IOS Press 2005 pp. 114–128

Book Title Book Editors IOS Press, 2003 1

Knowledge Management to Support Situation-aware Risk Management in Autonomous, Self-managing Agents

Martin Lorenz ^a Jan D. Gehrke ^a Joachim Hammer ^b Hagen Langer ^a Ingo J. Timm ^a

^a University of Bremen, {mlo|jgehrke|hlanger|i.timm}@tzi.de ^b University of Florida, jhammer@cise.ufl.edu

Abstract. We present a novel approach to enable decision-making in a highly distributed multiagent environment where individual agents need to act in an autonomous fashion. Our architecture framework integrates risk management, knowledge management, and agent deliberation to enable sophisticated, autonomous decision-making. Instead of a centralized knowledge repository, our approach supports a highly distributed knowledge base in which each agent manages a fraction of the knowledge needed by the entire system. Our approach also addresses the fact that the desired knowledge is often highly dynamic, context-sensitive, incomplete, or uncertain. Thus risk management becomes an integral component which enables context-based, situation-aware decision making, which in turn supports autonomous, self-managing behavior of the agents. A prototype system demonstrating the feasibility of our approach is being developed as part of an ongoing funded research project.

Keywords. Intelligent Agents, Multiagent Systems, Knowledge Management, Risk Management, Decision Support, Logistics

1. Introduction

In this paper we describe a new approach to enable robust decision-making in a highly distributed, multiagent environment where agents need to act in an autonomous fashion. Our application is the logistics domain where autonomous agents are seen as a promising and effective approach to represent the different planning, scheduling, and controlling processes in an enterprise. For example, we can envision a scenario in which agents are used to represent real-world entities such as truck and container, abstract objects such as weather or traffic service, or even human decision makers, such as the ramp agent at a loading dock. In this scenario, the agents need to make decisions about which containers to transport, what the fastest route to a specific destination is given current road or weather conditions, or what to do with goods damaged during unloading, for example. We believe this kind of autonomous, decentralized decision-making can help make the operational processes more efficient, cost-effective, and allow the participating enterprise to stay competitive. It is also a major improvement over traditional centralized approaches in which individual agents are ill-equipped to deal quickly with sudden events

since control usually resides with the entities that are removed from the scene of the event and thus have only delayed access to the relevant information.

Enabling this type of autonomous decision-making is challenging given the potentially large number of agents that could be involved, the dynamic and sometimes even competitive environment in which the agents operate. In principle, enabling a technical system, e.g., an autonomous agent, to make decisions that are designed to impact realworld entities delegates the assessment of consequences of the decisions to the agent. To the same extent as the future is perceived as decision-dependent, any decision to be made by the agent must be regarded as risky [18, p.77]. In the context of this work we define risk as uncertainty about the future state of the world which implies that any decision by the agent might turn out wrong. The goal of risk management (RM) is to attempt to optimize the agent's decisions in the presence of incomplete, imprecise, or debatable information by reducing the uncertainty about future events.

Knowledge management (KM) is an important means to achieve this. Here KM is regarded in terms of a knowledge consumer role presupposing knowledge providers but more complex interactions are possible. Our approach to KM aims at finding a rational basis to obtain needed information and to choose an appropriate provider. Furthermore, the agent is challenged by the fact that the knowledge that is needed is often highly dynamic, context-sensitive, incomplete, or uncertain. Thus the integration of risk and knowledge management enables context-based, situation-aware decision-making, which in turn supports autonomous, self-managing behavior of the agents.

Current agent architectures are not designed to model this complex decision-making process which requires agents to process knowledge about internal structures and organizations, show awareness of other agents and communicate or even cooperate with them, and perceive changes in their environment. A common approach in the agent community is to discriminate the steps percept, reason, and do as a basis for decision-making (deliberation cycle). In more sophisticated approaches, logical reasoning behavior is described. For example, in the *BDI* (belief, desire, intention) approach as introduced in [21], the strategic layer of agents may be modeled within desires, operational aspects within beliefs, and tactical features within intentions or plans. The BDI approach also attempts to closely mimic human decision-making ([3]) and represents the dominant approach for modeling intelligent behavior within the agent research community [6].

The major shortcoming of current agent deliberation cycles is the relatively simple discovery and evaluation of alternatives. The standard approach to creating consistent subsets (goals) for action selection is not sufficient for dynamic environments, as the agent must often conduct multi-criteria optimization, which may also be based on competing goals. Hence an important challenge for this project is to augment the agent's deliberation cycle with the ability to identify and assess the underlying risks that are associated with the options that determine the next course of action. If necessary, the agent must be able to augment its knowledge base with missing or updated knowledge, for example, from other agents, to be able to properly assess and evaluate the feasible options.

In the remainder of the paper we introduce our architecture framework, which integrates risk management, knowledge management, and agent deliberation to enable sophisticated, autonomous decision-making.

2. Framework

Our framework is depicted in Fig. 1. It includes explicit *risk* and *knowledge* management, termed decision-support in the figure, which may work in an inter-leaved fashion to augment the deliberation cycle of the agent. Generally speaking, we use risk management to identify and assess the risks associated with one or more options, and knowledge management to acquire missing knowledge, for example, to improve risk assessment or to generate additional options. Our decision-support system can be integrated into any intelligent agent that utilizes some form of deliberation with separate option generation and selection phases.



Figure 1. Conceptual overview of the framework depicting the interaction between agent deliberation, risk management, and knowledge management.

We realize that not all deliberation and subsequent option selection involves sophisticated risk management. In fact, many important actions are the result of a trained response (e.g., to avoid imminent danger). However, in this paper, we are focusing on agents in complex decision situations, such as the ramp agent wondering whether it is safe to start unloading paper rolls from a ship given the possible threat of a rain storm. We will elaborate on the use of risk and knowledge management in this scenario later in Sec. 5.

Starting with the deliberation cycle at the top of Fig. 1, we assume that some perceptions are leading to a situation, where the agent has to decide on its next action. Before making a decision, the agent invokes risk management to help with the assessment of the option(s) (e.g., unloading the paper rolls immediately or delaying it until the next morning). We are envisioning that all components have access to a common repository or knowledge base (not shown in the figure) containing the options currently under review. For the remainder of the paper, we will use the term "beliefs" to refer to this knowledge base. By invoking the risk management module, the agent also passes along a pointer to the option(s) currently under deliberation as input. The first step in risk management is the identification of potential risks associated with each option. For example, in our scenario, a risk of unloading immediately could be that the rolls get wet if it starts raining. Each identified risk must be evaluated to assess the magnitude of the risk and its probability of occurrence. In the ideal case, the agent has sufficient knowledge to arrive at a meaningful risk assessment. Upon completion, the result of the assessment is returned to the deliberation process which uses the information to aid in the selection of the best possible option.

Due to incomplete or uncertain knowledge (e.g., weather information has only a limited life-span and must be updated frequently), risk management may be unable to decide on risk. The exact approach for estimating both is described in more detail in Sec. 3. This triggers knowledge management to acquire the missing information or detailed information on the current situation — including alternative actions. Knowledge acquisition may retrieve knowledge from other agents (e.g., weather service) or directly from external sources/sensors (e.g., a barometer).

A central component of our approach is the representation of *decision-support parameters* which govern the RM and KM processes as well as the interactions between them. For example, when RM invokes KM to acquire missing knowledge to help assessment of risk, it communicates the *importance* of obtaining the missing knowledge to KM. This helps KM to select the proper strategy, which could be to obtain weather information from a free Internet service in case the importance is low to trading or even purchasing weather information from a reputable broker in case the perceived importance is high. Another parameter used by KM is *availability* which expresses the probability that an item of knowledge is available from any known source at this time. Availability of knowledge is based on prior experiences and used by KM, for example, in deciding which knowledge items should be acquired (in case there are choices). So far we have identified a total of eight parameters for RM and KM which will be described in more details in the following sections.

As we mentioned above, RM and KM are closely intertwined and can be invoked multiple times during a single decision-support cycle. For example, during knowledge acquisition, there could be a need to decide between different knowledge brokers both offering similar information. Based on the importance of the intended knowledge (importance parameter), the perceived trust of each knowledge broker (confidence parameter), the cost of the information (cost parameter), as well as the perceived value of the offered information compared to what is expected (similarity parameter), the RM module can be invoked by KM to help assess the "risk" of using one broker over the other.

In the following sections we describe the RM and KM modules including the shared decision-support parameters in more detail. A short scenario illustrating the use of our framework in a logistic environment is presented following the overviews in Sec. 5.

3. Risk Management

As stated in the introduction we associate uncertainty with risk. Thus the acquisition of facts that can reduce uncertainty is one strategy to handle risk. In this section we present an approach to assess the the amount of uncertainty and a strategy to reduce it by invoking knowledge management.

Risk management is a continuous process that will trigger further deliberation as soon as a fact is added to the knowledge base, which makes the situation risky. As already mentioned in the introduction, risk arises whenever a subsequent decision must be based on incomplete knowledge and thus might turn out wrong.

Our concept of risk management is heavily depending on knowledge. Therefore it can only function in close collaboration with a knowledge management infrastructure. In the following we will describe the mechanisms of this collaboration and subsequently describe the core task of knowledge-based risk assessment.

3.1. Interplay Between RM and KM

The correlation between risk and knowledge management is at least threefold. First of all knowledge of risk is one part of an agent's beliefs. Thus it can be communicated by our approach of knowledge management.

Secondly, an agent can use its knowledge of the world to identify risks. From that point of view—the *knowledge-based risk identification* view—knowledge is needed when the agent wants to reason about the possible risks it will face. Without knowledge risks degrade to a *threat*, i.e., to some incalculable future state of the world the agent will tumble into.

And thirdly, the act of communicating knowledge is in itself a risk to the agent because it can fail in various ways and of course the agent has to incorporate the possibility of false information into its reasoning about the value of intended information items.

As sketched out in Fig. 1 and mentioned earlier, the first element of the risk management process is to evaluate each incoming perception as to whether or not it adds evidence to the beliefs which would make the next decision risky.

3.2. Risk Identification

The initial task and most important prerequisite for successful risk management is its ability to identify risk and evaluate its potential consequence. Risk identification in an autonomous knowledge-based system can be achieved by matching fractions of the beliefs with situation patterns.

In the situation analysis phase of an agents deliberation cycle (see Fig. 1) incoming perceptions are integrated with the current beliefs B. Subsequently the agent generates a set of options $\Omega = \bigcup_{i \in \{1...n\}} \{O_i\} |\Omega| = n$ that are accessible given the current situation (for details and a formal specification of this process we refer to recent work by Timm [25]). Following the formalization each option O_i contains a desire and a plan to achieve it. Furthermore, the agent is able to generate an assumption on a future state of affairs based on its beliefs which is again part of its beliefs¹ (referenced as B^+). For risk identification a generalization B^{O_i} of B^+ is created, which contains only those beliefs that might be affected or are required by the plan. Risk identification will then work on B^{O_i} and the option itself to search for incidents that may impact the execution of O_i .

Following the approach presented by Lattner et al. [16] we define a risk pattern as a formal description of a situation where certain occurrences may be dangerous for the agent. A risk pattern $\mathscr{P} = \langle S, \chi \rangle$ is defined by a situation description sentence S and a gravity value χ . χ is a value for the possible outcome of the incident described by that pattern. A risk pattern is marked as identified if S subsumes B^{O_i} .

Consider the following simplified example:

¹For the sake of simplicity we omit all modal or temporal operators in the following.

Lorenz et al. / KM for RM in Agents

 $\langle (\exists x, y, z \text{ agent}(x) \land carries(x, y) \land water(z) \land damage(z, y)), 0.8 \rangle$

This pattern expresses the fact that water damages the load an agent carries. This situation is associated with a gravity of 0.8.

In the next step of risk assessment the agent evaluates all evidences $E = \{E_1, ..., E_j\}$ (i.e., beliefs) that lead to a match of the risk pattern according to the degree of uncertainty it has about each evidence E_j . The agent's attempt to acquire additional knowledge is then triggered by a threshold on the combined gravity and uncertainty.

3.3. Risk Parameters

To order the acquisition of new knowledge from KM, risk management assigns a parameter *importance* to every item k it requests.

Importance Imp(A, k, t) denotes the importance of an agent A's intention to obtain the item of knowledge k at time t. It is a measure for k to contribute to the achievement of strategic goals. The parameter value ranges from 0 (irrelevant) to 1 (maximum significance).

To determine $\text{Imp}(\mathcal{A}, k, t)$, the risk identification process interprets every perception in E as a belief k that supports a given risk hypothesis. Every belief k is associated with two probability values which denote the subjective confidence the agent has. In this we follow the basic idea of the Dempster-Shafer theory of evidence (c.f. [8] or [14]).

Support and plausibility Support Supp(k) for a hypothesis k indicates the probability mass given to sets of evidence that are enclosed by it. In other words, it gives the amount of belief that directly supports a given hypothesis. Plausibility Pl(k) is 1 minus the masses given to sets of events whose intersection with the hypothesis results in an empty set. Again, in other words, it gives an upper bound on the belief that the hypothesis could possibly happen, i.e., it "could possibly happen" up to that value, because there was no evidence that would contradict that hypothesis. Hence $Supp(k) = 0 \equiv Pl(k) = 1$ denotes total ignorance concerning k.

Based on this support values the agent can express its need for new evidences. More precisely, it calculates the difference $\psi_i = Pl(k_i) - Supp(k_i)$ for all relevant facts in its set of beliefs. Relevant are those facts k_i that are present in a risk pattern i.e., in E.

- **Ignorance factor** The *ignorance factor* $\psi = \sum \psi_i$ denotes the agents lack of crisp knowledge to be able to soundly evaluate the risk in question. Together with a *gravity value* χ these define the Imp parameter for knowledge management. A first approach to derive Imp from ψ and χ is given by Imp = $\chi * \psi$.
- **Gravity value** χ is part of an agents knowledge base. It is tied to the risk pattern and provides a value for the possible damage caused by the incidence of that risk. Motivation for this approach is taken from the classical formalization of risk as function of the occurrence probability of an incident and the severity of its consequence. As a part of the knowledge base χ is in itself subject to knowledge management and can therefore be communicated. The gravity value is considered to be an a-priori set constant here. We will consider a dynamic gravity function which includes several factors like e.g. occurrence probability, amount of loss/damage, damage classification, etc. in an ad-hoc manner.

Based on a threshold for the ignorance factor ψ , risk management decides whether the evidence that is already present in the beliefs is sufficiently crisp to assess the risk which was identified. If ψ is below the threshold, the option O_i is annotated with a risk value composed from all gravity values assigned to it in E and returned to option selection (see Fig. 1). If the ignorance is too high to give a concise risk estimation for O_i , new evidence is necessary to support or contradict the risk hypothesis. Therefore a request consisting of the belief that needs to be updated and the importance derived from ignorance and gravity is sent to KM.

4. Knowledge Management

If the risk management component identifies the need for additional information, knowledge management is invoked as depicted in Fig. 1.

Our approach to knowledge management consists of three main components: **conceptual knowledge**, **roles**, and **parameters**. The conceptual knowledge is represented as an OWL ontology. For the purpose of our logistic application domain this ontology includes a representation of the transportation or production network, 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), and the properties of 'inactive' objects, such as highways, depots, etc.

4.1. Roles

In contrast to previous approaches to agent-based knowledge management [28] we do not presuppose a one-to-one correspondence between agents and knowledge management functions, such as *providing* knowledge or *brokering* knowledge. In our approach these functions are implemented as *roles*. A knowledge management role includes certain reasoning capabilities, a visibility function on an agent's beliefs, a deliberation pattern (i.e., a plan how to accomplish the KM task), and a communication behavior with interacting roles. The aim of KM roles is to provide a formal description of KM tasks that eases the development of agents and reduces computational complexity by means of a minimum set of processed knowledge and applied reasoning capabilities.

One agent can assume different roles and may change them over time. The minimum role model includes the roles of a *provider* offering information and a *consumer* being in need of information. The next extension would be a *broker* mediating between the two [29]. In [15] we introduced an extended role model of eight roles. It incorporates for instance a *translator* between different knowledge representation formalisms.

4.2. Parameters in Knowledge Acquisition

The deliberation pattern of roles is primarily determined by decision parameters as a rational basis to allow for a design of logical foundations, and for analysis and verification of decision processes. Each role has a set of parameters characterizing aspects of decisions to make. In case of the knowledge consumer role, i.e., the rational process of selecting an item of knowledge from a set of possible items and their providers, we introduce a set of four decision parameters. Together with *importance* (see Sec. 3.3) these are correlated and balanced by an agent during a knowledge transfer process in order to reach a rational decision. The parameter set is considered necessary but not definitely complete. Except for *similarity*, all parameters are agent-specific, i.e., they need not to be objective.

- Availability Avail (\mathcal{A}, k, t) is the probability supposed by agent \mathcal{A} of an item of knowledge k to be available from any source at time t in principle. Due to its independence of specific sources the parameter does not help chose an appropriate source. Availability is consulted to select the most intended item of knowledge.²
- **Cost** $Cost(\mathcal{A}, \mathcal{B}, k, t)$ determines the costs resulting from the knowledge transfer of k between consumer \mathcal{A} and provider \mathcal{B} at time t. This includes costs arising in the communication process and possible costs to obtain k as payment to the knowledge provider.
- **Confidence** $\operatorname{Conf}(\mathcal{A}, \mathcal{B}, k, t)$ describes the confidence of the knowledge-consuming agent \mathcal{A} at time t that knowledge-providing agent \mathcal{B} will answer the request for the intended item of knowledge k correctly. The parameter value ranges from -1 to 1. -1 means \mathcal{A} feels certain that \mathcal{B} is lying or just has incorrect beliefs, whereas a confidence of 1 corresponds to absolute confidence in \mathcal{B} 's answer. 0 stands for neutral confidence, i.e., agent \mathcal{A} has no clue whether \mathcal{B} 's answer will be rather right or wrong.
- Similarity Sim(k, k') denotes the similarity of two items of knowledge k and k' to compare the intended answer with the one actually obtained. The value ranges from 0 (no similarity) to 1 (exact match). The obtained item k' may differ in terms of integrity and accuracy. Integrity concerns missing knowledge, whereas accuracy concerns deviations, e.g., spatial, temporal, or precision of measurement. The consumer needs to evaluate kind and scale of a potential deviation in order to plan and execute appropriate actions (intentions) to finally get the knowledge needed. Similarity should be calculated based on information provided by the ontology, e.g., information on deviations and partonomies of spatial concepts.

The parameters discussed above are closely interrelated and determine the impact of the knowledge transfer on the behavior of the consuming agent. The first decision is what knowledge item should be obtained at first. Here the agent tests each item's *importance* (specified by RM) and *availability*. The availability estimation may be based on background knowledge and experiences made. Next the provider of the intended knowledge has to be selected. Influencing decision parameters are *cost* of transfer and *confidence* in the provider. When the transfer is finished the providers answer to the send query is evaluated using the *similarity* function. If similarity is low *confidence* in the provider will decrease or increase otherwise. Thus, successful knowledge transactions with a particular provider agent will strengthen the connection between the involved agents and increase the likelihood of future transactions.

If *similarity* and *confidence* are too low in relation to the intended item's *importance* the agent may try to get an answer from another agent provided that there is enough time and money left. The obtained knowledge item is added to the agent's beliefs if it is considered credible and valuable. After this KM informs RM of the belief update (cf. Fig. 1) enabling a more precise risk assessment.

²Figure a case of an information item being extremely useful (in terms of importance) but with little or no chance of obtaining it, e.g., the winning numbers of tomorrow's lottery drawing.

5. Case Study

We now illustrate the RM and KM concepts introduced in the previous two sections using a simple case study which expands on the logistics examples used so far. Let us assume the existence of a ramp agent at a dock who is responsible for unloading a shipment of paper rolls which must be brought to a warehouse on the dock before being transported to their customers by truck.

What makes this simple scenario interesting for our decision-support framework 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 example, water damage caused by rain during the unloading process or by excessive moisture in the storage rooms (e.g., as a result of a sudden temperature change) can render all affected rolls unusable. Considering that the cost of a roll ranges from \in 1300 to \in 2000, that delivery schedules are specified down to the desired hour of the day, and that rolls have to be subjected to weather several times during the unloading process our scenario represents a very challenging decision problem that requires careful planning and on-the-fly re-planning by the ramp agent. For example, in case of a severe weather threat, the ramp agent must decide quickly whether unloading should be interrupted risking delays in the delivery of rolls to customers as well as additional docking fees for the ship, or if it should continue, risking damage to some of the rolls.

For the rest of this case study, we will focus on the ramp agent and its decision to continue the unloading process. Let us assume the ramp agent is represented as an intelligent agent \mathcal{A}^3 . Let us further assume that $B(\mathcal{A}, t)$ and $I(\mathcal{A}, t)$ represent *beliefs* and *intentions* of agent \mathcal{A} at time t. For simplicity, we assume that B and I as well as items of knowledge (k, k', etc.) are represented as (sets of) definite clauses. For example, on the day of this fictitious example, \mathcal{A} has intention I_1 , "take paper rolls to the warehouse for quality control and redistribution to the delivery trucks." Another intention I_2 could be to "identify all rolls that do not meet the quality requirements of the recipient before the rolls leave the warehouse." These intentions could have been formed by the agent's situation analysis given one or more desires. In addition, at the same time t, \mathcal{A} 's beliefs contain items k_1 , "each ship carries approximately 5,000 paper rolls" and k_2 , "our fork lift crew can unload a new roll every 60 sec.", as well as weather-related items such as k_3 , "it is late spring and weather conditions change quickly".

The agent beliefs k_4 = "Heavy rain within 20 minutes" with $Supp(k_4) = 0.2$ whereas $Pl(k_4) = 0.95$ such that $\psi = 0.75$ which is far to high given that contact with water will cause a total loss of approx. $\in 1000 \ (\chi \text{ might be } 0.9 \text{ in that case})$. So the agent requests a current short term weather forecast with $Imp(\approx 0.83)$. Recall that Imp(k) is the importance RM assigns to a knowledge item k. RM invokes KM with the request to obtain knowledge item k containing temperature readings, air pressure, wind speed, etc. for the next 12 hours for the specific geographic location of the port.

Obtaining the needed weather information k is the task of knowledge management. It involves the ramp agent A in the role of a *knowledge consumer* and two weather service agents, A_{w1} and A_{w2} , in the role of *knowledge providers*. We assume for simplicity that A has prior experience about knowledge providers for weather information. Oth-

³In more complex scenarios requiring interactions between the shipment and the personnel on the dock, individual paper rolls may also be represented as agents.

erwise, \mathcal{A} would have to consult a *knowledge broker* (a different type of role) first. \mathcal{A} intends to add k to its beliefs $B(\mathcal{A})$, where k is a fully instantiated clause subsumed by another clause q, the query. We further assume that k is not already included in $B(\mathcal{A})$ and that k cannot be inferred from $B(\mathcal{A})$, given \mathcal{A} 's current inferential abilities. In a knowledge transfer, k' is part of an informative communicative act directed to \mathcal{A} by some other agent, for example, \mathcal{A}_{w1} .

In communicating with A_{w1} and A_{w2} , A learns that there are two weather information packages k_{w1} and k_{w2} available with different level of detail and at different cost to the agent. In order to judge the usefulness of each package (e.g., a package may include incomplete or inaccurate data), in addition, to the importance value obtained from RM, A uses the availability, confidence, and cost measures described in Sec. 4 to decide which of the two weather providers should be used. In our scenario, A_{w1} may be a free Internet weather service with unknown reputation and whose information contains relatively little detail for small regions such as the port where the unloading is going on. On the other hand A_{w2} may be the weather service at the nearby airfield providing highly detailed weather information for the desired area. A_{w2} , which is known for its reliable data and high availability, requires a fee.

Despite the fact that weather agent A_{w2} requires a fee, A decides to obtain k_{w2} given the high value for importance (provided by RM), as well as confidence (based on prior experience with the two weather services). Availability of information does not enter into the decision-making process at this time since the information was readily available from both services.

When the requested information from A_{w2} arrives, agent A computes the similarity measure between the intended information k and the actually obtained information k_{w2} . Sim (k, k_{w2}) is then used to update the confidence measure for A_{w2} using the following rule of thumb: the greater the similarity, the more confidence one has in the service of the agent and vice versa. In addition to k_{w2} , Sim (k, k_{w2}) will be returned to RM, which uses it to evaluate the quality of the acquired information. Presumably, if the similarity is small, RM may decide to continue to request additional information.

Risk management uses the new information k_{w2} to compute the risk values as follows. k_{w2} adds new evidence to the beliefs thus $Supp(k_4)$ and $Pl(k_4)$ can be recalculated. The new evidence reduces A's ignorance concerning k_4 such that now a risk value can be assigned to the option for unloading paper rolls and it is forwarded to option selection.

Finally the agent decides to postpone the unloading, because the new evidence k_{w2} has increased the support of k_4 "heavy rain" to 0.7. The option to delay and avoid rain damage to the rolls has been favored over the option to proceed and avoid any late fees.

In this case study we showed an example for a straightforward invocation of KM by RM. However, other scenarios are possible. RM could be called by KM as well to determine the risk of acquiring new knowledge, e.g., lying or ill-informed knowledge providers.

6. Related Work

In the multiagent literature a variety of decision-making strategies has been described. Most multiagent systems however employ rather simple decision strategies and concentrate on phenomena emerging from the collaboration and communication mechanisms. An interesting theoretical approach to decision mechanisms in collectives and complex systems is presented by Tumer and Wolpert[27].

One of the most important concept for deliberative decision-making in autonomous agents has been developed by Rao and Georgeff[21] based on the belief-desire-intention theory of human rational action by Michael Bratman [3]. Kakas and Moraitis consider argumentation depending on the particular context that the agent finds himself[11].

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., [31]), or the internal structure is defined by a set of equations (e.g., [2]). Scholz et al. [22] apply MAS to shop floor logistics in a dynamic production scenario. It aims at flexible and optimal scheduling of production plans in a heterogeneous shop floor environment. Hofmann et al. [9] 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. Smirnov et al. [24] present a prototype of a multiagent community implementation and a constraint-based protocol designed for the agents' negotiation in a collaborative environment.

According to Dieng-Kuntz and Matta [5] "Knowledge Management [...] 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. [19]) as well as to KM within pure IT driven systems like the autonomous logistics scenario we proposed in this paper. Agentbased or agent-mediated knowledge management (cf. [28], [29]) is a relatively young but currently very active field of research. Van Elst et al. [29] 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 Maes and Henry Lieberman [17], the anticipatory knowledge mediator "Watson" [4], and others. [20] 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. Our proposed approach employs deliberative agents for which Timm [25] introduces a formal model.

Serafini and Tamilin [23] use a P2P architecture to define a sound and complete algorithm for global subsumptions based on local knowledge. Borgida and Serafini [1] investigate the issue of integration of information from multiple sources in a cooperative information system.

Risk assessment as scientific topic is basically known from management science, finance, environmentalism and health care. Therefore a number of methodologies for organizational risk identification and management can be found in the literature [7,26]. Risk identification is thereby described as the ongoing risk management task of identifying the significant risks to the success of an endeavor. All proposed techniques are of organizational nature, i.e. checklists of risks and their factors, brainstorming of risks and their factors, cross functional teams, interviews with stakeholders and domain experts, etc. In the later literature much attention is paid to software engineering risk management (cf., e.g., [13]) which tends to adapt existing methodologies to the special needs of software development projects.

Lorenz et al. / KM for RM in Agents

An upcoming field is the development of computer-based tools to assist in the risk management process. Zoysa and Russell [32] give a thorough overview on "computerized knowledge-based methodologies [...] to capture and reuse risk-related knowledge". An additional interesting approach which fits in this category is proposed by Kim [12].

Explicit, knowledge-based risk identification based solely on autonomously acquired data (in contrast to specific software-guided user input), i.e., a fully automated knowledge-based risk management system, has not yet been proposed in the considered literature.

7. Status and Conclusion

We have described our conceptual framework for enabling autonomous decision-making in agents. Our approach, which integrates risk and knowledge management, allows an agent to evaluate decisions/options based on the likelihood of certain beliefs that the agent uses as supporting evidence. In case, the supporting evidence is weak, i.e., below a certain threshold, or even missing, knowledge management attempts to provide the missing information. Instead of a centralized knowledge repository, our role-based approach supports a highly distributed knowledge base in which each agent manages a fraction of the knowledge needed by the entire system. Our approach also addresses the fact that the desired knowledge is often highly dynamic, context-sensitive, incomplete, or uncertain. Thus the integration of risk and knowledge management enable context-based, situationaware decision-making, which in turn supports autonomous, self-managing behavior of the agents.

7.1. Benefits and Contributions

The approach described in this paper has the following three important benefits: (1) Our approach augments agent deliberation with sophisticated decision-making capabilities not found in current architectures. (2) By using risk management to also support the acquisition of knowledge, our approach is better equipped to manage the highly dynamic, context-sensitive, and uncertain information needed to make autonomous decisions in realistic environments. This is of particular importance, since we do not presuppose benevolent behaviour. (3) Our role-based knowledge management enables the distribution of knowledge and knowledge management functionality which eliminates the need for a centralized knowledge repository. On the other hand, it provides the necessary flexibility to allow existing KM infrastructure to co-exist with our approach.

As such we expect that our project 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 risk and knowledge in large-scale multiagent systems. Other important contributions include the development of a formal representation that is powerful enough to represent agents, their roles, and the underlying decisionsupport mechanism, 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.

7.2. 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.⁴ 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 case study in Sec. 5 serves as our starting point for describing knowledge management needs in the logistics 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 decision-support system will be realized as part of an agent's behavior. Knowledge management 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) the assessment of risk can enhance the deliberation cycle; (3) our approach to distributed knowledge management is robust; 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.

8. Acknowledgments

This research is funded by the German Research Foundation (DFG) as the Collaborative Research Centre 637 "Autonomous Cooperating Logistic Processes: A Paradigm Shift and its Limitations" (SFB 637). Additional Information can be found at http://www.sfb637.uni-bremen.de/.

We would like to thank the reviewers for valuable hints and suggestions which helped to improve this paper.

References

 A. Borgida and L. Serafini. Distributed description logics: Directed domain correspondences in federated information sources. In R. Meersman and Z. Tari, editors, On the Move to Meaningful Internet Systems 2002: Confederated International Conferences CoopIS, DOA, and ODBASE 2002, volume 2519 of Lecture Notes in Computer Science, pages 36–53. Springer-Verlag, 2002.

⁴http://jade.tilab.com/

Lorenz et al. / KM for RM in Agents

- [2] A. Bos, M. M. de Weerdt, C. Witteveen, J. Tonino, and J. M. Valk. 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.
- [3] M. E. Bratman. Intentions, Plans, and Practical Reason. Harvard University Press, Cambridge, MA, USA, 1987.
- [4] J. Budzik and K. Hammond. Watson: Anticipating and contextualizing information needs. In Proceedings of the Sixty-second Annual Meeting of the American Society for Information Science, 1999.
- [5] R. Dieng-Kuntz and N. Matta, editors. *Knowledge Management and Organizational Memo*ries. Kluwer Academic Publishers, July 2002.
- [6] M. d'Inverno, M. Luck, M. P. Georgeff, D. Kinny, and M. Wooldridge. The dMARS architecture: A specification of the distributed multi-agent reasoning system. *Autonomous Agents* and Multi-Agent Systems, 9(1–2):5–53, 2004.
- [7] S. Figlewski and R. M. Levich, editors. *Risk management: the state of the art*. Kluwer Academic Press, 2002.
- [8] J. Y. Halpern. Reasoning about uncertainty. The MIT Press, 2003.
- [9] O. Hofmann, D. Deschner, S. Reinheimer, and F. Bodendorf. 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, page 8028. IEEE Computer Society, 1999.
- [10] M. C. Horsch and D. Poole. An anytime algorithm for decision making under uncertainty. In G. F. Cooper and S. Moral, editors, UAI '98: Proceedings of the Fourteenth Conference on Uncertainty in Artificial Intelligence, July 24-26, 1998, University of Wisconsin Business School, Madison, Wisconsin, USA, pages 246–255. Morgan Kaufmann, 1998.
- [11] A. Kakas and P. Moraitis. Argumentation based decision making for autonomous agents. In AAMAS '03: Proceedings of the second international joint conference on Autonomous agents and multiagent systems, pages 883–890, New York, NY, USA, 2003. ACM Press.
- [12] D. Kim. Structural risk minimization on decision trees using an evolutionary multiobjective optimization. In M. Keijzer, U.-M. O'Reilly, S. M. Lucas, E. Costa, and T. Soule, editors, *Genetic Programming. 7th European Conference, EuroGP 2004, Proceedings*, volume 3003 of *Lecture Notes in Computer Science*, pages 338–348. EvoNet, Springer-Verlag, 5-7 Apr. 2004.
- [13] J. Kontio. Software Engineering Risk Management: A Method, Improvement Framework, and Empirical Evaluation. PhD thesis, Helsinki University of Technology, Sept. 2001.
- [14] R. Kruse, E. Schwecke, and J. Heinsohn. Uncertainty and Vagueness in Knowledge Based Systems. Springer-Verlag, 1991.
- [15] H. Langer, J. D. Gehrke, J. Hammer, M. Lorenz, I. J. Timm, and O. Herzog. Emerging knowledge management in distributed environments. In Workshop on Agent-Mediated Knowledge Management (AMKM 2005) at the Fourth International Joint Conference on Autonomous Agents & MultiAgent Systems (AAMAS 2005), 2005. To appear.
- [16] A. D. Lattner, I. J. Timm, M. Lorenz, and O. Herzog. Knowledge-based risk assessment for intelligent vehicles. In *International Conference on Integration of Knowledge Intensive Multi-Agent Systems (KIMAS '05): Modeling, Evolution and Engineering*, Waltham, MA, USA, Apr. 2005.
- [17] H. Lieberman. Letizia: An agent that assists web browsing. In C. S. Mellish, editor, Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI-95), pages 924–929, Montreal, Quebec, Canada, 1995. Morgan Kaufmann Publishers Inc.
- [18] N. Luhmann. Soziologie des Risikos. Walter de Gruyter, 2003.
- [19] H. Maurer. Important aspects of knowledge management. In R. Klein, H.-W. Six, and L. Wegner, editors, *Computer Science in Perspective: Essays Dedicated to Thomas Ottmann*, volume 2598 of *Lecture Notes in Computer Science*, pages 245–254. Springer-Verlag, July 2003.

14

- [20] T. Nabeth, A. A. Angehrn, and C. Roda. Enhancing knowledge management systems with cognitive agents. Systémes d'Information et Management, 8(2), 2003.
- [21] A. S. Rao and M. P. Georgeff. Decision procedures for BDI logics. *Journal of Logic and Computation*, 8(3):293–342, 1998.
- [22] T. Scholz, I. J. Timm, and P.-O. Woelk. Emerging capabilities in intelligent agents for flexible production control. In B. Katalinic, editor, *Proceedings of the International Workshop on Emergent Synthesis (IWES 2004)*, 2004.
- [23] L. Serafini and A. Tamilin. Distributed reasoning services for multiple ontologies. Technical Report DIT-04-029, Informatica e Telecomunicazioni, University of Trento, 2004.
- [24] A. Smirnov, M. Pashkin, N. Chilov, and T. Levashova. Multi-agent knowledge logistics system "KSNet": Implementation and case study for coalition operations. In V. M. et al., editor, *CEEMAS 2003*, pages 292–302, 2003.
- [25] I. J. Timm. Dynamisches Konfliktmanagement als Verhaltenssteuerung Intelligenter Agenten. PhD thesis, Universität Bremen, Bremen, Germany, Apr. 2004.
- [26] M. J. S. Timothy McDaniels, editor. Risk analysis and society : an interdisciplinary characterization of the field. Cambridge University Press, 2004.
- [27] K. Tumer and D. Wolpert, editors. *Collectives and the design of complex systems*. Springer, 2004.
- [28] L. van Elst, V. Dignum, and A. Abecker, editors. 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.
- [29] L. van Elst, V. Dignum, and A. Abecker. Towards agent-mediated knowledge management. In Agent-Mediated Knowledge Management: International Symposium AMKM 2003, Stanford, CA, USA, March 24-26, Revised and Invited Papers [28], pages 1–30.
- [30] J. Vassileva, S. Breban, and M. C. Horsch. Agent reasoning mechanism for long-term coalitions based on decision making and trust. *Computational Intelligence*, 18(4):583–595, 2002.
- [31] S. Zice, L. Zhengping, Q. Runtao, and S. Mansoor. Agent-based logistics coordination and collaboration. Technical Report SIMTech (AT/01/011/LCI), Singapore Insitute of Manufacturing Technology, 2001.
- [32] S. D. Zoysa and A. D. Russell. Knowledge-based risk identification in infrastructure projects. *Canadian Journal of Civil Engineering*, 30(3):511–522, 2003.