuse process in a state of incompleteness to gain practical results as the basis for incremental improvement is similar to our approach. The information stored in our repository becomes the more detailed, the more EEs are being reused, caused by the growing number of relations (e. g., the `uses/used_in` relation) between the EEs. However, no predefined relation structure between the components is offered, which is an essential part of our repository structure. On the other hand, the SFB-EB is currently missing the ability of direct measures concerning the usage of the stored EEs in the SFB-EB that is an integral part of Henninger's repository. These measures would help us, for instance, in defining the maturity levels of the SDL patterns, and therefore, we plan to integrate them in the SFB-EB in the near future.

The ***ASSET Reuse Library WSRD*** [38] is an example of a web-based implementation of an object repository. It is a domain-oriented reuse repository that contains more than 1,000 EEs, dealing with topics such as software reuse practice or the Y2K problem. The repository is organized according to certain domains and collections that offer a mixture of different EEs, like lessons learned, process models, or code fragments. Therefore, a certain domain (or collection) that stores EEs from similar, but yet different projects in one entry can be compared to an experiment documentation of our experiment-specific section. However, complete project documentations are not an integral part of WSRD. Cross-references that interrelate the entries in the collections are offered, but they are much more general and unstructured than the relations defined in the SFB-EB.

4.2 Technical Realizations of Repositories

In the Arcadia project [36] several ***object management system***s [18, 35, 40] have been developed to support their process-centered software engineering environment. ***Triton*** [18] is one of those object managers. The whole system is based on the Exodus [8] database system toolkit to avoid the cost of new construction from scratch, just like

we are using the standard functionality of a commercially available ORDBMS and extend it, in accordance with our needs. Heimbigner describes Triton in [18] as follows: "*It is a serverized repository providing persistent storage for typed objects, plus functions for manipulating those objects*". Whereas we, on demand, link the stored EEs from the ORDBMS to the file system, so that different (commercial) tools can use them, Triton uses Remote Procedure Calls (RPC) for communication between client programs (tools) and the repository. Therefore, the server offers a procedural interface to its clients, acting as a kind of library of stored procedures. They are accessible from programs written in a variety of programming languages (such as Ada, C++, or Lisp). But this is a limitation that we can not accept, since most of the (commercial) tools are not offering the needed functionality to use RPC's and can not be modified because their source code is mostly not available. Mediators as described in [39] can be used to bridge the gap between a repository like Triton and such tools. But this would call for a wide variety of mediators. To avoid this construction effort, we consider EEs as BLOBs to keep the original storage formats of the different tools. Triton, on the other hand, uses a homogeneous storage schema to provide efficient representation for the wide variety of software artifacts. Consequently, all data that is shared by two or more tools have to be converted to the Triton schema, even if a direct data exchange via a format like RTF is available. Again, this is acceptable in an environment, where most of the tools have been developed from scratch, but it is not suitable for environment where commercial tools are heavily in use.

In [2] a ***hybrid system for delivering marketing information in heterogeneous formats via the Internet*** is described. An object-oriented client/server document management system based on an RDBMS is used for storing the documents of the repository. Standard attributes, like creation date, author, or a version number are automatically assigned to the repository entries by the document management system. Additional attributes, which further characterize a document, have to be included as HTML metatags. Web servers are running search engines to index the repository with the help of these tags. Since we are not using HTML metatags in our repository, all describing attributes are summarized in the CVs assigned to EEs. Based on these CVs, we offer tailored search mechanisms, including similarity-based search functions, for the different tasks, whereas the system described in [2] offers the standard web functionality for search and retrieval.

5 Conclusions

Exploiting ORDBMS features, we succeeded in conceptualizing and realizing a reuse repository that fulfills the major requirements of software development processes following the Quality Improvement Paradigm. We described the database schema, the interface functions, the flexible possibilities of integrating tools, and outlined the software architecture of the SFB-EB. To the knowledge of the authors, this is the first approach based on ORDBMS technology for such a kind of repository.

Although we consider ORDBMS to be the best choice w. r. t. our purposes, we have to admit that there are still open questions. One important requirement of software development applications is the support of an adequate versioning model for resulting products and experience elements. Current ORDBMS do not support versions at all. Consequently, we plan to realize an appropriate version model on top of the ORDBMS by exploiting its extensibility infrastructure. First steps into this direction are described in [25, 30].

Acknowledgments

The authors would like to thank Prof. Dr. R. Gotzhein, Prof. Dr. T. Härder and Prof. Dr. H. D. Rombach for their support. Part of this work has been conducted in the context of the Sonderforschungsbereich 501 “Development of Large Systems with Generic Methods” (SFB 501) funded by the Deutsche Forschungsgemeinschaft (DFG). Last but not least, we would like to thank Sonnhild Namingha from the Fraunhofer Institute for Experimental Software Engineering (IESE) for reviewing earlier versions of this paper.

References

[1] A. Aamodt, E. Plaza. Case-based reasoning: Foundational issues, methodological variations, and system approaches. In *AICom - Artificial Intelligence Communications*, 7(1):39–59, March 1994.

[2] V. Balasubramanian, A. Bashian. Document Management and Web Technologies: Alice Marries the Mad Hatter. In *Communications of the ACM*, 41(7):107–115, July 1998.

[3] V. R. Basili, G. Caldiera, H. D. Rombach. Experience Factory. In J. J. Marciniak (ed), *Encyclopedia of Software Engineering, Volume 1*, John Wiley & Sons, 1994, 469–476.

[4] V. R. Basili, G. Caldiera, H. D. Rombach. Goal Question Metric Paradigm. In J. J. Marciniak (ed), *Encyclopedia of Software Engineering, Volume 1*, John Wiley & Sons, 1994, 528–532.

[5] V. R. Basili and H. D. Rombach. Support for comprehensive reuse. In *IEE Software Engineering Journal,* 6(5):303–316, September 1991.

[6] A. Birk, C. Tautz. Knowledge Management of Software Engineering Lessons Learned. In *Proc. of the 10th Int. Conference on Software Engineering and Knowledge Engineering (SEKE'98)*, (San Francisco, CA, June 1998), Knowledge Systems Institute, Skokie, Illinois, USA.

[7] M. J. Carey, L. M. Haas, J. Kleewein, B. Reinwald. Data Access Interoperability in the IBM Database Family. In *Bulletin of the Technical Committee on Data Engineering*, 21(3): 4-12, September 1998.

[8] M. Carey, D. Dewitt, G. Graefe, D. Haight, J. Richardson, D. Schuh, E. Shekita, S. Vandenberg. The EXODUS Extensible DBMS Project: an Overview. In S. Zdonik, D. Maier (ed), *Readings in Object-Oriented Databases*, Morgan Kaufman, 1990.

[9] D. Cisowski, B. Geppert, F. Rößler, M. Schwaiger. Tool Support for SDL Patterns. In *Proc. of the 1st Workshop of the SDL Forum Society on SDL and MSC (SAM98)*, Berlin, Gemany, 1998, 107–115. ISSN: 0863-095.

[10] R. L. Feldmann, B. Geppert, F. Rößler. Continuous Improvement of Reuse-Driven SDL System Development. In *Proc. of the 11th Int. Conference on Software Engineering and Knowledge Engineering, (SEKE'99)*, (Kaiserslautern, Germany, June 1999), Knowledge Systems Institute, Skokie, Illinois, USA.

[11] R. L. Feldmann, B. Geppert, F. Rößler. First Results from an Experimental Evaluation of SDL-Pattern based Protocol Design. SFB 501 Technical Report 03/99, Sonderforschungsbereich 501, Dept. of Computer Science, University of Kaiserslautern, 67653 Kaiserslautern, Germany, 1999.

[12] R. L. Feldmann, J. Münch, S. Vorwieger. Towards Goal-Oriented Organizational Learning: Representing and Maintaining Knowledge in an Experience Base. In *Proc. of the 10th Int. Conference on Software Engineering and Knowledge Engineering (SEKE'98)*, (San Francisco, CA, June 1998), Knowledge Systems Institute, Skokie, Illinois, USA.

[13] K. Geiger. Inside ODBC. Microsoft Press, Redmond, Washington, 1995.

[14] B. Geppert, R. Gotzhein, F. Rößler. Configuring Communication Protocols Using SDL Patterns. In A. Cavalli, A. Sarma (ed), *SDL'97 - Time for Testing*, Elsevier Science Publishers, Proc. of the 8th SDL Forum, SDL '97, Paris/Evry, France, September 1997.

[15] B. Geppert, A. Kühlmeyer, F. Rößler, M. Schneider. SDL-Pattern based Development of a Communication Subsystem for CAN. In *Proc. of the IFIP Joint Int. Conference on Formal Description Techniques & Protocol Specification, Testing, and Verification, FORTE XI / PSTV XIII '98*, Paris, France, 1998.

[16] B. Geppert, F. Rößler. Generic Engineering of Communication Protocols - Current Experience and Future Issues. In *Proc. of the 1st IEEE Int. Conference on Formal Engineering Methods (ICFEM'97)*, Hiroshima, Japan, 1997.

[17] B. Geppert, F. Rößler, R. L. Feldmann, S. Vorwiger. Combining SDL Patterns with Continuous Quality Improvement: An Experience Factory Tailored to SDL Patterns. In *Proc. of the 1st Workshop of the SDL Forum Society on SDL and MSC (SAM98)*, Berlin, Gemany, 1998, 97–106. ISSN: 0863-095.

[18] D. Heimbigner. Experiences with an object manager for a process-centered environment. In *Proce. of the 18th VLDB Conference*, Van-

couver, British Columbia, Canada, August 1992.

[19] S. Henninger. Supporting the Construction and Evolution of Component Repositories. In *Proc. of the Eighteenth Int. Conference on Software Engineering*, IEEE Computer Society Press, March 1996, 279–288.

[20] S. Henninger. An Evolutionary Approach to Constructing Effective Software Reuse Repositories. In *ACM Transactions on Software Engineering and Methodology*, 6(2):111–140, April 1997.

[21] Informix Universal Server Documentation, Informix Software Inc., http://www.informix.com/, 1997.

[22] ITU-T Recommendation Z.100 (03/93). CCITT Specification and Description Language (SDL). International Telecommunication Union (ITU), 1994.

[23] ITU-T Recommendation Z.120 (10/96). Message Sequence Chart (MSC). International Telecommunication Union (ITU), 1996.

[24] W. Kim. Object-Oriented Database Systems: Promises, Reality and Future. In *Proc. 19th Int. Conf. on Very Large Databases*, Dublin, Irland, 1993, 676-687.

[25] W. Mahnke, N. Ritter, H.-P. Steiert. Towards Generating Object-Relational Software Engineering Repositories. In *Proc. Datenbanken in Büro, Technik und Wissenschaft (BTW'99)*, Freiburg, Germany, March 1999.

[26] I. Narang, C. Mohan, K. Brannon. Coordinated Backup and Recovery between DBMSs and File Systems. Technical Report, IBM Almaden Research Center, 1996.

[27] I. Narang, R. Rees. DataLinks - Linkage of Database and File Systems. In *Proc. of the 6th Int. Workshop on High Performance Transaction Systems*, Asilomar, September, 1995.

[28] R. Orfali, D. Harkey. Client/Server Programming with Java and CORBA. Wiley Computer Publishing Group, New York, 1997.

[29] N. Ritter, B. Mitschang. Capturing Design Dynamics - The CONCORD Approach. Data Engineering, 1994.

[30] N. Ritter, H.-P. Steiert, W. Mahnke, R. L. Feldmann. An Object-Relational SE-Repository with Generated Services. In *Proc. of the 1999 Information Resources Management Association International Conference (IRMA99)*, Hershey, Pennsylvania, USA, May 1999.

[31] J. Rumbaugh, I. Jacobson, G. Booch. The Unified Modeling Language Reference Manual. Addison-Wesley, 1999.

[32] M. Stonebraker, M. Brown. Object-Relational DBMSs - Tracking the Next Great Wave. Morgan Kaufman, 1999.

[33] M. Stonebraker, P. Brown, M. Herbach. Interoperability, Distributed Applications and Distributed Databases: The Virtual Table Interface. In *Bulletin of the Technical Commettee on Data Engineering*, 21(3): 4-12, September 1998.

[34] C. Szyperski. Component Software, Beyond Object-Oriented Programming. Addison-Wesley, Harlow, 1998.

[35] P. Tarr, L. A. Clark. PLEIADES: An object management system for software engineering environments. In D. Notkin, (ed), *Proc. of the 1st ACM SIGSOFT Symposium on the Foundations of Software Engineering*, ACM Press, December 1993, 56–70. Published as ACM SIGSOFT Software Engineering Notes 18(5), December 1993.

[36] R. N. Taylor, F. C. Belz, L. A. Clarke, L. Osterweil, R. W. Selby, J. C. Wileden, A. L. Wolf, M. Young. Foundations for the arcadia environment architecture. In P. Henderson (ed), *Proc. of the 3rd. ACM SIGSOFT/SIGPLAN Symposium on Practical Software Development Environments*, November 1988. Appeared as ACM SIGSOFT Software Engineering Notes 13(5), November 1988.

[37] Telelogic. Tau 3.4 SDT User's Manuals. Telelogic, Sweden, 1998.

[38] The ASSET Reuse Library, http://www.asset.com/WSRD/indices/domains/REUSE_LIBRARY.html. December 1998.

[39] G. Wiederhold. Mediators in the architecture of future information systems. In *IEEE Computer*, 25(3):38–49, March 1992.

[40] J. C. Wileden, A. L. Wolf, C. D. Fisher, P. L. Tarr. PGraphite: An Experiment in Persistent Typed Object Management. In *3rd Symposium on Software Development Environments (SDE3)*, 1988.