Submitted to I-KNOW'03 Workshop on (Virtual) Communities of Practice within Modern Organizations,
Graz, Austria, July, 2003

Organic Perspectives of Knowledge Management

Jonathan Ostwald[1], Koichi Hori[2], Kumiyo Nakakoji[2,3], Yasuhiro Yamamoto[2,3]

[1]Department of Computer Science, University of Colorado at Boulder
[2]Research Center for Advanced Science and Technology, University of Tokyo
[3]PRESTO, JST

ostwald@cs.colorado.edu, hori@ai.rcast.u-tokyo.ac.jp, kumiyo@kid.rcast.u-tokyo.ac.jp, yxy@acm.org

1 Introduction

The contribution of the paper is a framework for integrating organization and community perspectives on knowledge management and its computational support.

Our research on knowledge management is rooted in the community perspective. We believe that knowledge systems should serve primarily to help people create and share new knowledge. But we also acknowledge the role of stable, structured and reliable information, both as a component of our systems and as a component of the organizations within which we work.

This paper will introduce, compare and contrast two prototype knowledge systems as it presents a more general framework for thinking about how community of practice and organizational perspectives to knowledge management might be integrated. The first prototype is called the *Knowledge Nebula Crystallizer (KNC).* The KNC is a demonstration of concept developed at the University of Tokyo. It is used in this paper to describe basic knowledge system functionalities and components. The second prototype is the *livingOM* (for Living Organizational Memory), a web-based knowledge system developed at the University of Colorado at Boulder. The livingOM shares the basic functionalities of the KNC, but differs in its implementation and application. The role of the livingOM in this paper is to illustrate how aspects of both community of practice as well as organization perspectives can live together within a knowledge system.

2 Knowledge nebula crystallizer (KNC)

The Knowledge Nebula Crystallizer is a computational prototype that demonstrates several fundamental knowledge management processes including knowledge reuse, creation, and evolution. In the scenario below we describe how an idealized version of the prototype might work.

2.1 Scenario

Jane is a graduate student who is developing her initial ideas for her master's thesis. Her first step is to look, from her own viewpoint, into the work her research group has done in the past. She logs into the Knowledge Nebula Crystallizer (KNC), and tells the system of her interests. She also uploads a 2-page description of her thesis topic to further describe what information she is after.

The screen goes blank except for a small logo signifying that the KNC is working to collect and *crystallize* (compile in various possible contexts) information relevant to Jane's research interests. After about 10 seconds the screen is filled with several different editions of reports summarizing the relevant work selected from among the entire body of work done by the research group; each edition shows her a different possible context or a viewpoint discussing the same topic. The reports are organized from the papers her group has produced that are relevant to her interests. The KNC also provided a list of outside references that were most often cited by her group as well as a list of the most prominent researchers in her field, according to the possible different contexts.

The KNC has not written Jane's thesis, but it has certainly given her a good start on the related work section, and stimulated her to consider possible different contexts to discuss the same topic. It not only saved her substantial time in collecting the material she needs for her own research, but also triggered creating new knowledge.

We now jump ahead 6 months. Jane has finished her thesis. She once again logs into the KNC only this time she is contributing knowledge rather than consuming it. The KNC provides a simple form into which Jane uploads her thesis as a Word document. The screen KNC again goes blank except for the "working" logo. After about 30 seconds a message appears on the screen thanking Jane for her contribution. The KNC has successfully integrated Jane's thesis into its knowledge base, and is now better able to help future users who request information related to Jane's topic. The KNC may now give the new user the parts of information Jane gave combined with other information under possibly new contexts which are different from the context of Jane's master thesis.

2.2 Components and Operations of the KNC

The above scenario was a fictionalized account of how the Knowledge Nebula Crystallizer might contribute to knowledge sharing and knowledge construction in a particular organization or community. The scenario depicted two interactions with the user. In the first it output reports based on Jane's query and in the second it accepted Jane's finished thesis. These interactions involve the two essential operations of the KNC, *Crystallization* and *Liquidization*, and the repository of the KNC, called *the Knowledge Nebula*.

The knowledge nebula is an unstructured collection of small information elements. During Crystallization, elements from the nebula are selected and structured according to a particular context, resulting in a new information artifact. During Liquidization, an information artifact is chunked into elements which are added to the knowledge nebula. These operations are depicted in Figure 1.

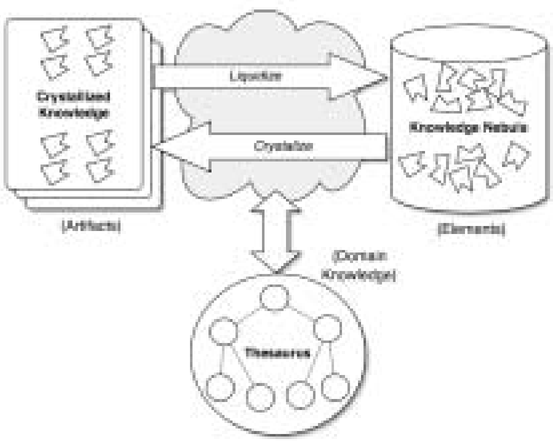

Figure 1: Components and Operations of the KNC

The third system component shown in Figure 1 is the *Thesaurus*, which serves as a solvent for liquidation (helping to chunk artifacts into elements) and as a catalyst for crystallization (helping to retrieve and structure elements into artifacts). Because the thesaurus determines the relation between elements and artifacts we call it *Domain knowledge.*

In the scenario, the knowledge nebula evolved as Jane input her thesis into the system. In this sense, the KNC is able to grow and improve over time as the result of the work of users. The scenario does not tell us, however, how the knowledge nebula came to be in the first place, nor how the thesaurus was created and by whom it was created.

3 Knowledge Management Process Framework

To extend the scope of our discussion to consider the long-term evolution of knowledge management systems such as KNC, we will use the seeding, evolutionary growth, reseeding (SER) Model (Figure 2) [Fischer, 1998]. The SER model provides a conceptual framework to understand how decentralized evolution, such as that resulting from Jane's interaction with the KNC, can be initiated and then sustained over time. The model was developed to understand the balance between centralized and decentralized evolution in sustained development of large systems. Its goal is to apply lessons learned from success cases, such as the open source software, to domains and communities, such as Knowledge

Management, that have not traditionally been viewed in from this perspective [Fischer & Ostwald, 2001].

The SER model describes three phases of evolution in terms of the stakeholders involved and their activities. The *seeding* phase creates the initial conditions for the use of a knowledge system, such as the KNC. The use of the system is the driving force of the *evolutionary growth* phase. *Reseeding* is a periodic effort to organize and tune the knowledge system. In this paper we will consider only the first two phases of the SER Model, which are now described in more detail.

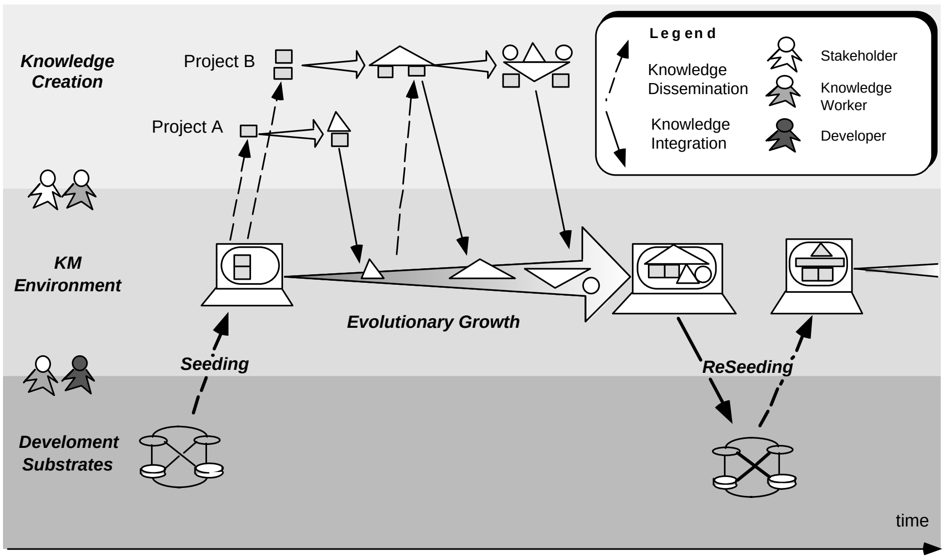

Figure 2: The Seeding, Evolutionary Growth, Reseeding (SER) Model

3.1 Seeding

As the name suggests, the seed is considered as a starting point for ongoing growth. Rather than chasing the impossible goal of complete coverage, the seed can be initially *underdesigned* [Brand, 1995], meaning that the system developers do not create final solutions but rather design spaces that can be changed and modified by knowledge workers at use time. Although the SER model acknowledges that the initial seed cannot be complete, the seeding process still requires a substantial up-front investment.

The KNC seed consists of a thesaurus and an initial nebula. In the current prototype, the nebula seed is created by first collecting artifacts from the target domain, such as existing documents or transcripts of interviews with workers about their domain. This collection of artifacts is then chunked into elements that are at approximately the paragraph level (about four sentences long), each containing a single idea. The thesaurus is a list of keywords with weights for calculating contexts . Currently, both the chunking operation and building the initial set of keywords are performed by domain experts, though the weights are automatically calculated. A future research issue would be to assess whether these manual operations could be automated.

The final step of the seeding phase is to compute a vector for each chunk in the nebula. These vectors represent the "semantic distance" calculated according to the thesaurus, and provide the KNC a way to determine relations between chunks, as well as between chunks and a query, during the crystallization process [Hori 1994]. As such, the vectors represent the system's domain knowledge [Hori 1997].

3.2 Evolutionary Growth

During this phase, the knowledge system simultaneously informs work, and accumulates the products of work. These roles are depicted as arrows in Figure 2, and correspond to the crystallization and liquidization processes depicted in Figure 1.

An essential aspect of the evolutionary growth phase is that the user community is responsible for driving the evolution of the seed. In the KNC, for example, the artifacts created as a result of crystallization are not considered to be complete, but rather they are starting points for further refinement. Because the knowledge nebula is not assumed to hold all the chunks necessary to form a new artifact, nor is it assumed to have the intelligence to assemble them as a human would, authors are expected to modify the artifacts, and create new chunks to better suit their purpose.

It is such contributions of new information by users that drives the evolution of the knowledge nebula. In the scenario, Jane contributed to the growth of the nebula by contributing her finished thesis. This contribution increased the size and quality of the nebula, but we are not sure from the scenario if her contribution resulted in any change in the underlying domain knowledge in the thesaurus.

The nebula and thesaurus of the KNC are two levels of knowledge that might evolve in the evolutionary growth phase. Another level is the functionality of the system. Making contributions to the knowledge system should be a part of everyone's job. But formalization of information and modification of system functionality may require significant programming knowledge and therefore will be the responsibility of *power-users* [Nardi, 1993], who are technically inclined and motivated to do this work.

The SER model assumes that some elements of an unselfconscious culture of design [Alexander, 1964] will emerge in the user community, and depending on the strength of this culture, the evolutionary growth phase may last for an extended period of time. However, there are limits to such decentralized evolution [Fischer & Ostwald, 2001], and eventually the usefulness and usability of the knowledge system will suffer. When this happens, it is time for developers to come back into the picture to reseed the KM environment.

4 Living Organizational Memory

The KNC a demonstration of concept belonging more to the CoP perspective of knowledge management than the perspective of traditional organizations. In this section we present another prototype, named livingOM, which shares many of the fundamental aspects of the Knowledge Nebula Crystallizer, but which also contains elements that are more oriented toward the organizational perspective.

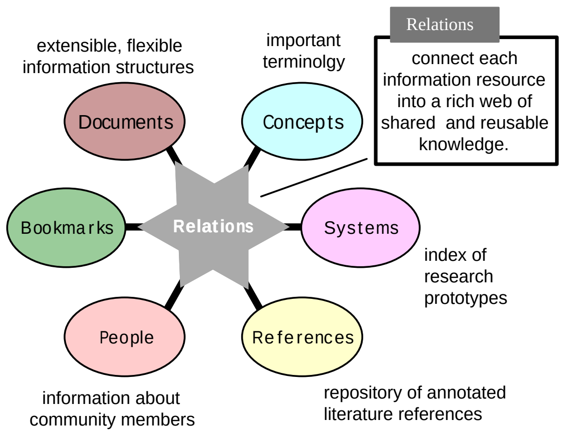

Figure 3: Domain Knowledge in the livingOM

The livingOM (living Organizational Memory) supports creative work within a community of knowledge workers who create, collect, and reuse information resources. It is living in the sense that it grows and improves as it is used to do work and solve problems. The livingOM is a follow-on to the DynaSites system [Fischer & Ostwald, 2001; Fischer & Ostwald, 2002], and has been conceptualized and developed as an organizational memory serving a mid-sized research group. It is designed to both support work practices and accumulate the products of work, such as research papers, glossaries of terms and information about group members and research prototypes.

4.1 Comparing the livingOM to KNC

Like the KNC, the livingOM manipulates information at both *artifact* and *element* levels of granularity. Artifacts in the livingOM are hypertext documents, whose elements (such as chapters, sections and paragraphs) are structured hierarchically. A typical way to create a document in the livingOM is to import a document written in MS-Word. The Word document is automatically liquidized into elements and stored in the livingOM. Unlike the KNC the elements are not through of as a nebula, although cross-document searches are supported which return lists of elements.

The livingOM's domain knowledge plays a similar role to the domain knowledge in KNC. Namely, it allows the system to collect and structure information in response to a specific situation, or query. The representation of domain knowledge, however, is different. Where the KNC represented domain knowledge by high-dimensional vectors, the livingOM represents domain knowledge as typed objects connected by relations as illustrated in Figure 3.

Relations can be either explicitly created by users or automatically created by the system. For example, a livingOM document contains relations to resources such glossary terms, citations, prototypes, and people that are displayed as hypertext links embedded in the text. When users traverse one of these links they access detailed information about the particular resource (such as the definition of a glossary term) as well as links to other related resources throughout the livingOM information space.

Table 1: Comparing KNC and livingOM

	KNC	**livingOM**
Artifacts (crystallized Elements)	Documents stored by users and ones dynamically created	Document stored by users and views that link related parts of documents and domain knowledge
Elements (liquidized Artifacts)	Nebula – information chunks created from documents and interview transcripts (but they could be created in anyway)	Word documents are automatically decomposed to the paragraph level (but they could also be created piece by piece)
Domain Knowledge	Thesaurus – Hand crafted set of keywords with automatically calculated weight values. There is also knowledge of people	Several types of information, including Glossary, references, people, research prototypes

4.2 LivingOM and the SER Model

4.2.1 Seeding

The data is highly structured and of high quality (meaning complete and probably trustworthy). But the point of seeding is not just to collect and represent get a bunch of good data — it is to lay the ground work for creating more knowledge.

We have found it effective to begin with a community's existing information repositories and tools, and to incrementally create prototypes that help developers and users to understand how their old information and technology can be cast into the new framework. This approach creates boundary objects that are familiar to users, allowing them to participate fully in the seeding process [Ostwald, 1996].

4.2.2 Evolutionary Growth

Ideally the contents of the livingOM are accumulated as a natural consequence of doing work. This is a community perspective of work as a knowledge-generating activity (the organizational perspective is that knowledge is only an *input* into work).

5 Discussion

As Table 2 illustrates, organizations and community of practice seem to have different perspectives on knowledge.

Table 2: What is knowledge? Two perspectives

Organization	Community of Practice
Static	Dynamic
Information	Process
Context Independent	Context Sensitive (situated)
Produced by doing work	An input into work

By looking at the SER model from the two perspectives, such as what drives growth, what it would mean to seed and reseed, or who is doing what, we have come to a conclusion that organizations and communities are not always so far apart.

Modern organizations need to see themselves not only as the keepers of knowledge, but also as creators of knowledge. And communities of practice, likewise, have a real interest in sustaining their practice and providing some centralized resources to their users.

What both types of organization need is a medium to help their members communicate and manage their information.

6 References

Alexander, C. (1964) The Synthesis of Form, Harvard University Press, Cambridge, MA.

Brand, S. (1995) How Buildings Learn: What Happens After They're Built, Penguin Books, New York.

Fischer, G. (1998) "Seeding, Evolutionary Growth and Reseeding: Constructing, Capturing and Evolving Knowledge in Domain-Oriented Design Environments," Automated Software Engineering, 5(4), pp. 447-464. Available at: http://www.cs.colorado.edu/~gerhard/papers/ase-093097.pdf.

Fischer, G., & Ostwald, J. (2001) "Knowledge Management — Problems, Promises, Realities, and Challenges," IEEE Intelligent Systems, January/February 2001, pp. 60-72. Available at: http://www.cs.colorado.edu/~gerhard/papers/km-ieee-2001.pdf.

Fischer, G., & Ostwald, J. (2002) "Seeding, Evolutionary Growth, and Reseeding: Enriching Participatory Design with Informed Participation," Proceedings of the Participatory Design Conference (PDC'02), Malmö University, Sweden, pp. 135-143.

Fischer, G., & Ostwald, J. (in press) "Transcending the Information Given: Designing Learning Environments for Informed Participation," Proceedings of ICCE 2002 International Conference on Computers in Education, Auckland, New Zealand.

Hori, K. (1994)"A System for Aiding Creative Concept Formation," IEEE Transactions on Sytems, Man, and Cybernetics, 24(6), pp.882-894.

Hori, K. (1997) "Concept Space Connected to Knowledge Processing for Supporting Creative Design," Knowledge-Based Systems, 10(1), pp.29-35.

Nardi, B. A. (1993) A Small Matter of Programming, The MIT Press, Cambridge, MA.

Ostwald, J. (1996) Knowledge Construction in Software Development: The Evolving Artifact Approach, Ph.D. Dissertation, University of Colorado at Boulder. Available at: http://www.cs.colorado.edu/~ostwald/thesis.