

Integrating Knowledge Management Systems with Everyday Work: Design Principles Leveraging User Practice

Dick Stenmark
Göteborg University
Dept of informatics, P.O.Box 620,
S-40530 Göteborg, Sweden
Stenmark@informatik.gu.se

Rikard Lindgren
Viktoria Institute
Telematics group, P.O.Box 620,
S-40530 Göteborg, Sweden
Rikard@viktoria.se

Abstract

Much research argues that information technology can have a positive influence on knowledge application. However, practical results from research on knowledge management systems indicate that such systems often fail when implemented in contemporary organizations. Whilst maintenance of knowledge management systems has been recognized as an important research area, imbalance between additional workload and accurate content still appears to be a critical factor, resulting in systems of little use for organisations in their knowledge application processes. The objective of this paper is to demonstrate how knowledge management systems can be designed to better support knowledge application in organizational knowledge work processes. Building on lessons learned from three knowledge management systems, this paper contributes general design principles describing how knowledge management systems can be integrated with everyday work to leverage user practices.

1. Introduction

There has been much debate about concepts such as knowledge-based organizations [2], knowledge-creating companies [17], organisational knowledge [21], and knowledge work [18]. Consequently, knowledge management has been widely recognised as an important approach for organisations to achieve competitive advantage [11]. Knowledge management is often regarded as the generation, representation, storage, transfer, transformation, application, embedding, and protecting of organisational knowledge [20]. Whilst processes of knowledge generation, storage, and transfer do not necessarily result in improved organisational performance, effective knowledge application does [1]. As outlined by the knowledge-based theory of the firm, the source of competitive advantage resides in the ability of an organisation to turn knowledge into action and less on knowledge itself [7].

Even though an organisation's ability to apply its knowledge depends heavily on social factors, many researchers have argued that information technology (IT) can have a positive influence on knowledge application (e.g., [1]). Practical results from knowledge management systems (KM

systems) research, however, suggest that such systems often fail when implemented in the everyday practice of modern organisations [19]. Although KM systems maintenance has been acknowledged as an important research issue [10, 12], imbalance between additional workload and accurate content still appears to be a critical problem, resulting in systems of little use for organisations in their knowledge application processes [15]. In this context, a significant challenge is to develop design principles intended to keep KM systems alive – updated, current, maintained – by encouraging use [16].

It is widely recognised that contributions from all organisational members are an important prerequisite for successful KM systems [10]. Concurring with this assertion, we argue that KM systems must be designed so that the technology itself actively affords user participation. Drawing upon empirical findings from three KM systems implemented and evaluated at Volvo Information Technology in Göteborg, Sweden and Grudin's eight challenges for developers of groupware applications [9], this paper contributes general design principles describing how KM systems can be integrated with everyday work to leverage user practices. These design principles are particularly important for designers of KM systems intended to bridge the knowing-doing gap in large and/ or geographically dispersed organisations where the organisational members do not know or know of each other and the organisation as a whole does not know what it knows. The objective of this paper is to demonstrate how KM systems can be designed to better support knowledge application in organisational knowledge work processes.

The remainder of this paper is organised as follows. Section 2 discusses knowledge application and KM systems as a theoretical backdrop. Section 3 outlines empirical findings generated from the evaluation of the three KM systems under study. In section 4, we apply Grudin's CSCW principles to the realm KM systems. Section 5, finally, outlines the conclusions of this research in terms of five principles developed for designing KM systems leveraging user practices.

2. Background

Our study is geared towards *applied* knowledge as opposed to knowledge itself, and, consequently, on how to design IT

support for knowledge workers. In this section, we account for previous work in these areas.

2.1 Knowledge application

As formulated by several commentators, the knowledge-based theory of the firm postulates that services rendered by knowledge resources such as organisation culture and identity, routines, policies, systems, documents, and individual employees form the basis for achieving competitive advantage, cf. [7, 17, 21, 23]. Viewing the firm as an institution for knowledge application, Grant emphasises that the competitiveness of an organisation depends on its ability to effectively apply the existing knowledge and to take action rather than on the existing knowledge *per se* [7]. Consistent with all theories of the firm acknowledging the efficiency gains of specialisation, Grant suggests that the principal task of organisation is to coordinate the efforts of many specialists. Integration of knowledge, either explicitly or implicitly, of many different people to facilitate knowledge application, Grant argues, is the motivation for organisations comprising multiple individuals.

In this paper, we find such distributed, unusual, and unstructured tasks and work processes requiring personal and communication-intensive forms of integration particularly interesting. This type of work is characterised by variety rather than routine, is problematic to describe in manuals, job descriptions and charts [5], and is performed by professional or technical workers with high level of skill and expertise, e.g., researchers, product developers, advertisers, and consultants. Knowledge work, unlike service work, defies routinisation and requires the use of creativity in order to produce idiosyncratic and esoteric knowledge [2]. Knowledge work is thus untidy in comparison with operational or administrative business processes, in which tangible inputs are acted on in some predictable, structured way and converted into outputs. The inputs and outputs of knowledge work, i.e. ideas, interruptions, and inspirations, are often less tangible and discrete, and in knowledge work there are no predetermined task sequences that, if correctly executed, guarantee the desired outcome [3, 6]. Summarising the characteristics of a knowledge-intensive emergent process, Markus *et al.* define such a process as an “organisational activity pattern characterised by (1) an emergent process of deliberations with no best structure or sequence, (2) an actor set that is unpredictable in terms of job roles or prior knowledge, and (3) knowledge requirements for general and specific distributed expertise” [16: 184].

2.2 System support for emergent knowledge work processes

Recognising that emergent knowledge work processes differ qualitatively from semi-structured decision making processes, Markus *et al.* argue that existing types of systems and their associated design theories do not adequately serve the unique requirements of this class of design situations [16]. More specifically, they assert that the development literature on decision support systems, executive information systems, expert systems, organisational communication systems,

organisational knowledge repository systems, and organisational memory systems does not provide sufficient guidance for how to build systems that support emergent knowledge work processes.

Markus *et al.* argue that a new IS design theory for systems supporting emergent knowledge processes is needed, and they develop a theory intended to assist system developers in their efforts to design effective emergent knowledge processes support systems [16]. On the basis of characteristics of emergent knowledge processes and requirements for IT support of such processes, this theory matches principles guiding the selection of system features and principles guiding the development process with the unique user requirements of emergent knowledge processes. Markus *et al.* suggest a new set of research challenges. One such concern is the development of design principles intended to keep KM systems alive – updated, current, maintained – by encouraging use. We see this research challenge as critical for designing KM systems, in particular so when the ambition is to bridge the knowledge application gap in organisations.

Whereas current research on KM systems maintenance has mainly focused on motivational aspects and incentives such as monetary reimbursements, norms and social issues, or organisational arrangements such as knowledge librarians (see e.g. [12, 10, 13]), little attention has been given to design-specific approaches. While acknowledging the importance of cultural, social, and motivational factors, this paper addresses the technical aspects of the knowledge maintenance dilemma by providing design principles describing how KM systems can be integrated with everyday work to leverage user practices.

Grudin’s influential work within the field of Computer Supported Cooperative Work (CSCW) shows that in situations where one party does the work and someone else receives the benefit often leads to failure [8, 9]. Although KM systems as organisational-wide technologies have been discussed in terms of groupware applications, Grudin’s findings seem to be overlooked in the knowledge management literature. As activities on top of their ordinary responsibilities, we cannot expect organisational members to spend time and efforts feeding a “knowledge database” or maintaining a “knowledge system” for the benefit of the organisation only. Yet, there must be mechanisms to express or represent knowledge in ways that enable the individual employees, and therefore the organisations as a whole, to make better use of their knowledge. By exploiting a user’s everyday actions in an unobtrusively manner, e.g., by capturing intranet activities a user already performs in form of web server log files, published documents, or submitted search engine queries, these traces can be aggregated and turned into an organisational benefit revealing otherwise invisible patterns of knowledge application.

3. Stories from the field

In the following sub-sections, we shall account for the setup of and lessons learned from two organisational experiments orchestrated by us.

3.1 Experiment #1 - setup

Volvo Information Technology AB is the competence centre for IT services within the Volvo Group and has currently some 4,300 employees world-wide. The Watson project took place at Volvo IT's Göteborg office in Sweden during the autumn of 1998 when one of the authors spent four months implementing and studying the use of an agent-based recommender system. For this experiment, we invited some 80 users, of which 48 agreed to participate. These users were invited to a 2-hour introduction meeting, where we explained the purpose of the research, the concept of agent-based systems, the design of the application and how to operate it, how to register and login, and how to set up and run individual agents.

A recommender system may be seen as performing personalised information delivery, i.e., such a system typically anticipates (based on e.g., the user's previous actions, other users' previous actions, or some mathematical similarities) what information a user is likely to be interested in and recommends such documents. The aim of the initial research project was to examine how agent-based retrieval technology could be used to help organisational members deal with information overload by providing an awareness of relevant intranet information. The assumption was that Watson (read what's on) would be able to provide employees with more targeted information for a low user effort. At the time of this study (i.e., August to November 1998), Volvo's intranet consisted of some 450 web servers and contained little less than half a million documents.

Watson was built on top of Autonomy's AgentWare software: a commercially available tool that uses neural networks and advanced pattern-matching techniques to find similarities between texts. The rationale to use software agents was to off-load the users from having to search the intranet themselves. Instead, the individual agents offered by the system could be set to find intranet documents based on an implicit profile, i.e. a richer representation of an interest than merely a keyword-based query. To achieve personalisation, the users were required to identify themselves by logging in. Once given access, the users could create agents, name the agents, and assign them tasks. A task corresponded to a search engine query, but was expressed in natural language.

To support community building, the users were supposed to create a user profile in which they were to describe their job role or work responsibilities in a free text fashion. The profile was then converted to a digital signature and the "Community" feature enabled users to locate colleagues with similar assignments and organisational roles by matching these signatures. A list of users with matching profiles was displayed and the user could now display the email address or the profile of any found user by clicking the corresponding hyperlink, and had the opportunity to contact him or her. The intention with this feature was to make the users aware of each other's presence and thus facilitate communication and collaboration.

The "Similar Agents" feature was initially meant to allow the users be able to search for and find similar agents in order to have them cloned. In this way, new and inexperienced users would receive help to get their agents to a decent quality level

more quickly. However, the cloning service was not implemented in time for the study and the only feature offered to the users during the test was the option to find other users with similar agents.

Our methodological approach can best be described as an interpretative case study [24]. User experiences and hard data were collected from several sources including interviews, questionnaires, and web server log files. Seven semi-structured interviews were conducted covering different categories of respondents such as department managers, information staff members, system developers, and technicians. During the analytic phase we strived to preserve openness to the field data in order not to stifle potential new insights, letting the data itself suggest concepts and categories according to Walsham's suggestions [24]. The data was thus categorised, conceptualised, and interpreted and the concepts derived were analysed and evaluated in an iterative fashion where the initial categories were revised and refined until they sufficiently explained all data.

3.2 Experiment #1 – lessons learned

We soon realised that we had underestimated the difficulties involved in agent training. The users conceived setting up and training of agents as non-trivial and many users had experienced mainly negative actual results. A majority of the users reported "strange" or "unexpected" document matches. The most interesting results, however, came from the Community and Similar agent usage.

The Community feature was intended to enable users with similar job profiles to learn of each other's existence. However, not many users exploited the Community feature. Those who actually did try the Community feature used it only once or, in one case, twice. The low interest was attributed to the fact that the result was just not very exciting since the users already knew the people doing similar jobs. Many users with similar profiles already worked at the same departments and the respondents were not too interested in finding like-minded colleagues. As a substitute, they suggested that one should be able to search for people with profiles other than one self. Since this was a design implication apparently shared by many users and one that seemed to be adding value, it was implemented in a subsequent prototype system (see experiment #2). The low utilisation of the Community feature was an implicit critique of the underpinning principles of explicit profiles. One user actually explicitly complained about the Community feature, claiming to have been connected to people he did not know. This was not what he had expected and he concluded that "this was clearly a bug".

People are often viewed as performing their jobs according to their formal job descriptions though everyday practice provides evidence of the opposite (cf. [4]). The organisational structure and the department descriptions are not only already known to the members but also experienced as fictitious and depicting an espoused theory of work. The Community feature was built on static profiles provided by the users themselves to mirror the official responsibilities placed upon them by the organisation. A later inspection of the users' profiles suggested that the user invested a minimum amount of time on these

profiles, resulting in very short and sketchy profiles, sometime only containing the name of the department and the job title.

Although the Similar Agent feature was implemented using the same pattern matching mechanisms, and generated exact the same output as did the Community feature, the Similar Agents feature was much more frequently used. Several respondents reported that they were surprised to find certain people sharing their interests, and they were intrigued by the fact that the Similar Agents feature returned users whom they had not expected. Users appreciated this opportunity to see in what areas other users applied their knowledge, and they considered these results useful new insights.

The discrepancy between the explicit descriptions of who we are and the implicit profiles of what we do taught us that profiles based on practice are considered more trustworthy than espoused theory-based descriptions.

3.3 Experiment #2 – setup

The second episode covers the period from June 1999 to December 2000. This research project was also carried out at the Göteborg office of Volvo IT. Like many other large and dispersed organisations, Volvo IT had recognised the major problem regarding knowing who knows what. Accordingly, large investments have been made in both organisational arrangements and IT to strengthen their competence management capabilities. One such activity was the Tieto Persona/Human Resource (TP/HR) project, initiated in June 1999. This project had two main objectives: Firstly, to identify a competence structure for Volvo IT that could serve as a foundation for the mapping of employees' competencies. Secondly, to implement the identified competence structure in TP/HR and to define a maintenance organisation that on a regular basis keeps TP/HR's structure updated and relevant.

TP/HR system was a commercial off-the-shelf module-based client/server system developed by Tieto Datema AB in Sweden. We focused on the Education/Competence module so when we hereafter refer to TP/HR, we mean this module only. The system itself was implemented in February 2000 through a top down strategy where the competence structure was defined by management alone. Volvo IT's organisational structure can be described as hierarchical and this was reflected in TP/HR's closed system structure. While managers were authorised to see competence data about all their subordinates, employees in other positions could see their own competence descriptions only.

In Volvo IT's implementation of TP/HR, competence was divided into functional (i.e., tasks such as development or support) and technical skills (e.g., programming), which in turn had sub-levels. Competence was rated using five levels, ranging 1 to 5. The search feature in TP/HR made it possible for management to search for employees holding a particular competence on a certain level, e.g., a java programmer on level 3 or above. TP/HR also had features for measuring employees' competencies status and for competence gap analyses. Volvo IT planned to use these analyses as support for organisational activities such as resource and availability planning, internal and external recruiting, goal and personal development

discussions, forming teams of employees, finding competence when manning assignments, and mission steering.

Recognising the similarities between the formalised competence descriptions of TP/HR and the unused explicit profiles in Watson, we wanted to offer action-driven competence profiles, and we therefore implemented a second recommender system – Volvo Information Portal (VIP). VIP was again an agent-based recommender system built on Autonomy's AgentWare platform. The intranet had now grown to over 700 web servers and approximately 750,000 web pages. Figure 1 illustrates the similarities and differences between the Watson and the VIP prototypes.

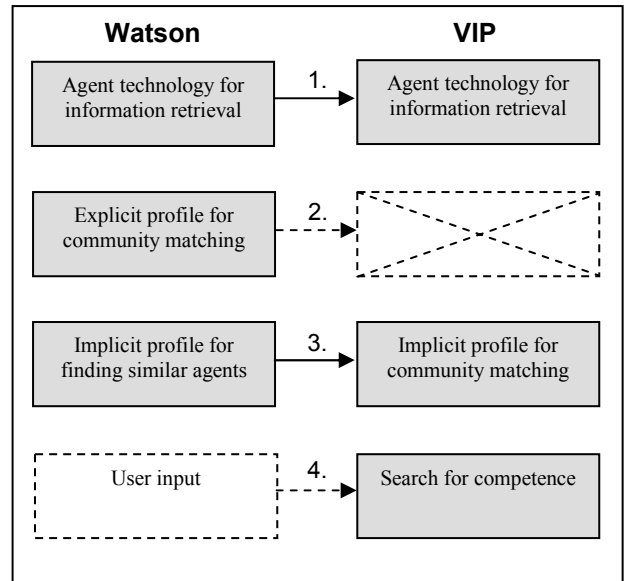


Figure 1: Design relationship between experiment #1 (Watson) and experiment #2 (VIP).

Information overload was still an issue, so we continued the agent-based information retrieval theme as the primary motivating factor (see 1 in figure 1). Since users found explicit profiles untrustworthy, we simply abandoned this idea (see 2 in figure 1). Instead, we kept (but re-labelled) the implicit profiles from Watson and introduced them as the new "Community Matching" feature (see 3 in figure 1). The "Search for Competence" feature, finally, was implemented as the result of user input from the Watson evaluation, where many users expressed their need to be able to find not just people similar to themselves but people with any arbitrary skill (see 4 in figure 1). This feature enabled the VIP users to enter a natural-language text describing a specific topic. VIP would then list all users with matching agents, i.e., all users who had agents actively searching for information related to the specified field.

Whereas the Community matching feature returned the names of those who shared your interest, the Search for competence feature could be used to find a person with an arbitrary interest. Viewing competence as knowledge applied in practice, creating information seeking agents in an area of interests was to us a practical application of existing or emerging knowledge. We realised that this view was not necessary shared by the members of the researched

organisation, so to introduce this as a competence search feature was a deliberate provocation intended to cause the organisational members to reflect. In contrast to general competence systems such as TP/HR, which rely on in beforehand-codified database records of formal competence, the VIP prototype based its results entirely on interest-driven and dynamically detected actions of organisational members. Based on the results from experiment #1, we formed a working hypothesis that these interests and actions of organisational members must have an impact on the design of KM systems. Experiment #2 was therefore designed to test this theoretical assumption in practice.

User viewpoints from the TP/HR system were collected through 10 semi-structured interviews, which lasted between 45 minutes and one hour. The interviewees were selected to represent different organisational roles and positions and included management consultants, systems programmers, and personnel from the human resource (HR) department. Another important source of TP/HR data was archival records and project documentation (including strategy plans for competence management within Volvo IT and written material about technical aspects). We also conducted 16 semi-structured, one-hour interviews with VIP users. The interviewees again occupied different positions within the organisation, ranging from non-technicians such as HR staff members, project managers, department managers, and financial controllers to technology watchers and systems programmers. As with experiment #1, we analysed the data using an interpretative approach although the action orientation was more pronounced in experiment #2.

3.4 Experiment #2 – lessons learned

Volvo IT tried to implement a competence structure that was common to and accepted by the entire organisation. The project team soon realised that to produce such a map was a non-trivial task that required much more work and consideration than they anticipated.

TP/HR was primarily designed to support management in activities such as recruiting, resource planning, and project steering. Organisational position determined how the system could be used and only managers were authorised to see their subordinate's competence description. Ordinary employees could only see their own profiles and could hence not use TP/HR to find people with a specific competence. These ordinary employees, presumed to regularly provide accurate information about their competencies, did not get much in return and had little incentive for participating.

When Volvo IT decided to implement the TP/HR system they did not foresee the problematic aspects above described. Instead, these emerged during the system implementation and while evaluating the system. Based on the troublesome work with creating a competence structure and keeping the structure relevant and updated in combination with the problems regarding competence data input and lack of commitment among the employees, the organisation realised the potential danger of the TP/HR system becoming an archive that would passively store increasingly inaccurate competence descriptions [15].

A significant finding from this study was that knowledge applied in practice was what attracted organisational members. While the problematic aspects related to the TP/HR system that arose in this study are typical for many KM systems (cf., [14]), a KM system where the content is based on daily application of individual knowledge seems to be much more appealing. Traditional competence systems (like TP/HR) have laborious maintenance processes and consequently infrequent updates, which inevitably leads to "static" systems with a more or less historic view of the organisation's competencies. The major advantage with the action-based approach, as illustrated by the VIP prototype, is that the organisation can begin to find competencies as soon as employees start to apply their existing or emerging knowledge

4. Design principles leveraging user practices

Grudin's influential research in the field of CSCW has taught us that the design of groupware applications introduce new challenges. In his early work, for example, Grudin highlights the differences between groupware applications and information systems (IS). While a large organisational IS can be considered successful if it meets the needs of most users, the majority of groupware is only useful if a high percentage of group members use it [8]. Although KM systems differ in significant ways from CSCW or groupware applications, we believe there are analogies, and that there are lessons to be learned from importing Grudin's findings to the KM systems realm.

Elaborating on his early observations, Grudin presents eight challenges for developers of CSCW applications that we argue are applicable also to KM systems [9]. These challenges are 1) Disparity in work and benefit, 2) Critical mass and Prisoner's dilemma problems, 3) Disruption of social processes, 4) Exception handling, 5) Unobtrusive accessibility, 6) Difficulty of evaluation, 7) Failure of intuition, and 8) Adoption processes. Drawing upon these eight challenges for developers of groupware applications, we shall here analyse experiences from our research effort at Volvo IT. This analysis is intended to generate general design principles describing how KM systems can be integrated with everyday work to leverage user practices.

4.1 Disparity in work and benefit

Grudin notices that groupware applications expected to provide a collective benefit still means that some people will have to adjust more than others, and CSCW applications therefore often require additional work from individuals who do not directly benefit from the application [9].

This approach was taken also when developing Watson, where users were expected to supply their own explicit profile descriptions. The sole purpose of this task was to enable other users to find them when engaging the community feature. Hence, the profiles had to be created for someone else's benefit, resulting in predictable and uninteresting profiles. The TP/HR system was also based on these principles in that employees were supposed to create and maintain their own competence database entries without even being able to use the

system. The expected benefit was on an organisational level only.

In addressing this problem, it is suggested that making the additional work required someone's explicit job might be a workaround. Such a solution seems most likely when large organisational IS are involved and when consequently strong management incentives are present. Another approach perhaps more feasible is to design the application with an accompanying process that ensures that usage creates benefits for all stakeholders.

Consequences for KM systems: When KM systems depend on making knowledge, competencies, or experiences explicit by requiring the organisational members to verbalise, rank, or document their skills, it is clearly for the benefit of the collective and not for the individual. Thus, the extra effort needed from the person interacting with the KM system, does not result in any perceived value and the interest in the KM system can be expected to decline quickly.

4.2 Critical mass and prisoner's dilemma

A groupware application requires a high percentage of all group members to interact with it in order to be truly useful. Depending on individual role or status, one or two defections may be enough to thwart an otherwise successful deployment. The problem is often to induce early adopters to stay on and not abandon the tool until a critical mass of users is achieved and they all can start to benefit [9].

With only 50 or so users in the Watson case and only some 30 in the VIP study, there were significant risks that individual users would create agents for which there were no matches. The community feature would in such cases result in zero hits and therefore generate no additional value. However, since the primary incentive for signing up with the applications was not to find community members but to receive targeted information, the lack of community members may not have had a negative impact on the overall use.

Grudin suggests that for large IS, management can force a critical mass by removing alternatives or mandate system usage until users are starting to experience the benefits and thus voluntarily continue to use it. This was the intended strategy in the TP/HR case, but the benefits were never planned to occur on the individual level, and reaching a critical mass did not help the system to survive. An alternative strategy would be to lower system thresholds by minimising the amount of additional work required, and to build in incentives for use by making salient both individual and collective benefits.

Consequences for KM systems: Critical mass means that a telephone is useful only when there is someone else to ring up. In isolation, the telephone has no value for the individual. TP/HR provided every employee with a telephone (a competence record), but removed the ability to phone other colleagues. In contrast, the VIP system provided every individual user with targeted information as a result of them training an agent: the more accurate agents, the better results. The incentive to participate was already there and a critical mass was not required to receive the primary benefits. The secondary benefit, i.e., being able to find colleagues, depended

on the number of users in the system, but this critical mass is presumably reached sooner when there is an incentive to participate even as a single user.

4.3 Disruption of social processes

Group activities are highly dependant on implicit social, motivational, economic, and political factors that change over time. If developers of groupware applications do not understand these factors, their tools may inscribe behaviour that is at odds with the subtle social dynamics of the organisation and thus hinder acceptance. If the tools violate social taboos, upset existing power structures, or reduce financial motivation, organisational members are likely to put up resistance [9].

Since knowledge is an increasingly valuable resource in today's organisations, one can expect group members to be reluctant to make explicit their knowledge and allow it to be captured by some KM system for the good of the collective. Such a process may ultimately result in them losing not only power and money but their jobs. TP/HR was a top-down system, designed for managers and from a management perspective. It is quite obvious that social factors with high influence on grass-root level were not considered.

When discussing possible solutions, Grudin reminds us of the importance of avoiding the assumption that work is carried out in a "rational" fashion. Obviously, some rationality is involved, but it has more to do with individual actors' hidden agendas than with some agreed-upon organisational goal.

Consequences for KM systems: There are at least three reasons why employees are unlikely to make explicit their knowledge: they are not fully aware of it, there is no personal need for it, and they risk losing power and competitive advantage [22]. In contrast, advertising yourself as a knowledgeable person may increase your status and salary. However, to avoid internal recruiting of experts and key individuals, TP/HR was closed to all but senior managers thus effectively removing the possibility for individuals to market themselves. When interrupting social processes in this way, KM systems as well as CSCW applications tend to fail.

4.4 Exception handling

When groupware applications are designed and implemented based on official office work handbooks and other readily available work specifications, the resulting tools may end up supporting the way things are supposed to work rather than the way they do work. Realising that descriptions of standard procedure often are post hoc rationalisations, we may recognise that what makes possible efficient performance is the ad hoc problem solving capacity of man [9].

The industrial organisation of the 20th century has been preoccupied with structures and standards, and this for good reasons. However, the breakdown of bureaucracy occurs when exceptions start to outnumber the routine. Knowledge must be renewed and find novel paths continuously to remain valuable. When yesterday's knowledge is no longer a prerequisite for tomorrow's work, old knowledge does not only become obsolete – it may actually be harmful to the company. As we

saw in the previous section, TP/HR was implemented on the basis of formal work manuals and corporate strategy policies. Many of the competencies needed in and work situations encountered during an ordinary office day were not covered by the system.

Instead of supporting rational myths, we must carefully study how work is actually done, suggests Grudin [9]. Systems must be tailorable and provide flexibility, although these requirements present challenges in themselves.

Consequences for KM systems: The databases used for storing and archiving old knowledge in today's competence systems are too rigid structures to be able to accommodate the need for ad hoc and flexible updating. To support and facilitate knowledge application, which is closely related to individual work practice, new approaches are required. Exception handling and ad hoc problem solving are the birthplaces of knowledge, and without ability to facilitate these situations the systems are useless. In daily work, exceptions are handled by people applying their knowledge in innovative ways. To be attractive, KM systems must be able to leverage these activities.

4.5 Unobtrusive accessibility

Even in groupware applications, the bulk of the work is carried out as individual tasks performed by individual group members, who mainly use groupware features to co-ordinate and communicate the result. Consequently, groupware features are typically used less frequently than many of the features supporting individual activities [9]. Grudin derive two important implications from this observation: For less frequently used feature to catch on, they must be tightly integrated with features that most users engage, and such integration must be unobtrusive not to obstruct the use of the more frequently used features.

The individual Volvo IT employee had no reason to enter the TP/HR system, except to once in a while update his or her profile in order to comply with corporate policy. In contrast, Watson and VIP rewarded the user by serving targeted information and monitoring the indicated field of interest on their behalf. If we assume information handling is something organisational members engage in on a daily basis, information agents would probably be a welcomed and relatively often used resource. The competence profiles derived from agent usage would then be maintained both frequently and unobtrusively.

Striking the right balance between being unobtrusive and yet accessible is otherwise indeed a challenge. Grudin suggests that infrequently used features should be added to and incorporated in existing and already successful applications rather than being launched as separate systems [9]. With such an approach, Grudin argues, the system can over time educate the users and slowly make them aware of the beneficial spin-offs.

Consequences for KM systems: It seems that KM systems should not be introduced as explicit stand-alone applications that user intentionally must interact with in addition to their other job responsibilities. Obviously, TP/HR suffered from this approach. KM systems should instead be invoked when

knowledge is applied in practice by exploiting spin-off from activities the organisational members are already engaged in.

4.6 Difficulty of evaluation

Whereas interaction with single-user applications can be sufficiently covered during an hour's observation, groupware interactions involve many different users and unfold over much longer periods of time. This makes evaluation of groupware applications more complex and less precise. Determining whether the application is a success or a failure may be easy, but not so to identify the factor(s) responsible for the result [9].

We were able to evaluate the Watson and VIP prototypes by studying single-users attending the primary objective of receiving relevant corporate information. We were less successful evaluating the organisational impact of the systems, since this would have required a much larger test population. The lack of historical data and ephemeral nature of the implicit profiles further added to the difficulties. However, TP/HR was even more difficult to correctly evaluate. Obviously, only three explicit competence profiles would have been a failure, but the existence of 30,000 profiles would not necessarily have indicated success.

Who should decide whether a KM system is successful: the organisation or the individual? As argued previously, there must be a benefit on the individual level before there can be a positive organisational effect. Yet, if return on investment is noticeable only at the individuals level, organisational sponsors may decide to abandon the system in lack of tangible proofs of success. Grudin's advice for how to deal with the problem of evaluation is to ensure the right mix of skills, i.e., both technical, sociological, and organisational, is allocated for the development task, and to disseminate the results actively to all stakeholders. His experience from the CSCW community is that too little accumulated learning is taking place due to the inability to learn from experiences.

Consequences for KM systems: Knowledge is an intangible resource that often affects the organisation indirectly. When evaluating IT systems in real organisational settings, it is very difficult to isolate the single factor contributing to the result. It may in fact not be one single contributing factor but a chain of concurring factors. Such evaluations are even more complicated when dealing with KM systems due to the nature of knowledge itself.

4.7 Failure of intuition

When software is constructed by the same people who are going to use it, intuition can be a reliable input to the design process; at least as far as single-user applications are concerned. Each of us has informed ideas about what we would require to get the job done. However, individual intuition is less likely to be able to predict the intricate demands on groupware tools that are to be used by a number of different users. Often, the unwelcome extra work required of other users to get the application to work is underestimated [9].

Developers typically rely on feedback from a few potential users (or sponsors); often those expected to benefit the most. A parallel from TP/HR is that mostly HR staff and managers (typical stake holders) were involved in the evaluation, whereas Watson and VIP were designed by a knowledge worker for other knowledge workers. An interesting observation is that managers on average were less impressed with the VIP approach than other employees.

Relying less on (stake holder's) intuition and more on user participation is the way forward, according to Grudin [9]. This may lead to fewer projects being run, but hopefully also to more realistic design goals and higher success rate amongst those that are actually started.

Consequences for KM systems: While systems designers capable of identifying managers' needs should be engaged when building KM systems to support management, entirely different developers should be brought in when designing for other user groups. This indicates that there should probably not be one large KM system solving everything but many small applications handling more specific aspects of knowledge management.

4.8 Adoption processes

Due to the critical mass problem mentioned earlier, groupware applications require more careful introduction in the workplace than developers may appreciate, and hence, they must pay more attention to the adoption process than product developers have in the past. The lower visibility of groupware features, which in turn generates less management support, also means that CSCW developers face more difficult acceptance problems than large-scale IS developers [9].

In our field studies, we noticed how the number of volunteering test users decreased from prototype #1 to prototype #2, something that typically happens when the group's curiosity wanes and people's attention returns to their ordinary work. Although our prototypes were built around information seeking – a process familiar to most employees – the tools themselves were new and unknown and obviously suffered from adoption problems. The small scale of our project, and consequently limited managerial attention, is likely to have contributed to the death of the prototypes.

Grudin's medicine is to sidestep the introduction problem as much as possible by adding features to existing applications, as discussed above. Building on the success of established systems and functions would, if not guarantee, at least substantially increase the likelihood of survival. As people continue to use the system, they will eventually discover the benefits of the added features and system usage will be further reinforced.

Consequences for KM systems: When KM systems depend on input from and interaction with many organisational members, the adoption process problems associated with CSCW tools apply, particularly so if the input needs to be explicit. This suggests that familiar applications used by many employees should be selected as hosts for the KM features to be added, e.g., email applications, word processors, web browsers or printer spooling systems.

5. Conclusions

A significant area of KM systems research is the development of systems with the potential to bridge the knowledge application gap in organisations. In this context, an important challenge is to develop design principles intended to keep KM systems alive – updated, current, maintained – by encouraging use. In addressing this challenge, this paper reports lessons learned from evaluating three different KM systems. The main contribution of this research is five general design principles describing how KM systems can be integrated with everyday work to leverage user practices:

- KM systems should not be introduced as explicit stand-alone applications that user intentionally must interact with in addition to their other job responsibilities. KM systems should instead be invoked when knowledge is applied in practice by exploiting spin-off from activities the organisational members already engage in. This indicates that there should probably not be one large KM system solving everything but many small applications handling more specific aspects of knowledge management.
- When KM systems depend on input from and interaction with many organisational members, familiar applications used by many employees should be selected as hosts for the KM features to be added, e.g., email applications, word processors, web browsers or printer spooling systems.
- Tomorrow's KM systems must be able to adapt to rapid changes in what sort of knowledge is being managed and to which field the knowledge is applied. KM systems based on rigid and well-defined structures are less likely to be able to do such adjustments and may therefore fail.
- To be perceived as attractive KM systems should provide organisational members with a natural incentive not only to participate but to provide as updated and as accurate information as possible. The most plausible way for this to happen is to have the system reward the contributor with direct and tangible benefits.
- KM systems must acknowledge and co-exist alongside existing social processes and organisational culture. Ignoring such issues and over-estimating the power of rational thinking is likely to lead to failure.

6. References

- [1] Alavi, M. and Leidner, D. E., "Knowledge Management and Knowledge Management Systems: Conceptual Foundations and Research Issues", *MIS Quarterly*, 25(1), 2001, pp. 107-136.
- [2] Blackler, F., "Knowledge, Knowledge Work and Organizations: An Overview and Interpretation", *Organization Studies*, 16(6), 1995, pp. 1021-1046.
- [3] Boland, R.J. and Tenkasi, R.V., "Perspective Making and Perspective Taking in Communities of Knowing", *Organization Science*, 6(4), 1995, pp. 350-372.
- [4] Brown, J. S., "Internet Technology in Support of the Concept of 'Communities of Practice': The Case of Xerox", *Accounting, Management and Information Technologies*, 8, 1998, pp. 227-236.

- [5] Brown, J. S. and Duguid, P., "Organizational Learning and Communities of Practice: Toward a Unified View of Working, Learning and Innovation", *Organization Science*, 2(1), 1991, pp. 40-57.
- [6] Davenport, T.H., Jarvenpaa, S.L. and Beers, M.C., "Improving Knowledge Work Processes", *Sloan Management Review*, summer 1996, pp. 53-65.
- [7] Grant, R.M., "Towards a Knowledge-Based Theory of the Firm", *Strategic Management Journal*, 17 (Winter Special Issue), 1996, pp. 109-122.
- [8] Grudin, J., "Social Evaluation of the User Interface: Who Does the Work and Who Gets the Benefit?", in Bullinger, H.-J. and Shackel B. (eds.) Proceedings of INTERACT 1987, Elsevier Science Publishers, Amsterdam, pp. 805-811.
- [9] Grudin, J., "Groupware and Social Dynamics: Eight Challenges for Developers", *Communications of the ACM*, 37(1), 1994, pp. 92-105.
- [10] Hahn, J. and Subramani, M. R., "A Framework of Knowledge Management Systems: Issues and Challenges for Theory and Practice", in Proceedings of ICIS 2000, pp. 302-312.
- [11] Hedlund, G., "A Model of Knowledge Management and the N-Form Corporation", *Strategic Management Journal*, 15, 1994, pp. 73-90.
- [12] Holtshouse, D., "Knowledge Research Issues", *California Management Review*, 40(3), 1998, pp. 277-280.
- [13] Jennex, M. E. and Olfman, L., "Development Recommendations for Knowledge Management/ Organizational Memory Systems", in Sein *et al.* (eds.) Contemporary Trends in IS Development, Kluwer Academic, 2001, pp. 209-222.
- [14] Lindgren, R. and Henfridsson, O., "Using Competence Systems: Adoption Barriers and Design Suggestions", *Journal of Information & Knowledge Management*, 1(1), 2002, pp. 65-77.
- [15] Lindgren, R. and Stenmark, D., "Designing Competence Systems: Towards Interest-Activated Technology", *Scandinavian Journal of Information Systems*, 14, 2002, pp. 19-35.
- [16] Markus, L.M., Majchrzak, A. and Gasser, L., "A Design Theory for Systems that Support Emergent Knowledge Processes", *MIS Quarterly*, 26, 2002, pp. 179-212.
- [17] Nonaka, I., "A Dynamic Theory of Organizational Knowledge Creation", *Organization Science*, 5, 1994, pp. 14-37.
- [18] Schultze, U., "A Confessional Account of an Ethnography about Knowledge Work", *MIS Quarterly*, 24(1), 2000, pp. 3-41.
- [19] Schultze, U. and Boland, R.J., "Knowledge Management Technology and the Reproduction of Knowledge Work Practices", *Journal of Strategic Information Systems*, 9, 2000, pp. 193-212.
- [20] Schultze, U. and Leidner, D., "Studying Knowledge Management in Information Systems Research: Discourses and Theoretical Assumptions", *MIS Quarterly*, 26(3), 2002, pp. 213-242.
- [21] Spender, J. C., "Organizational Knowledge, Learning and Memory: Three Concepts in Search for Theory", *Journal of Organizational Change Management*, 9(1), 1996, pp. 63-78.
- [22] Stenmark, D., "Leveraging Tacit Organisational Knowledge", *Journal of Management Information Systems*, 17(3), 2001, pp. 9-24.
- [23] Tsoukas, H., "The Firm as a Distributed Knowledge System: A Constructivist Approach", *Strategic Management Journal*, 17, 1996, pp. 11-25.
- [24] Walsham, G., "Interpretative case studies in IS research: nature and method", *European Journal of Information Systems*, 4, 1995, pp. 74-81.