In: Proceedings of the Workshop on Agent-Mediated Knowledge Management (AMKM 2005) at the Fourth International Joint Conference on Autonomous Agents & Multi Agent Systems (AAMAS 05), pp. 14-26. Universiteit Utrecht, Utrecht, NL, 2005.

Emerging Knowledge Management in Distributed Environments

Hagen Langer, Jan D. Gehrke, Joachim Hammer,
Martin Lorenz, Ingo J. Timm, and Otthein Herzog

TZI - Center for Computing Technologies, Universität Bremen,
Am Fallturm 1, 28359 Bremen, Germany
{hlanger|jgehrke|hammer|mlo|i.timm|herzog}@tzi.de

Abstract. We present a framework for role-based knowledge management in a multiagent environment. Our approach focuses on roles which are carried out by rational agents. The use of roles for knowledge management, which is orthogonal to the organizational entity represented by an agent, reduces the computational cost of reasoning and simplifies the agent model. Our approach and illustrative examples are couched in the context of the logistics domain.

1 Introduction

Continued and strong demand for increased customization of products and their delivery has brought about a sea change in the economical landscape: from markets that were predominantly controlled by sellers to markets that are now driven by buyers and their demands. To meet the resulting requirements in the *logistics domain*, participating enterprises investigate ways to restructure their business processes to allow for more autonomy in order to provide flexibility and rapid response when reacting to customer requests. Such restructuring away from the traditional centralized way of doing business is made possible by an emergence of hardware technologies including GPS-based telematics for trucks, more reliable and longer ranging wireless communication as well as low-power sensor devices.

In addition, innovative software is being developed to support autonomous decision-making and to provide the right information to the right processes when it is needed. Software systems implementing autonomous logistic processes (e.g., agents) need to share information on a continuous basis, for example, product specifications, manufacturing capabilities, delivery schedules, etc., and are required to make decisions which are consistent with the policies and overall economical situation of the enterprise they represent. In addition, agents must be able to negotiate, form coalitions, and thrive in the presence of competition, and are also subject to unpredictable changes in their environment.

Furthermore the dynamics of logistic processes require the ability to plan (re-plan) even in light of uncertainty, incomplete, or false knowledge. Standard scenarios of logistic processes typically have been modeled on the basis of static graph-theoretic representations. The well-known *traveling salesman problem* (TSP), the *vehicle routing problem* (VRP), or the *pickup & delivery problem*

(PDP) reduce the complex task of transportation to a route optimization problem. They neglect both the important role of knowledge and communication in real-world logistic processes and the fact that relevant parameters, e.g., traffic flow, incoming orders, etc. change over time.

Thus, in order to migrate from a centralized decision-making to one that is carried out by multiple, distributed processes acting autonomously, the traditional communication and data management infrastructures must be augmented with a sophisticated yet flexible knowledge management system to support the requirements described above.

The goal of our project, which is part of the Collaborative Research Center (CRC) *Autonomous Cooperating Logistics Processes – A Paradigm Shift and Its Limitations*,[1] is threefold: (1) To investigate the effects of different degrees of autonomy on the flexibility and robustness of logistic processes. We are specifically focusing on the role of knowledge and the flow of information in such processes. (2) To develop an agent-based distributed knowledge management system for the logistics domain. (3) To conduct and analyze experiments from large-scale simulations in order to assess how the accuracy, precision, and promptness of knowledge influences the quality of decision making in complex and dynamic environments.

In this initial overview paper, we report on the framework for our proposed knowledge management system (Sec. 3 and 4). We begin by introducing a concrete scenario from the logistics domain that illustrates some of the challenges mentioned above and conclude with related work (Sec. 5) and a summary and outlook (Sec. 6).

2 A Scenario

Consider a scenario in which a shipping company manages the shipping, intermediate storage, and distribution of paper rolls. The paper is produced by paper mills in North America, Sweden, or Russia and sold to newspapers, publishers, and manufacturers of paper products in Europe. In order to help inventory management as well as to reduce the price of paper (e.g., through high volume discounts), the shipping company combines and brokers orders from the consumers to the manufacturers.

Each order typically includes the number of desired rolls, the delivery date (including possible late penalties) and for each roll, the required dimensions and quality of the paper. To fulfill an order, the shipper brokers it to one or more of the paper mills The rolls are shipped to the buyer via ship, rail, or truck or a combination thereof.

What makes this scenario interesting from a logistics point of view is the fact that despite their weight (e.g., a roll typically weighs between 1,000 and 2,000 lbs.), paper rolls are very sensitive to shock, temperature changes, and moisture, and thus require special handling and care during loading and transport. For

[1] http://www.sfb637.uni-bremen.de

example, an accidental scrape by one of the forks from a fork lift during unloading at the dock can easily tear several inches of paper on a roll. In the best case, the damage can be controlled by unrolling and discarding the torn layers. In the worst case, the contact has caused the paper to misalign on the core of the roll. This renders the roll unusable for a high-speed printing press which unrolls paper at speeds in excess of 200 km/h during the printing process. Other problems include water damage caused by rain (e.g., during loading/unloading) or excessive moisture in the storage rooms (e.g., as a result of a sudden temperature change). Considering that the cost of a roll ranges from €1300 to €2000, that delivery schedules are specified down to the desired hour of the day, and that rolls have to be handled several times on their way from the mill to the consumer, our scenario represents a very challenging transportation problem that requires careful planning and on-the-fly re-planning. For example, in case of a damaged roll, the shipper needs to decide whether to sell the roll at a reduced price to the intended recipient or to a new recipient (who has to be identified first) who uses less sophisticated printing presses.

In addition, many decisions are made based on incomplete knowledge, for example, whether or not a roll sustained damage during transport even though no visible marks exist[2]. Furthermore, many decisions require communication and possibly negotiations among customer, buyer and possibly the manufacturer, for example, to establish a new price for a damaged roll or a new delivery time for a shipment that cannot be unloaded due to bad weather.

Our proposed framework for agent-based, distributed knowledge management described in Sec. 3 and 4 below enables this type of dynamic, autonomous decision-support.

3 Framework

Our framework consists of three main components: **agents**, **knowledge**, and **roles**. Agents represent process-owners (e.g., decision-maker) or real-world entities (e.g. vehicle, package, paper roll) in the logistics domain as described in Sec. 2. In addition, each agent has specific *properties* (e.g., weight, speed, enterprise affiliation), *capabilities* (e.g., transportation capabilities, sensors for measuring humidity), *desires* (e.g., minimizing delay of a shipment), and *intentions* (i.e., tactical plans). The set of beliefs forms an agent's knowledge base and is associated with specific inferential capabilities. We envision that these agents, which must act in a rational and autonomous fashion, can be implemented as *BDI (belief, desire, intention) agents* as introduced in [1]. The BDI approach is well suited for this purpose since it provides the appropriate concepts and structures for representing our agents. For example, the strategic layer of agents may be modeled within the desires, operational aspects within beliefs, and tactical features within intentions or plans. Furthermore, the BDI approach attempts to closely mimic human decision-making [2] and represents the dominant approach

[2] There are plans to outfit rolls with sensors that can detect certain damage during transport and store the information for later usage. See also Sec. 4.2

for modeling intelligent behavior within the agent research community [3]. The applicability of BDI to represent rational, autonomous agents has been proved in [4].

The second component of our framework provides knowledge management functionalities including knowledge representation, storage, and manipulation. In our framework, the terminological domain knowledge is organized as an ontology which includes

- a representation of the transportation network as a graph, together with a two-dimensional map-like representation (similar to geographic information systems) enabling spatial reasoning (e.g. inferring properties of proper subregions using a part-of relation)
- the basic types of agents and their properties (e.g., for a vehicle, its average and maximum speed, the types of routes in the network it can use, and its load capacity)
- the properties of 'inactive' objects, such as highways, depots, etc.

The visibility of the ontology is determined by an agent's predefined capabilities. For example, in contrast to a shipment agent, an agent representing a navigation system has to have complete access to all relevant details of the transportation network part of the ontology. Dynamic aspects of the scenario (weather, traffic density, etc.) are not part of the ontology, but triggered by an event generator.

Knowledge management enables agents to request new or missing knowledge, or update existing knowledge. Intuitively, our approach to knowledge management is similar to peer-to-peer knowledge management. Agents have the abilities to communicate with available data networks or form ad-hoc networks with other agents in their *neighborhood*. In these networks agents can interact on an operational level to coordinate activities with respect to their primary (logistical) task. Challenges include the handling of physical mobility of actors, building partially redundant systems (logistic and knowledge management), and providing adequate levels of interaction (collaboration, cooperation, competition, etc.).

The third component of our framework integrates the multiagent approach with knowledge management functionalities using *roles*. Examples of these roles are knowledge acquisition, retrieval, and mediation. They are described in more detail in Sec. 4. Depending on their functionality and task in the logistics domain, agents may assume any one of the roles, which may change over time. For example, an agent representing a ship may assume the role of a knowledge provider reporting weather information to other ships. At a different point in time, the same agent may also assume the role of a knowledge consumer requesting information about its cargo and destination from a dock agent after loading is complete.

In contrast to conventional knowledge management approaches, our framework does not depend on centralized knowledge repositories. Communication among the roles is carried out over the already existing agent communication infrastructure. We are tacitly assuming the existence of standard information technologies to provide the proper support such as networking, document storage, retrieval, metadata annotation, etc. In a sense, knowledge management *emerges* from the interaction of agents by virtue of implementing specific roles autonomously and in dynamic change.

It is important to note that emerging knowledge management is restricted by various sociological and technological boundaries. For example, on a sociological level, agents may represent competing enterprises, which may lead to inconsistent or even incompatible desires. In addition, there is the important issue of trust. Low trust levels could prevent agents to assume certain roles (e.g., that of a knowledge mediator or provider). High trust levels strengthen the connections between certain agents, causing an increase in traffic over time. As far as technological boundaries are concerned, the presence of embedded computational entities, which are partially moving in the physical world leads to hard restrictions on network availability and computational ability.

4 Roles of Agents in Emerging Knowledge Management

Recall that our approach to emergent knowledge management maps the organizational entities involved in the knowledge management process onto agents. KM usually involves *three* main roles: the *provider* who offers information, the *consumer* who needs information and the *broker* who mediates between the two. In contrast, the agent-oriented approach, which advocates decomposing problems in terms of autonomous agents that can engage in flexible, high-level interactions [5], employs a multitude of agents to solve the knowledge management problem. Agentification of classical KM (cf. [6]) has focused almost exclusively on knowledge or information brokering and on modeling KM tasks by specialized agents. Taking the agent-based approach, our claim to fully automate knowledge management raises new reasoning demands especially on the brokering and maintenance of knowledge, which have not been addressed so far. For example, in classical KM approaches, knowledge brokering and maintenance are performed by human actors (cf. [7]).

4.1 Roles

Figure 1 depicts a conceptual overview of the eight different roles in our framework together with the most important possible communication acts between them. Our approach extends the ideas of classical knowledge management by further dividing the organizational entities and knowledge management roles. This helps us to reduce computational complexity while gaining flexibility. We briefly describe the resulting role set and their respective tasks in our proposed framework:

knowledge consumer: an agent acts as a knowledge consumer the moment it discovers a lack in its own local knowledge repository. It then poses a question to a broker as for who would be able to provide the lacking piece of information.

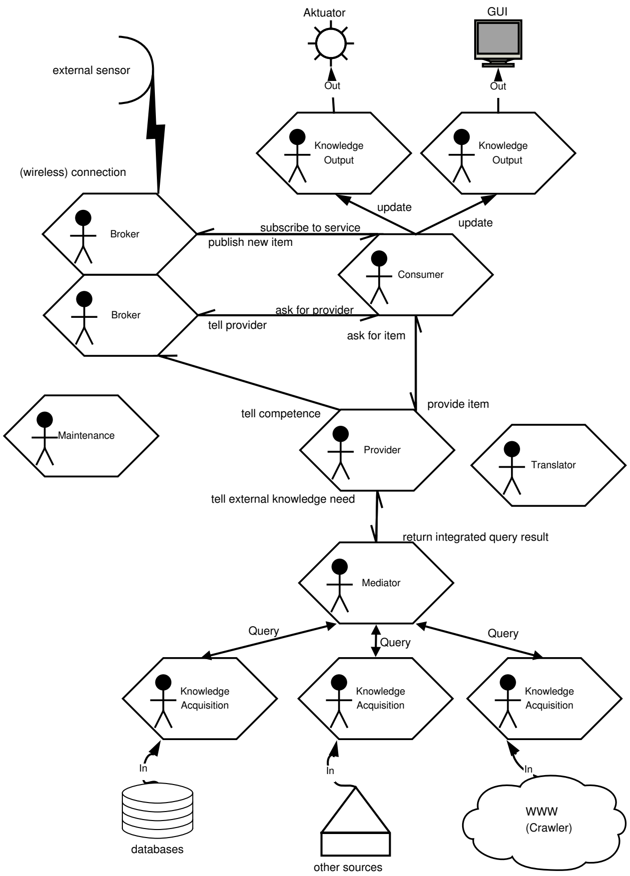

Fig. 1. Agent Roles

knowledge provider: provides items from its internal repository either on demand or as part of pro-active behavior.

knowledge acquisition: this role is intended to provide an interface to external data sources. It therefore needs the capability to query a specific source and build up an internal representation of it. Changes in the source might trigger the generation of new knowledge items.

knowledge maintenance: this role incorporates tasks which are needed to keep the distributed knowledge reasonable. Generalisation and learning capabilities keep the knowledge base (KB) as compact as possible. Surveying the KB for inconsistencies and checking these against external sources enables the deletion of obsolete or false items. It can also inform others of its findings.

This is the most complex role and may be further decomposed into subroles if it turns out necessary. It demands very sophisticated reasoning and learning capabilities.

knowledge broker: the knowledge broker acts as a yellow pages service within the system. It collects meta-knowledge from providers and points a knowledge seeking consumer to the right provider. The broker also implements a publisher/subscriber mechanism e.g. for values coming in from external sensors or the generation of specific items by an acquisition agent.

knowledge mediatior: whenever a consumer has more than one possible source for its current demand it can decide to employ a mediator to integrate and filter the responses it gets in order to reduce its own reasoning effort. The mediator acts as an intermediate between the consumer and the providers proposed by a broker and decides which items fit best to the given query or can be integrated into one, more general or more special answer. The mediator also maintains a reputation list and can therefore rule out answers from unreliable partners upon request (QoS enforcement).

knowledge output: knowledge output agents provide an interface to active elements of the external environment through actuators and user interfaces. The communication for this role is unidirectional toward the external interface. Possible responses from the environment are handled by knowledge acquisition or broker roles (through their external sensor capability) which can of course be implemented within one agent.

knowledge translation: to handle the possibility of different knowledge communication languages or dialects an optional intermediate can be engaged in every communication. The knowledge translation role acts as a wrapper agent who is able to translate e.g., between two differing domains based on their respective ontologies.

It is important to reiterate that one instantiated agent can incorporate more than one role (e.g., an agent representing a ship can first act as a knowledge provider and later as a knowledge consumer). Hence the incorporation of roles is a decomposition of the KM problem, which is in essence orthogonal to the mapping of organizational entities to agents. Furthermore, since different roles of agents need different reasoning capabilities, encapsulation of roles reduces the complexity of tasks which have to be performed by an agent at any given time. This encapsulation is realized in our framework by a *visibility* function as introduced in [8].

4.2 Roles and Knowledge Management - An Example

Let us clarify the role concept and its relevance to knowledge management using a simple example involving two roles, a *knowledge consumer role* and a *knowledge provider role*. Let us further assume that items of knowledge (k, k', etc.) are represented as definite clauses, ignoring the fact that a multi-modal logic might be more appropriate to formalize the roles. Moreover, $B(\mathcal{A}, t)$ and $I(\mathcal{A}, t)$, representing *beliefs* and *intentions* respectively, are finite sets of clauses associated with an agent $\mathcal{A}$ at time t. For example, in the context of our scenario described in Sec. 2, $\mathcal{A}$ could represent the loading/unloading foreman on the dock whose set of beliefs B during unloading of paper rolls from a ship on the 12^{th} of Dec. could include items k such as "the humidity is low and current temperature is 2 C" and one of his intentions i may be to "unload all paper rolls before 5pm." Another intention could be to "identify all rolls that do not meet the stringent quality requirements of the recipient before the rolls leave the dock." For the sake of simplicity we assume a discrete time line and ignore spatial and other parameters of the agents' environment.

The knowledge consumer role presupposes an agent $\mathcal{A}$ which intends to add k to the set of its beliefs, where k is a fully instantiated clause subsumed by another clause q which we will refer to as the *query*. We further assume that k is not already included in $B(\mathcal{A}, t)$ and that k cannot be inferred from $B(\mathcal{A}, t)$, given $\mathcal{A}'s$ current inferential abilities. $\mathrm{Rel}(k)$ is the relevance $\mathcal{A}$ assigns to k. In a knowledge transfer, k' is part of an informative communicative act directed to $\mathcal{A}$ by some other agent $\mathcal{B}$, which in this situation executes the role of *knowledge provider*.

Continuing the example, the dock agent $\mathcal{A}$ is the knowledge consumer who would like to add accurate knowledge (referred to as k above) about the status of each paper roll that is being unloaded to its set of beliefs (presumably to improve its decision making when identifying damaged rolls). In our example, this knowledge could be provided by sensors which are attached to each paper roll and monitor parameters such as temperature, humidity but also the occurrence of shocks or any kind of jarring that occurs during transport[3]. Each sensor agent $\mathcal{B}$ attached to a paper roll assumes the role of a knowledge provider which can be queried during unloading by the dock agent. Note that sensors measuring the condition of a paper roll are important knowledge providers since even small tears or scrapes that leave no visibl marks can render a roll useless to some printing companies. Thus, the dock agent will favor the sensor data coming from the roll sensor (referred to as k' above) over its own beliefs formed as a result of a visual inspection (assuming the sensors have been properly calibrated).

In the simplest case, k is identical to k', i.e., the sensor data from agent $\mathcal{B}$ provides accurate readings about the status of a paper roll. However, since k and k' can differ (e.g., k' is empty or contains inaccurate data), we need a similarity measure $\mathrm{Sim}(k, k')$ in order to be able to estimate $\mathrm{Rel}(k')$ on the basis of the initial $\mathrm{Rel}(k)$. $\mathrm{Sim}(k, k')$ returns the maximum value iff $(k' \cup B(\mathcal{A}, t))$ implies $(k \cup B(\mathcal{A}, t))$[4], and returns 0, if $B(\mathcal{A}, t)$ implies $\neg k'$ (not consistent) or $B(\mathcal{A}, t)$ implies k' (not informative). Since truth and appropriateness of k' might depend heavily on the competence and credibility of $\mathcal{B}$ (note, that we neither presuppose omniscience of knowledge providers nor benevolent behavior), we also introduce

[3] We know of at least one company who is experimenting with RFID and sensor technologies to monitor the condition of paper rolls during transport.

[4] If we presupposed that k is identical to k', we were able to derive, as a corollary, that $\mathcal{A}$ and $\mathcal{B}$ are not identical, given our assumption $k \notin B(\mathcal{A}, t)$.

a measure of confidence $\text{Conf}(\mathcal{A}, \mathcal{B}, k', t)$ ranging from -1 ($\mathcal{A}$ believes that k' is definitely false if provided by $\mathcal{B}$ at time t) to 1 (absolutely reliable)[5].

In the case of the sensor agent attached to a paper roll, mechanical failure or calibration problems are the most likely reasons for inaccurate or inconsistent knowledge k'. Regular calibration would therefore allow for higher confidence values. On the other hand, if one considers a scenario in which a truck is providing traffic updates to another truck competing for the same business, confidence of k' might be low. Additionally, the relevance weight assigned to an item of knowledge k might be influenced by the probability $\text{Avail}(k', t)$ of being available at time t. Finally, each act of knowledge consumption is associated with a cost function $\text{Cost}(\mathcal{A}, \mathcal{B}, k', t)$. Analogously to confidence the costs depend on both involved agents, the item of knowledge, and time.

The five parameters, discussed above, are closely interrelated and determine the impact of the knowledge transfer on the behavior of the consumer agent. In combination with the corresponding parameters of knowledge provider(s) and other agents, e.g., agents which participate as mediators[6], these parameters also provide the basis for the emerging structure of the KM system as a whole: successful knowledge transactions with a particular provider agent will strengthen the connection between the involved agents and increase the likelihood of future transactions.

5 Related Research

According to [9] "Knowledge Management (KM) [. . .] aims at capturing explicit and tacit knowledge [. . .] in order to facilitate its access, sharing out and reuse." This rather organization-centered view can be applied to information technology (IT) as supplementary technology (cf. e.g. [7]) as well as to KM within pure IT driven systems like the autonomous logistic scenario we propose in this paper.

Agent-based or agent-mediated knowledge management (AMKM, cf. [10], [6]) is a relatively young but currently very active field of research. [6] give a comprehensive overview of approaches, that use agent concepts for knowledge management. They hereby distinguish three areas: single agent systems, homogeneous MAS and heterogeneous or society-oriented MAS. Single agent approaches to KM usually are personal assistants like the well-known seminal works by Pattie Maes [11] and Henry Lieberman [12], the anticipatory knowledge mediator "Watson" [13], and others. [14] explore how cognitive agents can be used to design systems that implement their vision of knowledge management and that in particular support the knowledge management processes in social, organizational and individual dimension.

[5] $\text{Conf}(\mathcal{A}, \mathcal{B}, k', t)$ gives the reliability solely from the point of view of $\mathcal{A}$ and thus may differ significantly from the 'true' reliability of k' or some reliability assigned to k' by $\mathcal{B}$ as part of its communicative act.

[6] A detailed discussion of these 'corresponding parameters' goes beyond the scope of the present paper.

Applying the peer-to-peer (P2P) paradigm to personal information assistants (cf. [15]) gradually leads over to multiagent approaches, where Edutella [16], DIAMS [17], and relatives are mainly collaborative and not necessarily designed as agent platforms. The multiagent plattform FRODO aims at building a formal, logic-based organisational memory framework, implemented by an Intranet-enabled agent platform [18] and [6]. The SWF project [19] uses an ontology reasoning capable multiagent framework for semantic web retrieval and traversal.

Our proposed approach employs deliberative agents for which [4] introduces a formal model based on $\mathcal{LORA}$ and $\mathcal{VSK}$ ([20, 21]). Reasoning with multiple ontologies connected by semantic mappings is a problem present in many distributed KM approaches. [22] use a P2P architecture to define a sound and complete algorithm for global subsumptions based on local knowledge. [23] investigate the issue of integration of information from multiple sources in a cooperative information system. They propose a distributed description logic to solve a problem also seen in our presented work.

In the logistic domain several MAS approaches have been described in the literature (e.g., [24]). [25] aim at replacing conventional tracking and tracing in the logistics domain based on sending (i.e. pushing) EDIFACT messages by an agent-based pull mechanism. [26] present a prototype of a multi-agent community implementation and a constraint-based protocol designed for the agents' negotiation in a collaborative environment.

Previous research on MAS in the logistics domain has put a strong emphasis on price negotiations and auctions. In these approaches the inter-agent communication often reduces to bidding (cf., e.g., [27]), or the internal structure is defined by a set of equations (e.g., [28]). [29] applies MAS to shopfloor logistics in a dynamic production scenario. It aims at flexible and optimal scheduling of production plans in a heterogenous shopfloor environment.

6 Conclusions and Future Work

We have proposed an initial conceptual framework for knowledge management in support of autonomous processes in dynamic environments such as the presented logistic domain, for example. Our framework assumes an agent infrastructure based on the BDI approach and emphasizes the importance of knowledge and its efficient management. We propose a role concept which covers all relevant aspects of knowledge management, including knowledge acquisition, pro-active knowledge transfer, and knowledge maintenance.

The use of roles is motivated by a twofold advantage over common agent-based KM approaches. An agent incorporating different roles may partition its knowledge base accordingly and hence reduce the complexity of its reasoning task, and secondly roles enable a much clearer modeling of the agents at design time.

6.1 Contributions and Benefits

Our framework provides the following four important benefits: (1) It integrates the operational and knowledge management levels in decision-support applications through the allocation of roles. (2) It eliminates the need for a centralized knowledge repository. Existing KM infrastructure can easily co-exist with our approach. (3) The use of the different BDI tiers to organize knowledge allows for a natural mapping to the organizational, strategic, and tactical information layers that exist in the real world. (4) It supports collaboration and coalition forming among agents while also allowing for competitive behavior.

As such we expect that our project, when fully implemented, will not only contribute to a better understanding of the use of autonomous agents in the logistic domain but also provide new theories and algorithms for the efficient management of knowledge in large-scale multi-agent systems. Other important contributions include the development of a formal representation that is powerful enough to represent agents, their roles, and the underlying knowledge, as well as an efficient implementation of agents to allow experimental validation of the accuracy, precision, and promptness of autonomous decision making in complex and dynamic environments.

6.2 Plans for Prototype Development

We have conducted an initial feasibility study of the concepts proposed here using a simplified model of our logistic scenario. We are currently developing a proof-of-concept prototype system to help validate our approach. Specifically, we are developing a distributed multiagent system based on the FIPA compliant agent platform JADE [30]. This platform is aimed to be a testbed for various applications of autonomous agents in logistic scenarios. Basically those scenarios consist of a number of active objects modeled as agents and a traffic network of nodes and edges. For example, agents could model packages to be shipped as well as trucks, that want to maximize their utilization. Nodes may be logistic sources and sinks, or traffic junctions. Edges represent roads, railroads, waterways and the like that connect nodes. Manually or stochastically triggered *world events*, e.g., a traffic jam or a breakdown of a truck, force agents to reconsider their plans.

The scenario in Sec. 2 serves as our starting point for describing knowledge management needs in the logistic domain. In order to make use of such a scenario in our prototype, we defined a *logistics ontology* which forms a common ground for all KM-related tasks within the simulated world. Currently, the agents in our prototype are already able to exchange knowledge in a rudimentary way. The proposed KM approach consisting of roles and decision parameters is still under development. The parameter-based decision process will be realized as part of an agent's behavior. Agent roles follow from agent communication that will be specified as FIPA interaction protocols for each interacting role pair.

We will use the prototype as a means for validating whether (1) our set of decision parameters is complete and minimal; (2) our approach to distributed knowledge management is robust; (3) the role concept is appropriate to model communication and decision-making processes; and (4) the use of roles will reduce the computational costs of reasoning within the agents. Our long term objective is to evaluate possibilities and limitations of autonomy in logistics.

References

1. Rao, A.S., Georgeff, M.P.: Modeling rational agents within a BDI-architecture. Technical Report 14, Australian AI Institute, Carlton (1991)
2. Bratman, M.E.: Intentions, Plans, and Practical Reason. Harvard University Press, Cambridge (MA) (1987)
3. d'Inverno, M., Luck, M., Georgeff, M.P., Kinny, D., Wooldridge, M.: The dMARS architecture: A specification of the distributed multi-agent reasoning system. Autonomous Agents and Multi-Agent Systems **9** (2004) 5–53
4. Timm, I.J.: Dynamisches Konfliktmanagement als Verhaltenssteuerung Intelligenter Agenten. PhD thesis, Universität Bremen, Bremen, Germany (2004)
5. Jennings, N.R.: On agent-based software engineering. Artificial Intelligence **177** (2000) 277–296
6. van Elst, L., Dignum, V., Abecker, A.: Towards agent-mediated knowledge management. [10] 1–30
7. Maurer, H.: Important aspects of knowledge management. In Klein, R., Six, H.W., Wegner, L., eds.: Computer Science in Perspective: Essays Dedicated to Thomas Ottmann. Volume 2598 of Lecture Notes in Computer Science. Springer-Verlag, Heidelberg (2003) 245–254
8. Wooldridge, M., Lomuscio, A.: A computationally grounded logic of visibility, perception and knowledge. Logic Journal of the IGPL **9** (2001) 257–272
9. Dieng-Kuntz, R., Matta, N., eds.: Knowledge Management and Organizational Memories. Kluwer Academic Publishers, Netherlands (2002)
10. van Elst, L., Dignum, V., Abecker, A., eds.: Agent-Mediated Knowledge Management: International Symposium AMKM 2003, Stanford, CA, USA, March 24-26, Revised and Invited Papers. Volume 2926 of Lecture Notes in Computer Science., Springer-Verlag (2004)
11. Maes, P.: Agents that reduce work and information overload. Commun. ACM **37** (1994) 30–40
12. Lieberman, H.: Letizia: An agent that assists web browsing. In Mellish, C.S., ed.: Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95), Montreal, Quebec, Canada, Morgan Kaufmann publishers Inc.: San Mateo, CA, USA (1995) 924–929
13. Budzik, J., Hammond, K.: Watson: Anticipating and contextualizing information needs. In: Proceedings of the Sixty-second Annual Meeting of the American Society for Information Science. (1999)
14. Nabeth, T., Angehrn, A.A., Roda, C.: Enhancing knowledge management systems with cognitive agents. Systémes d'Information et Management **8** (2003)
15. Tsui, E.: Technologies for personal and peer-to-peer (P2P) knowledge management. Technical report, Financial Services, Computer Sciences Corporation (CSC), North Sydney, Australia and School of Business Information Technology, RMIT University, Melbourne, Australia (2002)

16. Nejdl, W., Wolf, B., Qu, C., Decker, S., Sintek, M., Naeve, A., Nilsson, M., Palmér, M., Risch, T.: Edutella: A P2P networking infrastructure based on RDF. In: WWW '02: Proceedings of the eleventh International Conference on World Wide Web, ACM Press (2002) 604–615
17. Chen, J.R., Wolfe, S.R., Wragg, S.D.: A distributed multi-agent system for collaborative information management and sharing. In: CIKM '00: Proceedings of the Ninth International Conference on Information and Knowledge Management, ACM Press (2000) 382–388
18. van Elst, L., Abecker, A.: Domain ontology agents for distributed organizational memories. [9] chapter 13 147–158
19. Stollberg, M., Keller, U., Zugmann, P., Herzog, R.: Semantic web fred - agent cooperation on the semantic web. Demontration at the 3rd International Semantic Web Conference ISWC 2004, Hiroshima, Japan, 7 - 11 November 2004 (2004)
20. Wooldridge, M., Lomuscio, A.: Multi-agent VSK logic. In Ojeda-Aciego, M., de Guzmán, I.P., Brewka, G., Pereira, L.M., eds.: Proceedings of the Seventh European Workshop on Logics in Artificial Intelligence (JELIA-2000), Heidelberg, Germany, Springer-Verlag (2000)
21. Wooldridge, M.: Reasoning about Rational Agents. Intelligent robotics and autonomous agents. The MIT Press, Cambridge, Massachussetts/London, England (2000)
22. Serafini, L., Tamilin, A.: Distributed reasoning services for multiple ontologies. Technical Report DIT-04-029, Informatica e Telecomunicazioni, University of Trento (2004)
23. Borgida, A., Serafini, L.: Distributed description logics: Directed domain correspondences in federated information sources. In Meersman, R., Tari, Z., eds.: On the Move to Meaningful Internet Systems 2002: Confederated International Conferences CoopIS, DOA, and ODBASE 2002. Number 2519 in Lecture Notes in Computer Science, Springer-Verlag (2002) 36–53
24. Timm, I.J., Schleiffer, R., Davidsson, P., Kirn, S., eds.: Agent Technology in Logistics. Number 20 in Workshop Notes, Lyon, France, 15th Biannual European Conference on Artificial Intelligence (ECAI-02) (2002)
25. Hofmann, O., Deschner, D., Reinheimer, S., Bodendorf, F.: Agent-supported information retrieval in the logistics chain. In: HICSS '99: Proceedings of the Thirty-second Annual Hawaii International Conference on System Sciences-Volume 8, IEEE Computer Society (1999) 8028
26. Smirnov, A., Pashkin, M., Chilov, N., Levashova, T.: Multi-agent knowledge logistics system "KSNet": Implementation and case study for coalition operations. In et al., V.M., ed.: CEEMAS 2003. (2003) 292–302
27. Zice, S., Zhengping, L., Runtao, Q., Mansoor, S.: Agent-based logistics coordination and collaboration. Technical Report SIMTech Technical Report (AT/01/011/LCI), Singapore Insitute of Manufacturing Technology (2001)
28. Bos, A., de Weerdt, M.M., Witteveen, C., Tonino, J., Valk, J.M.: A dynamic systems framework for multi-agent experiments. In: European Summer School on Logic, Language, and Information, Foundations and Applications of Collective Agent Based Systems workshop (1999)
29. Scholz, T., Timm, I.J., Woelk, P.O.: Emerging capabilities in intelligent agents for flexible production control. In Katalinic, B., ed.: Proceedings of the International Workshop on Emergent Synthesis (IWES 2004). (2004)
30. Bellifemine, F., Poggi, A., Rimassa, G.: JADE: A FIPA 2000 compliant agent development environment. In: Proceedings of the Fifth International Conference on Autonomous Agents (Agents01), ACM Press (2001)