K. Kuutti, E. H. Karsten, G. Fitzpatrick, P. Dourish, and K. Schmidt (eds.). *Proceedings of the Eighth European Conference on Computer-Supported Cooperative Work, 14-18 September 2003, Helsinki, Finland*

Pruning the Answer Garden: Knowledge Sharing in Maintenance Engineering

Volkmar Pipek* and Volker Wulf+

*ProSEC, Inst. for Computer Science III, University of Bonn, Roemerstr. 164, 53117 Bonn, and International Institute for Socio-Informatics (IISI), Heerstr. 148, 53111 Bonn, Germany, volkmar.pipek@iisi.de

+Institute for Information Systems, University of Siegen, Hölderlinstr. 3, 57068 Siegen, Germany; and: Fraunhofer Institute for Applied Computer Science (FhG-FIT), Schloß Birlinghoven, 53754 Sankt Augustin, Germany, wulf@fit.fhg.de

Abstract. The Answer Garden supports knowledge sharing in two intertwined ways: by making relevant information retrievable and by mediating access to people with knowledge. We present a case study in which the Answer Garden approach was applied to encourage knowledge sharing in maintenance engineering of a steel mill. The results show that the sheer amount of drawings and the long history of changing classification schemes challenge the Answer Garden approach as well as domain-specific needs for technically mediated communication. Moreover, the given division of labor and organizational micro-politics prevent the Answer Garden approach from encouraging knowledge sharing. Based on these experiences, design directions for knowledge management systems are pointed out. Finally, the results of the study are related to a recent controversy on technology support for expertise location.

Introduction

Knowledge is typically distributed among different actors and embodied in various artifacts (cf. Hutchins 1995; Ackerman and Halverson 1998; Davenport and Prusak 1998). There are mainly two - often interrelated - ways to share knowledge among human actors. In the direct way, human actors with different kinds of expertise can communicate and help each other to construct new knowledge. In the mediated way, a knowledgeable actor can create artifacts which may facilitate knowledge construction processes of others.

People who need to learn face usually the problem of finding either the appropriate material or the right expert. Groupware applications can play an important role in tackling these problems. They can support the locating of experts as well as the communication with them. Moreover, groupware may stimulate joint creation and sharing of artifacts which capture aspects of human knowledge.

Within the field of CSCW, the sharing of knowledge in organizations have been studied both empirically and with regard to the design of information systems (cf. Bannon and Kuutti 1996, Ackerman and Halverson 1998, Trigg et al. 1999, Groth and Bowers 2001, Lutters and Ackerman 2002). Looking at the design-oriented approaches, the Answer Garden (AG) approach by Ackerman (with Malone 1990, 1994, with McDonald 1996, 1998) has been extremely influential and is widely referred to (e.g. Stahl and Herrmann 1999, Fagrell et al. 1999, McDonald 2000). Beyond the author's own realizations, there are several third party implementations available (Smeaton and Neilson 1995).

The AG supports knowledge management in two intertwined ways: by making relevant information retrievable and by making people with knowledge accessible. Ackerman (1994) evaluated the first version of the AG with groups of users who carried out different software engineering projects related to the X-windows system. The results of this study suggest that the basic design assumptions turned out to be supported, while certain design issues had to be reconsidered. Ackerman (1996) provides the results of another evaluation study. However, it focuses on the concept of "organizational memory" rather than on the AG system itself (see Bannon and Kuutti 1996). Due to the evaluations, the AG system was redesigned (cf. Ackerman and McDonald 1996). However, an empirical evaluation of the redesigned system (AG 2) has not yet been published.

The core innovation of the AG is the tighter coupling of information and communication spaces, and their integration into one architecture. Though we concentrate on the AG architecture here, we estimate our results to be relevant for other architectures and tools which aim at supporting knowledge management in a similar way, as well.

In contrasting with earlier success stories, in this paper we report about a failed attempt to introduce the AG approach. We tried to support maintenance engineering of a major German steel mill with the AG architecture. In the following we briefly describe the AG approach, and explain why it seemed promising for our application field. Then we report on problems which the steel mill encountered in knowledge sharing. We present the experiences made when applying the answer garden to the problems of the steel mill. Based on these findings, we evaluate the AG approach and suggest further research directions.

We finally relate our findings to the actual discussion on the nature of expertise sharing and tools for its technical support.

The Answer Garden Approach

The Answer Garden (AG) is an integrated approach which makes recorded knowledge retrievable as well as individuals with knowledge better accessible. Ackerman (1994) does not explicitly mention for which types of knowledge domains the approach fits best. However, the following criterias were relevant for him to select the domain of his field study: software-engineers in need of support in programming with the X-windows system; (cf. Ackerman 1994, p. 246):

- the complex nature of the knowledge domain,
- the interactive nature of the problem solving process,
- sufficient technological infrastructure (computers with network connections and e-mail) at the workplaces,
- lack of a body of commonly accepted knowledge,
- dynamic changes with regard to the relevant knowledge.

For these kinds of knowledge domains, the AG combines an information retrieval system with a communication system in an innovative way (cf. Figure 1). A user who has a problem and seeks information first of all checks in an information retrieval system where relevant artifacts about the knowledge domain are stored. If he can find the necessary information, he will continue working on the problem. If the user can not find the appropriate information, he can ask a question via e-mail to a human expert. By means of a location mechanism, the AG system locates an expert and routes this question to her. The expert may then answer the question (via e-mail) and if she thinks the answer is one of common interest, she can insert it into the knowledge database. In response to a question, the expert may also want to update existing material in the database or change the classification scheme of the database. So the database grows in relevant and retrievable information. Ackerman (1994, p. 245) believes that it pays for the experts to improve the AGs database, because they save time in answering frequently asked questions.

As a result of the first evaluation study, a second version of the AG system was developed. Besides a modularized technical infrastructure, the location mechanism to route the mails to the appropriate experts was improved. Moreover, the expert’s publishing process to the database received increased support.

In the first version of the AG system, the location mechanism was based on the idea that there was a fixed list of external experts to whom the questions could be routed. Indeed this group was divided into two subgroups: first and second level experts. All questions were first routed to the first level experts, who forwarded them to the second level experts whenever they felt that they were unable to answer the questions appropriately (cf. Ackerman 1994, p. 246). Such a location mechanism is based on the assumption that the group of experts is different from the other users. In the second version of the AG this assumption was abandoned and the collaborative help among all users became the design rationale of the location mechanism. An “Escalation Agent” was implemented which incorporated a hierarchy of different user groups or media (e.g. local chat system, bulletin board, human experts). Whenever a user cannot find the relevant information in the database, his question is routed to the first level of the hierarchy (e.g. chat system). After a certain period of time the escalation agent asks the user whether his question was answered. In case it is not, the agent sends the question to the next level in the support hierarchy. The escalation agent can be used to navigate through different levels of locality, assuming that this way the information seeker is investigating appropriate contexts first. (Ackerman and McDonald 1996, pp. 101 ff.).

Concerning the publishing process, in the first version of the AG it was assumed that the (edited) answers given by the experts would also be helpful to others. They were published to the shared database. In the second version of the AG, additional functions to support the authoring process of the experts were developed. These functions allow collecting, culling, and classifying, and finally distilling cooperatively those materials which have been authored to answer questions from support seekers (cf. Ackerman and McDonald 1996, pp. 102).

Figure 1 gives an overview of the AG approach for sharing knowledge within an organization. In the following we look at the case of knowledge sharing in maintenance engineering to find out whether the AG approach is applicable in that domain.

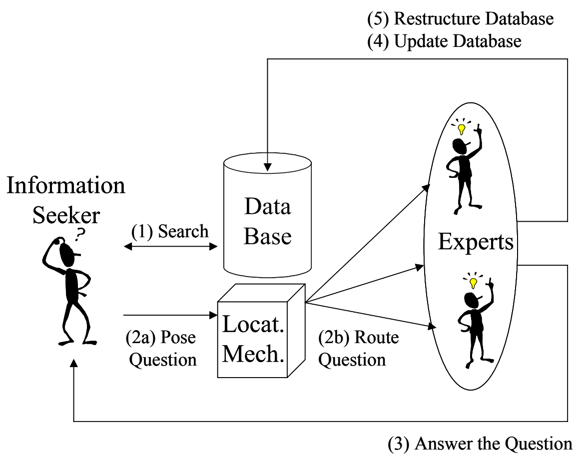

Figure 1: The Answer Garden Approach

Maintenance Engineering – Setting and Research Approach

We have investigated knowledge sharing in maintenance engineering processes of a Major German steel mill in the Ruhr area. The investigations took place in the context of the ORGTECH project (Wulf et al. 1999). The project’s goal was to support cooperative work processes within the steel mill as well as between the steel mill and two engineering offices. The two engineering offices take on subcontractual work from the steel mill in the field of maintenance engineering, e.g. the construction and documentation of steel furnace components. A construction department inside the steel mill coordinates the planning, construction and documentation processes, and manages the contracts with the external offices.

Research Method

The OrgTech project follows an action research approach, the Integrated Organization and Technology Development (OTD) framework (Wulf and Rohde 1995). The OTD process is characterized by: a parallel development of workplace, organizational and technical systems; the management of (existing) conflicts by discourse and negotiation; and the immediate participation of the organization members affected.

The project started with on-site visits, semi-structured interviews, a market overview of relevant applications, and participative workshops to build a shared understanding (project establishment phase). The project proceeded in an evolutionary way, i. e. completing cycles of problem analysis, interventions and their evaluation (cf. Wulf et al. 1999). Within this change process, the question of how to support the sharing of knowledge among the different experts involved in the maintenance of the steel mill became a focus of concern.

Our decision to apply the AG approach was motivated by several issues. First, our application field carried the properties that have been recommended in the literature for the application of the AG approach (see section above). Second, the engineers basically showed a culture of cooperation at a personal level that made it likely that technological support for knowledge sharing would be appreciated. So we developed an architecture for an AG system for the steel mill that involved and extended the existing technological infrastructures, and demonstrated prototypical implementations of some of its

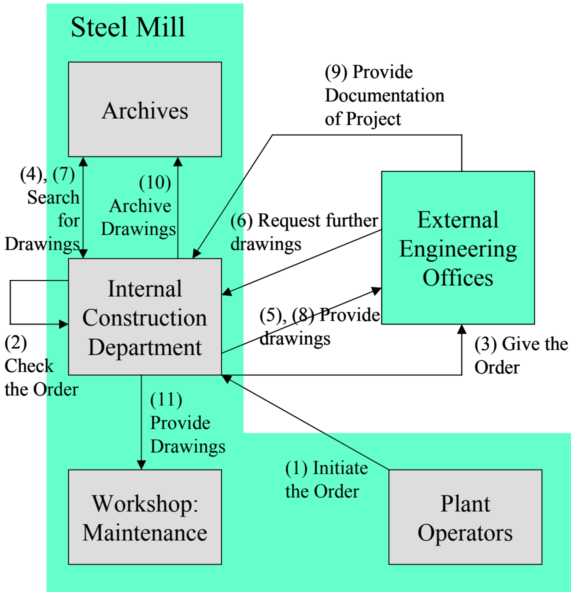

Figure 2: The Process of Maintenance Engineering

functionality in workshops. In the discussions following these presentations and workshops, it became clear that the AG architecture was rejected. These discussions and observations in that context helped us to reassess the AG approach as well as to broaden our understanding of organizational dynamics of knowledge sharing.

The results presented in this paper stem from a variety of different sources:

Analysis of the work practice: By means of about 25 semi-structured interviews, further workplace observations and additional open-ended interviews, the given work practice was examined.

Analysis of the documents available: By looking at the documents relevant for the maintenance engineering process, especially the drawings and the system descriptions (user handbooks), the relevant artifacts were investigated.

System evaluation: On the basis of task-oriented examinations like usability tests and expert reviews, the given database system was examined, especially with regard to its suitability for the maintenance engineering tasks.

Project workshops: During different workshops, organizational and technological interventions were presented and discussed to improve the maintenance engineering processes.

Field of Application

The Maintenance Engineering Department of the steel mill deals with repairing and improving the plant. Maintenance Engineering is a distributed process in which different organizational units of the steel mill and the external engineering offices are involved. Figure 2 gives a schematic overview of the maintenance engineering process.

In general, the starting point for a maintenance order is the plant operator. The steel mill consists of several organizationally independent plants, e.g. coke chambers, blast-furnace. The operators of each plant control the production equipment and machinery in their plant. Maintaining the plant also involves the repair, replacement or even redesign of outdated or deteriorated parts of the machinery, pipelines or buildings. When this kind of maintenance is necessary, the maintenance department of the plant operator asks the internal construction department for further processing. Depending on the type of order and the measures required, the transaction is handled internally or passed on to the external engineering offices. An external order will be prepared and surveyed by the responsible contact person in the internal construction department. For this reason, the necessary drawings and documents are compiled from the archives and passed on to the engineering office for further processing. Usually, the order specifications contain errors and need further clarification right from the very beginning. So discussions among the different actors and extensive re-ordering of drawings often become necessary. The additional requirements for documents are to be expressed in a comprehensive way and have to be returned to the construction department of the steel mill. Once again drawings and documents have to be found, coordination work has to be done and contacts with other departments have to be initiated. This process of renegotiating the order and reordering additional drawings requires a high level of work and expenditure of time for all participants involved.

After the external offices finish their engineering task, the internal construction department has to check it, to include the modified drawings and new ones into the archives, and to initiate the production process of the required spare parts. After being either produced by the internal workshop or ordered externally, the spare parts are assembled into the plant. While this is the general process schema of maintenance engineering, various sorts of informal communication and self-organized variations of the process can be found and add to the complexity of the problems we will describe now.

The Current State of a Plant – A Problem of Knowledge Sharing

In the following we will investigate the problem of knowledge sharing concerning one specific aspect of the process of maintenance engineering. In this case, the problem to be solved is to find out the "actual state" of those parts of a plant that are relevant for a maintenance engineering problem. "State" addresses e.g. details of the assembly of a machinery, the materials it consists of and their age, or the production process it is involved in. But it also addresses information on its location and related conditions that might interfere with necessary maintenance construction efforts, e.g. strength of the ground or where old pipelines run in the walls. This information should be available in the official technical drawings referring to that machine or location, but as we describe further below, the state described in the drawings is not necessarily the actual state of the plant.

The steel mill has a history of more than 100 years. During that time the different plants of the steel mill have been continuously modified, destroyed and replaced with other plants. The knowledge about this process is distributed among different actors in the plant and several archives containing drawings of the plant. These drawings are stored on various media. The central drawing archives contain about 300,000 technical drawings, about 2,500 files with technical descriptions, part lists, statics information and calculations, and about 500 files with plans of electronic and hydraulic devices.

A large part of these documents is filed in conventional paper form and saved on microfilm. The electronic drawing data consists of scanned drawings, which are saved in raster format and specific CAD formats. The electronically archived document stock contains about 5,000 CAD-drawings, about 20,000 raster format drawings and 30,000 scanned drawings on microfilm files, and about 90,000 documents describing the plant, maintenance processes and drawings.

In order to be able to handle the large number of drawings and documents, in 1995 an electronic archiving system was implemented to archive and provide the technical documentation. This system allows one to find drawings by numbers or keywords. Drawings are identified by two types of numbers. The first type is the drawing number which is given in a rather chronological order to all newly created drawings. It is specified by the filing clerk and does not contain any semantic information concerning the content of the drawing. For classification reasons the drawings get a second number, the so-called basic number. These basic numbers classify the drawings according to the plants and plant components they refer to. This classification scheme stems from accounting and controlling demands. It classifies the components of the plants in a way which does not fit well with the concepts and needs of engineers.

The electronic documentation is stored on a data jukebox, which is equipped with magneto-optical disks. Descriptions of the documents are stored on an Oracle data base and may be retrieved via the archiving system (called ADOS) programmed in Microsoft Access. At present, conventional and electronic archiving methods are being used in parallel, since the conversion from conventional to electronic archives is expensive and takes a long period of time to be implemented. A continuous conversion of all relevant data is aimed for. The central archives suffer from a couple of problems (cf. Hinrichs 2000):

- Approximately 20% of the drawings stored in the archives do not have any classification; a direct assignment to plants or their location is not stated. Their categorization can only be processed with in-depth system knowledge.
- A large number of drawings is old, of bad quality or has to be reconstructed in order to provide the information.
- Approximately 25% of the drawings are saved in the archive system without the correct basic or drawing number or are stored without keywords. Such drawings can only be found by the description used in the ADOS system or are accessible by search via indirect paths only (e.g. asking colleagues).
- The existing archive system does not offer extended search functions. Only drawing numbers and basic numbers (both of which are not self-explanatory) can be used as search attributes.

Due to the problems of the central archives there exist a couple of local archives maintained by the different actors involved in plant maintenance.

- Additional drawings are distributed among individuals in the maintenance department of the different plants. These individuals often have built up their private paper-based archives of those aspects of the plant they are responsible for. These archives contain up to 500 sketches and often occupy several shelves in their offices.

- Individuals of the internal construction department store documents relevant for construction, such as plant specific compiled drawing lists, in a decentralized way. They are typically filed in local paper-based archives. Some of the engineers prefer using these documents instead of the "official" ones when searching for drawings, since they contain important information in a compact form. Such documents are not within the range of electronic search functions and electronic access.

Further aspects add to the complexity of the problem to find out about the actual state of a plant in the steel mill. While already the handling of the different types of documents and archives requires some knowledge of the actors involved, certain modifications of the state of the plant cannot be found in the drawings at all. When handling accidents, plants may be modified instantly without prior planning and documentation by means of drawings. At the end of a budget year, certain work is carried out instantly to use still available funds of the respective plants. This work is typically not documented in drawings. Finally, even well planned and documented modifications of the plant may have been finally realized in a slightly different way than they are documented in the drawings. This can be caused by inadequate plans, which have to be adapted to the given environment. Sometimes the realization is carried out completely disregarding the given plans. So the knowledge about the actual state of the steel mill is distributed between different drawing archives and human actors. The workers of the maintenance departments of the different plants within the mill typically kept the best record about the actual state of their plant.

So, the problem of finding out the actual state of a plant has similar properties to the problem discussed in Ackerman's study (1994). As a reason for choosing his approach to support knowledge management for programmers of the X windows system he described specific attributes of the problem to be solved (cf. section 2). In our these attributes also apply:

- The knowledge domain is typically complex because the individual plants of the steel mill typically consist of a big variety of different components, which are interrelated in different ways. Moreover a huge variety of physical constraints has to be taken into account to understand the state of a plant.
- Like in the software engineering case, there is a lack of a body of commonly accepted knowledge. In the case of the steel mill, actors in the different organizational units maintain different drawing archives which again may document different states of the plant.
- The state of the plant changes dynamically. Maintenance activities are typically triggered by accidents, technological progress, capacity adjustments, or market changes. Some of these modifications of the plants

are not documented in drawings, because they had to be carried out instantly. This prevents actors in the steel mill from acquiring a lasting expertise.
- Maintenance engineering is often an interactive process because ideas to improve a plant have to be discussed and reversed several times.

So, the AG approach should be well applicable to the problem of finding out about the actual state of a plant.

Applying the Answer Garden Approach to Maintenance Engineering

Finding out about the actual state of a certain plant in the steel mill is a central problem in maintenance engineering. We investigated how the different elements of the AG approach could be applied to this problem.

The Steps of the Answer Garden Approach

We will discuss our experiences along the different steps of the AG approach and describe underlying organizational problems.

Search the database: The first element of the approach is the provision of electronic artifacts about the knowledge domain in a retrievable form. These artifacts exist in the case of the steel mill; the central archives contain different types of technical documents in three different media. The internal construction department and the plant operators maintain their own local paper archives. The paper artifacts are not easily retrievable as they are distributed among various locations of the steel mill. Their owners regard them as their „private" property. By contrary, the central electronic archives are easily accessible by any engineer. But the artifacts are not easily retrievable in the database. The common way to retrieve documents is via their basic number. The classification scheme according to basic numbers is not very intuitive to engineers because it was set up by the accounting department to allocate costs. Additionally, many documents are not included in the database or are not classified by basic numbers. Finally, the specification of queries to the archives is not very intuitive. So, retrieving documents in the archive database already requires quite some experience.

We have discussed different approaches to improve the electronic support for the retrieval of drawings with the engineers. First, all the drawings could be included into the database and could be classified correctly according to the basic numbers. However such an approach would require a huge amount of time of human experts. Due to the high costs and the fact hat the long-term future of the steel mill is unclear, this option is not viable from the point of view of the steel mill's management. Another approach could be to make more attributes of the drawings available for retrieval. For instance, one could apply pattern recognition or optical character recognition (OCR) algorithms on the scanned drawings to make the legend on each of the drawings available for key word search. The information given in drawing lists could be used to find drawings of artifacts which have been created or modified within the same project. All these extensions of the database scheme would require considerable input of labor because it cannot be implemented automatically. Due to the high costs involved it was not acceptable for the management, either. A similar scenario and the related problems have been described by Trigg et al. (1999).

Finally, the user interface of the drawing database could be improved. For instance, users could document and share successful search inquiries to help each other in finding relevant search results. This approach was agreed upon and is currently implemented in a prototype.

In conclusion one can say that the retrieval of the drawings is a complex problem in databases which have a long history of usage. As classification schemes change over time, information retrieval requires considerable expertise. In these cases information retrieval in the database can not be seen as an isolated first step as Ackerman (1994) describes it. Only if the required expertise to retrieve the information is available, can the drawing database serve as a primary source of information.

Pose a Question: If the database retrieval does not lead to the desired information, the AG offers the chance to formulate a corresponding question in textual form. In our case, the textual explication of a question is not always a viable way. The engineers typically ask specific questions, which may be further refined in an ongoing discussion with each other and the plant operators. Within such a discussion, artifacts like drawings, sketches or even the parts of the plants help to improve the mutual understanding.

Route the Question: Next, in the AG system the questions are routed to human experts. The system incorporates location mechanisms to find the appropriate expert. Such a location mechanism requires explicit algorithms for how to locate human experts. Indeed, such an algorithm is given in the case of the steel mill. If the information seeker is able to refer to a certain part of a plant by means of the two leading figures of the basic number, the relevant experts in the internal construction department as well as in the maintenance departments of the plant operators can be located. Additionally, the members of the archives have gained relevant experience since they have always been responsible when drawings of the plant had to be found. So, the guidance to find appropriate experts could also be provided by the information systems.

We discussed an extension of the database, which would give hints to the users whom to ask concerning certain aspects of a plant. This was assessed by the plant staff as an improvement which was mainly relevant for newcomers. Engineers who work for a longer time in maintaining the plant would typically know whom to ask.

Answer the Question and Update the Database: As a next step, in the AG approach the experts answer the questions of the information seekers and they incorporate this answer into the database with little extra effort. By updating the database they benefit from being less often asked by information seekers lateron. As already mentioned, in our case it may be difficult to share knowledge by one-way text-based communication channels. Typically an artifact-supported discussion is required. To document a given advice, multimedia recording and editing tools need to be applied. The complexity in handling these tools may well exceed the technological abilities of the experts.

Moreover, the efforts being made to help the information seeker can not always be represented within the database. The database offers a structure to store drawings. Indeed some of the advice given to information seekers will lead to either an update or a better categorization of certain drawings. However most advice will just be given verbally. The current structure of the database does not allow this type of content to be represented.

Organizational Issues

Even if these problems could be overcome, the existing division of labor within the steel mill would prevent the AG approach from being applied. The maintenance of the drawing archives and especially of the database is the responsibility of the archives group. This group is, like the internal construction department, a part of the central support division of the steel mill. Between the central support division and the different plant operating divisions there is an ongoing rivalry for power and resources (despite good relations among the engineers at a personal level). The competition for resources has led to a strict division of labor between these organizational units. Only the archives group has the right to modify the central database. The construction department has to send the drawings to the archives group after their job is finished. Afterwards, they have only read access to the central database. They cannot modify missing or incorrect classifications or update documents. On the other side, they don't have a strong interest in the material stored in the central archives. The workers in the local maintenance departments, who have built up their own local archives based on paper drawings, do not use the electronic drawing database very much. This is also due to the poor user interface for searching (Hinrichs 2000).

The restrictive access rights make it difficult to gain the benefits of the AG approach in our case. Most of the persons responsible for the inquiries of external engineers concerning the current state of a plant, cannot document their answers or at most parts of them in the shared database. While the given technological equipment at the workplaces does not allow for multimedia explanations, the reclassification and the update of certain parts of the drawings would be possible.

In case of a more flexible division of labor and corresponding access rights, the maintenance departments of the plant operators and the internal construction department could both update and improve the content of the database. When we discuss this issue with the workers of the different plants, they were not willing to improve a database they have been neither responsible for nor using much. The given division of labor and the existing conflicts between the organizational units prevented the actors from activities which would have improved the quality of the central database.

Discussion

We discuss our experiences along two lines. First, we abstract from our concrete experiences and comment on possible future directions of the AG approach and similar functionality. Second, we add to a recent discussion on the usefulness of expertise location algorithms.

Growing an Answer Garden

Our experiences showed several obstacles for realizing the AG approach. We believe that the AG is based on the following implicit assumption which were not all given in our case:

- The information seekers know exactly what their problem is.
- The categorization scheme of the database is understandable by information seekers. Moreover, they are able to handle the retrieval mechanism of the database in an appropriate way.
- Information seekers are able to use computer mediated communication (CMC) and are able express their questions in plain text.
- The location mechanism incorporates appropriate assumptions about the location of human expertise within the organization.
- If an expert is found, she is able to understand the question and give an understandable answer by means of written language.
- The experts (group of persons) have an interest in investing additional work in updating the common database.
- The experts are able to understand the given categorization scheme of the database.

Given these assumptions underlying the AG approach, we now can explicate criteria for settings in which the AG is especially applicable:

Settings in which CMC is a usual way for organizational communication: In distributed settings, like virtual organizations, there is a chance that members have built a culture of organization- and work-related CMC, which also eases the communication between expert and information seeker.

Settings with a complex expertise structure: The basic problem the AG approach addresses is: "What do I do when expertise is not at hand?". In some situation this relates to the fact that it is hard to get an overview on the competences and expertise available in complex organizational settings. Again, that problem is more likely to occur in distributed settings, when there is less opportunity for a peripheral awareness of others' expertise. The AG helps to find knowledgeable people in these settings (cf. Groth and Bowers 2001).

Settings with a strong tradition in using digital storage and collaboration media: Obviously, the problems surrounding digital artifacts made the application of the AG difficult in our case. However, one has to realize that the AG approach relies on algorithms which use metadata (e.g. who worked on what project?) to operate. This data has to be available. So the reality the data represents needs to be formalized appropriately, and that there should be cost-effective ways to maintain the metadata and the associated models. This was surely given in the AG's first application domain, i. e. the work group of programmers of the "X windows" system.

Settings with an open help culture: Several aspects related to the organizational culture may lead to restrictive practices of passing knowledge to others. Our example showed - despite good relations at a personal level - a conflict between two departments, but knowledge ownership might also be a relevant issue for individual experts. Establishing and maintaining an open help culture is important for many approaches to support knowledge transfer, but to systems like the AG it is crucial.

Referring to our experiences, we can also comment on some ways to improve the approach:

Opening up personal acres for the Answer Garden: We saw in our examples, that the engineers created private archives with own information artifacts. These often manifest a personal perspective on a problem, item or project, and often represent information which is only interpretable with the background of the private archive's creator. Yet, private archives are the first information experts refer to when explaining things to others, and they support the experts' re-contextualization process related to a problem they solved maybe years ago. Relating this personal information landscape to the public information base (e.g. as retrievable, contextualized "second opinion" on information available in public artifacts, or as a secondary stage for information retrieval before asking the expert), and supporting the use of this material in expertise communication seem to be an important direction of future work.

Growing answers using representations as seed: In the context we described, the drawings are not purely information "containers", but also focal points for given explanations. They represent an abstraction of the reality which is, when related to a specific problem, usually enriched by further information from other sources (bookkeeping information, project descriptions, co-workers) to complete the picture accordingly. These kind of representation can be valuable to clarify the communication between an expert and the information seeker, and to use representation in these communications should be supported. Using representations from the database would add the opportunity to collaboratively improve the information stored there (Buckingham Shum 1997, Pipek and Won 2002).

Harvesting in the neighbor's garden: Not being able to retrieve the information needed from own resources, the employees of the steel mill sometimes ask external engineering offices involved in related projects for material (drawings and other information) to approach a problem. This practice shows that networking AG systems could also improve the overall quality of the information infrastructure. "Distributed AGs" have been part of the original concept, but our case would also call for a concept of interconnectivity for different implementations of the AG concept which works across organizational boundaries.

On technological support for expert finding

A recent controversy on a specific part of the AG approach again posed the question, to what extent technology support may be appropriate for knowledge management. McDonald and Ackerman (2000) described an architecture (ER-Arch) which used concepts from the field of recommender systems for the problem of expertise location. In this architecture they used metadata derivable from information artifacts or other sources (groupware, email clients, browsers, organization models) to support the identification and the selection of experts.

Groth and Bowers (2001) challenged the architecture and the relevance of these algorithms with a study on expert finding in a consultancy company. Consultants there showed patterns of behavior (i.e. choosing accessibility and availability for prioritizing which expert to contact) when searching expertise which were seemingly not covered by the ER-Arch (the ER-Arch prioritizes according to heuristics which are based on data concerning expertise and concerning personal/organizational relations between the information seeker and potential experts). The study's authors further claim that, since expert location in practice is always situated, an architecture for building expertise locators would pose undesired restrictions on the information seeker. Rather than supporting expertise location they suggest to increase the awareness of experts' activities and availability.

We will now revisit our case study and provide further empirical evidence to that debate. Expertise location follows very explicit heuristics in our case. After the identification of a problem in maintenance engineering, the experts are rather easy to be found when asking concerning the current state of (one part of) the plant. Many of the workers are very confident with the machines and plants they work with, and they are experts for the functioning and current state for them. But the process of maintenance engineering is a specialized activity in which only few actors (plant operators, maintenance staff) are involved. So, the choice among the experts is rather limited. A heuristic can be easily explicated and implemented. Often the given organizational micro-politics determine whom to ask in which order (see section 5).

Contrary to Groth and Bower's (2001) observation, the current availability or a notion for activity/workload patterns is not a key factor in selecting somebody to ask in our case. Although steel mill workers devote a considerable part of their time to meetings, contact with them is not the major problem. But even in case it would become a major problem, there is little chance to ask somebody else. Awareness mechanisms may facilitate the interaction between information seeker and expert here.

Our case study suggests that history data concerning artifacts which represent knowledge may be a source for further heuristics for expertise location. In case more external service provider supports the maintenance engineering process, the relevant knowledge will be distributed among more actors. The simple matching algorithm between the experts and plants would not apply anymore. Due to the incomplete documentation provided by the drawings up to now, such a heuristic may just partially solve the problem.

A closer examination of McDonald's work sheds new light on the controversy discussed here. In his dissertation in which the ER-Arch is worked out, McDonald (2000, p. 176) describes a use scenario where he explicitly describes how recommendations brought up by the system are discarded because the user has additional information about the recommended experts' availability. The realizations of the architecture obviously are open to users' choices on how to proceed in the overall expert contacting process. So, it is possible (and potentially helpful) to work with heuristics which seem to cover a significant amount of strategies encountered in an organization. However, one should allow the choice among different strategies (e.g. with the "escalating" functionality by McDonald, or with different or no technological support). On the other hand, even from the case given by Groth and Bowers (2001) it is not credible that information seekers do not consider the competences and knowledge of their colleagues before thinking about their availability. The steps of the ER architecture reflect these considerations and give the interactive choices to navigate in recommendations.They do not restrict the strategies and options users choose from.

This controversy shows a common discussion pattern among the disciplines involved in CSCW. Designers of technology are criticized for working with inappropriate abstractions of the reality. But there is an inherent need to abstract when computer support is implemented. We have to remember that technological artifacts always face the challenge of practice some day, and that their appropriation in practice decides merciless on the quality of these abstractions.

Given McDonald's (2000) design approach which is grounded in a profound field study, the complete inappropriateness of the resulting artifacts is rather implausible. However, his work does not reflect the “total cost of modeling” (meaning model generation as well as model maintenance) to which his approach leads. In his approach, huge amounts of metadata have to be generated and frequently updated to make the expertise identification and expertise selection algorithms work. Whether the benefit of the concept justifies its costs is still to be evaluated, and repositioning the concept with regard to appropriate organizational contexts might be required.

Conclusion

The understanding of the dynamics between technical tools and organizational aspects is still a challenge for knowledge management in organizations (cf. Ackerman, Pipek, and Wulf 2003). In this contribution we discussed the failure of an introduction of the Answer Garden (AG) architecture in a steel mill.

We mentioned reasons for the introduction of the AG when describing the application field (see above). We now want to briefly comment on traditional approaches from the field of knowledge management we considered and rejected. Looking back at the experiences presented in the section “Search the Data Base”, it becomes clear that the traditional approach of repository-based support for knowledge management would have failed in this field of application. Approaches such as improving the interoperability of the distributed databases, imposing a general categorization scheme or establishing a universal access point (e.g. a web portal) to all documents could have helped improving the knowledge exchange processes. However, inaccurate and incomplete data, a complex classification problem, cost considerations, and organizational rivalries made such options unviable.

Extending traditional repository-based approaches by combining information and communication spaces into one architecture, the Answer Garden offered an innovative architecture. Therefore we applied it to a knowledge management problem which could not have been solved by traditional repository-based approaches: retrieving the current state of the plant within the general maintenance engineering process. However, our proposal has been rejected by the different actors although the scenario in the application field seems to fulfill the preconditions of the AG. So the case study adds to the discussion of appropriate organizational settings and possible improvements of the AG, or similar architectures. From revealing some implicit assumptions of the AG, we derive four conditions for organizational settings where the AG approach might be especially appropriate to use. These relate to the acceptance of computer-mediated communication, to the complexity of the expertise structure, to the tradition of digital collaboration media, and to the organization's culture.

Extensions of the concept our experiences suggest are

- the consideration of the experts' strategies to maintain private archives with information artifacts they can use for re-contextualization and explanation processes connected to a request,
- the use of representation in communication, the option to collaboratively improve representations from the database,
- and the idea to connect available archives and databases to improve the quality of the overall information infrastructure.

We also showed that in our case the given division of labor and intra-organizational rivalries also hindered an AG implementation, processes of organizational development may be required. Additionally, actors which provide information to others may need some additional qualification in documenting their knowledge (e.g. in editing the corresponding documents) and indexing their input appropriately. If knowledge cannot be easily expressed via electronic media or if the size of the knowledge domain requires sophisticated classification schemes, it is questionable whether experts are willing to take the additional efforts required. All our experiences stress the importance of an integrated perspective on technical tools and organizational issues in the context of knowledge management support.

Along the second line of argumentation, we referred to a recent discussion on the appropriateness of heuristics generated for expert finding support. This is a problem which is also central to the AG approach. Groth and Bowers (2001) challenged McDonald's and Ackerman's (2000) proposal questioning the usefulness of their heuristics. According to them, the situatedness of expertise finding strategies prevents implemented heuristics for expertise location from being effective. Due to the distribution of expertise in our field study, appropriate heuristics could easily be generated. However, they were not regarded to be useful.

We believe that the ER Arch approach (like the AG and other expertise management systems) may be problematic due to the high costs involved in building and maintaining appropriate computer models on expertise location. So the "total costs of modeling" have to be considered carefully.

Like our study with the AG, we need more in depth evaluations of expertise management systems in different organizational settings. So insights about the match between technical artifacts and organizational settings can be gained. Inappropriate design assumptions become obvious and directions for future research can be derived.

This controversy around the ER-Arch shares some aspects with the one between Suchman and Winograd concerning the Coordinator (Flores et al. 1988, Winograd 1988, Suchman 1993, JCSCW 1995). In both cases, the technological system (there: Winograd's "Coordinator") which derived design aspects from technically manifesting empirical concepts and results (there: speech acts used to coordinate work) was challenged because the resulting artifacts imposed major restrictions on users' ability to act. We do believe that it is important to observe these design trajectories carefully and critically, but we want to point out that there is a difference whether or not an architecture or technological concept is meant to replace other practices in the work setting it addresses. “Replacement architectures” (like the “Coordinator” since its concepts address a task – coordination – which inherently affects all work group members) have to be challenged much harder than “Supplement Architectures”. With regard to Supplement Architectures, practice will show whether or not the ideas of the designers work out. The AG approach and McDonald’s ER-Arch were never meant to replace other knowledge finding strategies. So other choices are still available to work group members.

Acknowledgements

We would like to thank Gerry Stahl for inspiring and valuable discussions on earlier versions of this paper. We would also like to thank our colleagues Joachim Hinrichs, Markus Klann, Matthias Krings, Bernhard Nett, Tim Reichling, Markus Rohde, Gunnar Stevens, and Markus Won for their comments and contributions to the field work.

References

Ackerman, M.S.; Malone, T.W. (1990): *Answer Garden: A tool for growing Organizational Memory*; in: Proc. of the ACM Conference on Office Information Systems, pp. 31 - 39

Ackerman, M.S. (1994): *Augmenting the organizational memory: a field study of Answer Garden*, in: Proc. of the CSCW'94 Conference, ACM Press, pp. 243 - 252

Ackerman, M. S. (1996): *Definitional and Contextual Issues in Organizational and Group Memories*, in: Technology and People, Vol. 9, No. 1, pp. 10 - 24

Ackerman, M.S.; McDonald, D.W. (1996): *Answer Garden 2: Merging Organizational Memory with Collaborative Help*, in: Int. Conf. on CSCW'96, ACM Press, pp. 97 - 105

Ackerman, M.S. (1998): *Augmenting Organizational Memory: A Field Study of Answer Garden*, in: ACM Transactions on Information Systems (TOIS), 16 (3). pp. 203 - 224

Ackerman, M. S.; Halverson, C. (1998). *Considering an Organization's Memory*. Int. Conf. on CSCW'98, Seattle, WA, ACM Press, New York, pp. 39 – 48

Ackerman, M. S.; Pipek, V.; Wulf, V. (eds) (2003): Sharing Expertise: Beyond Knowledge Management, MIT Press, Cambridge MA

Bannon, L.; Kuutti, K. (1996): *Shifting Perspectives on Organizational Memory: From Storage to Active Remembering*, in: Proceedings of HICSS-29, pp. 156 - 167

Buckingham Shum, S. (1997): *Negotiating the Construction and Reconstruction of Organisational Memories*; in: Journal of Universal Computer Science (Special Issue on IT for Knowledge Management), 3 (8), 899-928, http://www.jucs.org/

Davenport, T.H. and Prusak, L. (1998): *Working Knowledge: How Organizations Manage What They Know*. Harvard Business School Press, Boston, MA, USA.

Ehrlich, K. (2003): *Locating Expertise: Design Issues for an Expertise Locator System*, in: Ackerman, M.; Pipek, V.; Wulf, V. (eds): Expertise Sharing: Beyond Knowledge Management, MIT-Press, Cambridge MA, pp. 137 - 158

Fagrell, H., Ljungberg, F. and Kristofferson, S. (1999): *Exploring support for knowledge management in mobile work*. in: Proceedings of the 6th European Conference on CSCW (Copenhagen, Denmark, 1999), Kluwer, Dordrecht, The Netherlands, pp. 259 - 276

Flores, F., Graves. M., Hartfield, B., Winograd, T. (1988). *Computer Systems and the Design of Organizational Interaction*. Tran. on Office Information Systems. 6(2), pp.153-172.

Groth, K.; Bowers, J. (2001): *On Finding things out: Situated organizational knowledge in CSCW*, in: Proceedings of the 7th ECSCW, Kluwer, Dordrecht, pp. 279 - 298

Hinrichs, J. (2000): *Telecooperation in Engineering Offices - The problem of archiving*. in: Designing Cooperative Systems (COOP 2000), IOS Press, Amsterdam, pp. 259 - 274

Hutchins, E. (1995): Cognition in the Wild. MIT Press, Cambridge MA

JCSCW (1995), Several contributions in the Commentary Section, Int. Journal on CSCW 3(1), Kluwer, pp. 29-95

Lutters, W. G.; Ackerman, M. S. (2002): Archiving Safety : A Field Study of Boundary Objects in Aircraft Technical Support, in: Int. Conf. on CSCW'02, New Orleans, LO, ACM Press, New York, pp. 266 – 275

McDonald, D. W. (2000): *Supporting Nuance in Groupware Design: Moving from Naturalistic Expertise Location to Expertise Recommendation*, PhD-thesis, University of California, Irvine 2000

McDonald, D. W.; Ackerman; M. (2000): *Expertise Recommender: A Flexible Recommendation System and Architecture*, in: Int. Conf. On CSCW, ACM Press, New York, pp. 231-240

Pipek, V.; and Won, M. (2002). *Communication-oriented Computer Support for Knowledge Management*. Informatik/Informatique - Magazine of the Swiss Informatics Societies 2002(1): pp. 39-43.

Smeaton, C.; Neilson, I. (1995). *The Answer Web*. Fourth International World Wide Web Conference 1995, Boston, MA, USA, available upon request from *http://www.orbitalsw.com*.

Stahl, G.; Herrmann, T. (1999): *Intertwining Perspectives and Negotiation*; in: Int. Conference on Supporting Group Work (GROUP'99), ACM Press, New York, pp. 316-325

Suchman, L. (1993). *Do Categories Have Politics? The Language/Action Perspective reconsidered*. Europ. Conf. on Computer Supported Cooperative Work (ECSCW'93), Milan, Italy, pp. 1-14.

Trigg, R.H., Blomberg, J. and Suchman, L. (1999): *Moving document collections online: The evolution of a shared repository*. In: 6th European Conference, Kluwer, Dordrecht, NL, pp. 331-350.

Winograd, T. (1988): *A language/action perspective on the design of cooperative work*. In: Greif, I. (ed.), CSCW - A book of readings, Morgan Kaufman, San Mateo, CA, pp. 623-656.

Wulf, V.; Krings, M.; Stiemerling, O.; Iacucci, G.; Maidhof, M.; Peters, R.; Fuchs-Fronhofen, P.; Nett, B.; Hinrichs, J. (1999): *Improving Inter-Organizational Processes with Integrated Organization and Technology Development*, in: JUCS, Vol. 5, No. 6, 1999, pp. 339 - 365

Wulf, V.; Rohde, M. (1995): *Towards an Integrated Organization and Technology Development*. in: Symposium on Designing Interactive Systems (DIS'95), (Ann Arbor, MI, USA, 1995), ACM Press, New York, pp. 55 – 64