

Multiagent-based Knowledge Networks

José Luis Aguirre, Ramon Brena, Francisco J. Cantu*
Center for Artificial Intelligence, Monterrey Institute of Technology
Ave. Eugenio Garza Sada 2501, Monterrey, N.L., 64849, Mexico

jlaguirr@campus.mty.itesm.mx, rbrena@campus.mty.itesm.mx, fcantu@campus.mty.itesm.mx*

* corresponding author

Abstract

In this article, we present an approach for the design and development of knowledge networks and corporate memories based on multi-agent systems (MAS) technology. A corporate memory is conceptualized as a network of agents which collaborate to provide the users with knowledge services for both intranets and the internet. Lessons learned from the introduction of knowledge management (KM) practices into organizations are presented. These lessons have influenced and driven the MAS approach to knowledge networks and corporate memories described in this paper. An implementation of the RICA system (Knowledge and Information Networks with Agents) which incorporates these ideas is presented.

1. Introduction

Knowledge Management and Multi-Agent System Technology are two disciplines that share in common the fact of having integrated concepts, ideas, methods and techniques from various disciplines which include Management Science, Computer Science (CS) and Artificial Intelligence (AI). KM is defined as the process of creating value from an organization's intangible assets (Liebowitz, 1999). MAS technology integrates many concepts and methods from AI and CS such as knowledge-based systems, automated reasoning, machine learning, cooperative problem solving, rational decision making, object technology, computer networks, graphical interfaces, among the main ones. Our research takes problems and methods from both areas and develops the RICA (Redes Informaticas de Conocimiento mediante Agentes, in Spanish; Agent-Based Knowledge and Information Networks, in English) system to address some of the issues which emerge from such an integration.

This paper is organized as follows: Section 2 describes KM as a discipline that emerges from ideas developed in the AI field and shaped by disciplines such as Management Science, Organizational Behavior, and from trends in the globalization of international markets. Section 3 presents three case studies in introducing KM practices into Mexican organizations and the way the lessons learned in these projects has influenced the RICA system; section 4 describes the RICA system; section 5 explains the architecture of the RICA system; section 6 gives an example of the use of the RICA system for knowledge collection and distribution among users of an intranet; section 7 presents conclusions and future work.

2. Artificial Intelligence and Knowledge Management

2.1 Knowledge Management

Knowledge management is an area that has attracted the interest of many organizations, companies, government, and academicians from universities and research institutions (Wiig, 1993; O’Leary, 1998; Beckman, 1999). Various methodologies have been proposed for introducing KM practices into organizations (Wiig, 1995; Liebowitz&Beckman, 1998; Schreiber et al, 1999). Although these methodologies emphasize different aspects of the KM enterprise, all of them have in common the facts that knowledge can be treated as a commodity, that knowledge services are a key competitive factor for an organization’s survival, and that the power resides in sharing knowledge rather than in the knowledge in its own (Wiig, 1999a; Wiig, 1999b).

KM has been nourished with contributions of researchers and practitioners from many disciplines which include Management Science, Psychology, Social Science, Cognitive Science, Informatics and Artificial Intelligence among others. As a result, KM has been enriched with ideas, methodologies, approaches, and technologies from such diverse sources, but at the same time, it has made it difficult to integrate all of the diversity of concepts in a coherent and uniform framework for the scientific and methodological analysis of the problems posed by KM.

It has been documented elsewhere that KM comprises perspectives that include cultural, organizational, methodological, and technological aspects (Liebowitz&Beckman, 1998). It has been our experience that cultural, organizational, and technology-fear considerations are the first barriers that we have encountered when introducing KM practices into an organization. There is an extensive literature on these issues (Davenport, 1999; Wiig, 1999b) and the case studies we present in section 3 confirm this belief.

2.2 Knowledge in Artificial Intelligence

The importance of knowledge in modern complex problem solving was first discovered and studied by the Artificial Intelligence (AI) community in the middle 60s. In particular, Edward Feigenbaum and his group at Stanford University conducted the seminal research that led to the establishment of the Expert and Knowledge Based System paradigm with the classical DENDRAL and MYCIN projects (Feigenbaum et al, 1971; Buchanan, 1984). One of the many key contributions of this research was a paradigm shift from designing AI systems for general problem solving which had been the main concern of the AI community, into designing AI systems for complex, knowledge-based problem solving. This paradigm shift allowed AI researchers to analyze and solve industrial-strength, real world problems, instead of the toy, over-simplified problems that had been faced by AI researchers until the late 60s. The concepts, methodologies, and techniques around Knowledge Engineering were developed and established during the 70s (Feigenbaum, 1977). The Knowledge-based Systems industry was developed during the 80s (Feigenbaum, 1988; Liebowitz, 1991) and the Fifth Generation Computer Project was launched in 1982 by the Japanese government with the aim of developing during a 10 years period, intelligent, knowledge-based computers as vehicles for extending the

capabilities of the human brain and the tools for starting a knowledge industry revolution (Feigenbaum&McCorduck, 1983). In perspective, and after 8 years of the end of the Fifth Generation project, we can conclude that the main hypotheses and objectives were right, but the computer, the electronic, the network, and the software technologies and management / business theories were not ripe enough to support the goals. In other words, the predictions for developing smart, knowledge-based computers were underestimated by at least 10 years. Now, with the development of more powerful computers and of the internet, and the involvement of brilliant researchers from other disciplines such as economics, management and business, the technological and methodological bases are more robust to support the announced knowledge revolution.

2.3 Knowledge is power

The adage coined by the Fifth Generation project was the one used by the by Sun Tzu during the Chou dynasty, around the fourth century B.C. in a brief treatise called “The Art of War” which made much of knowledge for the successful conduct of war: “Knowledge is Power, -says Sun Tzu, and permits the wise sovereign and the benevolent general to attack without risk, conquer without bloodshed, and accomplish deeds surpassing all others” (Feigenbaum&McCorduck, 1983). Feigenbaum and McCorduck write on page number 6 of their book *The Fifth Generation*: “The New York Stock Exchange recently published its own treatise which says, less poetically, the same thing: increased productivity derives from more capital, from better capital, but most important of all, from “working smarter” with the capital at hand. American business leaders are as concerned with the art of war as Sun Tzu and his legion of international disciples have been, but in this century the battlefield has changed. Instead of the mountains and valleys of ancient China, the vital battlefield has become the international marketplace”. Then, on page number 11 of the same book they conclude the chapter with the following conclusion wrote in 1983: “The world is entering a new period. The wealth of nations, which depended upon land, labor, and capital during its agricultural and industrial phases – depended upon natural resources, the accumulation of money, and even upon weaponry – will come in the future to depend upon information, knowledge, and intelligence. This is not to say that the traditional forms of wealth will be unimportant. Humans must eat, and they use up energy, and they like manufactured goods. But in the control of all these processes will reside a new form of power which will consist of facts, skills, codified experience, large amounts of easily obtained data, all accessible in fast, powerful ways to anybody who wants it –scholar, manager, policymaker, professional, or ordinary citizen. And it will be for sale”.

As we can see, the main concepts and ideas around KM, whereby knowledge is at the core of them, were defined and studied in the AI community. The term coined and used as a banner of this new paradigm was the word “Expert System”. Another key idea was the separation of knowledge and control. The knowledge was packaged in knowledge bases and the control/inference mechanism were packaged separately in the form of expert system shells. One unfortunate result of this separation was the fact that the industry which arose concentrated in the commercialization of expert system shells with empty knowledge bases. In this way, the “knowledge is power” adage was hidden and relegated by the algorithmic part of the newborn systems. Edward Feigenbaum resumed this phenomenon in his AAAI 93 invited talk “A Tiger in a Cage”. Another unfotunate result was the overselling of the advantages and merits of the expert system technology. Many of the

promises advertised by the proponents of the technology were not met, resulting in disillusionment and skepticism from current and potential customers. However, this does not mean that the expert system technology is dead or disappearing. On the contrary, it has grown steadily, is present in many non-AI software packages and applications, and is quietly used in many intelligent systems application. The World Congress on Expert Systems series started by Professor Liebowitz in the late 80s to spread out the ES technology (Liebowitz, 1991, Liebowitz, 1994; Lee, 1996; Cantu et al, 1998), and the many conferences on intelligent systems, AI, and other non-AI conferences, as well as industry reports, witness these facts.

2.4 Sharing knowledge is power

Fortunately, researchers with non-computer science backgrounds came into the knowledge arena in the middle 80s, contributing with methods, enterprise approaches, and business strategies at which AI researchers do not excel. These academicians, with fresh ideas and systemic, organizational, human, and economic views of the organizations, assimilated and shaped the knowledge ideas, presenting the issues in such a way that CEOs, managers, business leaders, consultants, and decision makers, listened and paid attention to this new way to present the knowledge concepts that had arisen from the AI research in the last two decades. This process was enabled by both economic and market driven requirements of the enterprises created by customer demands and international competition, as pointed out by Feigenbaum and McCorduck in their Fifth Generation book. The development of the internet and the World Wide Web has made it possible to break the geographical barriers and put together groups of people dispersed around the world via intranets. Thus, the sharing of knowledge becomes possible, and is used as a banner by corporations and consulting companies whose main business is about providing knowledge services. One of the researchers who has contributed to this new trend is Karl Wiig, who coined the term Knowledge Management in a keynote address at a European management conference sponsored by the International Labour Organization of the United Nations in 1986. In his article, "Knowledge Management: Where did it come from and where will it go?", Wiig describes a biased story of KM which does not make justice to the origins and development of the knowledge paradigm (Wiig, 1997). The article ignores the efforts of the two previous decades which gave shape to the knowledge paradigm. It talks about some early KBSs like DEC's XCON, without mentioning how it came into existence, and presents a reductionist view of KBS implying that it consists solely of the visible end products which the user has access to, like technology-based networks with email groupware, knowledge repositories accessed through the internet. The knowledge paradigm is much more than this. Wiig's approach has a broader focus in an effort to include knowledge aspects that include the viability and success of the enterprise. This generality makes it difficult to formalize and to treat in a systematic way, leaving room to a narrative, qualitative description and modelling of an enterprise. This is valuable in its own, as it provides the CEOs with ways to find new business opportunities and keep the company wealthy, but these ideas would be impossible to achieve without the use of methodologies, engineering, and technology that emerges from CS and AI. Efforts to formalize general, qualitative, wide scope knowledge had been pursued by AI researchers through the development of knowledge taxonomies, hierarchies, and ontologies (Waterman, 1987; Lenat&Guha, 1990; Turban, 1996).

2.5 Intellectual capital

K.E. Sveiby and T. Loyd published in 1987 the first KM book (Sveiby&Loyd, 1987) which focuses on the management and measuring of intellectual capital possessed by people and by the organizations. The Swedish insurance company Scandia is one of the leaders in trying to make visible the non-monetary and intangible assets of a company. Scandia gives a decomposition of the market value of the company by dividing it into financial capital and intellectual capital. Intellectual capital is made of human capital, structural capital, and customer capital. Structural capital is what is left after the employees are gone: databases, knowledge repositories, software, network, etc. According to Scandia's perspective, intellectual capital is the most important element of the company. Scandia has developed a comprehensive approach to measure intellectual capital (Skandia, 1996). Although corporations realize the importance of knowledge assets and intellectual capital, they are still struggling to develop a comprehensive framework to identify, value and manage these type of assets and capital. J. Wilkins, B. Van Wegen, and R. de Hoog propose a framework for knowledge assets based on relevant differences between human resource assets and intellectual property (Wilkins et al, 1997).

2.6 Learning organizations

Peter Senge presents an approach for Learning Organizations (LO) (Senge, 1990). In his book "The Fifth Discipline" Senge argues that people continually expand their capacity to create the results they truly desire. In a LO, new and expansive patterns of thinking are nurtured, and people are continually learning how to learn together. The organizations that will truly excel will be the ones that discover how to tap people's commitment and capacity to learn at all levels in an organization. Learning organizations are fundamentally different from traditional authoritarian "controlling organizations." The Fifth Discipline is defined around Systems Thinking. Systems thinking is a conceptual framework, a body of knowledge and tools that has been developed over the past fifty years, to make the full patterns clearer, and to help us see how to change things effectively and with the least amount of effort –to find the leverage points in a system. Systems thinking makes understandable the subtlest aspect of the learning organization -the new way individuals perceive themselves and their world. At the heart of a learning organization is a shift of mind –from seeing ourselves as separate from the world to connected to the world, from seeing problems as caused by someone or something "out there" to seeing how our own actions create the problems we experience. A learning organization is a place where people are continually discovering how they create their reality. And how they can change it. Senge defines Systems Thinking in terms of positive and negative feedback loops, in terms of qualitative laws, and in terms of system archetypes (Santos, 2000). Although these goals sound idealistic, it has the merit of emphasizing the value and importance of the human factor to accomplish the goals of an organization. Online communities and communities of practice help toward building a learning organization.

2.7 Agents in KM

In most organizations, knowledge is distributed among the many individuals they are made of. This is why many efforts for building centralized "knowledge repositories" of enterprises (Breslin and McGann, 1998) failed to be comprehensive or even useful. A more natural approach, naturally distributed, could come from research in "intelligent agents"

(Weiss 1999, Sycara 1998, Wooldridge and Jennings 1995). The term “agent” has been overused, so it could mean many things, but we will consider mainly two: 1) An agent is a computational long-lived process acting on behalf of a user in a specific environment. This view emphasizes the traits of an individual agent, dealing with chores delegated by a master; 2) “Multiagent” systems work through the interaction with other agents, cooperating or competing to achieve their goals. There are other views of agents as well, like that of Marvin Minsky in his seminal work about the society of mind, which tries to explain how the mind works (Minsky, 1985). He argues that a mind can be built from many little parts, each mindless by itself.

The importance of automated intelligent agents in the organizations comes from their natural distribution, their ability to pre-process information for users, their continuous presence and availability and the interactions with other agents and users.

In summary, we have seen how the knowledge paradigm was developed and shaped during the 60s and 70s, especially through the work of E. Feigenbaun and his group at Stanford. This launched the knowledge revolution during the 80s led by the knowledge industry and the Japanese Fifth Generation project to create intelligent, knowledge-based computers and networks of them, as the building blocks of the new era. Results showed that matters were not ripe enough to accomplish the stated goals, although the hypotheses were right. The appearance of researchers from other disciplines and the international markets competition, gave new perspective to the knowledge paradigm, allowing corporations to redefine their objectives and find new niches and business opportunities around knowledge services, which were enabled by the parallel development of the internet, the web, multi-agent systems, and other technologies.

3. Case studies in KM

In this section, we describe a sample of efforts for introducing KM practices into an organization. We have been developing KBS solutions for Mexican corporations for the last 15 years in areas of diagnosis, control, scheduling, configuration, planning, and other application domains (Cantu, 1991; CIA, 2000). Many of these systems were developed, deployed, used for some time, and eventually abandoned, as the knowledge contained in the KB became outdated. Companies were not conscious of the need to maintain updated the knowledge bases, or if they were, the tools and methodologies to keep the KB up to date and consistent, were not available. If they were available, the updating process was painful and error-prone, especially for KB containing hundreds of rules. The effort to keep the KBs updated was justified in a few applications that were critical and indispensable for the operation of the company, as was the case of DEC’s XCON expert system for configuring computer equipment. With the advent of the internet and the Web in the second half of the 90s, the tools to gather and distribute various kinds of knowledge and information became available. Information technologies (IT) companies and consultants developed tools and environments to support the management of documents containing text, images and video, where the users took responsibility for the reasoning and interpretation of the contents of those documents. Standards such as the KQML language were developed to facilitate documents exchange. This new infrastructure provides KBSs a more gentle and powerful means for addressing again the problems that proved difficult during the 80s and early 90s.

Developing a knowledge-based organization demands the addressing of cultural, organizational, managerial, and technological aspects. Figure 1 shows steps that, according to our experience, a company may follow for introducing technologies to support KM.

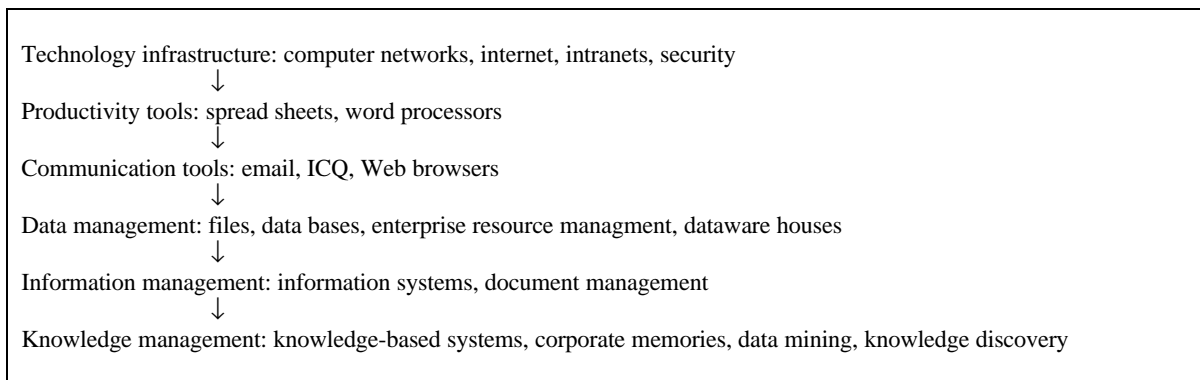


Figure 1. Steps for introducing KM technology into a corporation

3.1 CYDSA

CYDSA Corporation is an enterprise for which we have developed at least 20 KBS in the last few years. Most of these KBSs were stand alone, off-line systems, which were used to retain the knowledge and experience of people about to retire, train operators, standardized practices, and assist in trouble-shooting and problem solving in the production lines. Efforts were made to keep the systems alive, although these attempts were not always successful. Today, CYDSA has developed an IT infrastructure with an intranet that links more than 30 plants and offices distributed over the country, and email, browsers, and collaborative work facilities based on Lotus Notes.

Currently, we are working with CYDSA's Chief Knowledge Officers (CKOs) and IT managers to develop a KM company-wide plan which includes (a) repositories with best-practices in several production lines, (b) updating of existing KBSs to support trouble-shooting, customer support, maintenance of industrial machinery, production scheduling, and control of process variables, (c) development of new KBSs in the same areas as in (b). The idea is to develop a KM package that implements a corporate memory of a *knowledge pump* type, that is, a knowledge repository with active knowledge collection from authorized users and active knowledge distribution to users, according to user /job profiles and work interests specified in a knowledge map (Liebowitz & Beckman, 1998). This package will run over Lotus Notes, in order to take advantage of the existing facilities and the proficiency that the users have developed in this platform. We conducted a cost-benefit analysis comparing the internal development of the package versus acquiring an existing KM tool. The comparison favored the former alternative, although it may take a bit longer to develop. The RICA tool described in section 4 is a subset of the functionality devised for the proposed package.

3.2 Peñoles

Peñoles is a steel, chemical and mining corporation. The Mining Division which has 11 plants in the center and north of the country, started in 1998 an initiative to develop a

corporate memory containing best/worse, good/bad practices in the construction of mining plants. The company had plans to open new plants in the near future. A concern of the executives of the company was the fact that the same costly errors were made over and over again, each time a new plant had been built in the past. Reasons for this included that the experiences were not documented and that the people that participated in the construction projects had moved to other areas or had left the company.

Regarding technology infrastructure, the company was not quite prepared to introduce KM practices. The computer network was under development, the data and information systems were under conversion to data base systems, and ERP tools were being introduced. Regarding cultural aspects, the main concern of the potential participants was the fact that they would have to document their errors and make them public, at least for the intranet members. People were willing to talk about their successes, and sometimes, over-emphasize them, but not the other way around. We recommended establishing economic and recognition incentives for those sharing their good/bad practices through the corporate memory. At the end, the head of the division concluded that the organizational culture was not sufficiently developed to start a KM project of this type. We recommended to start a culture assessment project to address these issues, and a knowledge audit, to assess information and knowledge requirements.

3.3 The Monterrey Tech experience

The Monterrey Institute of Technology started in 1997 a university-wide process to redesign the teaching practices at the undergraduate and graduate levels. Monterrey Tech is a university system comprising 30 campus in the country and offices in South America, the USA and Europe. A satellite-based computer network links all those places. The redesign of the teaching practices had two main objectives: (a) use pedagogical and didactical methods and (b) use computer-based technology to support the learning process. The former looks to promote student self-learning, collaborative work, information search, problem/project based learning, practical experiences and some other methods. Under this view, the role of the teacher changes from the traditional scheme in which he/she is the center of the learning process which occurs in the classroom, into being a facilitator of the learning process, which is student-centered and may occur in the classroom, the cafeteria, the laboratory, the gardens, the office, or the student home. The latter takes advantage of the computer technology to support and facilitate the learning process. The technology includes the development of knowledge repositories for each course. The repositories contain the knowledge elements that the student must acquire during the semester. Knowledge elements include slides, exercises, references, articles, projects, electronic books, and a set of URLs of sites with relevant knowledge and information. The repositories also contain a schedule which guides the learning activities during the semester: homework assignments, project deadlines, exam dates, etc. Collaborative work activities performed by teams of students in the class is also included in the repository. The development of these knowledge repositories is the responsibility of the teachers of the courses. Each student owns a laptop where the knowledge repositories of each course he is enrolled in, is loaded at the beginning of the semester. The classrooms, the cafeterias, the gardens, the laboratories and the library are equipped with internet ports so that the students and the teachers can communicate. The knowledge repositories reside also in central computer equipments which are maintained by the university computer center. The

software supporting the knowledge repositories is Lotus Notes/Learning Space. The technology also includes the use of intelligent tutoring systems and computer-aided learning software.

There were three main challenges in this initiative. The first has to do with cultural aspects. As expected, the first reaction was a change resistance by the faculty members. This was exacerbated by the unfortunate use in public by management of expressions such as “what you have done in the past is wrong, this is the correct way to do things”. Actions were taken to facilitate the transition process. These included the establishment of economic incentives such as salary increases, bonus, free laptops, budgets to initiate pilot projects, training, traveling funds, special prices, and workshops, among the main ones. There are about 6,000 faculty members in the university system. The majority of them are now involved in the teaching redesign process. There are about 80,000 students in the university system. The majority of them are now taking courses with the new scheme. Each semester several thousand courses and groups are offered with the new scheme.

The second challenge was the scaling of the software Lotus Notes/Learning Space that supports the knowledge repositories, the bandwidth of the communication lines, and the capacity of the central computers. All of those elements were upgraded to provide the service required by such massive requirements of data, information, knowledge and communications.

The third challenge, which is the most fundamental one, is a demonstration that the new scheme is a more effective way of learning. It is clear the the new scheme is more expensive, as it requires an expensive infrastructure of computers, communications, and multimedia equipment in the classrooms. The first question is: “is the new scheme a better way of learning?”. If it is, the second question is: “by doing a cost-benefit analysis, does the better way of learning of the new scheme justifies the extra cost?”. There are several studies being conducted to answer these questions. Qualitatively speaking, and from the experience that we have had as teachers of those new courses, we think that both questions have a positive answer (ITESM, 2000).

These experiences have driven and influenced the design of RICA, a long-term research-oriented project to support KM through the use of multi-agent technology.

4. The RICA project

As we mentioned earlier, our focus is on knowledge sharing. In the RICA project we propose an architecture and a set of integrated tools for facilitating knowledge dissemination in a variety of contexts. We assume that a group of users has already some basic electronic facilities such as internet/intranet connection, web, email, electronic meeting, and so on. What we are adding is the possibility to automatically directing relevant pieces of knowledge to the right persons.

We think that automated knowledge handling within an organization requires the application of Artificial Intelligence tools, and further, the use of the novel Intelligent Agents technology (Weiss 1999). We take the meaning of Intelligent Agent in two ways:

one is a computer task that executes some tasks on behalf of its user – this is a kind of “electronic assistant” at the service of users who subrogates some chores to it. The other meaning of Intelligent Agents relates to Distributed Artificial Intelligence, and focuses on how one agent could relate in a rational way to other agents, building in this way a kind of artificial society.

Of course, the use of electronic assistants acting on behalf of their masters has been seen as dangerous by some potential users, as they lose control over some tasks. But the benefits of delegating tasks could easily overcome those dangers, because the human user can become much more efficient, just as a human boss needs to subrogate tasks to his/her subordinates to make the organization work efficiently.

In the RICA project users have a personal agent, which handles a *user profile* for selecting which pieces of knowledge could be relevant for him/her. This profile is very much application dependent, but in our pilot application (cooperation between computer science researchers) the profile includes a set of academic interests. These interests are any point in a standard tree-structured taxonomy of areas. For example, in our pilot project we used the ACM Computer Science Classification (ACM 1998) as a map for the possible user interest areas. Other aspects in a user profile include his/her personal identification, as well as his/her place in the organization, according to an enterprise dependent scheme. We will call those classifications generically “*user maps*”. Part of the entries in a user map is provided directly by the user (like his/her interests), while others are provided by the organization (like his/her responsibilities, the place in the enterprise, etc).

The advantage of using structured user maps is that they allow to identify automatically whether a piece of knowledge concerns a user or not. Further, we have been using automated deduction tools to fully exploit the information implicit in the user maps. For example, if a user includes in his/her profile an interest in computer hardware, he/she will receive information concerning any sub-area of hardware, like peripherals or memory. This gives the system a great flexibility for determining which users should receive a specific piece of information.

The pieces of information a RICA user can receive could be either generated by some human user (for example, a call for papers in a scientific community, or a report of a best practice in a quality control setting in an enterprise) or automatically triggered by a local or remote computer linked to RICA (for instance, if the information in a certain database meets some criteria, or if certain web page has changed).

Put in a condensed way, the role of RICA in knowledge management is the dissemination of opportune and relevant knowledge in the organization, by the use of agents and user maps. We think that the task of constructing and maintaining a collection of user maps is in itself of great value for the enterprise, as it helps in characterizing the human assets. RICA offers as well other knowledge-sharing tools, like a digital repository of documents. It is a facility for collaborative working with versions of a document.

5. Architecture of RICA

In order to provide the kind of services mentioned above, we devised an architecture able to fulfill the following goals:

- a) Users should connect from personal computers, but their agents should not stop working when that computer is turned off or disconnected;
- b) Administration of the system should be kept at a minimum, by centralizing as little as possible. Each department of the organization should be accountable for their use of computational resources and user administration.
- c) Agent components should run on top of conventional technologies such as client – server connections, relational databases, and so on. This goal is intended to reduce the development time, as well as increasing compatibility with existing systems in the organizations.

The resulting architecture is depicted in figure 2.

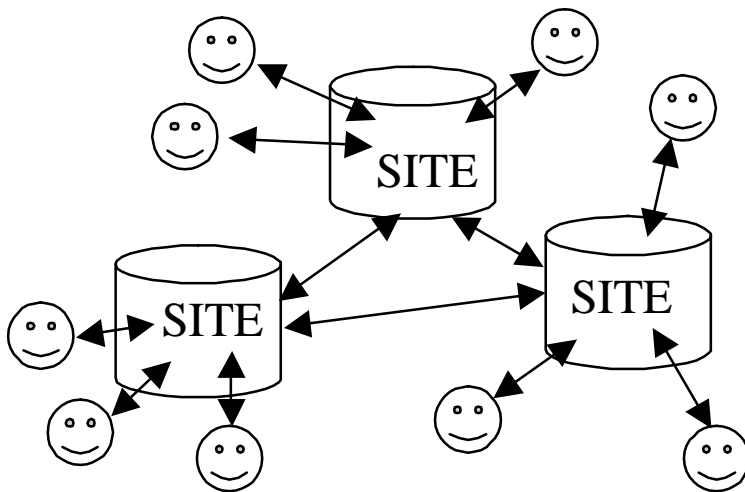


Figure 2. RICA architecture

Each user has a personal agent, represented by a smiling face in the diagram, running continuously in a server, called “SITE” in the diagram. The personal agent can be accessed through a personal computer connected to that site. Then, personal agents interact with a “server agent” –one by site- , which provides the general services of the system.

The user communicates with his/her personal agent through a web-based interface, where the web server runs on the server, and the web browser in the user's personal computer.

This means that the client side software is kept to a minimum, while the burden is passed to the server. Actually there were many technological options for this; the most extreme ones are presented in table 1.

***** Insert table No. 1 about here *****

We chose the first possibility as the development was faster, and also the server load has not appeared to be a bottleneck, especially taking into account that decentralization of the servers provides room for great scalability. Anyway we plan to experiment with intermediate schemes as well, for better distributing the load between the client and the server.

Typically there is one server agent per department, campus in universities, etc., depending on the size of each organizational entity. Given the heavy workload caused by the server agent, one agent per user, and the web server comprising our current architecture, it is not convenient to include many users in a single site -typically there will be less than 100 per site.

The server agent itself comprises the following elements:

- a) A database, storing on disk information about the services being offered, like the web pages that have to be continuously monitored to check for changes.
- b) An inference engine, which performs deductions from data stored in the fact base. Those deductions allow to propagate consequences of the information in very flexible ways.
- c) An object administrator, which loads informations from the database or external devices to main memory, and updates them as required.
- d) “Adapters”, which are interfaces to the external world. For example, there is a web monitor adapter, which continuously monitors specific web pages and generates an event to the inference engine when something happens.

The inference engine, a key component in our architecture, uses a “production rules” knowledge representation. They are pairs of the form “if <condition> then <action>”, where the action could be to manipulate internal variables or to trigger external actions, like sending a message.

Rules in the inference engine are classified in the following categories: APPLICATION RULES, which perform the RICA services themselves, CONNECTION RULES, for activating the deduction process when required, and MAINTENANCE RULES, for cleaning up useless information and the like. Two illustrative examples of specific rules are presented below.

Rule 1:

*If webPage(P,URL) and changeWebPage(P) and interestedInPage(P,E)
then notify(P,URL,E).*

This means that if there is a web page P that changed, and users in entity E (interest area or group or individual) are interested in monitoring that page, then E should be notified. Of course, if E is an area, this is not going to generate a message to a specific user; we still have to map the users with the areas in order to send individual messages. This is done in the following rule:

Rule 2:

If area(A) and notify(A,M) and subarea(S,A) then notify(S,M).

We encourage the reader to compare the ease of this transitive closure calculation of the “subarea” relation with the way it would have to be programmed in conventional programming languages.

The inference engine has been programmed in the Jess shell (Friedman-Hill 2000). For the database we used the MiniSQL system (Jepson and Hughes, 1998), and the object administrator and other components were directly programmed in Java (Sun 1999).

The server agents communicate among them to provide overall coordination within the whole system. We made an adaptation of the KQML agent communication protocol (Finin et al. 1992) for the inter-server communication.

So far we have just developed very basic tools for server-agent database replication. New or modified records in each local database are marked during normal operation, then they are extracted and transmitted to the other sites, so that these data are incorporated there. Possible conflicts concerning modification of the same records in different sites are minimized as we associate access rights to data so that when it originates at a different site is read-only.

In the future we plan to develop more sophisticated schemes to coordinate the servers concurrent operation.

Strict access policies have been developed for the shared-document repository as well, so that cooperative and concurrent access to documents is provided, while avoiding conflicts and inconsistencies (Eguia 1999).

6. An example in RICA

The example that we develop in this section is related to the use of RICA for the “Red de Desarrollo e Investigación en Informática” (REDII, in english “Network for Development and Research in Informatics”). REDII is supported by the “Consejo Nacional de Ciencia y Tecnología” of Mexico (CONACyT, in english “National Council for Science and Technology”).

The main knowledge management problem that we address is the cooperation and collaboration between the members of the different projects that compose REDII. The goal is to give tools to those members so they can actually distribute and share their knowledge related to their projects. Also, it can be used by the REDII administrator so she/he will be notified each time that an important change in one of the projects happens, and she/he can communicate with the other members of REDII about issues concerning all of them.

In the following, we describe first the REDII organization; next, we show a possible REDII project’s taxonomy that can be used for RICA; finally, we show how RICA can help the members of REDII in the distribution and sharing of the knowledge that arise from their projects and in facilitating their collaboration and cooperation.

6.1 What is REDII

REDII is a network of universities and research institutions in Mexico which receives funding from the CONACYT in order to conduct research and development in Computer Science and Informatics. Currently, it is composed of 15 institutions, of which our university is a member.

6.2 A REDII project’s taxonomy

At the time this paper is written, REDII is composed of several projects dealing with different computer areas, domains and problems. For the purpose of the RICA system, the projects can be arranged in a taxonomy, as it is showed in figure 3.

- A. REDII Project
 - A.1 Software Tools

- A.1.1 Symbolic object-based tool
- A.1.2 Distributed CASE
- A.1.3 Quality maps for software development
- A.1.4 Intelligent interfaces
- A.2 Data bases
 - A.2.1 Distributed data base management system
 - A.2.2 HYPERTEXTER (multimedia data bases)
 - A.2.3 Object-oriented and deductive data base management system
- A.3 Computer asisted education
 - A.3.1 Education
 - A.3.1.1 Remote education
 - A.3.1.2 EVA (Virtual learning spaces)
 - A.3.1.3 On-line education center
 - A.3.1.4 Collaborative learning environments
 - A.3.2 Tutoring
 - A.3.2.1 Intelligent tutoring simulating tool
 - A.3.2.2 Analytical geometry laboratory
 - A.3.2.3 Tutor for analysis and optimization of structures
- A.4 Digital libraries
 - A.4.1 PHRONESIS
 - A.4.2 A digital library for collaborative learning
 - A.4.3 A digital library in the domain of flowers
- A.5 Perception and action
 - A.5.1 Robotics and automation
 - A.5.1.1 Symbolic image's analysis and mobile robot's control
 - A.5.1.2 EJACOM: Java specification for mechanical remote control
 - A.5.1.3 Virtual robotics laboratory
 - A.5.2 Vision and speech
 - A.5.2.1 Image calculator
 - A.5.2.2 Tridimensional optical digitizer
 - A.5.2.3 Face recognition by indexing
 - A.5.2.4 Digital speech recognition
- A.6 Networks and communication
 - A.6.1 Digital teleconferences
 - A.6.2 A system for analysis and design of conmutation networks
 - A.6.3 Distributed object models for INTRANET architectures
- A.7 Knowledge management
 - A.7.1 Web pages text recovery
 - A.7.2 RICA (Agent based knowledge and information networks)

Figure 3. A taxonomy for REDII

As it can be seen, the root of the taxonomy represents the set of all REDII projects; the inner nodes represent application areas or subareas; the terminal nodes are specific projects, for example, RICA is a REDII project. Each project has associated with it a certain number of Web pages with information of several kinds about the project.

It has to be noted that the taxonomy can be defined by the REDII administrator. She/he can update the taxonomy with the help of a Web interface. The taxonomy can be arranged in the most useful way so it can serve the purposes for which the RICA system is proposed. For example, the REDII administrator can define nodes in the taxonomy for handling pages where REDII members can put announcements about seminars, conferences, important dates and events for REDII, etc.

6.3 RICA services for REDII

By using the taxonomy of the preceding section, REDII members can benefit from the services that the RICA system offers, such as the monitoring of Web pages and the digital repository of documents.

Traversing the taxonomy from the bottom to the upper nodes, the first documents of interest for monitoring are the Web pages that are directly associated with the projects, as each of these pages has information of several kinds about the projects. For example, RICA project has Web pages with information about the description of the project, the team members, publications, technical reports, manuals, and even an online version of the system. It is sure that all of these pages are of interest for the members of the RICA team, and it could be the case that the same pages are interesting for other members of REDII. Anyway, the pages are declared of interest for the RICA “area”. The same approach is followed for the other projects.

As it could be the case that there exist some other pages of interest for the member of a project, for some area or subarea of the taxonomy, or for all the members of REDII, the pages must be located and someone must declare that there are of interest for some node in the taxonomy. For example, the main page of the Robotics Laboratory of Stanford University can be declared as an interest point to the node “Perception and Action” in the REDII taxonomy. Next, the users of the RICA system activate the services of monitoring, by declaring its user profile. In fact, it consists in declaring that they are interested in some areas of the taxonomy. The REDII administrator can declare her/his interest in the root of the taxonomy, so any time that something changes in one of the pages that has been declared for monitoring, he/she will be notified. The changes may be caused by a new technical report, changes in the participants of a project, the releasing of new versions of some of the systems, etc. The main screen of RICA system is displayed in figure 4.

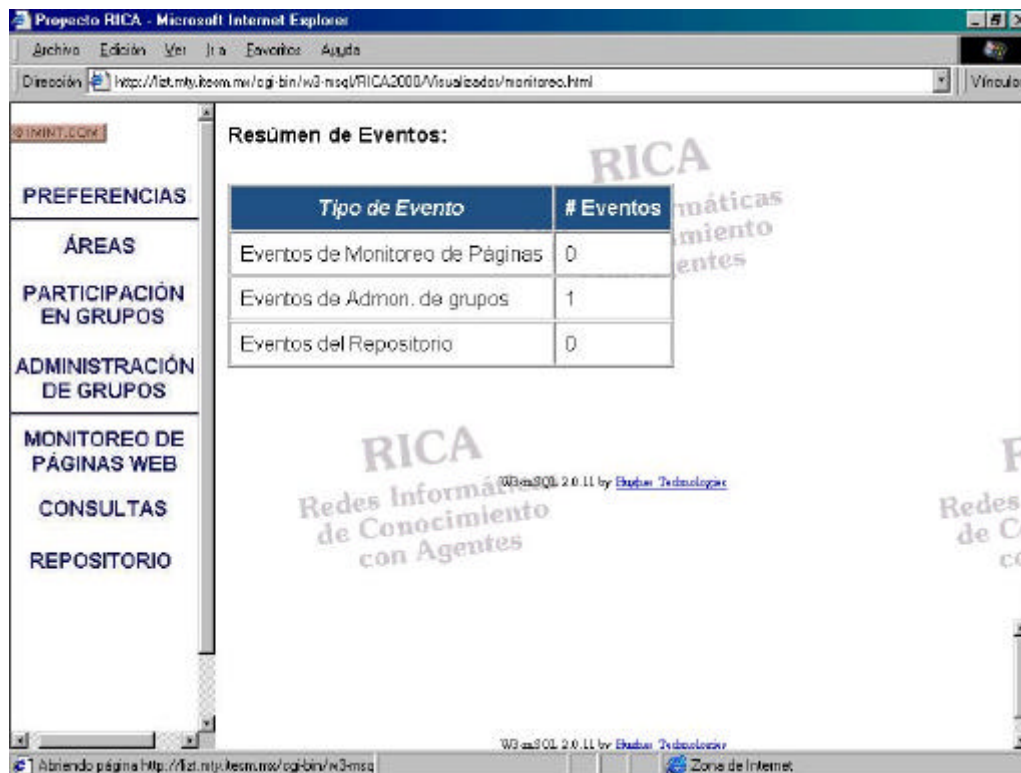


Figure 4. A RICA screen

From figure 4, we can see on the left the options that the user can choose, and on the right, we have the working area where RICA puts information of interest to the user, such as the events notified by RICA according to the user profile.

Another service that offers the RICA system is the digital repository of documents. As it was already mentioned, it is basically a facility for collaborative working with versions of a document. The digital repository works on the same taxonomy that RICA and share also the same users. In this case, a document has an owner, which is responsible of declaring the areas for which the document is interesting. She/he can also give access privileges to other users, mainly the ones that will collaborate in the preparation and reviewing of the document. For example, the group responsible for the Digital Libraries projects can decide to write a technical report about their projects, and can use the digital repository for that purpose. The notifying services work the same way as they do for Web pages.

7. Conclusions and future work

The ideas described in this paper have been partially implemented in the RICA system. The actual implementation of RICA offers some services that permits the creation of user collaborative networks installed on one site and working around a common domain of knowledge. By the moment, we are working in the integration of the distributed agents approach, so the interaction between users of different sites could be done. Another point of interest in our future work, is to enhance the reactive and proactive capacities of our agents, so they can take better decisions in behalf of the users, as searching and proposing new sources for monitoring. Finally, we are on the way to integrate with our system a tool called Clasitex. Clasitex is capable to analyze documents in order to recognize the main subjects and concepts that appear on them. With such a tool, several services can be refined, and new ones can be offered.

We plan to pursue the development of RICA to have an intelligent environment that supports the introduction of KM practices into an organization using MAS, AI, and KBS technologies. Experiences with institutions that are committed to becoming knowledge organizations will continue to guide our plans.

Acknowledgements

The authors gratefully acknowledge the financial support from CONACYT-REDII grant number REDII-CE-1 and from the Monterrey Institute of Technology.

References

- Association for Computing Machinery (1998). The ACM Computing Classification System, Technical Report, <http://www.acm.org/class/1998>
- Beckman, T. (1999). The Current State of Knowledge Management. In Knowledge Management Handbook, Jay Liebowitz (ed.), CRC Press, Boca Raton, FL.
- Breslin, J. & McGann, J. (1998). The Business Knowledge Repository, Quorum Books. Westport, CT.
- Buchanan, B. & Shortliffe E. (1984), Rule-Based Expert Systems: The MYCIN Experiments of the Stanford Heuristic Program project. Reading Mass: Addison-Wesley.
- Cantu, F. (1991). Operational Expert Systems Applications in Mexico, Pergamon Press, New York.
- Cantu, F. Soto R., Liebowitz, J. (eds) (1998). Proceedings of the Fourth World Congress on Expert Systems, Cognizant Communications, New York.
- CIA (2000). Center for Artificial Intelligence Web page, Monterrey Institute of Technology, <http://www.cia.mty.itesm.mx>
- Davenport, T.H. (1999). Knowledge Management and the Broader Firm: Strategy, Advantage, and Performance. In Knowledge Management Handbook, Jay Liebowitz (ed.), Boca Raton FL: CRC Press.
- Eguia, S. (1999) A distributed document repository for the RICA project (in Spanish), Master Thesis, ITESM, Mexico
- Feigenbaum, E. A., Buchanan, B.G., & Lederberg J. (1971). On Generality and Problem Solving: A Case Study involving the DENDRAL program. In Machine Intelligence, B. Meltzer and D. Michie, pp 165-190, New York: Elsevier.
- Feigenbaum, Edward (1977). The Art of Artificial Intelligence: Themes and Case Studies of Knowledge Engineering. In Proceedings of the the Fifth International Joint Conference on Artificial Intelligence.
- Feigenbaum, E. & McCorduck, P. (1983). The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World. New York: New American Library.
- Feigenbaum, E. McCorduck, P. & Nii, H.P. (1988). The Rise of the Expert Company. New York: Times Books.
- Finin, T., McKay, D. and Fritzson, R. (1992). An Overview of KQML: A knowledge query and manipulation language. Technical Report, U. of Maryland CS Department.

ITESM, (2000). Monterrey Institute of Technology Teaching Redesign Page: <http://www.sistema.itesm.mx/va/desedu.html> (Spanish)

Jepson, B. and Hughes, D. (1998). The Official Guide to Mini SQL 2.0. John Wiley and Sons.

Lee, J. K. (ed) (1996). Proceedings of the Third World Congress on Expert Systems, Cognizant Communications, New York.

Lenat, D.B.& Guha, R.V. (1990). Building Large Knowledge-Based Systems, Reading Mass: Addison-Wesley.

Liebowitz, Jay (ed.) (1991). Proceedings of the First World Congress on Expert Systems, Pergamon Press, England.

Liebowitz, Jay (ed.) (1994). Proceedings of the Second World Congress on Expert Systems, Cognizant Communications, New York..

Liebowitz J. & Beckman T. (1998). Knowledge Organizations: What Every Manager Should Know. Boca Raton FL: St. Lucie Press.

Liebowitz, Jay (ed.) (1999). Knowledge Management Handbook. Boca Raton FL: CRC Press.

Minsky, Marvin (1985). The Society of Mind. New York: Simon&Schuster Inc.

O'Leary, Daniel (ed.) (1998). Knowledge Management: Taming the Information Beasts. Special issue of the IEEE Intelligent Systems and their Applications Magazine.

Santos, Aldo (2000). Summary and comments about the Fifth Discipline Book: <http://home.nycap.rr.com/klarsen/learnorg/senge.html>

Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N., Van de Velde, W. & Wielinga B. (1999). *Knowledge Engineering and Management: The CommonKADS Methodology*. Boston Mass: MIT Press.

Senge, P.M. (1990). The Fifth Discipline: The art and practice of the learning organization. New York: Doubleday Currency.

Skandia (1996). Power of Innovation Intellectual Capital. Skandia Market Communications, Stockholm, Sweden.

Sveiby K. E., & Loyd, T. (1987). Managing Know-how, London: Bloomsbury.

Sycara, K. (1998). Multiagent Systems, Artificial Intelligence 19(2), pp79-92.

Turban, E. (1995). *Decision Support and Expert Systems*. Fourth edition. Englewood Cliffs NJ: Prentice Hall.

Waterman, D.A. (1986). *A Guide to Expert Systems*. Reading Mass: Addison Wesley.

Weiss, G. (1999). *Multiagent Systems: A modern introduction to Distributed Artificial Intelligence*, MIT Press.

Wiig, K. M. (1993). *Knowledge Management Foundations: Thinking about Thinking – How People and Organizations Create, Represent, and Use Knowledge*, Arlington, TX: Schema Press.

Wiig, K. M. (1995). *Knowledge Management Methods: Practical Approaches to Managing Knowledge*, Arlington, TX: Schema Press.

Wiig, K. M. (1997). *Knowledge Management : Where did it come from and where will it go?*. In *Expert Systems with Applications* 13(1), pp 1-14, Pergamon.

Wiig, K. M. (1999a). *Introducing Knowledge Management into the Enterprise*. In *Knowledge Management Handbook*, Jay Liebowitz (ed.), Boca Raton FL: CRC Press.

Wiig, K. M. (1999b). *Establish, Govern, and Renew the Enterprise's Knowledge Practices*, Arlington, TX: Schema Press.

Wilkins, J., van Wegen B. & de Hoog R. (1997). *Understanding and Valuing Knowledge Assets: Overview and Method*. In *Expert Systems with Applications*, 13(1), pp. 55-72, Elsevier.

Wooldridge, M., Jennings, N. (1995). *Intelligent Agents - Theories, Architectures and Languages*. LNAI 890, Springer Verlag.

Technology	Advantages	Disadvantages
Web browser as client, personal agent runs on server	Virtually no client software, non-stop operation	Additional load for server, slower response
Dedicated client, personal agent runs on client	Reduced load for server, faster response	Need for a dedicated client, personal agent stops when disconnected

Table 1. Technological options