An Overview of Agents in Knowledge Management

Virginia Dignum

institute of information and computing sciences, utrecht university
technical report UU-CS-2004-017
www.cs.uu.nl

An Overview of Agents in Knowledge Management

Virginia Dignum

Abstract. Current developments in Knowledge Management concern the sharing and usage of knowledge in dynamic environments. The need for systems that both react to and anticipate the needs and expectations of users calls for flexible and adaptable development and implementation frameworks. These are exactly the characteristics that identify software agents and agent societies, which make natural the application of the agent paradigm in KM. This paper attempts to identify both the advantages of agents for KM, and the aspects of KM that can benefit most from this paradigm. Furthermore, the paper describes several current KM projects that use agent technology and identifies open research areas.

1. INTRODUCTION

The main goal of Knowledge Management (KM) is to provide relevant knowledge to assist users in executing knowledge intensive tasks. KM is about facilitating an environment where work critical information can be created, structured, shared, distributed and used. To be effective such environments must provide users with relevant knowledge, that is, knowledge that enables users to better perform their tasks, at the right time and in the right form.

As it is often mentioned in the literature, knowledge tasks have a collaborative aspect, that is, an individual can best acquire and use knowledge by reusing information already collected and annotated by others or by making use of existing relations among people (communities) [16]. Furthermore, a KM system must be able to adapt to changes in the environment, to the different needs and preferences of users, and to integrate naturally with existing work methods, tools and processes. That is, KM systems must be *reactive* (able respond to user requests or environment changes) and *proactive* (able to take

initiatives to attend to user needs). These aspects also characterise intelligent software agents, what seems to indicate the applicability of agent technology in the KM area.

Intelligent agents are a new paradigm for developing software applications and are currently the focus of intense interest on the part of many fields of computer science and artificial intelligence [13]. Agents have made it possible to support the representation, co-ordination, and co-operation between heterogeneous processes and their users. Internet technology and software agents make thus possible to build sophisticated, well performing KM systems designed to deliver content, from multiple, heterogeneous sources, to each individual, in the individual's specific context and under the individual's own control. This ability contributes to improve the relationship between knowledge suppliers and consumers by providing both parties more precise control over the interaction.

In this paper, we intend to show the applicability of agent technology to the development of collaborative KM environments that address the problems highlighted above. The paper is organised as follows: in section 2, we briefly introduce collaborative KM environments and intelligent agents. Section 3 discusses the applicability of agents in the KM area. In section 4, some current KM projects that make use of agent technology are listed. In section 5, the Agent-Mediated Knowledge Management framework being developed at our organisation is presented. Finally, section 6 discusses some open issues and highlights aspects for further research.

2. KNOWLEDGE MANAGEMENT AND AGENT TECHNOLOGY

In this section, we introduce relevant issues of knowledge management and agent technology. Our aim is not to provide a precise definition of either area but to give an impression of the issues that relate both areas and can benefit most from cross-fertilisation.

2.1. Knowledge Management Environments

Business processes are often highly dynamic and unpredictable. This makes it difficult to give a complete a priori specification of all the activities that need to be performed, which their knowledge needs are, and how those activities should be ordered. Within organisations, there is a decentralised ownership of the tasks, information and resources involved in the business process. Different groups within organisations are relatively autonomous, in the sense that they control how their resources are created, managed or consumed, and by whom, at what cost, and in what period [11]. Furthermore, often multiple and physically distributed organisations (or parts hereof) are involved in the business process. Each organisation attempts to maximise its own profit within the overall activity. That is, there is a high degree of natural concurrency (many interrelated tasks and actors are active at any given point of the business process) which makes important to monitor and manage the overall business process (e.g. total time, total budget, etc.). These characteristics call for an environment that integrates the business

process aspects of knowledge work with active support for searching, using and adding heterogeneous knowledge sources [21].

The first attempts towards business integration were made at the data level, where distributed computing frameworks have been developed to support distributed computing in heterogeneous environments and provide an interface description language and services that allow objects to be defined, located and invoked across locations and applications. The most popular of such distributed object paradigms are **CORBA**, the Common Object Request Broker Architecture developed by OMG and **DCOM**, the Microsoft's Distributed Component Object Model. Such frameworks encapsulate the heterogeneity of legacy systems and applications within standard, interoperable wrappers. These frameworks are defined and are well suitable to the 'data' level of communication. They presuppose a relatively stable environment and some common grounds of understanding.

However, knowledge assets available in an organisation are more than data sources alone. Such assets include structured and unstructured information, multimedia knowledge representations and links to people (ex. through knowledge maps or yellow pages). Besides using existing knowledge sources, the environment should be able to create (and store) new knowledge based on its observation of the user's task performance [18]. Furthermore, there is a need to handle and combine formal and informal knowledge representations, as well as heterogeneous multimedia knowledge sources.

This means that the integration of knowledge and business processes requires a higher level of integration than the one provided by CORBA-like frameworks. At knowledge level, integration must be based on the semantics and the context of the problem at hand. A knowledge-level integration framework must be able to create dynamic relationships between knowledge-intensive business processes and knowledge sources that do not compromise the autonomy of the different parts. In order to be able to support the execution of knowledge-intensive tasks, using knowledge from heterogeneous sources, according to diverse user preferences, KM system must be able to provide a common knowledge description. In this way, integration and autonomy are achieved by separating the use of knowledge from the specific characteristics of the knowledge source. KM systems must therefore provide uniform access to a diversity of knowledge and information sources of different degree of formality. Furthermore, knowledge management environments must be able to adapt to the different needs and preferences of users, and integrate naturally with existing work methods, tools and processes. That is, such environments must be *reactive* and *proactive*.

In summary, KM systems must be able to [6]:

- Provide uniform and transparent access to a diversity of knowledge and information sources of different degree of formality and format, wherever it is situated in the organisation, and even outside (e.g. World Wide Web).
- Proactively identify and deliver timely, task relevant information which may not have been explicitly asked for (e.g. because the decision-maker is unaware of its existence).

- Create virtual, dynamic links between knowledge needs and knowledge sources according to the work context of the user.
- Inform users about changes that have been made elsewhere in the business process, which have consequences upon the current work context.
- Identify and contact parties that may be interested in the outcome and results of the activity of the user.

2.2. Intelligent Software Agents

Although there is not yet a real agreement to the core question of what exactly is a software agent, most researchers accept the following definition [12]:

> *A software agent is an encapsulated computer system that is situated in some environment and that is capable of flexible, autonomous action in that environment in order to meet its design objectives.*

A few of the notions introduced in this definition are worth further explanation. The term 'encapsulated computer system' indicates a clear distinction between the agent and its environment. Moreover, the definition implies that there is a well-defined boundary and concrete interface between the agent and its environment. The key aspect of the definition is **autonomy**, which refers to the principle that agents can operate on their own, without the need for human guidance. An autonomous agent has the control over its own actions and internal state. That is, an agent can decide whether to perform a requested action. The definition above situates an agent in a particular environment, which the agent can sense and effect. This indicates **responsive** behaviour. Furthermore, the definition implies that agents are problem-solving entities, with well-defined boundaries and interfaces, designed to fulfil a specific purpose, that is, having particular goals to achieve, and exhibiting **flexible** and **pro-active** behaviour. Furthermore, agents are often capable of **social** behaviour, that is, they can **communicate** and **co-operate** with each other and with users. Lastly, for agent to be truly intelligent, they must be able to **learn** as they react and interact with their external environment.

Current 'real world' agent applications, cover a large number of domains in industry, commerce, health care and entertainment, and range from comparatively small systems such as e-mail filters to large, open, complex, mission critical systems such as air traffic control.

3. AGENTS IN KNOWLEDGE MANAGEMENT

KM environments can be described as distributed system where different actors, acting autonomously on behalf of a user, and each pursuing its own goals, need to interact in order to achieve their goals. In such environments, the ability to communicate and negotiate is paramount. Furthermore, the number and behaviour of participants cannot be fixed a priori and the system can be expected to expand and change during operation, both in number of participants as in amount and kind of knowledge shared. The choice for multi-agent systems for KM is motivated by the following observations:

- KM domains involve an inherent distribution of data, problem solving capabilities and responsibilities (conforms to the ideas of autonomy and social ability of agents).
- The integrity of the existing organisational structure and the autonomy of its subparts need to be maintained (uses autonomous nature of the agents).
- Interactions in KM environments are fairly sophisticated, including negotiation, information sharing, and coordination (requires complex social skills with which agents are endowed).
- KM domains call for a functional separation between knowledge use and knowledge sources as a way to incorporate dynamic behaviour into information systems design (agents can act as mediators between source and application of knowledge).
- The solution for KM problems cannot be entirely prescribed from start to finish and therefore problem solvers are required that can respond to changes in the environment, to react to the unpredictability of business process and to proactively take opportunities when they arise (requires the reactive and proactive abilities of agents).

In KM environments, agents can check of the dynamic conditions of the environment, reason to interpret those perceptions, solve problems, draw inferences and determine actions, and finally, act accordingly. The use of agents in KM can be seen in two perspectives. In one hand, agents can be used to model the organisational environment where the KM system will operate and, on the other hand, software agents can be used to implement the functionality of KM systems. Most existing KM projects involving agent technology concentrate on using agents as implementation tool modelling primitives. Agents are used there to support and extend the activity of (human) users. However, more and more interest is arising about the advantages of agent-based modelling of KM environments. In the remaining of this section, we will describe both perspectives in more detail.

3.1. Agent-based models for KM environments

Agent-based models for knowledge management see agents as autonomous entities (like employees in a company) that are endowed with certain behaviours and the interactions among these entities give rise to complex dynamics. In this context, agents can be defined as ‘one that acts or has the power or authority to act’ or ‘one that takes action at the instigation of another’. The concept of agent in this sense is not new, nor restricted to software. In this perspective, agents are autonomous social entities that exhibit flexible, responsive and proactive behaviour.

An organisation can be seen as a set of entities and their interactions, which are regulated by mechanisms of social order and created by more or less autonomous actors to achieve common goals. Business environments must furthermore consider the behaviour of the global system and be able to incorporate collective characteristics of an organisation such as stability over time, some level of predictability, and clear

commitment to aims and strategies. While current research often takes individual agents as starting point and looks at interaction from the perspective of an individual agent, that is, how it affects and influences the goals and beliefs of the agent, agent models for organisations must take the perspective of the organisation as a whole. That is, multi-agent systems, or **agent societies**, must therefore be able to define the global aims as well as the roles and responsibilities of participants.

Agent societies represent the interactions between agents and are as such the virtual counterpart of real-life societies and organisations. Individual agents model specific roles in the society and interact with others as a means to accomplish the goals specified by those goals. This perspective makes the design of the system less complex since it reduces the conceptual distance between the system and the real-world application it has to model. Therefore, agent societies are an effective platform for virtual organisations because they provide mechanisms to allow organisations to advertise their capabilities, negotiate their terms, exchange rich information, and synchronise processes and workflow at a high-level of abstraction [19].

From an organisational perspective, the main function of an individual agent is the enactment of a role that contributes to the global aims of the society. That is, society goals determine agent roles and interaction norms. Agents are actors that perform role(s) described by the society design. The agent's own capabilities and aims determine the specific way an agent enacts its role(s). However, the society is often not concerned about which individual agent will actually play a specific role as long it gets performed. Several authors have advocated role-oriented approaches to agent society development, especially when it is manifest to take an organisational view on the application scenario [7, 22].

3.2. Using agents to implement KM systems

Knowledge Management Environments can be implemented as communities of different types of agents that collaborate to provide the required support to users on their knowledge intensive tasks. In agent-based implementations of knowledge management systems, software agents are employed as tools to manage loosely coupled information sources, to provide unifying presentation of distributed heterogeneous components and to personalise knowledge presentation and navigation. Possible agent-based services in a KM system are [15]:

- Search for, acquire, analyse, integrate and archive information from multiple heterogeneous sources,
- Inform us (or our colleagues) when new information of special interest becomes available,
- Negotiate for, purchase and receive information, goods or services,
- Explain the relevance, quality and reliability of that information,
- Learn, adapt and evolve to changing conditions.

These services are often specified in terms of the following types of agents:

Personal Assistants represent the interests of the user and provide the interface between users and the system. They are concerned with user preferences and needs, and will present information in the preferred format, at the right time. A proactive personal assistant agent will not only perform the tasks given to it by the user, but will also suggest knowledge sources or other resources that are not explicitly requested if they match the user's interests.

Cooperative Information Agents (CIAs) focus on accessing multiple, distributed and heterogeneous information sources. A CIA needs to maintain actively its information by communicating with others and reasoning about its own information.

Task analysts are agents that monitor a certain task in the business process, determine the knowledge needs of the task, and gather that knowledge by communicating with other agents. The agent can also monitor the execution of the task and evaluate the applicability of the knowledge provided. The lessons learned here are used to update its internal state and optimising task knowledge.

Source keepers are agents dedicated to maintaining knowledge sources and are responsible for describing the knowledge contained in the source and extract relevant information for a given request. Source keepers can also actively propose uses for its source to other agents based on its own knowledge of other agents' needs.

Finally, **mediators** are agents that can provide a number of intermediate information services to other agents. They may suggest collaboration between users with common interests, or provide information about the tools available. These agents contain specialised knowledge about the domain including where resources can be found.

4. APPLICATIONS OF AGENTS IN KM

In this section, we present some current projects that illustrate different possibilities for the applicability of agent technology to KM environments.

4.1. COMMA [4]

CoMMA (Corporate Memory Management through Agents) is a European that started in 2000 [1]. The main objective of the project is to implement and test a Corporate Memory management framework integrating several emerging technologies: agent technology, knowledge modelling, XML technology, information retrieval and machine learning techniques.

The use of multi-agents architectures in CoMMA is aimed at personalisation. That is, the resulting system must be able to adapt to the user, to the context, and support the retrieval of relevant information in the Corporate Memory. Agents in CoMMA can (1) communicate with others to delegate tasks, and (2) make elementary reasoning and decisions, supporting the choice between several documents. They use inference mechanisms exploiting ontologies that can be used to help authors of documents to annotate the documents, to perform technological monitoring on the Internet and to diffuse the acquired innovative ideas to the interested employees of the company. The

project focuses on the case where the corporate memory is materialised by XML documents and annotated by meta-information in RDF in order to offer intelligent search functionality and improve document retrieval. The project further aims at the exploitation of machine learning techniques in order to make agents adaptive to their users and the context. CoMMA makes use of pre-existing software framework for the development of agent applications, namely JADE, which is compliant with the FIPA specifications.

4.2. DIAMS

DIAMS is a system of distributed, collaborative information agents which helps users access, collect, organise and exchange information on the World Wide Web [3]. DIAMS was developed at the AIM Group, a research and development effort of NASA's Ames Research Centre and aims to encourage collaboration among users. Personal agents provide their owners with dynamic displays of well-organised information collections, as well as friendly information management utilities. Personal agents work closely with each other and with other types of information agents such as matchmakers and knowledge experts to facilitate collaboration and communication.

Object-based structure is used in information repositories to promote easy information sharing and exchange. A flexible hierarchical display is integrated with indexed query search to ensure effective information access. Automatic indexing methods are employed to support translation between user queries and communication between agents. Collaboration between users is aided by easy sharing of information and facilitated by automated information exchange. Connections between users with similar interests can be established with the help of matchmaker agents.

4.3. FRODO [10]

The FRODO (Framework for Distributed Organisational Memories) project, at the DFKI in Germany, aims at developing methods and tools for building and maintaining distributed Organisational Memories in a real-world enterprise environment [9]. Ontologies are often used as knowledge level middleware for distributed organisational memories. However, an ontology is not an eternal truth, but a socially constructed artefact that evolves over time. FRODO’s objective is to deliver a middleware for organisational memories based on a FIPA-compliant agent platform. Agent roles characterise ontology-related actors specified in terms of their respective goals and obligations in the ontology society. The suitability of the FRODO approach is being tested in an application scenario in the realm of KM for nuclear power engineering expertise. FRODO seeks to provide the following results:

- A flexible, scalable OM framework for evolutionary growth,
- A comprehensive toolkit for the construction and maintenance of domain ontologies, and
- Improve information delivery by the OM by developing more integrated and easier adaptable DAU techniques.

- A methodology and tool for business-process oriented knowledge management relying on the notion of weakly structured workflow.

4.4. KRAFT [17]

KRAFT supports the fusion of knowledge from multiple, distributed and heterogeneous sources [19]. Partners in the KRAFT project are the universities of Aberdeen, Cardiff and Liverpool and British Telecom. KRAFT has an agent-based architecture, in which all knowledge-processing components are realised as software agents. The architecture uses constraints as a common knowledge interchange format, expressed against a common ontology. Knowledge held in local sources can be translated into the common constraint language and fused with knowledge from other sources and is then used to solve some problem or deliver some information to a user.

The generic framework of the architecture can be reused across a wide range of knowledge domains and has been used in a network data services application as well as in prototype systems for advising students on university transfers and advising health care practitioners on drug therapies.

4.5. K-InCA [14]

In the project K-InCA under development at the Centre for Advanced Learning Technologies of INSEAD and sponsored by Xerox, agents are used to guide, monitor and stimulate managers towards the understanding of KM concepts and the adoption of KM practices in organisational contexts [20]. The K-InCA system behaves as a personal KM coach for its users. K-InCA agents can be seen as experts on organisational behaviour and change management assisting users in the transition from their current working habits to new habits that integrate some new behaviour (e.g. KM practices, entrepreneurial attitude, etc.). The system allows for different modes of interaction (practice and coaching) aiming at bringing the user to adopt a desired behaviour. In order to achieve this goal, agents react to the current user activity on the basis on information stored in a domain model and a user model, as well as through interaction with other agents. Target behaviours are described in the learning domain model.

4.6. Campiello [2]

The European project Campiello aims to experiment on the use of innovative ICT to develop new links between local communities and the visitors of historical cities of art and culture [16]. The objectives of the project are to better connect local inhabitants of historical places, to make them active participants in the construction of cultural information and to support new and improved connections with cultural managers and tourists. The system includes a recommender module, a search module and a shared data space. In order to facilitate the integration, tailoring and extensibility of these components, an agent model was chosen for the services in Campiello. The architecture supports interaction between distributed, heterogeneous agent and is built using a

Voyager platform extended with directory and broker services, administration tools and agent classes.

5. AGENT-MEDIATED KNOWLEDGE MANAGEMENT

We are developing a framework for Agent-Mediated Knowledge Management (AMKM) that uses agent concepts to analyse and model organisations and their knowledge needs, and to provide a reusable architecture to build KM systems. Different knowledge intensive tasks need knowledge from different sources and in different presentation formats. Therefore, the framework distinguishes between application, description and representation of knowledge and provides a common, uniform description of knowledge items (both sources and needs). A community of collaborative agents is responsible for the matching of knowledge supply and demand taking in account the user needs and preferences and the knowledge needs of a task. By collaborating with each other and with users, agents will learn and dynamically extend this framework by checking the current conditions of the environment. Agents will collaborate to interpret those perceptions, solve problems, draw inferences and determine actions, and finally, act accordingly. Information agents specialised in the different types of sources can provide this description.

5.1. AMKM model of organisations

The framework for agent societies we propose models the collective and interaction aspects of the society from an organisational perspective based on the notions of agent, role, norms, communication and goals. We propose a framework for agent societies consisting of three interrelated models each describing different aspects of the society that attempts to cope with the difference between desired order (from an organisational perspective) and actual behaviour (as actually realised by the participants) in dynamic environments [7]:

- The **organisational model** is the result of the observation and analysis of the domain and describes the desired behaviour of an agent society, as determined by the society 'owners' in terms of goals, rules, roles and interactions.
- The **social model** maps organisational roles to specific agents. Agreements concerning the role(s) an agent will play and the conditions of the participation are described in social contracts.
- The **agent model** specifies the interaction agreements between agents as interaction contracts. This model accounts for the actual (emergent) behaviour of the society.

A methodology to analyse a given domain and determine the type and structure of the agent society that best models that domain is described in [8]. Organisation theory shows that different organisations with exhibit different requirements for coordination and interaction. Coordination models are determined by transaction costs and reflect the balance between organisational objectives and activities. For example, the market model

fits well in an exchange situation, and the hierarchical model can be used in production settings. The methodology provides generic facilitation and interaction frameworks for agent societies that implement the functionality derived from the co-ordination model applicable to the problem domain. Standard society types as market, hierarchy and network, can be used as starting point for development and can be extended where needed and determine the basic norms and facilitation roles necessary for the society.

5.2. AMKM architecture

The AMKM architecture based on the model described in [8] and [7] consists of two layers: **operation** and **facilitation** as depicted in figure 1.

At production level, there are basically three types of agents: personal assistants, (business) process task analysts and knowledge source keepers. Depending on the application area, other agent types may be needed what can be determined by the application of the methodology introduced in [8]. Each of those agents provides a transparent access to its organisational background. That is, a personal assistant concentrates on the fulfilment and description of its user needs and does not need to know the type and format of knowledge sources, which are encapsulated by source agents, with whom the personal assistant can communicate. The facilitation level helps agents to locate each other, based on their needs and facilities. That is, at facilitation level, the 'norms' of the society are kept and enforced and interaction is ensured. Furthermore, facilitation agents ensure interaction by monitoring and supporting contract formation, take care of introducing new agents to the rules of the society and keep track of the reputation of trading agents. Typical facilitation agent roles are matchmakers, gatekeepers and reputation agents.

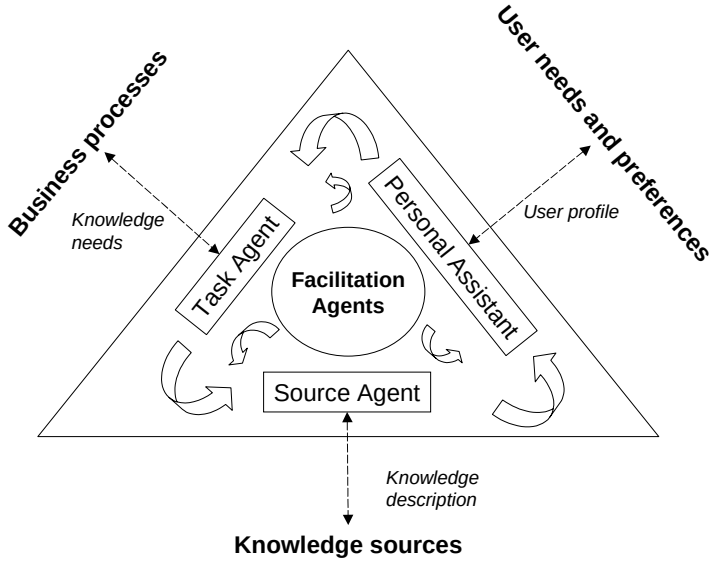

Fig. 1. Generic architecture for Agent-mediated Knowledge Management

5.3. An application of AMKM framework

Based on the AMKM architecture we are developing a Knowledge Market to support non-life insurance experts to exchange knowledge with each other, in a way that preserves the knowledge, rewards the knowledge owner and reaches the knowledge seeker in a just in time, just-enough basis [5]. Non-life insurance experts possess knowledge, experience and contacts of great value to each other. One of the main tasks of this system is to support and encourage collaboration and knowledge sharing. In this environment, both knowledge seekers as knowledge owners want to be able to decide on trade partners and conditions. Sharing is not centrally controlled but greatly encouraged by the management. The best-suited partner, according to each participant's own conditions and judgement, will get the 'job'. Furthermore, factors such as privacy, secrecy and competitiveness between brands and departments may influence the channels and possibilities of sharing and must thus be considered.

The architecture of the Knowledge Market is illustrated in figure 2. In this case, the operational roles deduced from the requirements, knowledge seeker and knowledge owner, are both specific aspects of personal assistants. This is because that the aim of this system is direct collaboration between human experts and not, as in the generic case described above, to integrate user descriptions, knowledge sources and business tasks. People seeking collaboration can initiate through the user interface of the Knowledge Market a personal agent that will act as their avatar in the system. This agent will use the preferences and conditions specified by the user to find appropriate partners and negotiate exchange terms. The 'goods' to be exchanged are knowledge descriptions.

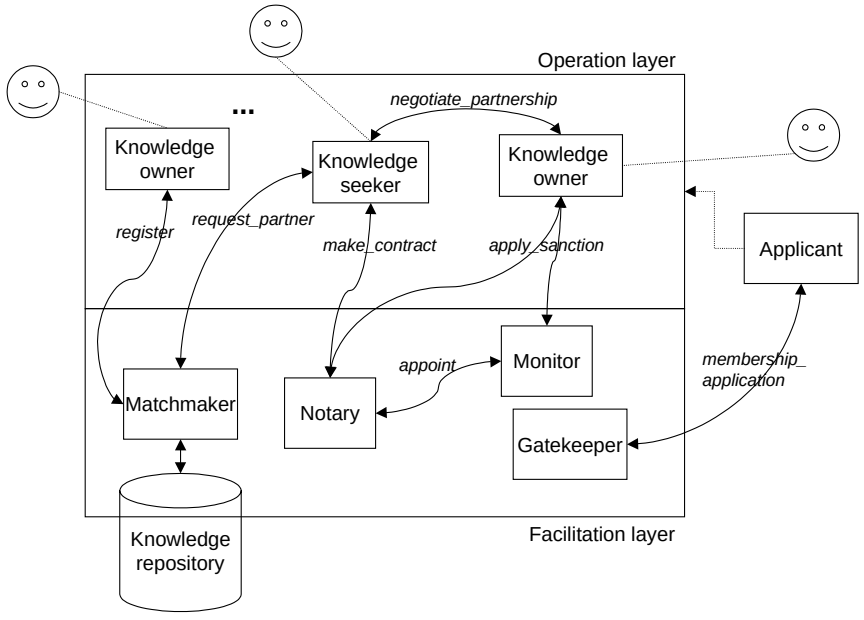

Fig. 2. Knowledge Market architecture

The following scenario illustrates the functionality of the knowledge market: Anne is working in the development of a new product for liability insurance. She requires an analysis of the products offered by competitors that is not directly available in the

knowledge repository. Via the user interface, Anne communicates her needs and conditions to her personal assistant. The conditions set by Anne refer for example to the deadline (when is it needed), type of source preferred (does she prefer to talk to someone, to read a document), and what is she willing to 'pay' for the result. Anne's assistant will then join the Knowledge Market where it assumes the role of *knowledge seeker* and contacts the *matchmaker* in order to find out potential partners. Using its own internal information about knowledge owners in the system, or possibly, referring to the knowledge repository to find out possible knowledge sources matching the request, the *matchmaker* will provide a list of *knowledge owners* to the *knowledge seeker*. Following its own strategy and the preferences specified by Anne, her personal assistant will then contact knowledge owners and try to get the best deal for Anne's request.

Matching and decision-making strategies differ for each agent and can be determined by its user. For example, an agent may decide to negotiate with all possible matches from the list supplied by the matchmaker and another will just contact the first one. In our example, Anne's personal assistant will make a commitment with the *knowledge owner* acting on behalf of Paul in order to receive the information requested by Anne. Such commitments are registered with the *notary* agent that will monitor its execution. Depending on the exact request and conditions, Anne and Paul will possibly need to contact each other in order for the requested knowledge to be exchanged.

6. OPEN RESEARCH ISSUES

Current developments in KM indicate a need for systems that are reactive and proactive in relation to the needs and expectations of its users. In this paper, we have discussed the role of agents in the design and functionality of knowledge management environments. In such environments, the flow of knowledge within an organisation (or organisations) must take in account not only the knowledge needs of business processes, but also the personal preferences and level of expertise of individual users.

Agent concepts, which originated in artificial intelligence but which have further developed and evolved in many areas of computing, hold great promise for responding to the new realities of knowledge management. While there are many conceptions of agents, most have embodied higher levels of representation and reasoning involving knowledge/belief, perception, goal, intention, and commitment. On the one hand, the technical embodiment of these concepts can lead to advanced functionality of KM systems, e.g. personalisation of knowledge presentation and matching supply and demand of knowledge. On the other, the rich representational capabilities of agents as modelling entities allow more faithful and effective treatments of complex organisational processes.

Although several projects are currently taking place that use agents for the design and implementation of KM systems, several research issues are still open that will provide a clear view of the applicability and consequences of using agents in KM. While the agent research area is very active, its concerns towards KM are not yet well covered. Agent concepts could fundamentally alter the nature of knowledge management both in the way KM systems are build as well as the way organisations are analysed and modelled.

Future research in agent-oriented approaches to knowledge management and collaborative systems must include:

- Methodologies are needed that support the analysis of knowledge management needs of organisations and its specification using software agents and agent societies
- Reusable agent-oriented knowledge management frameworks, including the description of agent roles, interaction forms and knowledge description
- Agent-based tools for organisational modelling and simulation that help determine the knowledge processes of the organisation,
- The role of learning in agent-based knowledge management systems, namely, how to use agent learning to support and extend knowledge sharing.

7. REFERENCES

[1] Bergenti, F., Poggi, A., Rimassa G.: Agent Architectures and Interaction Protocols for Corporate Management Systems. In: Proc. ECAI Workshop on Knowledge Management and Organizational Memories, Berlin, Germany, (2000), 39-47.

[2] Campiello: http://klee.cootech.disco.unimib.it/~campiello/

[3] Chen, J.R., Wolfe, S.R., & Wragg, S.D.: A Distributed Multi-Agent System for Collaborative Information Management and Sharing. In: Proc. 9th ACM International Conference on Information and Knowledge Management, (2000).

[4] CoMMA: http://www.si.fr.atosorigin.com/sophia/comma/Htm/HomePage.htm

[5] Dignum, V.: A Knowledge Sharing Model for Peer Collaboration in the Non-Life Insurance Domain. In: Proc. German Workshop on Experience Management, Lecture Notes in Informatics, German Society for Informatics, Berlin, (2002).

[6] Dignum, V.: Towards a People-Oriented Knowledge Management Environment. In: Proc. 11th Workshop on Database and Expert Systems Applications – DEXA2000, IEEE Computer Society, (2000), 1134-1138.

[7] Dignum, V., Meyer, J.J. and Weigand, H.: Towards an Organizational Model for Agent Societies Using Contracts. In: Proc. 1st International Joint Conference on Autonomous Agents and Multi-agent Systems, Bologna, Italy, (2002). To appear.

[8] Dignum, V., Weigand, H., Xu L.: Agent Societies: Towards framework-based design. In: Wooldridge, M., Ciancarini P., Weiss, G. (Eds.): Agent-Oriented Software Engineering II, LNCS 2222, Springer-Verlag, (2002), 33 – 49.

[9] van Elst, L., Abecker, A.: Domain Ontology Agents in Distributed Organizational Memories. In: Proc IJCAI'2001 Workshop on Knowledge Management and Organizational Memories, Seattle, Washington, (2001), 39-48.

[10] FRODO: http://www.dfki.uni-kl.de/frodo/

[11] Jennings, N., Faratin, P., Johnson, M., Norman, T., O'Brien, P., & Wiegand M.: Agent-based business process management. International Journal of Cooperative Information Systems, 5 (2&3), 1996, 105 – 130.

[12] Jennings, N.: On agent-based software engineering. Artificial Intelligence, 117, (2000), 277-296.

[13] Jennings, N. & Wooldridge, M.: Applications of Intelligent Agents. In: Jennings, N. and Wooldridge, M. (Eds.): Agent Technology: Foundations, Applications and Markets, Springer-Verlag, 1998, 3 – 28.

[14] K-InCA website: http://www.calt.insead.edu/project/K-InCA/

[15] Klusch, M. (Ed.): Intelligent Information Agents: Agent-based Information Discovery and Management in the Internet. Springer-Verlag, (1999).

[16] Koch, M.: Knowledge Management and Knowledge Agents in Campiello. In: Lees, B., Müller, H., Branki, C. (Eds.): Proc. Workshop on Intelligent Agents in CSCW, Dortmund, Germany, (1998), 44-52,

[17] KRAFT: http://www.csc.liv.ac.uk/~kraft/

[18] Leake, D., Birnbaum, L., Hammond, K., Marlow, C. & Yang, H.: Task-Based Knowledge Management. In: Proc. AAAI Workshop Exploring Synergies of Knowledge Management and Case-Based Reasoning, Orlando, Florida. AAAI Technical Report WS-99-10. (1999).

[19] Preece, A., Hui, K.-Y., Gray, W., Marti, P., Bench-Capon, T. Jones D. & Cui Z.: The KRAFT Architecture for Knowledge Fusion and Transformation. 19th SGES International Conference on Knowledge-based Systems and Applied Artificial Intelligence, (1999).

[20] Roda, C., Angehrn, A. and Nabeth, T.: Matching competencies to enhance knowledge sharing: An Intelligent Agent Approach. In: Proc. 7th International Netties Conference, September 2001, Fribourg, Switzerland.

[21] Staab, S. and Schnurr H.: Knowledge and Business Processes: Approaching an Integration. In: Proc. Workshop Knowledge Management and Organizational Memory (OM '99). Stockholm, Sweden, 1999.

[22] Zambonelli F., Jennings, N., Wooldridge, M.: Organizational Abstractions for the Analysis and Design of Multi-Agent Systems. In: Ciancarini P., Wooldridge, M. (eds.): Agent-Oriented Software Engine