

A Distributed Intelligence Paradigm for Knowledge Management

M. Bonifacio

Knowledge Services Group
Arthur Andersen Business Consulting
P.zza S. Vito, 37 - 31100 Treviso (Italy)
matteo.bonifacio@it.arthurandersen.com

P. Bouquet and A. Manzardo

Dept. of Computer and Management Sciences
University of Trento
Via Inama, 5 - 38100 Trento (Italy)
{bouquet,manzardo}@cs.unitn.it

Abstract

In the last ten years, *knowledge management* (KM) has become a new fashioned managerial practice. Though KM theories seem to benefit from a “contamination” with cognitive and social sciences, which emphasize a subjective, contextual, and distributed approach to knowledge representation and integration, current technologies support what we may call a “god’s eye” paradigm, in which knowledge is viewed as an objective resource. In this paper we discuss artificial intelligence theories and technologies that can support a shift to a new paradigm, called the “distributed intelligence” paradigm, in designing KM systems. Using the evolution of KM systems within Arthur Andersen Consulting as a motivating case study, we propose the framework of *MultiContext Systems* as a specification language for distributed intelligence KM systems, and sketch an *agent-based* architecture as an example of a KM system which embodies the assumptions of the distributed intelligence paradigm.

Introduction

In the last ten years, *knowledge management* (KM) has become a new fashioned managerial practice. Under this name, managers foresee the opportunity to control the processes of producing, distributing and using a “new” valuable resource: *knowledge* (Drucker 1994). This complex matter cannot be reduced to the traditional concepts of land, labor and financial. As a matter of fact, we need a “contamination” process between managerial practices and theoretical disciplines that are historically devoted to the study of knowledge and cognition, such as philosophy, artificial intelligence (AI), social sciences, psychology.

From a managerial point of view, this process of contamination is not linear, and somehow contradictory. On the one hand, cognitive and social disciplines increasingly regard *subjectivity*, *contextuality* and *distribution* as intrinsic features of knowledge. On the other hand, following a traditional approach to organizational life, current KM systems seem to implicitly rely on an

objective view of knowledge. In this situation, managers are facing the substantial failure of KM systems and have the intuition that this failure is somehow related to inadequate assumptions about the nature of knowledge (see e.g. (Nonaka & Takeuchi 1995)). However, though managers are attracted by the emerging approaches to knowledge, they are unable to accept the underlying paradigmatic shift. The result is a situation where KM systems are nominally consistent with emerging issues on the nature of knowledge, but are still substantially based on traditional assumptions.

Current KM systems are inspired to a traditional paradigm in organizational life based on the idea that the purpose of a KM system is to represent and organize knowledge into a single, shared, and coherent, structure (e.g. a taxonomy, an ontology), independently of when, how, where, and why it was originally produced. We call this paradigm, that descends from a traditional approach to cognition in social systems, the “god’s eye” paradigm (GEP). The emergent paradigm in KM, which we call the “distributed intelligence” paradigm (DIP), is based on the idea that knowledge is always and irreducibly distributed into multiple contexts of knowledge production (individuals, groups, time periods and spatial locations, and so on), and therefore cannot in general be straightforwardly organized into a single, shared and coherent structure. This dichotomy between paradigms cannot be discharged as a purely philosophical issue. Adopting one view or the other inevitably leads to completely different conceptual and architectural choices in designing a technology enabled environment for KM, and these choices can be the reason for the success or the failure of a KM system. In this paper we discuss artificial intelligence (AI) theories and technologies that can support a shift from the GEP to the DIP in designing KM systems. We first discuss the basic assumptions of the GEP and the DIP using the evolution of KM systems within Arthur Andersen Consulting as our case study. Then we propose the framework of *MultiContext Systems* as a specification language for distributed intelligence KM systems, and sketch an *agent-based* architecture as an example of DIP based KM system for representing and integrating knowledge.

The evolution of the KM in AA

Arthur Andersen Consulting (AA) is a global consulting firm organized in four main divisions that offer tax and legal, business consulting, financial and auditing services to medium and large size target clients. This services are provided by nearly 120.000 professional consultants that operate across 74 countries, and are organized in a complex matrix given by the intersection of several target industries (e.g. telecommunications, automotive) and service categories. Country specificity, industries and service categories are the basis of a work content that completely relies on the use of knowledge as the primary resource of value. That is to say, a consultant's work is primarily given by the ability to apply business knowledge to solve a client's business issues or problems. AA consultants are thus facing a particularly and increasingly complex, dynamic and differentiated environment. As a result, knowledge used by consultants becomes increasingly complex, dynamic and differentiated to effectively respond to emerging business needs.

The information and communication technology (ICT) revolution and the related explosion of business complexity has influenced the *knowledge needs* of consultancy and, as a consequence, the nature of a consultant's work. The revolution raised new issues and questions:

- the necessity of continuously revise and renew knowledge raised the issue of managing knowledge;
- the ability of consultants to effectively understand their business environments through their daily work slowed consultants and their work as a primary source of knowledge;
- the impossibility for a consultancy entity (an office, a team, an individual and so on) to retain enough knowledge to deal with environmental complexity posed the issue of making organizational knowledge (the knowledge produced in the overall organization) available to each consultant.

The opportunity of ICT to store and communicate information at a low cost represented the natural path to solve this knowledge managing issues. AA gave different concrete answers in terms of managing knowledge within a technology enabled environment. Briefly we describe some main approaches, which we interpret from an evolutionary perspective.

In a first approach, named the *Knowledge Base (KB) approach*, knowledge is seen as a set of formal representations of given knowledge domains. The latter are given by industries and service categories while the former by documents such as methodologies, presentations, articles and engagements descriptions. Knowledge is collected through a centralized process of gathering, organizing and formalizing context specific experiences that through generalization can be used in other contexts. Formal repositories of knowledge are called *Knowledge Bases*, initially distributed to consul-

tants on a CD support and later made directly accessible through a corporate Intranet. Specialized organizational entities played the role of knowledge producers, while consultants played the role of knowledge consumers.

A second approach, the *AAOnLine approach*, arose from some serious limitations displayed by the KB approach. On the one hand, the attempt to produce generally valid "knowledge packages" generated oversimplified contents that were practically useless, except as high level descriptions of consultants' knowledge domains. From this point of view, knowledge consumers didn't consume centrally packaged knowledge. On the other hand, it became evident how consultants' daily experiences and tacit knowledges weren't completely packageable. Rather than collecting and formalizing these experiences, AA decided to make them accessible at a source level. The basic idea was to give consultants a global discussion forum where each consultant could discuss, ask and give information about any topic related to work. This global forum, named AAOnLine, was organized in *communities* of interest, each community being a virtual space where consultants with a common interest could collaborate and share their individual knowledge. Each community was organized by a *Knowledge Manager* who had the role of supporting and enabling the interaction within the community.

However, the AAOnLine approach displayed some limitations as well. After a first period of enthusiasm, the number of participants and contributions begun to decrease. This reduced the quality of available contents, which in turn discouraged participation and contribution. A possible explanation of this partial failure is that AAOnLine was a virtual space where people who speak different languages (not just intended as country specific) could met without having enough technological and conceptual tools to understand each other. Context specificity, in the terms of linguistic and semantic differentiation, was represented by the specificity of individual contributions, each contribution being characterized by a particular use of language. Nonetheless, this diversity remained implicit and unexpressed, so that *accessibility to individuals contributions didn't correspond to accessibility to individuals' interpretive perspectives*. In other words, within AAOnLine consultants were asked to read other consultants' contributions without having the "linguistic key" to interpret the intended meaning of used terms. Within a physical environment, this opportunity is ensured by the dialectic process of *meaning negotiation* that is hardly reproducible within an asynchronous communication environment like AAOnLine. As a result AAOnLine, rather than solving the problem of linguistic and semantic heterogeneity, became the expression of its apparent irreducibility.

The KB approach made evident the need to consider knowledge as a tangible resource, organized and represented in a way that fosters its accessibility and, as a consequence, the value of its replication. The limit was an oversimplified view of the process of knowledge pro-

duction that underestimates – or perhaps ignores – the problem of context specificity. On the other hand, the AAOOnline approach made manifest the need to consider contexts as a constitutive dimension of knowledge rather than a “noise” that interferes and negatively influences knowledge production. The limit was the lack of both technological and conceptual tools to sustain meaning negotiation and knowledge exchange by making contexts accessible through an appropriate representation and interaction process. As a result AAOOnline recognized the value of contexts without providing tools to transform this potential in actual value. We may say that the KB approach focused on knowledge as a product without considering the nature of its generative process, while the AAOOnline approach focused on the constitutive elements of knowledge production without providing enough tools to configure it as a reusable product.

The need to consider knowledge both as a product characterized by accessibility and replication and as a process characterized by context specificity and meaning negotiation is the current process of AA knowledge systems. Nonetheless, although the issue has now been made explicit, the solution is still to be found. Currently, AA is experimenting a “third way” that tries to merge the two issues, called the *KnowledgeSpace approach*. KnowledgeSpace is a corporate Intranet that is characterized by the following features:

- both KBs and communities are integral part of the system. KnowledgeSpace is a virtual space organized in interest communities, each community being a complex space of knowledge repositories, interaction areas, roles and processes to ensure the vitality of the community;
- KBs tend to become the expression of each community. That is to say, each consultants community is somehow related to a KB that is intended as the “tangible” and “accessible” expression of a community knowledge and context. A community knowledge repository tends to become the outcome of a community contribution process;
- communities are experimenting tools that enable them to exploit their knowledge within their contexts. At the moment, these tools are given by structured contribution processes, increasingly complex navigation knowledge maps, and categorization structures. Knowledge Managers play the role of keeping the community knowledge and context accessible and understandable to visitors belonging to other communities.

AA is our case study in the development of an innovative KM system. The paper focuses on new conceptual and technological tools that can support *knowledge representation* within a community and *knowledge integration* through meaning negotiation across different communities.

The “god’s eye” paradigm

The paradigm of knowledge representation and integration that inspired the KB approach – and perhaps is still dominant in most KM systems – is based on the following general assumptions:

1. knowledge is independent from any subjective element, and thus the process of knowledge production is independent from the identity of the knower. The same knowledge could be produced by one person or more, in different groups, in different places, and so on;
2. knowledge can be represented in systems of general and abstract propositions, organized in a single, coherent, and sharable structure;
3. linguistic heterogeneity and semantic differentiation, which derive from context and specificity, are seen as limitations to the potential value of knowledge, and therefore must be wiped out through a process of homogenization (abstraction, generalization, ...).

This leads to a fundamental consequence in terms of how a knowledge system is to be organized, namely the possibility to use economies of scale in the process of knowledge production. Sub-consequences are: the opportunity to centralize knowledge production processes in order to minimize unitary costs; the opportunity to specialize people in the knowledge production process in order to gain efficiency; the opportunity to maximize the value of knowledge by replication and use in any context.

Concepts such as sociality and identity, though widely used in KM, are considered as human needs that are to be satisfied without interfering with the pure process of cognition. It is recognized that knowledge is often embedded into a context (tacit and implicit are terms frequently used (Polany 1966)). But its exploitation requires a sort of pre-processing, that is to say generalizing what is specific and abstracting away what is contextual. The result of this pre-processing should be a corpus of knowledge that can be shared within an organization. In conclusion, knowledge is related to the concept of “absolute truth”, learning to the passive process of acquiring what is already given, and inquiring to the attempt to see “the true reality” through the “god’s eye”. That’s why we call this paradigm, the “god’s eyes” paradigm (GEP). Here we present some elements that characterize a KM system consistent with the GEP:

- knowledge is stored in Knowledge Bases, repositories of formal representations about knowledge domains;
- knowledge is represented with a language that tends to be context independent: there is a univocal relation between meanings, words, and objects;
- knowledge is produced by team of experts deductively (by articulating methodologies from theories) and inductively (by extracting best practices from experiences);

- learning occurs when new truths are discovered (methodologies and practices) and acquired through use and study by knowledge consumers.

A critique of the GEP

The GEP has some obvious advantages, as it promises of eliminating from the outset problems such as conceptual or semantic heterogeneity between different representations. However, our intuition is that the GEP simply begs the question of knowledge integration, without facing the real complexity of the problem. Indeed, it does not address a number of crucial issues, such as the following:

- first of all, it underestimates the complexity of knowledge within a big organization. Designing a single structure (e.g. a taxonomy, an ontology) where every piece of knowledge produced within an organization can be coherently placed is very hard work, and it is not obvious at all that it is feasible in practice, when the approach is scaled up to large and complex organizations (such as AA);
- second, the GEP is not elaboration tolerant, namely it is not adequate to cope with highly dynamic environments. Small changes in the organization may require a very complex revision or even a complete re-design of the entire knowledge structure, and this is not desirable;
- third, it overrides the categories, concepts, terms that a professional group uses in its everyday activity. This could be a reason for a sub-optimal usage of the system, or even a complete failure;
- finally, perhaps the most serious issue is the assumption that knowledge representation can be completely de-contextualized by mapping local (idiosyncratic) conceptual schemas into a general (shared) conceptual schema.

The fact that knowledge (and its representation) is irreducibly contextual has been advocated by several leading researchers both in AI (McCarthy 1993; Lenat & Guha 1990; Guha 1991; Giunchiglia 1993) and in other knowledge related disciplines (Fauconnier 1985; Dinsmore 1991; Perry 1998) (see (Bouquet *et al.* 1999) for an interdisciplinary collection of papers on this topic). The common intuition is that the content of any linguistic representation depends on a collection of assumptions, on the purpose for which it was produced, on the intentions of an agent, and so on (a similar argument can be made for non linguistic representations, but here we are not concerned with this aspect of the problem). Whatever the right explanation for this irreducibility (the non existence of a so called reality, its inaccessibility through our perception, of its complexity as a constraint to understanding), the immediate consequence is the impossibility to establish a non ambiguous relation between words, meanings and objects. The meaning of a word is not always self evident. Thus it

is perfectly conceivable (and we see it happen in everyday life) that speakers who share the “same” language do not necessarily share its semantics, and so associate different content to the same word.

This is what happened within the AAOnline approach. Different professional communities produce (and represent) knowledge autonomously, use idiosyncratic lexicons, different concepts and implicit taxonomies. Managing knowledge is not just a matter of creating a shared repository, or having a team of Knowledge Managers that impose on it a general structure. We need new concepts that allow us to deal with conceptual and semantic heterogeneity. In short, we need a different paradigm.

The “distributed intelligence” paradigm

The paradigm we sketch in this section aims at overcoming some of the drawbacks of the GEP. It is inspired to recent work in distributed AI, and its basic assumptions can be summarized as follows:

- knowledge is context dependent, and its representation depends on contextual factors such as purpose, background assumptions, available resources, and so on;
- linguistic heterogeneity and semantic differentiation are not accidents that must be eliminated, but essential properties of knowledge that must be dealt with in the process of knowledge integration.

The basic intuition is that locally produced knowledge cannot be represented by mapping it onto a universal structure, because we cannot assume that this structure is shared and understood in the same way by different knowers (or groups of knowers). This does not mean, however, that knowledge cannot be integrated. Only, a different mechanism should be used. If there is no single, objective perspective from which knowledge can be represented, knowledge integration can only be the result of a process of *meaning negotiation* between autonomous entities. Integrating knowledge is therefore a mechanism of social agreement. As a consequence, knowing appears as the intersubjective process of negotiating relations between symbols, and knowledge as the system of negotiated symbols and meanings at a given time and place. Given this general assumptions, knowledge becomes intrinsically social, as it cannot be separated from the social agreement that sustains its constitutive elements.

Many are the consequences of these assumptions. We will emphasize only those that are relevant for our purposes.

- First, since knowledge “exists” in the context of a negotiation process, knowledge has no existence when considered apart from its context. That is to say, contextual elements become constitutive of knowledge itself. As a consequence, any knowledge, when intended as a system of negotiated meanings and symbols, is to be considered as a two face entity of sym-

bolic and mental representation within an interpretive context.

- Second, it is possible to imagine as many “knowledge sets” about the same domain as the possible interpretive contexts through which this domain is readable. As a consequence, the abstract concept of knowledge can concretely be viewed as a system of possible *local* knowledge sets.
- Third, learning appears, on the one hand, as the social process of generating, criticizing and reviewing interpretive contexts and, on the other hand, as the process of generating propositions within interpretive contexts. The problem of knowledge transfer and exchange involves the transfer and the accessibility of both symbolic propositions, and interpretive contexts through which symbols are able to gain a particular meaning.
- Fourth, diversity, redundancy and heterogeneity are seen as intrinsic qualities of a knowledge system. Basically, the many to many relationship that links a symbolic to a meaning system is the direct outcome of the many possible ways to read and represent the same domain. That is to say, a given symbolic system is able to represent as many thought worlds as the possible social perspectives through which a world can be read.
- Fifth, for our purpose, a context is an unspecified expression of social subjectivity. In other words, whether we assume a context to be a collective identity, a social paradigm, a task oriented use of a world or a mental model, the fundamental consequences of our description remain unchanged.

Within this framework, a KM system, viewed in its cognitive aspects, appears as a “distributed intelligence system”, and the emerging paradigm becomes a distributed intelligence paradigm (DIP). With this term we refer to a system where knowledge is distributed between social actors, each social actor being involved in the process of constituting knowledge by considering knowledge domains through negotiated and subjective interpretive perspectives.

The development of KM systems based on the DIP poses a great variety of challenges for researchers and practitioners. First of all, we need to take into account the social aspects of cognition and investigate how knowledge is generated and represented through social structures, processes and roles. In particular, assuming the community as a lens to read organizations as social learning systems, a distributed intelligence system appears as a constellation of communities, each community being a social context and an interpretive perspective (Orr 1990; Brown & Duguid 1991; Wenger 1998).

From a cognitive point of view, a second challenge is explaining how groups of individuals can communicate without sharing the interpretive contexts. In particular, assuming mapping processes as a mean to meaning

making, a distributed intelligence system appears as a constellation of interpretive maps and of mapping processes between different individuals and communities.

Third, from a technological point of view, the representation of interpretive context and the analysis of meaning negotiation processes can be a valid basis for research in the intelligent agents field. In particular, a distributed intelligence system within a technology enabled communication environment can benefit of technological agents able to “socially negotiate” information under the light of locally and contextually represented “knowledge”.

In this paper we do not address the first issue, as we are not personally involved in that research area (though we should mention the fact that the study of social aspects of cognition is part of the project within which the research presented in this paper was developed). Instead, we concentrate on the description of a formal framework for context-based knowledge representation, and sketch a high-level architecture of a multi-agent system for implementing a distributed intelligence KM system.

MCS for knowledge representation and integration

In this section, we show how to use *MultiContext Systems* (MCS) as a formal framework in which knowledge representation and integration in a distributed intelligence system can be specified. MCS were introduced in (Giunchiglia 1993). Their semantic counterpart, *Local Models Semantics*, was presented in (Giunchiglia & Ghidini 1998). A foundational account of MCS can be found in (Benerecetti, Bouquet, & Ghidini 2000). MCS have also been used as a specification language for multi-agent systems (Sabater *et al.* 1999).

MCS can be described as a logic of relationships between local, independent representations. The logic is based on two very general principles (Giunchiglia & Ghidini 1998). The first, named *principle of locality*, says that reasoning is intrinsically local, namely happens within a given context. The second, named *principle of compatibility*, says that reasoning in a context is partly constrained by its relationship with reasoning processes that happen in other contexts. These principles are given both a proof and a model theoretical formalization.

Locality. *Proof theoretically*, each context is associated with a logical theory, finitely presented as an axiomatic formal system $\langle L, \Omega, \Delta \rangle$ (where L is a formal language, Ω is a set of axioms, and Δ is a correct and complete set of inference rules defined over L). Notationally, we write $c_i : \Phi$ to express the (metalinguistic) fact that Φ is a formula of the context c_i (i.e. $\Phi \in L_i$). *Model theoretically*, each context c_i is characterized as a set of models (called *local models*) of the language L_i (for the moment, we require that these models satisfy at least the theorems of the theory associated with c_i).

Compatibility. *Proof theoretically*, compatibility is

formalized as a collection of *bridge rules*, namely inference rules whose premisses and conclusion belong to different contexts, for example:

$$\frac{c_i : \Psi}{c_j : \Phi}$$

where $c_j \neq c_i$. *Model theoretically*, this correspond to impose a compatibility relation R over sets of local models of c_i and c_j (if M_i and M_j are the set of all local models of c_i and c_j respectively, then $R \subseteq 2^{M_i} \times 2^{M_j}$). The bridge rule above would correspond to the following relation: a set of local models of c_i that satisfies the formula Ψ is compatible only with those sets of local models of c_j that satisfy the formula Φ (where satisfiability is local satisfiability in the theories associated with the two contexts).

The proof theoretical effect of the principle of compatibility is that it increases the set of theorems which are locally derivable in a context (with respect to the theorems that can be derived in the associated theory taken in isolation). The model-theoretical effect of the principle of compatibility is that it cuts off the set of local models that satisfy a context (again, with respect to the models that satisfy the associated theory taken in isolation), as it eliminates the (sets of) local models that are not compatible with (sets of) local models of other contexts.

Letting aside the technicalities, MCS are a highly flexible and modular way to formalize a collection of local representations and to model the process of knowledge integration (using compatibility relations) without resorting to the GEP. As an example, consider the interaction between a user and a technician who has been called to repair a photocopier. Each of them has a representation of the machine which is *partial* (e.g. the user knows very little of what is inside the machine, whereas the technician has no information about the “history” of the machine and the conditions in which it’s been used), *approximate* (the user and the technician have knowledge about the machine at very different level of detail), and *perspectival* (the user’s perspective on the machine is quite different from the technician’s perspective: the first has to use it, the second to repair it). In short, the user and the technician have local representations of the photocopier. In MCS, this means that we represent what the user and the technician know as two different contexts, each with its own representation language (they do not completely share the lexicon), its set of axioms (they have different information), its inference rules (namely we assume that they have the same reasoning abilities). Notice that the languages of both contexts are interpreted over a set of local models of the two contexts.

The crucial question now is: how do the user and the technician integrate what they know in order to communicate and cooperate in solving the machine’s problem? The intuition is that knowledge cannot be shared across contexts, but that the fact that they are talking of the same machines imposes a compatibility

relation between their representations. Some aspects of this relation can be known (e.g. the technician may know how to map part of what the user says into a more technical language), some need to be learned in the communication through a process of meaning negotiation. Whichever the case, the relation can be represented as a collection of bridge rules, namely rules that allow the technician to map onto his/her language what the user says (and vice versa). Notice that, in general, knowledge about this relationship can be incomplete (the user and the technician may have only partial knowledge about it), and even worse can be incorrect. Incompleteness means that the user or the technician (or both) lack some bridge rule; incorrectness means that they are not using the right bridge rules. Bridge rules (or, correspondingly, compatibility relations) are the way MCS formalize knowledge integration.

Notice that, as a consequence of accepting the DIP, both incompleteness and incorrectness cannot be eliminated *a priori* from the system, but must be detected in the communication process. In a conversation, there are some typical situations that allow us to realize that something is going wrong. For example, we can imagine that the technician uses a term that the user has never heard before; or the user describes a problem of the machine in such a way that the technician cannot make sense of it. In these situations, the two speakers start a process of meaning negotiation, whose goal is to establish new links between local representations (e.g. learning a new word, learning that a word has a different meaning, learning how to map a functional problem into a technical description, and so on).

These ideas, that in many respects recall ideas discussed from an organizational perspective in works such as (Weick 1993; Boland & R.V.Tenkasi 1994), have a direct application to KM. The case of AA described in section 4 is a paradigmatic example of how eliminating contextual aspects of knowledge (KB approach), or the possibility of meaning negotiation across context-dependent representations (AAOnLine approach), may lead to failures.

An Agent Based Architecture for Knowledge Integration

The high-level architecture of a multi-agent system (MAS) for knowledge integration described in this section is an attempt of characterizing a KM system in the spirit of the DIP. We use MAS because, in our view, agents are the most promising technology for realizing distributed systems whose components are to be autonomous and pro-active, and show social abilities. In the present section, we assume some familiarity with MAS; a recent handbook on theory and practice of MAS is (Weiss 1999).

The architecture, called KI-MAS, is described in figure 1. It has four macro-levels:

Knowledge Centers (KC). A KC is a knowledge source which represents knowledge that is produced and

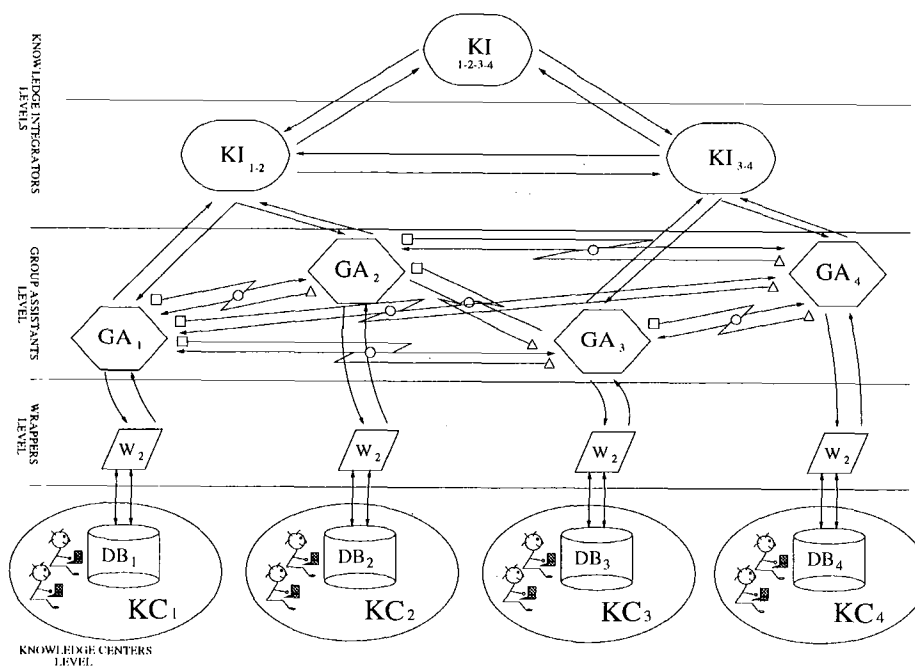


Figure 1: KI-MAS architecture

used by a professional community. In principle, it could be anything, from databases to human operators. In our case study, each KC is a taxonomy of documents, such as papers, internal reports, mail messages, transparencies, automatically generated by an application running on top of Lotus Notes. A KC is viewed as a context where knowledge is managed (i.e. produced, represented, organized, maintained, validated, updated).

Wrappers. The level of *wrappers* provides the service of software integration. According to FIPA's specification standards, wrappers are agents that interface non agent-based software with a MAS. Basically, they have two components: on one side, they know "how to talk" with the software they wrap; on the other side, they know how to translate I/O from the software in messages of an agent communication language (such as KQML, or FIPA's ACL). Though wrappers are a very specific type of agent, their role is very important, as they allow to integrate many KC without forcing people to learn how to use a new software they are not familiar with.

Group Assistants. *Group Assistants* are agents that act on behalf of a KC in the system. As such, they have a complete representation of what is available in the KC on whose behalf they act. If we use MCS as specification framework, it can be a first order theory that formalizes the taxonomy of a KC.

The main role of a GA is to interact with the other GAs in order to learn conceptual links between knowledge in its KC and knowledge in other KCs. This is to be done through a process of meaning negotiation. The aim of the negotiation process is to agree on the classification of knowledge about the same object (or topic)

in different KCs. Suppose, for example, that a member of KC1 seeks information about Linux. In KC1, this information is classified as "Operating Systems". Suppose that KC2 has information about Linux, but this information is classified as "Free Software". The goal of negotiation is to learn that, whenever KC1 needs information about Linux, a request should be sent to KC2 for data about "Free Software". This relation between the way KC1 and KC2 classify knowledge about the same topic is recorded by GA1, and used in the future as a defeasible rule of query translation with GA2.

An obvious way of doing what we are proposing from the perspective of the GEP is to define a common ontology agents can refer to. This has been proposed by many authors (see e.g. (Huhns, Singh, & Ksiezyk)). However, as we discussed in section , we don't believe that this top-down approach is scalable to really complex domains. Therefore we propose a bottom-up approach, where bridge rules (in this case, partial mapping between taxonomies of different KCs) are progressively learned by agents (though it is always possible for designers to pre-compile some rules in some GAs). See (Ghidini & Serafini 1998b) for a formalization of federated databases which uses bridge rules.

Finally, each GA observes his KC, in order to detect significant changes in knowledge organization. The idea is that a GA may decide to inform other GAs if a change has happened, so that they may update their mappings.

Knowledge Integrators. *Knowledge Integrators* (KI) form the highest level in our architecture. We can imagine that, for very complex organizations, there are several KIs, hierarchically organized (see figure), each of

which integrates information from the lower level. However, for the sake of simplicity, we will ignore this aspect.

KIs have the role of synthesizing complex mappings between KCs. We said that GAs learn one-to-one mappings with other GAs. When one of such mappings proves to be useful, GAs communicate such a mapping to their KI. Each KI receives information about mappings from many GAs. Its task is to infer new mappings starting from one-to-one mappings which have been negotiated between GAs. A trivial case is a transitive mapping: if GA1 has a mapping between P in its language and Q in GA2's language, and GA2 has a mapping between its Q and GA3's concept R, then a KI should be able to infer the mapping between GA1's P and GA3's R, even if GA1 and GA3 have never directly interacted.

When a GA is seeking information about something, the first thing to do will be to ask its KI whether it knows about mappings others that those the GA already knows. This not only helps saving time and communication exchanges between GAs, but also allows to establish connections that GAs could never find by themselves. This is particularly true in more structured organizations, where we assume that not all GAs know each other, or are allowed to communicate with each other.

Acknowledgments. This paper is part of a project on *Analysing Professional Communities as Distributed Intelligence Systems* within the *Center for Studies on Technologies in Distributed Intelligence Systems (TEDIS) – Wisdom and Intelligence Laboratory* at Venice International University (VIU). We thank *Arthur Andersen Business Consulting* and the *University of Trento* for their support to the authors.

References

- Benerecetti, M.; Bouquet, P.; and Ghidini, C. Contextual reasoning distilled. *Journal of Theoretical and Experimental Artificial Intelligence*. To appear.
- Boland, J., and R.V.Tenkasi. 1994. Designing information technology to support distributed cognition. *Organizational Science*.
- Bouquet, P.; Serafini, L.; Brezillon, P.; Benerecetti, M.; and Castellani, F., eds. 1999. *Modelling and Using Context – Proceedings of the 2nd International and Interdisciplinary Conference (9-11 September 1999, Trento, Italy)*, volume 1688 of *Lecture Notes in Artificial Intelligence*. Springer Verlag - Heidelberg.
- Brown, S., and Duguid, P. 1991. Organizational learning and communities-of-practice : Toward a unified view of working, learning and innovation. *Organization Science* 2(1).
- Dinsmore, J. 1991. *Partitioned Representations*. Kluwer Academic Publishers.
- Drucker, P. 1994. *Post Capitalist Society*. Cambridge University Press.
- Fauconnier, G. 1985. *Mental Spaces: aspects of meaning construction in natural language*. MIT Press.
- Ghidini, C., and Serafini, L. 1998b. Information Integration for Electronic Commerce. In Sierra, C., and Noriega, P., eds., *Proceedings of the Workshop on Agent Mediated Electronic Trading (AMET'98)*, volume 1571 of *LNAI*, 189–206. Minneapolis, USA: Springer Verlag.
- Giunchiglia, F., and Ghidini, C. 1998. Local Models Semantics, or Contextual Reasoning = Locality + Compatibility. In *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, 282–289. Morgan Kaufmann.
- Giunchiglia, F. 1993. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine* XVI:345–364.
- Guha, R. 1991. Contexts: a Formalization and some Applications. Technical Report ACT-CYC-423-91, MCC, Austin, Texas.
- Huhns, M., and Singh, M., eds. 1997. *Readings in agents*. S. Francisco (CA): Morgan Kaufmann Publishers, Inc.
- Huhns, M.; Singh, M.; and Ksiezyk, T. Glocal information management via local autonomous agents. In (*Huhns & Singh 1997*).
- Lenat, D., and Guha, R. 1990. *Building large knowledge-based systems*. Reading (MA): Addison-Wesley.
- McCarthy, J. 1993. Notes on Formalizing Context. In *Proc. of the 13th International Joint Conference on Artificial Intelligence*, 555–560.
- Nonaka, I., and Takeuchi, H. 1995. *The Knowledge Creating Company*. Oxford University Press.
- Orr, J. 1990. *Talking About Machines : An Ethnography of Modern Job*. Ph.D. Dissertation, Cornell University.
- Perry, J. 1998. Indexicals, Contexts and Unarticulated Constituents. In *Proceedings of the 1995 CSLI-Amsterdam Logic, Language and Computation Conference*. Stanford. CSLI Publications.
- Polanyi, M. 1966. *The Tacit Dimension*. Routledge & Keel Paul.
- Sabater, J.; Sierra, C.; Parsons, S.; and Jennings, N. 1999. Using multi-context systems to engineer executable agents. In Jennings, N., and Lespérance, Y., eds., *Pre-proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL'99)*, 131–148.
- Weick, E. 1993. Collective mind in organizations : Heedful interrelating on flight decks. *Administrative Science Quarterly* 38.
- Weiss, G., ed. 1999. *Multiagent Systems*. The MIT Press.
- Wenger, E. 1998. *Communities of Practice. Learning, Meaning, and Identity*. Cambridge University Press.