

A FRAMEWORK OF KNOWLEDGE MANAGEMENT SYSTEMS: ISSUES AND CHALLENGES FOR THEORY AND PRACTICE¹

Jungpil Hahn
Mani R. Subramani
Carlson School of Management
University of Minnesota
U.S.A.

Abstract

As the basis of value creation increasingly depends on the leverage of the intangible assets of firms, knowledge management systems (KMS) are emerging as powerful sources of competitive advantage. However, the general recognition of the importance of such systems seems to be accompanied by a technology-induced drive to implement systems with inadequate consideration of the fundamental knowledge problems that the KMS are likely to solve. This paper contributes to the stream of research on knowledge management systems by proposing an inductively developed framework for this important class of information systems, classifying KMS based on the locus of the knowledge and the a priori structuring of contents. This framework provides a means to explore issues related to KMS and unifying dimensions underlying different types of KMS. The contingencies that we discuss—the size and diversity of networks, the maintenance of knowledge flows and the long term effects of the use of KMS—provide a window into work in a number of reference disciplines that would enrich the utility of KMS and also open up fruitful areas for future research.

Keywords: Information in organizations, knowledge management, knowledge management systems, organizational behavior, organizational learning, system selection

1. INTRODUCTION

Organizational knowledge is now recognized as a key resource and a variety of perspectives suggest that the ability to marshal and deploy knowledge dispersed across the organization is as an important source of organizational advantage (Teece 1998; Tsai and Ghoshal 1999). Knowledge management initiatives in organizations are consequently increasingly becoming important and firms are making significant IT investments in deploying knowledge management systems (KMS). The primary focus of many of these efforts has been on developing new applications of information technology such as data warehousing and document repositories linked to search engines to support the digital capture, storage, retrieval and distribution of an organization's explicitly documented knowledge. KMS also encompass a variety of technology based initiatives such as the creation of databases of experts and expertise profiling and the *hardwiring* of social networks to aid access to resources of non-colocated individuals (Davenport et al. 1998; Pickering and King 1995).

While there is considerable energy, investment, and organizational activity to establish knowledge management initiatives, observers raise discomfoting questions related to the insufficient conception of what specific problems these solutions address within firms. For instance, in an insightful paper, Michael Zack (1999) wonders: "If managing knowledge is the solution, then what's the problem?" While this situation is problematic for practitioners, it is equally problematic for IT researchers studying these systems as the key issues that may advance our understanding of KMS are undefined and unclear.

¹We wish to thank Professor Gordon Davis for his valuable insights and assistance. We also wish to thank the members of the Twin Cities Knowledge Management Forum for their assistance in data collection.

To gain a rich understanding of phenomena related to this important emerging class of information systems, we conducted a series of semi-structured interviews with managers heading knowledge management initiatives in six leading firms in a large city in the Midwest.² The senior executives of these firms were members of an informal knowledge management forum and held periodic meetings to share experiences and learn from each other's experiences.

The intent of this paper is two-fold. First, based on our interviews with knowledge managers, we provide a framework that highlights that different approaches to knowledge management may differ fundamentally on two important attributes: the *nature of the knowledge* and the *locus of the knowledge* underlying knowledge management systems. Second, we highlight the tradeoffs involved in different KM implementations and the implications of differing choices on outcomes of KMS use in organizations. In discussing these issues, we also highlight avenues for research that may lead to a deeper understanding of the complex phenomena related to knowledge management systems. The rest of the paper is organized as follows: we provide a brief definitional overview of knowledge, knowledge management, and KMS. Section 3 presents the classificatory framework for KMS. Section 4 discusses contingencies influencing outcomes in KMS implementations and directions for research in KMS followed by Section 5, which concludes with a summary and a discussion of the contributions of the paper.

2. KNOWLEDGE MANAGEMENT, KNOWLEDGE MANAGEMENT SYSTEMS, AND KNOWLEDGE

Knowledge management is the systemic and organizationally specified process for acquiring, organizing, and communicating knowledge of employees so that other employees may make use of it to be more effective and productive in their work (Alavi and Leidner 1999). Knowledge management systems (KMS) are tools to effect the management of knowledge and are manifested in a variety of implementations (Davenport et al. 1998) including document repositories, expertise databases, discussion lists, and context-specific retrieval systems incorporating collaborative filtering technologies.

Knowledge is a complex concept and a number of factors determine the nature of knowledge creation, management, and sharing (Nonaka 1994). Drawing from prior discussions, we distinguish knowledge from data and information and view knowledge as “*fluid mix of framed experience, values, contextual information and expert insight that provide a framework for evaluation and incorporating new experiences and information*” (Davenport and Prusak 1997, p. 5). Prior research suggests the existence of different *types* of knowledge. Knowledge can be either *tacit* or *explicit* (Polanyi 1967), this attribute is also expressed as the distinction between *knowing* and *knowledge* (Brown and Duguid 1998; Cook and Brown 1999). Tacit knowledge refers to the knowledge that has a personal quality that makes it hard to articulate or communicate or, analogously, the knowing or the deeply rooted *know-how* that emerges from action in a particular context. In contrast, explicit knowledge refers to the codifiable component that can be disembodied and transmitted, a notion analogous to knowledge, the *know-what*, which can be extracted from the knowledge holder and shared with other individuals. Further, knowledge can be conceived as existing at multiple levels—not only at the individual level but also at the group and organizational levels. Organizational knowledge is created through cycles of combination, internalization, socialization and externalization that transform knowledge between tacit and explicit modes (Nonaka 1994). In light of this dynamic process of knowledge creation, linkages between individuals and groups sharing similar tasks—the *communities of practice* (Brown and Duguid 1991)—play an important role in communicating and sharing knowledge. However, communities have their own unique and context-specific vocabularies that, while facilitating knowledge exchange within the community, impede communication between them. The overlapping of understanding provided by boundary objects spanning multiple communities (Boland and Tenkasi 1995) provides a basis for communicating, sharing, resolving, and combining disparate perspectives. These issues thus have an important bearing on the choice of KMS to accomplish the access and deployment of knowledge in different contexts.

One widely adopted classification of knowledge management strategies recognizes two broad classes of strategies: personalization and codification (Hansen et al. 1999). The personalization strategy, recognizing the tacit dimension of knowledge, assumes that knowledge is shared mainly through direct interpersonal communication. On the other hand, the codification strategy assumes that knowledge can be effectively extracted and codified and thus uses a *document-to-person* approach where knowledge artifacts are then stored and indexed in databases that enable easy retrieval. Another influential source (Zack 1999; Zack 2000) casts organizational knowledge management strategies in a more abstract fashion. Zack suggests that knowledge management strategies need to be determined by the specific knowledge problems of uncertainty, equivocality, ambiguity, or complexity faced by the organization.

²Details of the methodology are provided in the appendix.

Despite the potential usefulness of previous strategies for knowledge management, our interviews with knowledge managers show that organizations still face difficulties in bridging strategy and practice. The difficulties arise because a clear link does not exist between knowledge strategy and various types of knowledge management systems. Indeed, the knowledge management literature has largely focused on general conceptual principles of knowledge management or case studies of knowledge management initiatives and systems in leading organizations without much guidance as to how to conceptualize the requirements of a knowledge management system. It is thus difficult for knowledge managers to plan knowledge management initiatives despite all the understanding from the literature about what knowledge is and how it is created, transferred, and shared. In the next section we present a framework of knowledge management systems that brings together the literature on knowledge management and the different approaches to knowledge management support via information technology.

3. A FRAMEWORK FOR KNOWLEDGE MANAGEMENT SUPPORT

The foregoing discussions suggest that two important considerations in managing knowledge are (1) where the knowledge resides and (2) the extent to which the knowledge is structured. The locus of the knowledge determines whether the KMS connects a user who has a problem or question to an artifact (e.g., a document) or directs her to a person. The level of *a priori* structure determines the extent to which KMS use imposes the burden of a *translation* or *transformation* of the problem or question to a form that corresponds to implicit logic underlying the *a priori* structure. We use these dimensions as the basis to categorize the different types of knowledge management systems currently used for knowledge management support. The framework is presented in Figure 1 with examples of currently used knowledge management systems represented in each category.

The horizontal dimension of the framework focuses on the location of the organizational knowledge resources managed by the KMS—whether the knowledge is embodied within individuals or whether it exists as externalized knowledge artifacts. The vertical dimension deals with the extent to which the KMS imposes or requires a structure, *a priori*. The location of commonly used IT based solutions for knowledge management (Davenport et al. 1998), as well as the systems adopted or developed by the organizations, are positioned in our framework (see Figure 1).

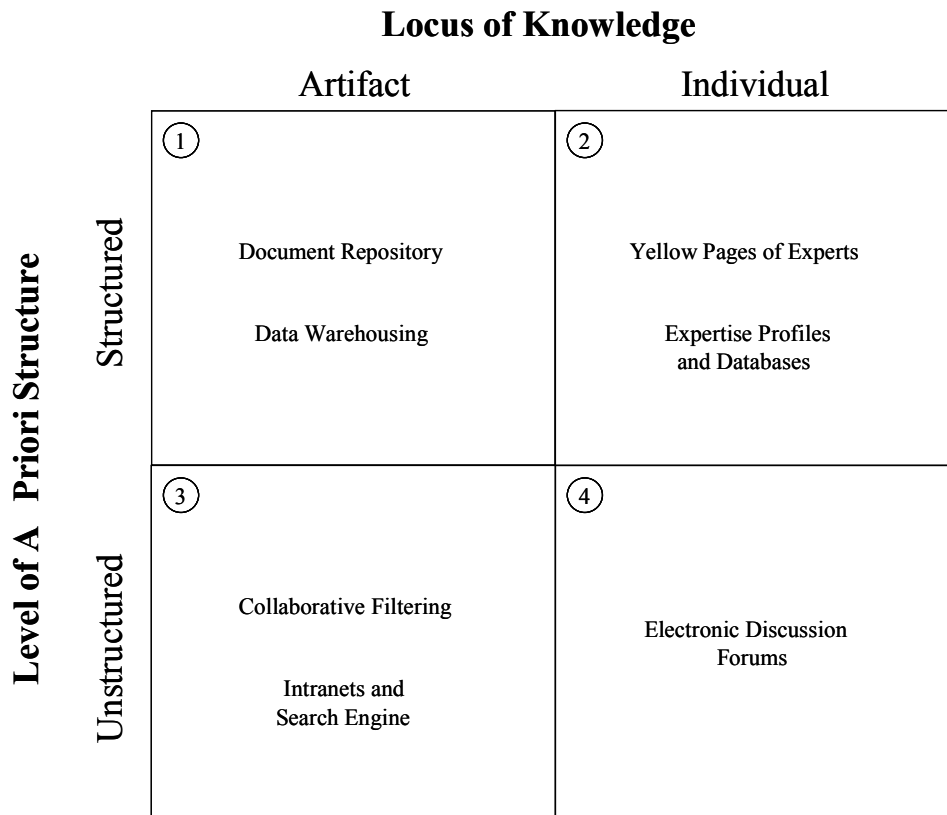


Figure 1. Framework for Knowledge Management Support

Cell 1 comprises KMS managing knowledge artifacts that have an inherent structure (e.g., enterprise wide data) or those where the KMS imposes a structure on the contents (e.g., consulting reports tagged with keywords). Essentially, the domain of these systems is restricted to the organizational knowledge that is or can be effectively codified. Document repositories and data warehousing systems fall into this category. These systems typically use database management systems (DBMS) designed to capture and store documents with predefined keywords and meta-data so that the contents can be accessed using the document categorization scheme.

Cell 2 comprises systems where the knowledge resides in individuals but the contents managed by the KMS are catalogued and structured employing *a priori* categorizing schemes. A database of experts is an instance of such a system. The contents of the KMS are created by employees filling out a questionnaire to describe their level of expertise in a predefined list of skill categories (e.g., Java programming, project management, vibration dampening, etc.). The experts database is intended to be useful for users to locate people with specific skills in domains where the user has a problem to contact them for help.

Cell 3 comprises systems where the knowledge is captured in artifacts but where the contents do not have *a priori* structures imposed on them. Instances include KMS systems incorporating document repositories that are fully indexed on the words they contain and KMS with document recommendation capabilities using *collaborative filtering* technology. The organization of the contents in such systems is dynamic and the systems aim to provide employees with relevant documents on the fly. For example, many organizations currently deploy corporate intranets so that important documents can be posted and accessed by other users browsing or searching through the site. The documents do not follow a predefined structure: search and retrieval is achieved via search engines that locate documents using full-text search. The use of collaborative filtering technology that *recommends* documents is an alternative approach to locating documents relevant to a user's question or problem without structuring contents *a priori*. The collaborative filtering system records the browse and search behaviors of users and recommends documents based on the past behaviors of other users when they performed similar searches.

Cell 4 comprises systems that provide means for users to access others who may be able to help and where the system imposes no *a priori* structure on the knowledge. In such systems, interpersonal contact enabled by the system results in knowledge sharing and transfer (Cross 1999). Instances of such systems include electronic discussion forums where employees may post questions to which other employees with answers or suggestions can post replies. Threaded discussions and e-mail distribution lists (Listservs) are typical technologies used in systems in this class.

The framework presents insights to the nature of assistance, the extent of assistance, and the limitations of different KMS. Implicit in this framework is the notion that different KMS implementations provide differing levels of support in locating knowledge and impose differing burdens on users of these systems. KMS that impose an *a priori* structure are useful as they enable contents to be categorized in a manner that is consistent with the vocabulary of the community or constituency they serve. For instance, a Class 1 KMS would return a highly relevant set of documents as the results of a query on *compiler construction* as this specific keyword would exclude irrelevant documents that dealt with compilers but not with compiler construction. However, such a KMS imposes a burden on users as retrieval of relevant content is often difficult for someone without access to the local vocabulary. A Class 3 approach that supports a free-text search of documents would result in a much larger number of documents being retrieved in response to the same search, including a number of irrelevant documents that contain *compiler* and *construction* but don't address the specific issue of compiler construction. Such a system, while providing minimal barriers to use, burdens users with the task of wading through irrelevant content to locate relevant content.

The framework thus also highlights the implicit assumptions regarding the nature of knowledge and sources of knowledge required for problem solving that underlie different KMS. A repository of information, say a Class 1 KMS, even if it contains all the information required to solve the user's problem, is unlikely to be useful if the information is not in the form that corresponds to the way the user views the problem. For instance, uploading all technical references manuals of photocopy machines into a searchable document repository is not likely to be effective when the *knowledge* required to solve problems is embedded in *war stories* chronicled within the community of practice of repair technicians (Brown and Duguid 1991; Orr 1990). Similarly, in a consulting firm that provides predefined systems solutions, the major source of knowledge would be knowledge artifacts such as project proposals and previous solutions to similar problems so that these may be reused for more efficient performance. Supporting the reuse of such knowledge artifacts through an electronic discussion forum would probably lead to simple discussions where employees post questions such as "*Has anyone recently conducted a project with XXX product within the YYY industry? If so, can you send me your proposal?*" A document repository with an efficient indexing mechanism would be more effective for users to obtain the inputs they need without burdening other employees with requests for help.

While our framework is valuable in focusing attention on the utility, limitations, and assumptions inherent in different KMS, it does not address the variations in the utility of similar systems in different contexts that we encountered in our study. For instance, why was a database of experts that was so successful in one division of an organization viewed as being less useful when

it was expanded to cover multiple divisions? Why did discussion forums that were vibrant and successful in one organization fail to generate any user interest in another? We explore some of the contextual factors that appear to play an important role in determining the utility of KMS in the following section.

4. ISSUES AND CHALLENGES

We identify issues and challenges related to the utilization of information technology for knowledge management support in three phases of deployment: the setup phase, the on-going utilization and maintenance phase and finally long-term effects of knowledge management support. In addition, we briefly discuss some issues related to the development of KMS.

4.1 KMS Setup: Balancing Information Overload and Potentially Useful Content

The reach of the knowledge management system, both in terms of size of the user group and diversity of the group, is an important consideration. This is an important issue because increases in the size and the diversity, while beneficial, involve an implicit trade-off between irrelevant content and potentially useful information (Hiltz and Turoff 1985). If usage of the system (e.g., a Class 4 e-mail list) is restricted to a small group, the users would collectively lose valuable input that may potentially be obtained by a more inclusive policy encouraging broader participation (Constant et al. 1996). However, if the reach of the system is too broad, then the system risks an overload of irrelevant information. Deciding on an appropriate size and scope is important in order to leverage the strength of weak ties (Constant et al. 1996; Granovetter 1982).

Size: When the knowledge sources are artifacts (Class 1 and 3), size may have a positive network effect. For example, the greater the number of documents in the document repository or the intranet, the higher the chances one would be able to find documents of interest. However, size may have negative network effects when the knowledge sources are individuals (Class 2 and 4). Even though greater network size may increase the potential reach of an electronic discussion forum, which in turn would increase the chances of generating an offer to help, greater network size would tend to increase information overload in the form of multiple asynchronous yet concurrent threads of conversations (Kerr and Hiltz 1982). This increases the amount of effort required to follow and participate in the discussions (Subramani and Hahn 2000), possibly reducing the overall incentive to participate. On the other hand, if the size of the network is too small, the network may offer insufficient resources to make participation worthwhile.

Diversity: When knowledge sources are highly structured (Class 1 and 2), greater diversity of content poses minimal problems as long as the indexing scheme is appropriate. However, when the knowledge sources are loosely structured (Class 3 and 4), diversity of content and domains may become problematic. In case of lower levels of diversity, the resources made available by the KMS would be limited, in some cases corresponding closely with that of the individuals in search of the target content. Hence, the knowledge management system would be less valuable. One of the managers interviewed pointed out that “*employees would not come [to the document repository] because they already knew everything that was in there ... they were wanting to see new information from other parts of the organization.*” On the other hand, diversity can also be overwhelming as it may become difficult to find the appropriate knowledge sources due to information overload and a lack of shared vocabulary. This situation is often apparent in the use of search engines on the Internet (Class 3). A simple search on the Internet can retrieve literally millions of documents from which only a very small proportion is relevant and valuable. As these observations suggest, the utility of a knowledge management system, which largely depends on its ability to provide novel and interesting knowledge resources, is likely to diminish in the presence of overwhelming size and diversity. In essence, the leveraging of the strength of weak ties need to be balanced against the downside of information overload and incongruent vocabularies.

Even though normative solutions to this problem do not yet exist, we may learn from reference disciplines (e.g., such as computer science, human-computer interaction, and information science) several ways in which this problem of information overload may be tackled. Recent research suggests innovative systems and techniques for dealing with the potential information overload associated with the use of KMS. For example, Van Dyke et al. (1999) describe the use of intelligent agent technology to filter out irrelevant content and locate potentially useful conversations on Internet relay chat (IRC). The agent builds a profile of the user through a keyword-based model and recommends current chat sessions by sampling IRC channels to find current conversations that match the user’s profiles. Another line of research employs collaborative filtering methods to recommend potentially useful Usenet news posts to users (Konstan et al. 1997; Sarwar et al. 1998). The system filters messages based on users’ previous preferences and recommends Usenet news articles a user may find useful—those that other users with similar preferences have rated as being of high quality. Finally, several techniques have been developed to visualize and represent the conversational interface of synchronous and asynchronous discussion forums to express social patterns in online conversations. Systems such as *Loom* (Donath et al. 1999) and *PeopleGarden* (Xiong and Donath 1999), provide a visual thumbprint of

asynchronous threaded discussion groups by representing social patterns of interaction such as who starts discussions, who provides responses to questions, and how lively the discussion is. This visualization of the conversation interface allows users to gain an overall portrait of the discussion forums before spending time examining any content. *Chat Circles* is a visual interface for synchronous online chat (Donath et al. 1999). This system conveys locality of conversations within a chat room by incorporating a space dimension. The users can only *hear* conversations that occur in her proximity, yet she may be aware of other peripheral conversations that occur simultaneously in the periphery. The common theme underlying these systems is to provide a technological solution to prevent users from being overwhelmed by information overload and yet allow the size and scope of the system to increase so that the probability of users locating potentially useful content within the system is enhanced. We suggest that the utility of knowledge management systems would be considerably improved by incorporating these features.

4.2 KMS Maintenance: *Balancing Additional Workload and Accurate Content*

In a knowledge management system, maintaining a steady stream of contributions to content, the knowledge flow (Holtshouse 1998), is an important issue. Thus a critical problem is one of motivation: motivating users to contribute. This may be more or less difficult depending on the level of structuring of the content in the KMS. When the knowledge content is highly structured (Class 1 and 2), a great deal of effort is required up front to ensure the appropriate structuring. Employees are required to append appropriate keywords and meta-data to their documents prior to the uploading to the document repository or to fill out extensive skills and expertise questionnaires. While this may appear to be a simple, straightforward task, serious difficulties of knowledge representation and problems of motivation may exist. When employees are appending keywords and meta-data tags to their documents or are evaluating themselves in terms of the skills or expertise they possess, they are in essence creating answers for questions that have not yet been posed and are implicitly addressing the issue *to what queries should this document (or user profile) be a result?* Furthermore, motivation is also problematic since extra effort and time required for structuring contributions need to be allocated in addition to their regular job tasks and employees do not know *a priori* (and in some implementations, even *ex post*) if their document (or profiles) are later retrieved and used.

This is less of a problem when the knowledge sources are loosely structured on an ongoing basis (Classes 3 and 4). Participants in an electronic discussion forum need only respond when questions are asked. Since response behavior is usually highly visible in such forums, participants may be greatly motivated to provide answers whenever possible since frequent responses provide greater visibility, engendering a higher social status (Subramani and Hahn 2000). Similarly, maintaining the knowledge resources is straightforward in the case of knowledge artifacts. Simply posting documents on the corporate intranet does not require extensive effort up front to ensure the appropriate description of the document so long as the search algorithm is effective.

One potential consequence of increasing access to domain experts through KMS is that the small set of recognized experts may end up executing a significant portion of the work for users who seek their advice. If this shifting of the burden strains the experts' time and effort, it may result in the experts' unwillingness or reduced motivation to respond to inquiries and contribute in future periods.

The issue of motivation and incentives recurred frequently in discussions of the utility of knowledge management systems and appears to be a critical issue needing further research. In many cases, using straightforward monetary incentives does not appear to be an effective solution. As one manager commented in our interviews, "*providing monetary benefits for uploading content was not effective ... we ended up with a large repository full of garbage.*" It is likely that alternative approaches combined with monetary incentive methods may be more effective. For example, we frequently observe in electronic forums, such as Usenet newsgroups or mailing lists, that there are people who contribute greatly to the group by providing a lot of high quality content even though these groups do not offer monetary benefits. One potential driver behind such altruistic response behavior may be social status or self-esteem (Constant et al. 1996; Sproull and Kiesler 1991). Research to understand such altruistic response behaviors would be beneficial so that it may be leveraged to effectively maintain knowledge flow in KMS. Technical approaches to solve the problem also show promise. For example, *Answer Garden*, an organizational memory system, archives answers to frequently asked questions so that domain experts need not provide answers to the same questions over and over again (Ackerman 1998). Norms evolved in such forums (e.g., the norm that users examine the FAQ before posting a question) may be another means to minimize overloading experts (Subramani and Hahn 2000). Without such support, domain experts may be reluctant to provide answers because providing an answer would attract many subsequent questions. Evolving devices in KMS to shield experts would motivate the experts to contribute because access to them could be reserved only for interesting and challenging questions to which prior answers are unavailable. This is an area for future research in KMS.

Current research in computer science suggests technological solutions to the problems that have been explored. For instance, Honda et al. (1997) describe a prototype virtual office environment that captures a user's current degree of concentration through keystrokes and chair movement and uses this as the basis to guard against interruption. In KMS, a similar scheme could be used to route requests for help only when an expert is not currently occupied.

Technological solutions are not the only possibility. Our interviews with knowledge managers inform us that social and management measures are also useful in dealing with the problem of motivation and increased burden on domain experts. For example, one company participating in our interview adopts a management approach whereby domain experts in different engineering groups take turns manning the Usenet forum. This way, individual experts are able to concentrate on their own work when it is not their designated shift for monitoring requests for help. When designated employees receive a request for advice they cannot handle, they forward that request to an employee that has the expertise, who can later respond during her designated duty shift. Anecdotal evidence shows that such managerial solutions have worked very well in motivating experts to share expertise while at the same time minimizing the burden on workload.

Another managerial practice to solve the knowledge repository maintenance problem without putting the burden on knowledge workers is the creation of a new role within the organization: the knowledge librarian. The knowledge librarian is responsible for transferring content into the knowledge repository by tagging user submissions with appropriate keywords or meta-data tags. While this solves the problem of increased burden on domain experts and users, this may also create new problems. Since knowledge librarians are not the individuals that actually create the knowledge objects (e.g., documents, reports, videos), even though they may have familiarity with the domain, they may not have an accurate first-hand understanding of the content. Hence, the keywords and meta-data tags appended by the knowledge librarian may be inappropriate and as a result subsequent queries searching for a knowledge resource may not retrieve the right document.

4.3 Long-Term Effects of KMS: Balancing Exploitation and Exploration

A final issue we raise is the long-term impact of the use of KMS on learning, innovation, and expertise development. The availability of existing solutions may bias employees to adopt existing solutions rather than search for or develop novel solutions that may be more effective. In the long run, reliance on existing solutions may result in competency traps that inhibit organizational learning and innovation (Levitt and March 1988). The literature on organizational learning suggests that experience plays a vital role in the learning process. In the case of knowledge management systems, experience will be tightly tied with the use of the system. Thus, in the long run, knowledge workers may gain extensive experience at *assembling* knowledge components to solve problems instead of actually *creating* the knowledge components themselves, an important consideration related to expertise development in knowledge intensive firms that tend to promote heavy use of KMS. Further, the presence of KMS may predispose users to the use of explicit and easily available information rather than tacit knowledge that may be more effort intensive to access on account of the stickiness of such knowledge (von Hippel 1998).

These arguments indicate that the use of KMS may lead to both positive as well as negative outcomes. An organization may gain great efficiency by streamlining problem solving processes through reuse of knowledge, yet may slowly become rigid and lose its capacity to learn locally and innovate. The literature on organizational learning illustrates the trade-off between exploitation and exploration for organizational learning (March 1991). Exploitation of old certainties may be effective for short-term benefits but self-destructive in the long run. The use and re-use of existing solutions encouraged by knowledge management systems may bias organizational processes toward exploitation in ways that may ultimately prove detrimental to the organization. While it may be too early to observe such long-term effects of knowledge management support on organizations, the repercussions on the organization's capacity to learn and innovate are too great to be ignored. We believe that researchers should pay careful attention to these long-term effects.

4.4 KMS Development

Our discussions of issues related to the use of KMS create the possibility that the development of KMS may require attention to some issues particular to this class of information systems. The uniqueness of the knowledge management context may make it difficult for traditional systems development methodologies to be directly applicable. The nature of adaptations required to the traditional systems development and deployment methods in the context of knowledge management systems is, therefore, an important issue.

One difference between KMS and traditional IS (e.g., transaction processing systems, management information systems, etc.) is the often fluid *ex ante* nature of the objectives, outcomes, and processes of the system to be developed. Within the knowledge management context, it is difficult to know a priori what information will be requested, who will request the information, who will supply the information, and when and how the information will be used. This makes determining requirements for KMS development extremely difficult.

First, it is difficult to define in advance the profile of a *typical* user. Implicit in systems development methodologies is the assumption that users of the system to be developed are homogeneous with respect to the objectives in using the system. Hence,

systems development methodologies generally sample typical or representative users to determine requirements and perform user-testing. However, this approach may be limiting in the knowledge management context.

Second, in the knowledge management context, the final outcome of development efforts needs to be flexible. Due to difficulties in defining a priori structures for knowledge and the need for ongoing refinement of these structures, traditional approaches to systems development, such as those adopted for organizational databases for instance, may not be appropriate. For example, with traditional systems, the domain data is structured into a conceptual schema that is translated into a data model. The resulting data model rarely changes and is shared throughout the organization. This may no longer be true since knowledge schemas constantly change over time and fundamentally different schemas may be used in different parts of the organization. Therefore, systems development in the knowledge management context may require the application of development methodologies used for dynamic/multiple schema databases (or knowledge bases). With constantly changing and multiple knowledge schemas, interoperability becomes a major issue.

Another implicit assumption of traditional systems development methodologies is that the system developed is a final product. Even though a developed system goes through significant maintenance, the basic philosophy is passive maintenance where errors are only corrected when highlighted. The difficulty in the knowledge management context is that, due to the ill-defined nature of the goals/processes of the KMS, a normative criterion with which to evaluate whether the KMS is being effectively used or means to identify factors that impede desired usage are not available. Hence, rather than a final product-oriented approach, an evolutionary approach to systems development may be required. With an evolutionary approach, the basic KM platform is initially developed and evolves proactively in an on-going fashion. Since a normative model of KMS use cannot be determined, experimentation with novel features will become extremely important.

Finally, in the knowledge management context, one of the most important factors for success is motivating users to use the system. A KMS can be regarded as analogous to a tool. A tool is only successful if the users of the tools succeed with the tool. Thus, system acceptance becomes critical. Previous research shows that an effective way to increase system acceptance is user involvement during systems development (Tait and Vessey 1988). In such a situation, participatory design (PD) (Kensing and Munk-Madsen 1993) may be an appropriate approach.

5. CONCLUSION

As the basis of value creation increasingly depends on leveraging firms' intangible assets, knowledge management systems (KMS) are emerging as powerful sources of competitive advantage. However, the general recognition of the importance of such systems seems to be accompanied by a technology-induced drive to implement systems with inadequate consideration of the fundamental *knowledge problems* that the KMS are likely to solve.

This paper contributes to the stream of research on knowledge management systems by proposing an inductively developed framework for this important class of information systems, classifying KMS based on the locus of the knowledge and the a priori structuring of contents. This framework provides a means to explore issues related to KMS and unifying dimensions underlying different types of KMS. The contingencies that we discuss the size and diversity of networks, the maintenance of knowledge flows, and the long term effects of the use of KMS provide a window into work in a number of reference disciplines that would enrich the utility of KMS and also open up fruitful areas for future research.

References

- Ackerman, M. S. "Augmenting Organizational Memory: A Field Study of Answer Garden," *ACM Transactions on Information Systems* (16:3), 1998, pp. 203-224.
- Alavi, M., and Leidner, D. E. "Knowledge Management Systems: Issues, Challenges, and Benefits," *Communications of the AIS* (1:Article 7), 1999.
- Boland, R. J., Jr., and Tenkasi, R. V. "Perspective Making and Perspective Taking in Communities of Knowing," *Organization Science* (6:4), 1995, pp. 350-372.
- 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.
- Brown, J. S., and Duguid, P. "Organizing Knowledge," *California Management Review* (40:3), 1998, pp. 90-111.
- Constant, D., Sproull, L., and Kiesler, S. "The Kindness of Strangers: The Usefulness of Electronic Weak Ties for Technical Advice," *Organization Science* (7:2), 1996, pp. 119-135.

- Cook, S. D. N., and Brown, J. S. "Bridging Epistemologies: The Generative Dance Between Organizational Knowledge and Organizational Knowing," *Organization Science* (10:4), 1999, pp. 381-400.
- Cross, R. "More than an Answer: How Seeking Information Through People Facilitates Knowledge Creation and Use," Working Paper School of Management, Boston University, 1999.
- Davenport, T. H., De Long, D. W., and Beers, M. C. "Successful Knowledge Management Projects," *Sloan Management Review*, Winter 1998 1998, pp. 43-57.
- Davenport, T. H., and Prusak, L. *Working Knowledge: How Organizations Manage What They Know*, Cambridge, MA: Harvard Business School Press, 1997.
- Donath, J., Karahalios, K., and Viegas, F. "Visualizing Conversation," *Proceedings of the Thirty-second Annual Hawaii International Conference on System Sciences (HICSS 32)*, Los Alamitos, CA: IEEE Computer Society Press, 1999.
- Granovetter, M. "The Strength of Weak Ties: A Network Theory Revisited," in *Social Structure and Network Analysis*, P. Marsden and N. Lin (eds.), New York: John Wiley and Sons, 1982, pp. 105-130.
- Hansen, M. T., Nohria, N., and Tierney, T. "What's Your Strategy for Managing Knowledge," *Harvard Business Review*, March-April 1999, pp. 106-116.
- Hiltz, S. R., and Turoff, M. "Structuring Computer-Mediated Communication System to Avoid Information Overload," *Communications of the ACM* (28:7), 1985, pp. 680-689.
- Holtshouse, D. "Knowledge Research Issues," *California Management Review* (40:3), 1998, pp. 277-280.
- Honda, S., Tomioka, H., Kimura, T., Ohsawa, T., Okada, K., and Matsushita, Y. "A Virtual Office Environment Based on a Shared Room Realizing Awareness Space and Transmitting Awareness Information," in *Proceedings of the Tenth Annual ACM Symposium on User Interface Software and Technology*, Banff, Canada, 1997, pp. 199-207.
- Kensing, F., and Munk-Madsen, A. "PD: Structure in the Toolbox," *Communications of the ACM* (36:4), 1993, pp. 78-85.
- Kerr, E. B., and Hiltz, S. R. *Computer-Mediated Communication Systems: Status and Evaluation*, New York: Academic Press, 1982.
- Konstan, J. A., Miller, B. N., Maltz, D., Herlocker, J. L., Gordon, L. R., and Riedl, J. T. "GroupLens: Applying Collaborative Filtering to Usenet News," *Communications of the ACM* (40:3), 1997, pp. 77-87.
- Levitt, B., and March, J. G. "Organizational Learning," *Annual Review of Sociology* (14), 1988, pp. 319-340.
- March, J. G. "Exploration and Exploitation in Organizational Learning," *Organization Science* (2:1), 1991, pp. 71-87.
- Nonaka, I. "A Dynamic Theory of Organizational Knowledge Creation," *Organization Science* (5:1), 1994, pp. 14-37.
- Orr, J. "Sharing Knowledge, Celebrating Identity: Community Memory in a Service Culture," in *Collective Remembering*, D. Middleton and D. Edwards (eds.), London: Sage Publications, 1990, pp. 168-189.
- Pickering, J. M., and King, J. L. "Hardwiring Weak Ties: Interorganizational Computer-Mediated Communication, Occupational Communities, and Organizational Change," *Organization Science* (6:4), 1995, pp. 479-486.
- Polanyi, M. *The Tacit Dimension*, Garden City, NY: Doubleday, 1967.
- Sarwar, B. M., Konstan, J. A., Borchers, A., Herlocker, J., Miller, B., and Riedl, J. "Using Filtering Agents to Improve Prediction Quality in the GroupLens Research Collaborative Filtering System," in *Proceedings of the ACM Conference on Computer Supported Cooperative Work (CSCW '98)*, Seattle, WA, 1998, pp. 345-354.
- Sproull, L., and Kiesler, S. *Connections: New Ways of Working in the Networked Organization*, Cambridge, MA: The MIT Press, 1991.
- Subramani, M. R., and Hahn, J. "Examining the Effectiveness of Electronic Group Communication Technologies: The Role of the Conversation Interface," MISRC Working Paper 00-03, Carlson School of Management, University of Minnesota, May 2000.
- Tait, P., and Vessey, I. "The Effect of User Involvement on System Success: A Contingency Approach," *MIS Quarterly* (12:1), 1988, pp. 91-108.
- Teece, D. J. "Capturing Value from Knowledge Assets," *California Management Review* (40:3), 1998, pp. 55-79.
- Tsai, W., and Ghoshal, S. "Social Capital and Value Creation: The Role of Intrafirm Networks," *Academy of Management Journal* (41:4), 1999, pp. 464-476.
- Van Dyke, N. W., Lieberman, H., and Maes, P. "Butterfly: A Conversation-Finding Agent for Internet Relay Chat," in *Proceedings of the 1999 International Conference on Intelligent User Interfaces*, Redondo Beach, CA, 1999, pp. 39-41.
- von Hippel, E. "Economics of Product Development by Users: The Impact of "Sticky" Local Information," *Management Science* (44:5), 1998, pp. 629-644.
- Xiong, R., and Donath, J. "PeopleGarden: Creating Data Portraits for Users," in *Proceedings of the Twelfth Annual ACM Symposium on User Interface Software and Technology*, Asheville, NC, 1999, pp. 37-44.
- Zack, M. H. "Developing a Knowledge Strategy," *California Management Review* (41:3), 1999, pp. 125-145.
- Zack, M. H. "If Managing Knowledge is the Solution, Then What's the Problem?," in *Knowledge Management and Business Model Innovation*, Y. Malhotra (ed.), Hershey, PA: Idea Group Publishing, 2000.

Appendix

Research Methodology³

The orientation of the research project was to enrich our understanding of knowledge management projects and the activities that are associated with success and the delivery of value through knowledge management systems in organizations.

We conducted semi-structured interviews of managers heading the knowledge management function in eight large organizations in a large city in the Midwest. The managers at these firms were members of an informal group of senior managers who met periodically to share experiences and leverage each other's expertise. The organizations in this informal group were large multidivisional firms with national or global operations who had created formal knowledge management functions within the organization. The interviews were conducted in the offices of the managers and, in most cases, the managers arranged for a demonstration of the system so that we could see the features of the KMS implementation. The length of the interviews varied between one and a half hours to about three hours. At least two researchers participated in the interviews and took detailed notes. All interviews were taped and later transcribed.

Details of the organizations and the knowledge management systems in the organization are provided below. The names of the organizations are masked to preserve confidentiality.

1. PowerCo: Division of a large national utility company engaged in both power generation and power distribution. KMS implementation in the power generation division that operates a mix of conventional coal and oil fired power generating plants as well as nuclear power plants. KMS consisted of a directory of about 500 technical employees who had voluntarily entered details of their expertise into the KMS. The KMS could be searched using keywords. An interesting feature of the search engine was a thesaurus of technical terms; the engine would automatically substitute synonyms of the keywords entered by the user and return a richer set of results than provided by the keywords. The system was viewed as being successful in lowering the *mean-time-to-repair* of generating equipments in power plants. However, the enlargement of the scope of the system to cover specialists in other divisions of the firm has been less successful.
2. FoodCo: Leading consumer food company, division of large global corporation, producing a variety of prepared goods sold through grocery chains. KMS implementation was built around a Lotus Notes platform and was intended to streamline the interface between the global research and development group and manufacturing plants worldwide. The KMS was originally designed to provide the ability to search technical reports produced by research and development groups as well as to attempt to link non-colocated researchers and technicians. A variety of later projects added a range of functionality to the original system, broadening the scope. The system was viewed as having enabled the company to improve *time-to-market* for new products.
3. SoftCo: Software development company headquartered in Boston that had grown rapidly, largely fueled by mergers and acquisitions of smaller firms to an organization with 5,500 consultants distributed in over 90 offices worldwide. Software consultants in offices all had a large variety of backgrounds, skills, and expertise and KMS was viewed as a means to enable employees to access expertise located outside the local offices as well as encourage re-use of software modules such as standard reporting modules developed for clients, etc. The KMS was built in-house around a discussion database backed by a document repository. The KMS was not adopted by a majority of consultants and failed to become a significant component of routine resources used by them. The system was abandoned 12 months after launch and the resources diverted to other projects.
4. CompCo: Global IT consulting company providing turnkey hardware and software solutions worldwide. The primary purpose of the knowledge management initiative at CompCo was to create a repository of content for globally distributed marketing groups to improve the chance of winning bids by creating teams comprising non-colocated team members to address customer requirements in different industry segments. The secondary purpose was to augment the internal learning and training programs of the training department by accessing resources located worldwide within the corporation. The KMS was built around a collaboration environment that was developed in-house and based on the metaphor of the "team-room." It provided facilities for collaboration among team members in such a manner that the content of their interactions as well

³Due to constraints of space, the discussion of our methodology is highly abbreviated. A detailed description is available on request from the authors.

as the product of their collaboration could be captured within the system. At the end of the project, each team was required to “clean-up” the materials that had been created over the life of the project and submit it to a central repository. The central repository was fully indexed and searchable and returned documents as well as details of individuals who had created the content. There was a “scorecard” created for each customer so that it allowed a coordinated picture of services provided to global customers through different offices of the firm. The KMS effort was considered only modestly successful; the repository grew rapidly but the quality of submissions varied considerably. User feedback indicated that information contained was often out-of-date or not usable without extensive effort, leading users to view the KMS as an information source of last resort.

5. DataCo: Leading software firm manufacturing and marketing a variety of software products. KMS implementation was in the consulting division aimed at reusing solutions created for specific customer projects across a wider customer base. KMS comprised an archive of project documentation and client reports accessible through a free text search. The implementation was considered unworkable by consultants because of the very large number of results returned by searches for commonly used terms. An attempt to structure the existing repository by hiring professional librarians met with limited success because of the highly specialized and technical nature of documents. Proposals to require the *a priori* structuring of submissions by consultants encountered considerable resistance by users. The KMS continues to be supported and maintained but is not viewed as being successful.
6. MainframeCo–Aviation Division: Division of large computer manufacturer providing specialized large-system solutions to vertical industry segments worldwide. KMS created in aviation division was built around the Lotus Notes platform to let non-colocated experts collaborate on project teams. Documents generated in the project were archived for retrieval in keyword searches. The project was considered highly successful in transferring best practices across different geographies and standardizing processes across projects executed by teams in different countries.
7. MainframeCo–Development Labs: Division of a large computer manufacturer providing hardware and software development services to other divisions. KMS implemented to create a database of specialists to improve staffing decisions for hardware and software projects. Expertise profiles of research staff were created based on voluntary self reported expertise. KMS rollout was part of ERP implementation and managed by the HR Division. The project is in the early phases and considered promising.