Journal of Universal Computer Science, vol. 3, no. 8 (1997), 929-954
submitted: 1/6/97, accepted: 10/8/97, appeared: 28/8/97 © *Springer Pub. Co.*

Corporate Memories for Knowledge Management in Industrial Practice: Prospects and Challenges

Otto Kühn
(DFKI Kaiserslautern, Germany
Otto.Kuehn@dfki.de)

Andreas Abecker
(DFKI Kaiserslautern, Germany
Andreas.Abecker@dfki.de)

Abstract: A core concept in discussions about technological support for knowledge management is the *Corporate Memory.* A Corporate or Organizational Memory can be characterized as a comprehensive computer system which captures a company's accumulated know-how and other knowledge assets and makes them available to enhance the efficiency and effectiveness of knowledge-intensive work processes. The successful development of such a system requires a careful analysis of established work practices and available information-technology (IT) infrastructure. This is essential for providing a cost-effective solution which will be accepted by the users and can be evolved in the future. The current paper compares and summarizes our experiences from three case studies on Corporate Memories for supporting various aspects in the product life-cycles of three European corporations. Based on the conducted analyses and prototypical implementations, we sketch a general framework for the development methodology, architecture, and technical realization of a Corporate Memory.
Key Words: corporate memory, knowledge management, case studies
Category: H.4.m, I.2.1

1 Introduction

The global marketplace of the future will be characterized—with respect to *products*—by shorter development cycles and harder fights on the quality battlefield, and—with respect to *organizational structures*—by lean organizations, Total Quality Management, and the advent of the Virtual Enterprise. In this situation, knowledge has been recognized as a company's most important asset for successful competition. Management strategies like *Organizational Learning* and *Corporate Knowledge Management* (KM) [Drucker, 1993; Senge, 1995; Wiig, 1993] find growing interest.

Ann Macintosh from the AI Applications Institute (AIAI, Edinburgh, UK) motivates the need for knowledge asset management as follows [Macintosh, 1997]: "Enterprises are realizing how important it is to 'know what they know' and be able to make maximum use of the knowledge. This is their corporate knowledge asset. These knowledge assets reside in many different places such as: databases, knowledge bases, filing cabinets and peoples' heads and are distributed right across the enterprise. All too often one part of an enterprise repeats work of another part simply because it is impossible to keep track of, and make use of, knowledge in other parts. Enterprises need to know: what their corporate knowledge assets are; and how to manage and make use of these assets to get maximum return."

Knowledge management has become a hot topic both in management and computer sciences, and numerous suggestions have been made concerning various aspects of the problem. We will give a brief review of the ongoing discussion and the current state of the art in the next section. We will then focus on the question: *"What kind of computer system is best suited to support knowledge management in industrial practice?"* Such a computer system we will call "Corporate or Organizational Memory".

In order to elucidate the requirements for a Corporate Memory and to investigate possible solutions, we will present three case studies which were conducted in cooperation with major European companies. We will then summarize the lessons learned from these case studies and suggest a generic reference architecture for a Corporate Memory together with some directions for a development methodology. We will finish by presenting some core problems in the realization of a Corporate Memory which require further research.

2 Knowledge Management and Corporate Memories

2.1 Current Knowledge Management Deficits

Ann Macintosh's analysis of the current deficits in knowledge management, which was quoted above, perfectly agrees with those of other authors as well as with our own experiences from industrial practice. The most serious impediments to more productivity in knowledge-based work processes can be summarized as follows:

- *Highly-paid workers spend much of their time looking for needed information:* This ubiquitous fact is in stark contrast with the considerable efforts which have been made in enhancing the productivity of lower-paid production workers by placing every tool and part they need within the reach of their hands. Even though it may be doubted that work processes which require knowledge and creativity can be rationalized to the same extent, it is more than obvious that a huge potential for improvement exists.
- *Essential know-how is available only in the heads of a few employees:* This lack of documentation is becoming ever more serious with changing work habits (e.g. frequent job changes, shorter and more flexible work hours, teleworking), which reduce the availability of individual know-how and impede communication between employees.
- *Valuable information is buried in piles of documents and data:* Even when relevant knowledge is explictly documented, its identification is becoming more and more difficult due to the continuously growing flood of mostly irrelevant information.
- *Costly errors are repeated due to disregard of previous experiences:* This is mostly a direct consequence of the forementioned deficits. It is mentioned here explicitly, since it highlights the costs of insufficient knowledge management and offers a tangible goal for improvements.
- *Delays and suboptimal product quality result from insufficient flow of information:* This is how lack of knowledge management shows at the bottom line, since the rapid development of new products with high-quality and low costs is becoming more and more essential for a company's successful competition in the global marketplace.

2.2 Computer Support for Knowledge Management

Most surveys of knowledge management (KM) concepts and their operationalization coming from the management sciences treat only roughly the aspects of computer support; the topic is usually discussed by giving an unstructured list of useful computer services, such as E-Mail, Groupware, or Hypertext Systems (see, e.g. [Wiig, 1996]).

On the other hand, a series of recent and upcoming scientific events concerning the use of Information and Communication Technology (ICT) for KM support (cf. [Dieng and Vanwelkenhuysen, 1996; Wolf and Reimer, 1996; Rose Dieng, 1997; Abecker *et al.*, 1997] shows that computer scientists have a growing interest in the topic. Whereas they discuss in detail specific technical topics of computer support for KM, they tend to ignore the specific requirements and constraints for successful KM support in industrial practice which can only be met by a combination of several ICT technologies. Nevertheless, there can be identified two main streams of work:

1. The *process-centered view* mainly understands KM as a social communication process which can be improved by various aspects of *groupware support*.
2. The *product-centered view* focuses on knowledge documents, their creation, storage, and reuse in computer-based *corporate memories*.

Groupware Support for Communication and Cooperation

Figure 1 gives an overview of techniques typically mentioned in the KM discussion. Interestingly, the core idea of "Improving knowledge management by applying innovative techniques for communication and cooperation" seems to be predominating in management-oriented papers.

In our opinion, these techniques are important *enabling* technologies; they form a substantial part of the necessary technical infrastructure for KM, but they are definitely not sufficient. Intranet technology or workflow management systems may still have an innovative touch for managers, but they have to be intelligently used and synergetically combined to overcome the current deficits mentioned in the previous section.

Using E-Mail or Video-Conferencing for consulting a knowledgeable colleague depends on a yellow-page system or a personal-skills profile for finding her; building better workflows based on previously performed projects relies on a thorough analysis of stored performance data; distributed authoring of hypertext documents makes only full use of its technical possibilities if the collaborative argumentation space built-up during the group-decisions processes can be exploited in other, similar projects and the document contents can comfortably be exploited, without "getting lost in hyperspace". All these examples show that the full benefits of an improved communication and cooperation by innovative technologies will only become effective, if they are based on a flexible and comprehensive information depository in the form of a Corporate Memory.

Knowledge Capitalization with Corporate Memories

This view on KM is being pursued mostly by computer scientists and psychologists who have been working in Artificial Intelligence. In analogy to human memory, which allows us to build on previous experiences and avoid a repetition

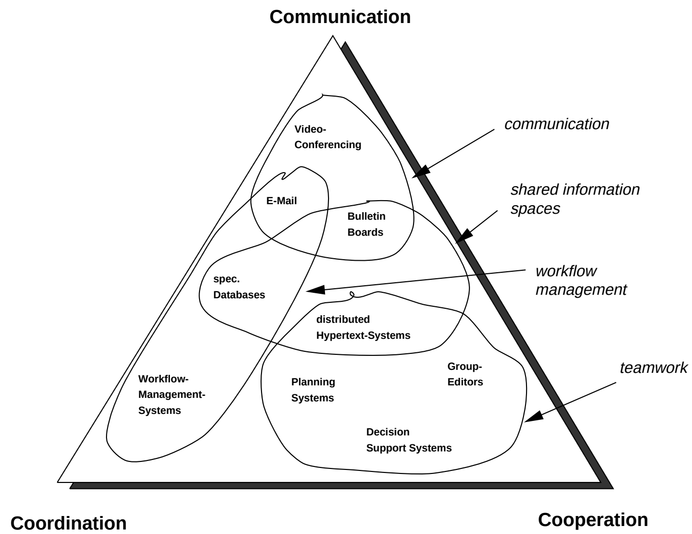

Figure 1: Dimensions of Groupware Support (taken from [Stammwitz, 1996])

of errors, a Corporate Memory is to capture information from various sources in an organization and make it available for different tasks at hand.

The vision of a such a computer system, which has also been called Organizational Memory Information System (OMIS) [Stein and Zwass, 1995], has so far been realized only in a very rudimentary way, but nevertheless, many interesting suggestions can be found in the literature:

1. Lessons-learned archives and best-practice databases typically rely on conventional Information Retrieval/Document Management techniques but introduce business process models and application domain ontologies as additional metadata categories for describing and indexing knowledge documents [van der Spek and de Hoog, 1994; van Heijst *et al.*, 1996; Hofer-Alfeis and Klabunde, 1996].
2. (Distributed) case-bases capture problem-solving expertise for specific application-problems, sometimes combined with sophisticated knowledge-based retrieval mechanisms or coupled hypertext information systems [Kamp, 1996; Kamp *et al.*, 1996; Aamodt and Nygard, 1995; Prasad and Plaza, 1996].
3. Expert systems and formal knowledge structures represent explications of formerly tacit corporate knowledge [Lukose, 1997; Wiig, 1990].
4. Formal representations of argumentation-structures record group-decision processes and embed the design documents created during theses processes (cf.: issue-based information systems, IBIS, [Reddy, 1993; Shum, 1997]).

The heterogeneity of this enumeration shows that it is not easy to distillate the main characteristics of OMIS technology which should be subject to further research and transfer into operational prototypes. One main goal of OMIS's is to *capitalize on existing knowledge*, i.e., making intangible knowledge assets tangible in order to manage them in a similar way as other resources a company deals with. For this purpose, we see two central mechanisms:

- The first mechanism is to explicate tacit knowledge and to capture it in some kind of (more or less formal) knowledge representation.
  The most typical examples are expert systems [Wiig, 1990]; lessons learned archives can also be seen in this category, but on the basis of an informal knowledge representation.
- The second mechanism starts with the knowledge sources which are already explicit in the company, e.g. in paper or electronic form, but are insufficiently used. A better exploitation of these resources can often be achieved by adding formal knowledge which describes and organizes them, eases to find them, and ensures their utilization whenever useful.
  A good example for the second mechanism is reported by Kamp [Kamp, 1994] who uses a case-based diagnosis approach for retrieving appropriate technical documentation in a helpdesk application. Another example is described in [Tschaitschian *et al.*, 1997] as "Information Tuning": the organization of software-requirement specification texts along several formal models (of the software product, the application domain, or the development process) allows to produce a much better quality of requirements documents because the formal models make possible views on the thousands of specification text fragments; such views select for several tasks exactly those information units which are relevant.

2.3 Towards the Realization of a Corporate Memory

As pointed out in the previous section, the vision of a computer-based Corporate Memory is on the one hand rather comprehensive and abstract, whereas on the other hand several specific solutions, which however only deal with isolated aspects of the whole problem, have been proposed. There still remain many open questions which require a careful analysis and evaluation.

In order to gain a better understanding of what OMIS technology can afford, we will therefore first give a rather detailed report of three case studies on OMIS support for product development which were conducted in cooperation with major European companies.

3 Three Case Studies

3.1 Study 1: Crankshaft Design

Motivation

The development department of a German company which produces motor-powered tools and vehicles inquired about the feasibility of a computer system for enhancing product development and product maintenance, i.e. the continuous improvement of products already on the market with respect to higher quality and lower production costs. The crankshaft, as one of the core components of their products, which affects and is affected by many other design decisions, was suggested as the target domain for a system prototype.

Requirements analysis

Interviews were conducted with two engineers: a designated crankshaft expert and a team leader. Product development was done in teams composed of specialists for the various components of the product (e.g. crankshaft, casing). The crankshaft expert, as the only designated specialist for this domain, was more or less engaged in all development projects at various stages from initial product conceptualization to serial production. For product maintenance, he had to answer numerous queries from manufacturing, quality assurance, and marketing departments concerning suggested modifications and the feasibility of cost reductions. The crankshaft expert suffered from heavy work overload which prevented him from keeping track of recent technical innovations which might be exploited by his company. The indispensability of the crankshaft expert became most obvious, when during his leave incorrect advice was given to a query from manufacturing which could be corrected just in time before causing a considerable loss of time and money.

As the interviews showed, the main problem in crankshaft design consists of finding an optimal solution within the quality-cost continuum. This requires the proper consideration of many incidental circumstances, such as availability of manufacturing facilities, previously encountered quality problems, etc. Such expertise can only be acquired in many years of on-the-job experience. Since this expertise was only available in the head of the over-worked crankshaft expert, we decided that the primary goal of the to-be-developed design support system should be to capture as much as possible of this expertise and make it available for answering the different types of queries which the expert is confronted with in his job.

Contrary to more general engineering knowledge, expertise for finding an optimal design under some given circumstances is subject to rapid evolution. For instance, it was found that the crankshaft expert changed a previously adopted design guideline in order to take advantage of a new supplier for blank parts who due to a currency devaluation had become more competitive. An essential system requirement was thus to support a continuous modification and update of the captured knowledge. This knowledge evolution should be accomplished by the system users with a minimal disruption of normal work processes.

System conceptualization

The system prototype which was developed based on these criteria was named KONUS, which is a German acronym for both construction support and (knowledge) conservation support. An overview of the KONUS system is shown in figure 2. A more detailed description can be found in [Kühn and Höfling, 1994] and [Kühn *et al.*, 1994].

The various users, such as the crankshaft expert, other engineers in the development team, and members of the quality assurance department interact with the KONUS system via a graphical user interface which integrates the different knowledge utilization and knowledge evolution services. These services can be grouped into three categories which were named according to [Green, 1992]:

The design aide provides direct support for the construction and modification of crankshafts, so that this activity can now also be performed by engineers who are not crankshaft specialists. Even though being unable to design a

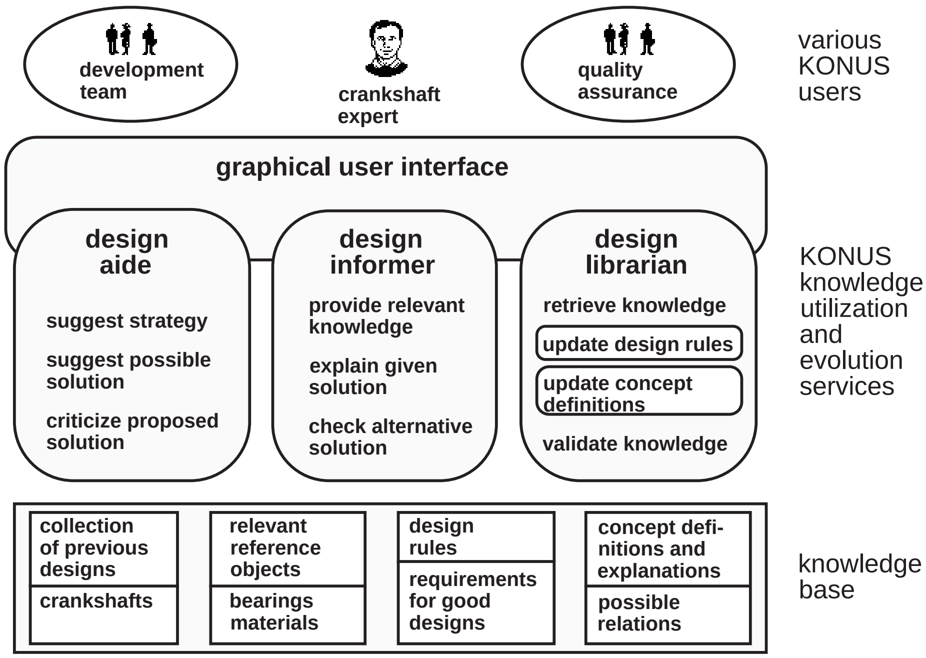

Figure 2: Overview of the KONUS Corporate Memory for crankshaft design.

crankshaft automatically, the design aide suggests a strategy, points out viable solution alternatives, and provides critique if design guidelines are violated.

The design informer provides answers to different types of questions about previous designs which frequently occur in everyday practice. Most important are "Why?" and "Why not?" questions which explain previously made design decisions and check the feasibility of design alternatives which are frequently proposed by the manufacturing or marketing departments.

The design librarian supports the management and evolution of the design knowledge by providing functionalities for knowledge retrieval, update, and validation. For knowledge update, a rule and a concept editor are provided which have been described in [Kühn *et al.*, 1994].

The captured expertise about crankshaft design is stored in a **knowledge base** in which four different types of knowledge are distinguished.

The collection of previous designs comprises descriptions of crankshafts and other crank-drive components (e.g. the connecting rod) which are designed together with the crankshaft. Each component is described by attribute-value pairs. Contrary to the technical drawings, the attribute-value representation provides an abstract and qualitative description which lists those design features which play a central role in the experts' reasoning when solving a design problem.

The relevant reference objects such as materials and bearings are also described by attribute-value pairs. Even though they are not designed but selected, their properties must be taken into account when designing the crankshaft and its components.

The design rules contain the essence of the company's design expertise. They indicate what criteria (attribute-value combinations) should be satisfied by good designs with respect to engineering, cost-efficiency, and manufacturing concerns. Each design rule may be given an informal explanation which further elaborates why that rule should be followed.

The concept definitions and explanations constitute a kind of ontology of all terms and structures which may occur in the previously mentioned types of knowledge. Besides a formal definition of concepts and relations, it also comprises informal explanations as well as information about how the various objects are to be displayed to the users. A detailed discussion of the KONUS ontology can be found in [Kühn, 1994].

System realization

With the crankshaft expert as the principal source of information, knowledge acquisition was performed in a sequence of increasingly structured interviews. The expert was asked to tell what he would like his colleagues to know about crankshaft design so that they might be able to answer for themselves many of the questions they now posed to him. In addition the expert was asked to explain some exemplary cases based on available CAD drawings.

The ontology, consisting of about 150 concepts, and the 126 design rules were extracted from the interview transcripts and presented to the expert for validation in a semi-formal notation. Concept explanations and rule justifications were taken directly from the interview transcripts and were attached as text strings to the formal knowledge units. This approach which may be seen as a combination of a formal knowledge-base with hypertext proved very effective in conveying background information to the users without having to formalize and later re-translate this information into a user-adequate form.

The core of the KONUS prototype was implemented in LISP for which knowledge representation and processing formalisms had previously been developed [Boley *et al.*, 1995]. User interaction occurred via a graphical user interface written in TCL/TK which communicated with the LISP process on a Unix machine.

Even though the utility of a Corporate Memory as exemplified by the KONUS prototype was not disputed either by prospective users or by the department management, the company could not be convinced to invest into the development of a fully usable system. On the one hand, it was hard to estimate how much effort would have to be invested for solving technical problems such as integration with currently used CAD systems and databases. On the other hand, the department was just in the process of reorganizing its outdated IT infrastructure and wanted to postpone investments into innovative AI technologies until the improvement potential of traditional database and CAD systems had been further explored.

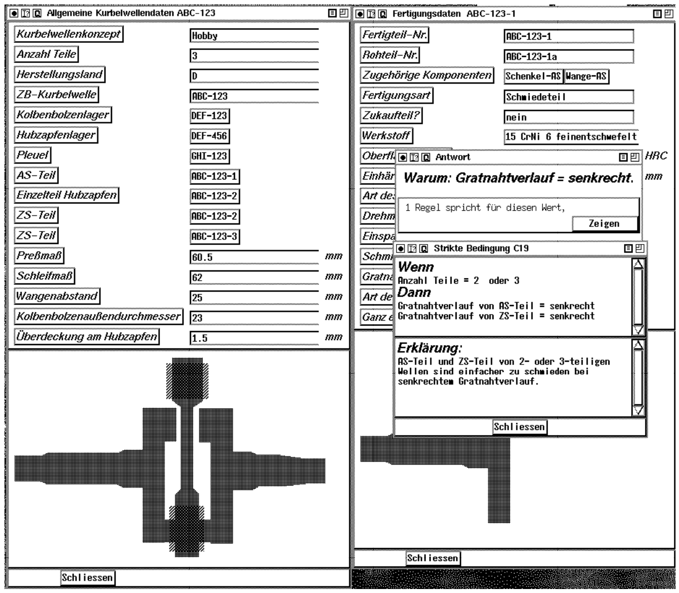

Figure 3: Snapshot of the KONUS User Interface

3.2 Study 2: Quality Assurance for Vehicle Components

Motivation

The head of the testing department of a German company which supplies components to the automobile and aircraft industry was impressed by a demonstration of KONUS, as it promised a solution to many problems he faced in his department. In particular, he sought computerized support to avoid costly errors, which were frequently committed due to disregard of previous experience, and to reduce the time that his engineers were spending looking for relevant information.

Requirements analysis

Interviews were conducted with 6 of the 20 members of the testing department. Quality assurance was done in parallel with product development and included roughly the following sequence of activities: 1) determine which tests have to be performed based on customer requirements, product features, legal prescriptions, etc.; 2) work out detailed testing instructions specifying test parameters, data to be collected, etc.; 3) calculate the testing costs depending on the number of patterns needed, work hours required for test set-up, execution and data analysis, etc.; 4) execute the tests according to the specifications and document the obtained results; 5) write test reports.

All of these activities were already performed with the help of computers (using word processors, spreadsheets, data analysis tools, etc.) and a lot of information about each project was already available in electronic form. The various data files were, however, poorly organized so that the relevant information was hard to find and was frequently overlooked. Furthermore, the most crucial information for building on previous experiences (e.g. reasons for test failures, customer complaints about employed test parameters, etc.) was not available on the computer network and its transmission depended mostly on verbal communication between the employees.

The following requirements were thus considered crucial for developing a Corporate Memory for quality assurance tests:

1. integrate all project data and documents in a systematic way so that relevant information can be accessed easily;
2. during the execution of a project capture all information which may be relevant for subsequent steps (e.g. test repetitions) in the same project or for similar projects in the future;
3. support all activities in a project by providing the specific data needed and relevant experiences from other projects in a comprehensive but concise way;
4. provide tools for capturing and continuously updating the general knowledge that may be helpful in future projects (e.g. standards required by the customer, legal regulations, technical and cost restrictions for tests);
5. provide interfaces to the other computer tools which are currently used to perform the various tasks.

System conceptualization

The global architecture of the projected system (called RITA) is quite similar to the KONUS architecture: The core of the system consists of an information depository containing all relevant information about the target domain. Project specific information is distinguished from general information. Information may be represented in the form of knowledge items (e.g. concept definitions, rules, constraints), data records, or documents (texts or drawings).

The information in RITA is entered and accessed via components which are tailored to the five above mentioned tasks. These components are to be interfaced with the other computer tools (e.g. word processor, CAD system) so that relevant information can be interchanged while the users perform their work processes in the accustomed way. Besides components supporting the five tasks in a testing project, the RITA architecture includes a knowledge acquisition and evolution component which is functionally similar to the design librarian in KONUS. Furthermore, a component for project control by the management was added which provides detailed information about the current state and encountered problems for individual projects, as well as summary statistics.

System realization

Since the RITA system requires a good deal of database and system interface programming, Visual-Works SMALLTALK was proposed as an implementation platform that provides all relevant building blocks and system interfaces for fast and efficient system development.

The effort required for a basic realization of RITA was estimated at about 4 person years. This investment into innovative information technology was deemed as rather high by the company management, in particular, since a return of investment could not be expected within one or two years. The RITA project was thus given low priority and was indefinitely postponed.

3.3 Study 3: Bid Preparation for Oil Production Systems

Motivation

An international engineering company was expanding its business in the domain of oil production system (PS) which heretofore had played only a marginal role. Besides hiring new experts, who are scarce and heavily contested among competitors, the company strove for a better capitalization on available know-how by means of better IT support. Financial benefits were expected primarily from shorter bid-preparation times and an improved quality of submitted bids, which would help the company to expand its share of a rapidly growing market.

Requirements analysis

Due to the large scope of the project, an elaborate requirements analysis was performed comprising some 30 hours of interviews with engineers, management, secretaries, and personnel from the IT department.

Work processes dealt with the preparation of bids and the execution of projects after winning a bid. The former was analyzed in more detail, since bid preparation was the primary target for computer support. The sequence of activities in bid preparation can be characterized as follows:

1. After receiving a request to submit a bid from a prospective customer, the requirements for the projected facility are collected in as much detail as possible. Relevant data are the chemical composition of the crude oil, location of the field, expected yield, prescribed security standards, etc. Most of these data can be extracted from documents supplied by the customer, which however are often enough incomplete, vague, or contradictory so that a further clarification is required.
2. A conceptual design of the facility is developed and verified with the help of a process simulator.
3. Based on the components specified in the design, equipment costs and mobilization costs are computed or estimated. For precise cost estimates, which are essential for winning a bid and not losing money, a lot of experience is needed. Many factors have to be taken into account such as worker productivity in different countries, quality and reliability of components bought from different suppliers, etc.
4. The bid document to be sent to the customer is compiled by integrating the information determined in the preceding steps together with general information about the company, standards and legal procedures to be applied, etc. The consistency and completeness of the bid document together with its overall quality is essential for winning a bid.

Concerning the current level of IT support and desired improvements, the opinions of management and staff differed considerably. Management judged that basic IT support for managing and archiving data and documents was well established and that intelligent tools for the collection of customer requirements (called 'automated questionnaire') and for the configuration of production facilities (called 'design assistant') were needed.

Staff members, however, agreed that current data and document management was a total mess. Due to the recent growth of the department nobody knew

Figure 4: Architecture and information flow of the PS-Advisor system

where to find relevant information and whom to ask. Instead of an automated questionnaire and a design assistant they desired a workflow management and document archiving system with an emphasis on support for cost estimation and compilation of bid documents. The latter was judged particularly important, since embarassing errors had occurred in bid documents which were created from old documents by a cut and paste process.

System conceptualization

In our system conceptualization, we tried to take into account all the voiced user requirements and to provide support for all work processes in bid preparation with a possible later extension to the execution of won projects. The scope of the Corporate Memory (called PS-Advisor) was delimited by the tasks for which shared product data and relevant know-how was already available.

An overview of the architecture for the PS-Advisor system is shown in figure 4. The information depository (called archive system) contains information about individual projects as well as general information. The figure also shows how information for a current project is successively added and retrieved by four tools, each of which is tailored to support one of the four main tasks in the bid preparation process.

Each of these tools was to be used by several user groups in different situations. For the automated questionnaire, for instance, three possible usage scenarios were considered: 1) a client accesses the questionnaire via WWW, 2) field people complete the questionnaire together with a client on a portable computer, 3) an engineer extracts client requirements from supplied documents. The support given by the automated questionnaire was of course to be tailored to the background knowledge and the information needs of the different users. Giving all of them access to the same information was however considered essential for encouraging mutual learning and collaboration.

The main function of the design assistant was to support the reuse of both available equipment and of previous design experiences. Relevant design parameters were to be acquired directly from a hydrocarbon simulator, which was used by the engineers to test the viability of various design alternatives.

The cost assistant was to select relevant information from previous cases based on equipment descriptions (obtained from the design assistant) as well as supplier and region data (obtained from a corporate database). In addition, it was to employ rules for cost calculations and estimations acquired from human experts.

The bid compiler was to employ rather conventional technology (e.g. text-processing macros) in order to facilitate the creation of bid documents and verify their consistency. Even though conceptually simple, it was considered crucial for the entire PS-Advisor system, since it made accumulated information from the bidding process not merely accessible but directly usable and thus facilitated the completion of a rather boring routine task.

System realization

The PS-Advisor system was planned to be realized in cooperation with a large software company mostly by customizing and enhancing software packages available on the market. An innovative approach was to be employed for the automated questionnaire, which was to be realized with HTML and JAVA and be made available via the World-Wide Web. Effort estimations for a basic realization of all systems components amounted to some 10 person years.

When the proposal for the PS-Advisor sytem was presented to the chief executive of the business unit, he was convinced of its usefulness and that a return of investment was to be expected within a few years. From past experiences, he considered it however 'philosophically wrong' to finance the development of any software just for his own department. He suggested to discuss the project together with other business units, many of which also sought better IT support for bid preparation, and the IT department. A final decision about a realization of the PS-Advisor system has not yet been reached. The IT department is currently busy with the introduction of new data and workflow management software on a company-wide scale, into which the PS-Advisor system or a more general bid-support system may be integrated in the future.

4 Lessons Learned from the Case Studies

4.1 Crucial OMIS Requirements

From the numerous interviews conducted with prospective users and the discussions with IT personnel and management, the following requirements were identified as crucial for the success of an OMIS project in industrial practice:

Collection and systematic organization of information from various sources

Knowledge needed in work processes is currently scattered among various sources, such as paper documents, electronic documents, databases, e-mails, CAD drawings, and the heads and private notes of employees. The primary requirement for an OM is to prevent the loss and enhance the accessibility of all kinds of corporate knowledge by providing a centralized and well-structured information depository.

Integration into existing work environment

In order to be accepted by the users, an OM has to tap into the flow of information that is already happening in an organization [Conklin, 1996]. At a technical level, this means that the OM has to be directly interfaced with the tools that are currently used to do the work (e.g. word-processors, spreadsheets, CAD systems, simulators). For example, in the KONUS project, a tight integration with a widespread CAD system would be necessary. As another example: the knowledge-based integration of several application systems was an important part of the PS-Advisor system design.

Minimization of up-front knowledge engineering

Even though the benefits of having an OM are generally recognized, organizations are reluctant to invest time and money into a novel technology the benefits of which will be far-off. Furthermore, prospective users have little or no time to spend for requirements and knowledge acquisition. An OM thus has to exploit readily available information (mostly databases and electronic or paper documents), must provide benefits soon, and be adaptable to newly arising requirements.

Active presentation of relevant information

In industrial practice, costly errors are often repeated due to an insufficient flow of information. This cannot be avoided by a passive information system, since workers are often too busy to look for information or don't even know that pertinent information exists. An OM therefore should actively remind workers of helpful information and be a competent partner for cooperative problem solving.

Exploiting user feedback for maintenance and evolution

For the same reasons as up-front knowledge engineering, maintenance efforts for an OM have to be kept at a minimum. At the same time, an OM has to deal with incomplete, potentially incorrect, and frequently changing information. Keeping an OM up-to-date and gradually improving its knowledge can only be achieved by collecting feedback from its users, who must be enabled to point out deficiencies and suggest improvements without causing a major disruption of the usual flow of work.

4.2 Corporate Memory vs. Expert System

On the one hand, the three case studies have shown that there exists a considerable industrial interest and application potential for Corporate Memories. On the other hand, none of the three projects went beyond the prototype stadium, which makes it obvious that companies shun the risks and costs of investing in novel technologies that have not yet found widespread application. This reservation is partly due to the unsatisfying results obtained from expert system technology with which all of the three companies had experimented in the past. It is therefore important to look closer at the differences (and commonalities) between a Corporate Memory and an Expert System.

The main objective of expert systems has been the automatic solution of a particular task, e.g. the generation of a lathe production plan. Why this is extremely difficult for a complex, real-world task is elucidated by figure 5. For an expert solution of any such task three kinds of knowledge are required: common-sense knowledge, domain knowledge, and task expertise. Capturing only the latter was attempted in first-generation expert systems, but was found to be insufficient for providing good-enough solutions under varying circumstances. Adding deeper knowledge was a primary concern of second generation expert systems, but the resulting knowledge acquisition, representation, and processing tasks were found to be enormous so that practical and cost considerations usually prohibited the development of such ambitious systems.

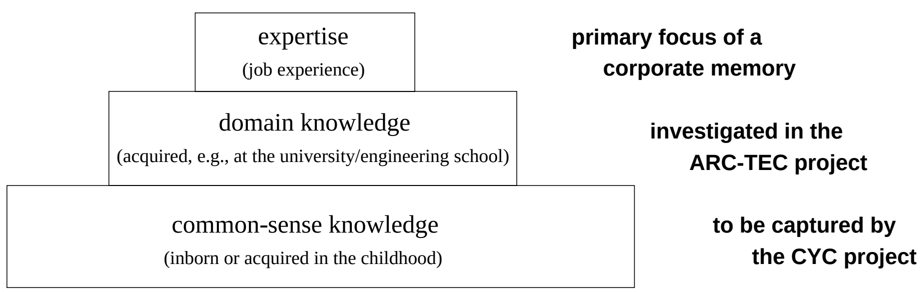

Figure 5: The knowledge pyramid.

A Corporate Memory therefore adopts a more moderate goal which has been motivated by the industrial success of database and hypertext systems (e.g. workflow and product-data management systems, Internet and Intranets). Such systems store and supply relevant corporate information but leave its interpretation and evaluation in a particular task context mostly to the user. On the other hand, a Corporate Memory enhances these technologies by knowledge processing in order to improve the quality of task support which can be provided to the user.

This approach is exemplified by the OMIS prototype developed in the first case study. The design process is monitored by a design assistant which provides critique and suggests solution alternatives, without the users having to be aware that a better alternative might exist. Thanks to knowledge processing, the design informer can directly answer practically important "Why?"- and "Why not?"-questions. In a mere database and hypertext system, users would have to look up relevant information for answering such questions themselves, having to sift through large amount of potentially relevant information, which however does not apply in the particular case.

4.3 Core Functionalities of an Organizational Memory

Let us summarize and move towards a vision of the next generation OMIS: A Corporate Memory/OMIS is an enterprise-internal application-independent information and assistant system. It stores large amounts of data, information, and knowledge from different sources of an enterprise. These are represented in various forms, such as databases, documents, and formal knowledge-bases. The OMIS will be permanently extended to keep it up-to-date and can be accessed enterprise-wide through an appropriate network infrastructure (cf. Figure 6).

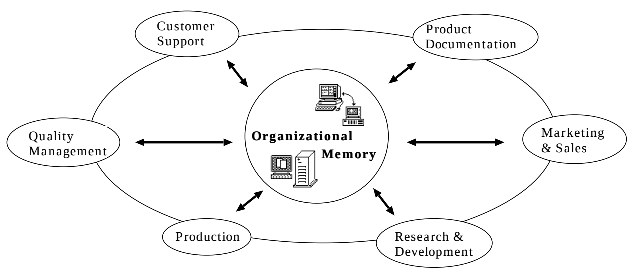

Figure 6: Enterprise-Wide Knowledge Management in an OMIS

The ultimate goal of an OMIS is to provide the necessary knowledge whenever it is needed. To assure this, the OMIS realizes an *active knowledge dissemination approach* which does not rely on users' queries but automatically provides knowledge useful for solving the task at hand. To prevent information overload, this approach has to be coupled with a highly selective assessment of relevance. The resulting system shall act as an intelligent assistant to the user.

As the design assistant of the KONUS system exemplifies, monitoring the execution of a task and providing critique when a suboptimal solution approach is detected is one technique which can be used for this end. User modeling techniques such as those used in information filtering systems, may be applied to enable not only a task-specific but also a user-specific dissemination of information.

Finally, an OMIS not only has to disseminate information but it also has to be always ready to accept new information from its users. Techniques for an integrated knowledge utilization and evolution [Kühn *et al.*, 1994], which collect usage statistics and encourage user feedback at various levels of detail can be exploited for this purpose.

4.4 OMIS Architecture

Figure 7 suggests the general architecture for a Corporate Memory which was derived by abstracting from the three case studies. Naturally, at the core of the system stands the *Information Depository*. The information contained can be structured wrt. two dimensions. Figure 7 also gives some examples for the several kinds of information we found in our case studies. According to the complexity and the representation format of information, we can distinguish data, knowledge, and documents. In an enterprise, *data* and *documents* usually are already stored in databases, paper-based or electronical document archives, or hypertext information systems. The *knowledge*—characterized by complex formal representations that can be processed with complex inferences (deduction, induction, abduction, e.g.), and not only with retrieval mechanisms—has to be captured by AI techniques to be developed together with the Corporate Memory system.

According to the level of abstraction, its role in the problem solving process, and the temporal stability of information, we can distinguish three information layers: *Case-Specific Information* describes process and control information used/generated during concrete executions of the operational task under consideration. *General Information* abstracts from specific task instances and describes how to use and to interpret the concrete case data. Such information can be employed to generate suggestions, critique, or explanations. Another form of general information can be data or documents of general importance relevant for several task executions and not depending on a concrete situation. The most abstract kind of information is managed in the *Ontological*, or the *Meta Information* layer. Generalizing the idea of a database schema, it represents the basic assumptions underlying the system development and representation. The ontology mainly describes the conceptual framework for the other system parts and is thus essential for controlling the knowledge-base evolution (see [Kühn, 1994]). A simple example is the automatic generation of structure-based rule and concept editors (cf. [Kühn *et al.*, 1994]) from ontological information that make syntactically incorrect or semantically meaningless inputs impossible.

The central information depository is accessed by a number of *Basic Information Management and Processing Services*. Principally, this generalizes the database management system in a conventional IT solution. However, because of the more heterogeneous information storage mechanisms and the more complex possible inferences, we can have a lot of different components here. Now, we can have combined *Knowledge Utilization and Evolution Components* employing the basic mechanisms for solving subtasks occurring in the work processes to be supported. The workflow, a result of requirements analysis, describes the respective subtasks, performed by the end-user(s), typically using already existing (standard) software, with specific information demands and support needs. Knowledge evolution should be a continuous activity performed by a Corporate Memory Administrator in close cooperation with the users who can make improvement/update suggestions tightly integrated into their work processes [Kühn *et al.*, 1994].

4.5 How to Develop an OMIS

As in every software project, the first step in developing a Corporate Memory has to be a careful requirements analysis in which the following questions should

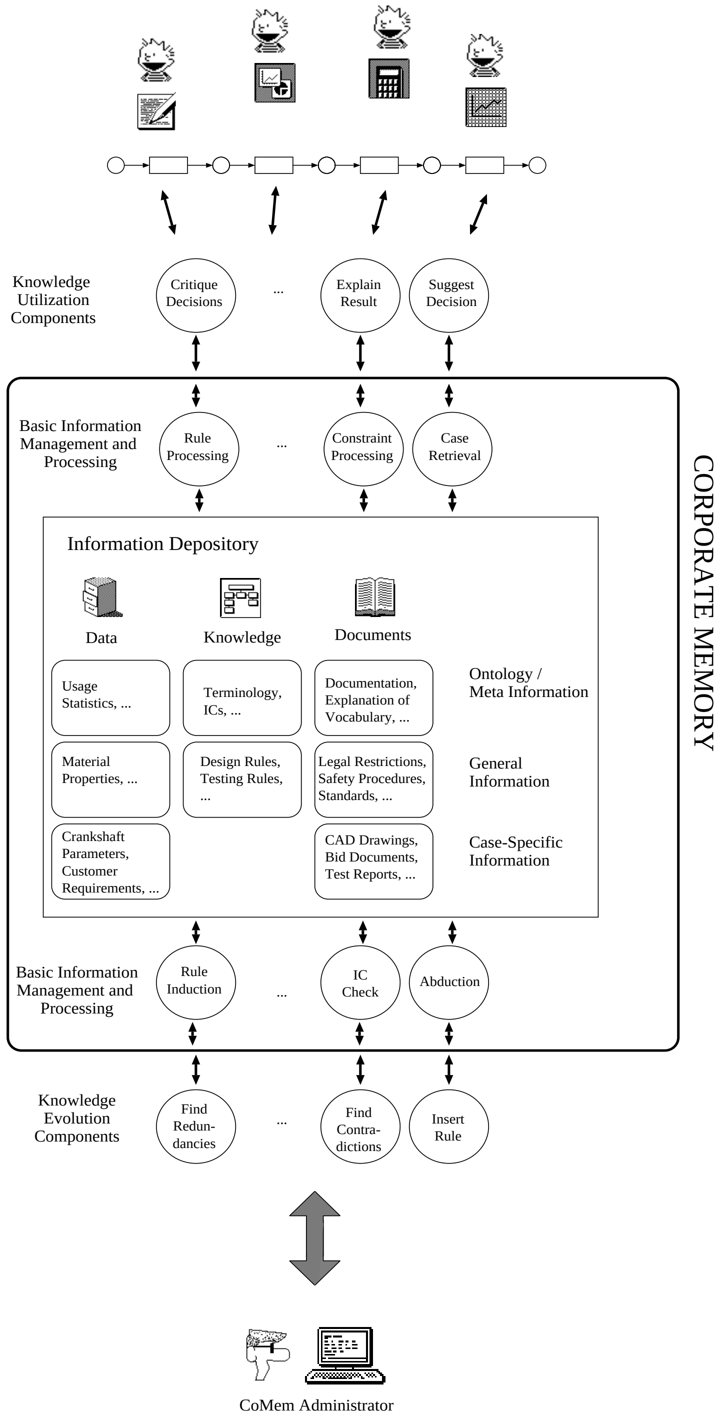

Figure 7: The Corporate Memory Reference Architecture.

be asked:

1. What are the *tasks* to be supported?
2. What is the *information needed* to perform these tasks?
3. What is the *current task environment*?
4. Which type of *support* is *desired by the users*?
5. What are the *costs and benefits* for providing different kinds of support?
6. What are the *changes to be expected* in the future?

While quite obvious, we repeat that each knowledge-based system development is first of all a software project, such that all observations valid for conventional software project management also hold true for the design and implementation of a Corporate Memory. Nevertheless, we want to draw attention to some topics that are of particular importance here:

Human Factors

A main reason for failure of early expert system projects was that the developers ignored the real needs, abilities, and goals of the intended system users (cf. [Malsch *et al.*, 1993]). Since our system is not intended to replace an expert, but to support him or her, the thorough analysis of human factors in the operational task under consideration becomes even more important [Kühn *et al.*, 1994]. The question how to motivate users for storing valuable personal expertise into the OMIS requires careful analysis in each concrete application.

Cost-Benefit-Analysis

Knowledge-based techniques are not yet an off-the-shelf product, rather they are hand-made, especially tailored to an application, which causes relatively high and not-easy-to-estimate development costs. Moreover, the benefit achievable by better IT support is very hard to quantify. Thus, we propose to minimize the risks of a Corporate Memory project by focusing it in a twofold way. First, a project introducing Corporate Memory technology in an enterprise should concentrate on a product or process central to the operational outcome, already suffering from the lack of support or exposed to a high risk of (costly) errors caused by insufficient information. This was a main characteristic of the crankshaft design example, described in section 3.1. Second, we should not overload the initial system with too many services, which may be desirable but do not promise a quick return of investment.

Knowledge Evolution

Corporate Memory maintenance is an important topic for several reasons. First, electronic support is especially valuable in areas being subject to rapid change in the enterprise itself and in its economic/political environment, because under such conditions it is extremely difficult to provide up-to-date, consistent information at any time. Second, for practical reasons, a Corporate Memory may often be put into action with an initial knowledge seed which is to be extended, improved, and refined during its use; so, we must care about how to keep the (heterogeneous) stored information consistent during this evolution process. Third, capturing knowledge in an electronic form that was previously distributed over the personal expertise of several experts and partially stored in separate, non-interacting systems, opens new possibilities for *Knowledge Discovery* and *Data Mining*. Here, the knowledge can both be subject to analyses in order to detect, e.g., inconsistent decisions or new decision alternatives, or it can be used as background knowledge for analysing and interpreting the data and documents (cf. section 4.4). We understand *Knowledge Evolution* as a lifelong quality-improving maintenance activity for knowledge bases, technically achieved by the combination of validation and exploration activities. In [Hinkelmann and Kühn, 1995], we discuss this topic in some more detail, which was also the subject of our basic research in the VEGA project [Abecker *et al.*, 1995].

Technical Realization

The knowledge-processing part of our case studies, as the most innovative part from the IT point of view, was surprisingly easy to manage. With state-of-the-art knowledge-based system technology, tailored to technical domains, one can already be notably useful in real-world applications. However, the real world provides enough other difficult problems. One has not only to incorporate and integrate existing data and information bases (which was relatively harmless in the described projects, but is still an active research area in the *Enterprise Integration* and *Knowledge Sharing and Reuse* communities) and to realize tightly coupled data and knowledge processing. One also has to closely interweave knowledge-based support with the already existing work processes typically performed using available standard software. For example, it should be possible to have a knowledge-based validation of the data actually processed within an Excel spreadsheet. To tackle this problem, we found the concept of client-server architectures as especially useful (which is already suggested by the logical system architecture as depicted in figure 7). Another contribution to Corporate Memories is to have commonly agreed upon standards for communication, product and process data, as well as knowledge and meta-knowledge interchange.

5 Future Work

As already pointed out in section 4.2, we regard an OMIS not only as an information system, but rather as an intelligent assistant to the user. It accompanies the execution of tasks, makes suggestions to the users, or provides relevant information that enables the employees to do their work better and more effective. Besides this knowledge supply, the OMIS makes it easy for the user to extend or refine it whenever he or she becomes aware of new knowledge.

In section 3, we have described several application projects as pragmatic approaches to implement handcrafted solutions that realize parts of our OMIS vision. All projects failed to jump into daily operational work. Main reasons were the costs of customer-tailored solutions with unpredictable return of investment, the insufficient experiences with OMIS applications, and the poor integration into the conventional IT landscape. This leads to the demand for further substantial application-oriented basic research with the ultimate goal to provide software modules necessary for building OMIS as off-the-shelf solutions.

In order to identify the main research goals to be attacked, let us first sketch once again the kinds of knowledge to be managed in an OMIS and the way how active user support can be achieved by synergetically exploiting formal and non-formal knowledge.

Kinds of Knowledge and Their Roles in an OMIS

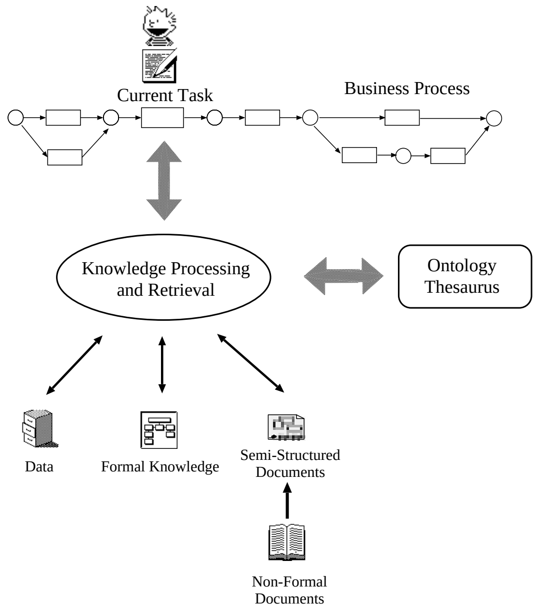

Figure 8: Active Support by Integrated Knowledge Processing and Retrieval

Our view of how a next generation OMIS could work is shown in Figure 8. The OMIS makes available individual knowledge across the organization, thus enabling organizational learning and supporting problem solving and decision making. This can be done by handling *formal* knowledge items, e.g. business rules or design guidelines which typically represent 'condensed' knowledge. An overwhelming part of individual knowledge, however, will not be available in this form. Instead, *semi-structured* representations (e.g. keyword-oriented memos, notes, e-mails) are widespread and suitable to capture the employees' experiences and lessons learned. They need not be processed automatically, but provided to colleagues when needed.

Paper-based or electronic *documents*, e.g. letters, manuals, technical documentation, are another important carrier of knowledge in an organization. The OMIS should ease the capitalization on this company legacy knowledge by context-dependent, active presentation within the human work process. By providing appropriate descriptions this material can be treated similar to the semi-structured items. Finally, the large amounts of *data* stored in databases or data warehouses can be regarded as some kind of formal knowledge. The resulting system intelligently assists the user on the basis of formal, semi-structured and non-formal knowledge. The role of the formal knowledge is two-fold:

On the one hand, it enables support of a user's decision making or automatically solving certain tasks (like a conventional knowledge-based system). However, if the limits of the formal knowledge are reached, the OMIS still offers support by presenting relevant non-formal knowledge: to explain formal inferences or to provide information instead of or in addition to formally derived answers.

On the other hand, the formal knowledge and the inferences enable a precise and effective retrieval. The goal of the inferences is to exactly determine the information needs and to map the particularities of the current domain and the task at hand to the information access structures. This mapping is supported by ontological and/or thesaurus background knowledge in order to be robust against changing vocabulary, and it is triggered by information-need specifications attached to the workflow model.

Research Questions

For realizing our visionary OMIS architecture, the following research questions must be attacked:

1. *Integrate Data Models, Thesauri, and Ontologies.* A basis for the combined exploitation of data, documents, and formal knowledge is a unified meta data model. First steps into this direction are already done in the KACTUS project [Ostermayer *et al.*, 1996]. Another interesting point in this area (wrt. the requirement of a minimum of up-front knowledge engineering) is to develop techniques for automatic thesaurus generation from existing document corpora. These techniques can profit from available ontological knowledge. The combined ontology/thesaurus can be used to improve retrieval, filtering, and routing of documents.
2. *Combine Inference and Information Retrieval.* Conjoint exploitation of formal and non-formal knowledge representations is a consequent extension of the logic-based view on Information Retrieval and the conceptual indexing approach [van Rijsbergen, 1989; van Bakel *et al.*, 1996]. For tackling this research goal, expressive document representation formalisms and feasible inference methods must be investigated [Rölleke and Fuhr, 1996; Meghini and Straccia, 1996; Abecker *et al.*, 1998]. Inference is used as a powerful and declarative retrieval mechanism which is able to take into account several kinds of background knowledge (e.g. retrieval heuristics or thesaurus knowledge).
3. *Bring together Business Processes and Knowledge Management.* Business process models can be content of the OMIS knowledge base as well as meta information for annotating other knowledge elements at the meta level. The formal representation and seamless integration of business process models into other knowledge-based processes in the enterprise is an important prerequisite. First steps are done in the knowledge acquisition community [Decker *et al.*, 1996] and in the field of Enterprise Ontologies [Fox *et al.*, 1995; Uschold *et al.*, 1995]. The ultimate goal is to detect an information need within an ongoing process execution and to determine knowledge relevant

in the particular task execution context. A first, pragmatic step into this direction is described in [Hinkelmann and Kieninger, 1997] where the authors propose to exploit task-context information for Information Filtering.

6 Summary

A Corporate Memory integrates product data with product legacy knowledge in order to bring previously accumulated experiences to bear on new tasks. Thereby, a repetition of errors can be avoided, know-how can be enhanced systematically, and knowledge-intensive work processes can be made more efficient and effective. In contrast to expert systems (and, to a lesser degree, also to intelligent assistant systems), the primary target of a corporate memory is not the support of one particular task but the better exploitation of the essential corporate resource: 'knowledge'.

We reported on three industrial case studies, all aiming at providing better IT support for crucial operational tasks by a Corporate Memory. A Corporate Memory uses state-of-the-art knowledge-processing technology, tightly coupled with the existing applications as well as data and document management systems in order to achieve a better exploitation of knowledge as a corporate resource.

A main insight of our studies was that developing a Corporate Memory essentially amounts to developing a complex software system and must thus be guided by conventional software engineering principles. The following issues are of particular importance: current organization of work processes and information interchange, human factors in cooperative problem solving and know-how sharing, cost-benefit considerations for desired system features and functionalities, and technical integration into the available IT infrastructure.

We sketched a general architecture for a Corporate Memory which has been evolved over several years based on experiences from various case studies. It is used as a general schema or template for approaching other Coprorate Memory projects in the domain of helpdesks and customer support on which we are currently working. Based on feedback from practice, the conceptual framework for Corporate Memories will be further developed towards a technology with generic modular system components and an established project methodology.

References

[Aamodt and Nygard, 1995] A. Aamodt and M. Nygard. Different roles and mutual dependencies of data, information, and knowledge–An AI perspective on their integration. *Data & Knowledge Engineering*, 16, 1995.

[Abecker *et al.*, 1995] A. Abecker, H. Boley, K. Hinkelmann, H. Wache, and Franz Schmalhofer. An Environment for Exploring and Validating Declarative Knowledge. DFKI Technical Memo TM-95-03, DFKI GmbH, November 1995. Also in: Proc. Workshop on Logic Programming Environments at ILPS'95, Portland, Oregon, Dec. 1995.

[Abecker *et al.*, 1997] A. Abecker, St. Decker, K. Hinkelmann, and U. Reimer, editors. *Workshop "Knowledge-Based Systems for Knowledge Management in Enterprises" at the 21st German Conference on AI.* September 1997. `http://www.dfki.uni-kl.de/km/ws-ki-97.html`.

[Abecker *et al.*, 1998] A. Abecker, M. Sintek, and H. Wirtz. From hypermedia information retrieval to knowledge management in enterprises. In *First Int. Forum on Multimedia & Image Processing (IFMIP-98), Anchorage, Alaska*, May 1998.

[Boley *et al.*, 1995] H. Boley, Ph. Hanschke, K. Hinkelmann, and M. Meyer. CoLab: A hybrid knowledge representation and compilation laboratory. *Annals of Operations Research*, (55), 1995.

[Conklin, 1996] E. Jeffrey Conklin. Designing organizational memory: Preserving intellectual assets in a knowledge economy. Electronic Publication by Corporate Memory Systems, Inc., 1996. `http://www.zilker.net/business/info/pubs/desom/`.

[Decker *et al.*, 1996] St. Decker, M. Erdmann, and R. Studer. A unifying view on business process modelling and knowledge engineering. In *[Dieng and Vanwelkenhuysen, 1996]*, 1996.

[Dieng and Vanwelkenhuysen, 1996] Rose Dieng and Johan Vanwelkenhuysen, editors. *Tenth Knowledge Acquisition for Knowledge-Based Systems Workshop KAW'96, Special Track on Corporate Memory and Enterprise Modeling.* November 1996. `http://ksi.cpsc.ucalgary.ca:80/KAW/KAW96/KAW96Abstracts.html`.

[Drucker, 1993] Peter F. Drucker. *The Post-Capitalist Society.* Butterworth-Heinemann Ltd, 1993.

[Fox *et al.*, 1995] M.S. Fox, M. Barbuceanu, and M. Gruninger. An organisational ontology for enterprise modelling–preliminay concepts for linking structure and behaviour. In *4th Workshop on Enabling Technologies - Infrastructures for Collaborative Enterprises, West Virginia University*, pages 1–15, 1995.

[Green, 1992] M. Green. *Conceptions and misconceptions of knowledge aided design*, volume 10 of *Knowledge-Based Systems*, chapter 1. Academic Press, 1992.

[Hinkelmann and Kieninger, 1997] K. Hinkelmann and Th. Kieninger. Task-oriented web-search refinement and information filtering. DFKI GmbH, 1997.

[Hinkelmann and Kühn, 1995] K. Hinkelmann and O. Kühn. Revising and Updating a Corporate Memory. In *EUROVAV-95, Proc. of the European Symposium on Validation and Verification of Knowledge-based Systems*, June 1995.

[Hofer-Alfeis and Klabunde, 1996] J. Hofer-Alfeis and St. Klabunde. Approaches to managing the lessons learned cycle. In *[Wolf and Reimer, 1996]*, October 1996.

[Kamp *et al.*, 1996] G. Kamp, P. Pirk, and H.-D. Burkhard. Falldaten: Case-based reasoning for the diagnosis of technical devices. Technical Report LKI-M-96/02, Labor für Künstliche Intelligenz, Universität Hamburg, June 1996. Also in Proc. 20th Annual German Conf. on Artificial Intelligence, Springer, LNAI 1137.

[Kamp, 1994] G. Kamp. Integrating semantic structure and technical documentation in case-based service support systems. Technical Report LKI-M-94/04, Labor für Künstliche Intelligenz, Universität Hamburg, 1994. Also in: Topics in Case-Based Reasoning - Selected Papers from EWCBR93, Springer, LNCS.

[Kamp, 1996] G. Kamp. Using description logics for knowledge intensive case-based reasoning. Technical Report LKI-M-96/03, Labor für Künstliche Intelligenz, Universität Hamburg, August 1996. Also in Proc. of the 4th European Workshop on Case -Based Reasoning - EWCBR96, Lausanne, Switzerland, Springer Verlag.

[Kühn and Höfling, 1994] O. Kühn and B. Höfling. Conserving corporate knowledge for crankshaft design. Proc. of the Seventh Int. Conf. on Industrial and Engineering Applications of AI and Expert Systems (IEA/AIE-94), 1994.

[Kühn *et al.*, 1994] O. Kühn, V. Becker, G. Lohse, and Ph. Neumann. Integrated Knowledge Utilization and Evolution for the Conservation of Corporate Know-How. 1994. ISMICK'94: Int. Symposium on the Management of Industrial and Corporate Knowledge.

[Kühn, 1994] O. Kühn. An Ontology for the Conservation of Corporate Knowledge about Crankshaft Design. August 1994. ECAI'94 Workshop on Comparison of Implemented Ontologies.

[Lukose, 1997] D. Lukose. Knowledge Management Using MODEL-ECS. In *[Rose Dieng, 1997]*, 1997.

[Macintosh, 1997] A. Macintosh. Knowledge asset management. *AIring*, (20), April 1997.

[Malsch *et al.*, 1993] Th. Malsch, R. Bachmann, M. Jonas, U. Mill, and S. Ziegler. *Expertensysteme in der Abseitsfalle? – Fallstudien aus der industriellen Praxis.* edition sigma, Reiner Bohn Verlag, Berlin, 1993.

[Meghini and Straccia, 1996] Carlo Meghini and Umberto Straccia. A relevance terminological logic for information retrieval. In *[?]*, 1996.

[Ostermayer *et al.*, 1996] R. Ostermayer, E. Meis, A. Bernaras, and I. Laresgoiti. Guidelines on domain ontology building. Technical Report Esprit Project 8145, Deliverable DO1c.2, KACTUS Consortium, 1996.

[Prasad and Plaza, 1996] M.V. Nagendra Prasad and E. Plaza. Corporate memories as distributed case libraries. In *[Dieng and Vanwelkenhuysen, 1996]*, 1996.

[Reddy, 1993] Y. Reddy. Computer support for concurrent engineering. *Computer*, 26(1), 1993.

[Rölleke and Fuhr, 1996] Th. Rölleke and N. Fuhr. Retrieval of complex objects using a four-valued logic. In *[?]*, 1996.

[Rose Dieng, 1997] Brian R. Gaines et al. Rose Dieng, editor. *AAAI Spring Symposium Artificial Intelligence in Knowledge Management.* March 1997. `http://ksi.cpsc.ucalgary.ca/AIKM97/AIKMDrafts.html`.

[Senge, 1995] Peter Senge. *The Fifth Discipline Fieldbook: Strategies and Tools for Building a Learning Organization.* Doubleday, 1995.

[Shum, 1997] Simon Buckingham Shum. Negotiating multidisciplinary integration: From collaborative argumenta tion to organisational memory. In *[Rose Dieng, 1997]*, 1997.

[Stammwitz, 1996] G. Stammwitz. The CIM Center Kaiserslautern CSCW Lab, 1996. Presentation Slides, WMK'96–1. Workshop Wissensmanagement Kaiserslautern, `http://www.cck.uni-kl.de/wmk/`.

[Stein and Zwass, 1995] E.W. Stein and V. Zwass. Actualizing organizational memory with information technology. *Information Systems Research*, 6(2), 1995.

[Tschaitschian *et al.*, 1997] B. Tschaitschian, A. Abecker, and F. Schmalhofer. Putting Knowledge Into Action: Information Tuning With KARAT. In *10th European Workshop on Knowledge Acquisition, Modeling, and Management (EKAW-97)*, October 1997.

[Uschold *et al.*, 1995] M. Uschold, M. King, S. Moralee, and Y. Zorgios. The enterprise ontology. Technical report, October 1995. Available from AIAI, Edinburgh, `http://www.aiai.ed.ac.uk/~entprise/papers/md31-v1.ps`.

[van Bakel *et al.*, 1996] B. van Bakel, R.T. Boon, N.J. Mars, J. Nijhuis, E. Oltmans, and P.E. van der Vet. Condorcet annual report. Technical Report UT-KBS-96-12, Knowledge-Based Systems Group, University of Twente, September 1996.

[van der Spek and de Hoog, 1994] R. van der Spek and R. de Hoog. Towards a methodology for knowledge management. In *ISMICK'94: International Symposium on the Management of Industrial and Corporate Knowledge*, October 1994.

[van Heijst *et al.*, 1996] Gertjan van Heijst, Rob van der Spek, and Eelco Kruizinga. Organizing corporate memories. In *[Dieng and Vanwelkenhuysen, 1996]*, 1996.

[van Rijsbergen, 1989] C.J. van Rijsbergen. Towards an information logic. In *Proc. of the 12th Annual Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*, 1989.

[Wiig, 1990] K. Wiig. *Expert Systems: A Manager's Guide.* Schema Press, 1990.

[Wiig, 1993] K.M. Wiig. *Knowledge Management: Foundations.* Schema Press, Arlington, 1993.

[Wiig, 1996] K. Wiig. Knowledge management is no illusion! In *[Wolf and Reimer, 1996]*, October 1996.

[Wolf and Reimer, 1996] M. Wolf and U. Reimer, editors. *PAKM'96, Proc. First Int. Conference on Practical Aspects of Knowledge Management, Basel, Switzerland.* October 1996.

Acknowledgements

Parts of the work reported here were performed in the projects ARC-TEC, IMCOD, VEGA, and KnowMore funded by the German Ministry for Education and Research (BMBF). The vision of the next generation OMIS heavily profited from discussions in the KnowMore project team: Bernd Bachmann, Ansgar Bernardi, Andreas Dengel, Knut Hinkelmann, and Michael Sintek.