

Role of acquaintance models in agent's private and semi-private knowledge disclosure

Michal Pěchouček *, Vladimír Mařík, Jaroslav Bárta

Gerstner Laboratory, Czech Technical University in Prague, Technická 2, 166 27, Prague 6, Czech Republic

Received 9 May 2003; accepted 4 October 2005

Available online 13 February 2006

Abstract

The organizational architecture of the multi-agent systems and the structure of social knowledge that the members of the community administer are critical factors for assuring such patterns of information exchange that keep agents private knowledge confidential. In this paper we will introduce the concept of agents' private and semi-private knowledge and we will explain the difference between the *alliance* – a semi-permanent organizational structure and a *coalition* – a goal-oriented, non-permanent organizational structure. We will provide the reader with an analysis on how does the agents' social knowledge, stored in the *tri-base acquaintance model*, contributes to permanent confidentiality of agents' private knowledge, preferences, decision making models, resources, etc. The study has been experimentally verified in the domain of planning for humanitarian relief operations within a high number of hardly collaborating and vaguely linked non-governmental organizations is a challenging problem.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Agents; Coalition formation; Private/semi-private knowledge; Alliances

1. Domain settings

The application domain of this research belongs to the area of *war avoidance operations* such as peace-keeping, peace-enforcing, non-combat evacuation or disaster relief operations. Unlike in classical war operations, where the technology of decision making is strictly hierarchical, *operations other than war* (OOTW) are very likely to be based on cooperation of a number of different, quasi-volunteered, vaguely organized groups of people, non-governmental organizations (NGO's), institutions providing humanitarian aid, but also army troops and official governmental initiatives.

Collaborative, unlike hierarchical, approach to operation planning allows greater deal of flexibility and dynamics in grouping optimal parties playing an active role in the operation. New entities shall be free to join autonomously

and involve themselves in planning with respect to their capabilities. Therefore any organization framework must be essentially “open”. OOTW have, according to [20], multiple perspective on plan evaluation as there does not need to be one shared goal or a single metrics of the operation (such as political, economical, and humanitarian). From the same reason, the goals of entities involved in a possible coalition may be in conflict. Even if the community members share the same goal, it can be easily misunderstood due to different cultural backgrounds.

The main reason why we can hardly plan operations involving different NGO's by a central authority results from their *reluctance to provide information* about their intentions, goals and resources. Consequently, besides difficulties related to planning and negotiation we have to face the problems how to assure sharing the detailed information. Many institutions will be ready to share resources and information within some well-specified community, whereas they will refuse to register their full capabilities and plans with a central planning system and to follow centralized commands. They may agree to participate in exe-

* Corresponding author.

E-mail address: pechouc@labe.felk.cvut.cz (M. Pěchouček).

cuting a plan, in forming of which they played an active role. In our interpretation, an agent is a complex, organized entity (representing a NGO, humanitarian organization, army troop, etc.) playing an active role in the OOTW planning. A multi-agent system consists of a number of agents that group themselves in various, temporary coalitions (each solving a specific mission/part of the mission).

The main ambition of our research has been to analyze the problem of OOTW coalition formation and to propose a novel approach that would (i) make the coalition formation process simpler in comparison to the “classical” methods, and thus more efficient and (ii) at the same time maintain confidentiality of the private information. In our case, we decided to sacrifice the total optimality of the formed coalitions as we found this is not the most important aspect in the OOTW planning. We have suggested a concept of *alliances* – a set of agents that agreed to share some of their private information and to cooperate eventually. The coalition formation complexity is reduced by splitting the whole community of agents into disjunctive alliances and by the attempts to create a coalition preferably within the single alliance. Agents within a single alliance share *social knowledge*, that is stored in the *acquaintance models* of the individual agents. Knowledge sharing and structuring the OOTW community in the alliances allows implementation of a coalition formation system with the following key functionality aspects:

- minimizes required *communication traffic* which influences the problem solving efficiency,
- keeps the *quality of the coalition* ‘reasonably good’ (e.g., the delivery time per task),
- minimizes the loss of *agents’ private information* (e.g., about services resources, intention), when negotiating the mission
- *minimize the amount of shared information* – information that the agents decided to make available to others in order to plan an optimal mission.

The developed approach has been tested on the CPlanT multi-agent system implementation (see Section 6) [11].

1.1. Structure of the paper

In the following, the theoretical concept of agents’ neighborhood (Section 2.1) and knowledge sharing (2.2) will be presented in order to provide definition of an alliance, a coalition and a team action plan (2.5) and set the problem of knowledge disclosure 2.4. Section 3 explains the concept of an acquaintance model and social knowledge and its application for inter-agent communication. Section 5 describes computational processes in the CPlanT multi-agent system. Section 7 provides a brief review of coalition formation approaches and their comparison with the knowledge-based approach studied in this paper.

2. Preliminaries

2.1. Agent’s neighborhood

Each agent may participate in one alliance of ‘friendly’ agents and at the same time it may be actively involved in several coalitions of agents cooperating in fulfilling specific shared tasks. Computational and communication complexity of forming such a coalition depends on the amount of pre-prepared information the agents administer one about the other and on sophistication of the agents’ capability to reason about the other agents’ resources, plans and intentions. The agents can allow others to reason about them and at the same time they can reason differently about the agents that belong to their different scopes of reasoning – neighborhood. Therefore, we distinguish among several types of agents’ neighborhoods:

- $\alpha(A)$ – agent’s *total neighborhood*, a set of all agents that the agent A is aware of (e.g., knows about their existence and is able to communicate with them)
 - $\mu(A)$ – agent’s *social (monitoring) neighborhood* that is a set of agents, which the agent A keeps specific information about (e.g., services they provide, status, load, etc.). This neighborhood consists of the set of the agents that the agent A reasons about $-\mu^+(A)$ and the set the agents that reason about the agent $A - \mu^-(A)$. Therefore
- $$\forall B \in \mu^-(A) : A \in \mu^+(B). \quad (1)$$
- $\varepsilon(A)$ – agent’s *cooperation neighborhood* that is a set of agents jointly collaborating (or committed to collaboration) in achieving one or more shared goals.

We assume that each agent belongs to its neighborhoods: $\forall A \in \Theta : A \in \alpha(A) \wedge A \in \mu(A) \wedge A \in \varepsilon(A)$.

2.2. Knowledge sharing

In order to reason one about the other, the agents must share some of their knowledge. Let us introduce the operator $(\text{Bel } A\varphi)$ that expresses the agent’s A awareness of the formula φ being true [21]. We say that the agent A_0 intentionally shares its knowledge $K_\delta(A_0)$ within a set of agents $\delta \subseteq \Theta$ ($A_0 \in \delta$) provided that:

$$K_\delta(A_0) = \{\varphi | (\text{Bel } A_0\varphi) \wedge \forall A_i \in \delta : (\text{Bel } A_0(\text{Bel } A_i\varphi)) \wedge \forall B_i \notin \delta : (\text{Bel } A_0 \neg (\text{Bel } B_i\varphi))\}. \quad (2)$$

From the previous follows, that if an agent B knows some of the shared information without the agent A_0 being aware of this fact, the agent B is not regarded as a member of the $\delta(A_0)$ set of agents, representing A_0 ’s knowledge sharing neighborhood. According to this classification, we suggest three levels of the agent’s knowledge sharing:

- *Public knowledge* is shared within the entire multi-agent community. If it is assumed that all the agents know one about the other (i.e., $\forall A, A \in \Theta : \alpha(A) = \Theta$), public knowledge $K_p(A_0)$ of an agent A_0 is defined as

$$K_p(A_0) = K_\delta(A_0), \quad \text{where } \delta = \alpha(A_0). \quad (3)$$

This class of knowledge is freely accessible within the community. As public knowledge we understand the agent's name, the type of the organization the agent represents, the general objectives of the agent's activity, the country where the agent is registered, agent's human-human contact (telephone, fax number, email), the human-agent type of contact (http address), the agent-agent type of contact (the IP address, incoming port, ACL) and, finally, available services.

- *Semi-private knowledge* is shared within agents' social neighborhoods. Semi-private knowledge $K_s(A_0)$ of an agent A_0 is defined as

$$K_s(A_0) = K_\delta(A_0), \quad \text{where } \delta = \mu(A_0). \quad (4)$$

As in the OOTW domain, we do not assume the knowledge to be shared within the overlapping alliances, we will require the social neighborhood to fulfill the property of *completeness*: $\forall A \in \Theta : \mu^-(A) = \mu^+(A) = \mu(A)$. Members of a social neighborhood share information about availability of their resources.

- *Private knowledge* is owned and administered by the agent itself. Private knowledge $K_p(A_0)$ of an agent A_0 is defined as

$$K_{pr}(A_0) = K_\delta(A_0), \quad \text{where } \delta = (A_0). \quad (5)$$

An important type of private knowledge includes agent's collaboration preferences, alliance restrictions, coalition leader restrictions and possible next restrictions, but also agent's planning and scheduling algorithms.

This concept can be generalized for multiple levels of knowledge semi-privateness. There may be different δ neighborhoods (and thus alliances) for each piece of knowledge that the agent is aware of. The question is how these alliances would be formed. One option is to assign an index to the set denoting agents semi-private knowledge $-K_s^i(A_0)$, where $K_s^0(A_0) = K_p(A_0)$. We would obviously require that

$$K_p(A_0) = K_s^0(A_0) \subseteq K_s^1(A_0) \subseteq \dots \subseteq K_s^n(A_0) = K_{pr}(A_0). \quad (6)$$

2.3. Knowing the same thing

Agents may not only to share knowledge intentionally but very often it happens that a collection of agents knows the same thing without being aware of this fact. Therefore, we need to introduce another knowledge structure R_δ , denoting a set of true formulae that are known by all agents, members of the environment δ :

$$R_\delta(A_0) = \{\varphi | \forall A_i \in \delta : (\text{Bel } A_i \varphi) \wedge \forall B_i \notin \delta : (\neg \text{Bel } B_i \varphi)\}. \quad (7)$$

In the following we shall see how these constructs are used in the agents' private and semi-private knowledge disclosure.

2.4. Disclosure of private and semi-private knowledge

Measuring the loss of information, that the agents may want to keep private, is an uneasy task. The revealed piece of information has got different value to agents with different meta-reasoning capabilities [9]. In order to vaguely categorize various types of information leaks, let us distinguish between two different types of leaks:

- *unintentional knowledge disclosure* (uKD) – If an agent looses the private knowledge by deliberately disclosing some of its knowledge to other agents (e.g., when sending an inform-type message).
- *intentional knowledge disclosure* (iKD) – if an agent looses some type of private (or semi-private) knowledge as a side effect of some proactive step (such as sending a request-type message) (see Figs. 1 and 2).

In situation calculus, we may define an action $iKD^\xi(A_0, \varphi)$ – agent's A_0 intentional disclosure of the formula φ in an environment ξ as:

$$\begin{aligned} \text{result}(iKD^\xi(A_0, \varphi), s_0) &= s_1, \text{ where} \\ \exists \delta : \varphi \in K_\delta(A_0) &\text{ is true in the situation } s_0 \text{ and} \\ \exists \delta' : \delta' = \delta \cup \xi \wedge \varphi \in K_\delta(A_0) &\text{ is true in the situation } s_1 \end{aligned}$$

Now, we may define an action $uKD^\xi(A_0, \varphi)$ – agent's A_0 unintentional disclosure of the formula φ in an environment ξ as:

$$\begin{aligned} \text{result}(uKD^\xi(A_0, \varphi), s_0) &= s_1, \text{ where} \\ \exists \delta : \varphi \in R_\delta(A_0) &\text{ is true in the situation } s_0 \text{ and} \\ \exists \delta' : \delta' = \delta \cup \xi \wedge \varphi \in R_\delta(A_0) &\text{ is true in the situation } s_1 \end{aligned}$$

Unlike in the case of iKD, the knowledge structure resulting from uKD has no direct relation to the agent

```
(inform
:sender hr-provider-003@cplant:1099/JADE
:receiver hr-provider-002@cplant@mas0:1099/JADE
:language XML
:ontology plan-task-ontology
:protocol fipa-request
:content
(<?xml version="1.0" encoding="UTF-8"?>
  <resource-available>
    <type>medical-assistant</resource=>
    <amount>10</amount>
    <from>20011026T00000000</from>
    <until>20011120T00000000</until>
  </resource-available>
)
```

Fig. 1. Example of an intentional knowledge disclosure – implemented by an inform communicative act. The agent hr-provider-003@cplant : 1099/JADE provides the agent hr-provider-002@cplant@mas0 : 1099/JADE with information about availability of 10 instances of its medical-assistant-type of resources in the specific time.

```

(cfp
:sender hr-provider-002@cplant@mas0:1099/JADE
:receiver hr-provider-003@cplant:1099/JADE
:language XML
:ontology plan-task-ontology
:protocol cfp
:content
(<?xml version="1.0" encoding="UTF-8"?>
  <resource-request>
    <type>medical-assistant</resource=>
  </resource-request>
)
)

(propose
:sender hr-provider-003@cplant:1099/JADE
:receiver hr-provider-002@cplant@mas0:1099/JADE
:language XML
:ontology plan-task-ontology
:protocol fipa-cfp
:content
(<?xml version="1.0" encoding="UTF-8"?>
  <resource-propose>
    <type>medical-assistant</resource=>
    <amount>10</amount>
    <from>20011026T0000000000</from>
    <until>20011120T0000000000</until>
  </resource-propose>
)
)

```

Fig. 2. Example of an unintentional knowledge disclosure – implemented by a cfp – call-for-proposal communicative act. The agent hr-provider-002@cplant : 1099/JADE initiates a contract-net-protocol (explained in Section 4.1) by sending a call for proposals to several agents. The agent hr-provider-003@cplant@mas0 : 1099/JADE replies with a proposal to deliver 10 instances of its medical-assistant-type of resources in the specific time.

A_0 . While in the case of the shared knowledge, the agent A_0 is expected to be aware of the fact that the knowledge of the formula φ is shared, in the case of uKD the agent A_0 may only reconstruct the ξ part of the δ' environment and the neighborhood of δ is unknown.

There are many different ways how one can measure the amounts of disclosed knowledge in the coalition formation process. We are interested in the amount of knowledge the others know about the given agent. Let us first analyze intentional knowledge disclosure. One possible way is a size of the part of the agent's social neighborhood $-\mu(A)^+$. Another alternative is to measure the size of the agent's semi-private knowledge, the knowledge shared within the alliance $-\mathcal{K}_s(A)$. Neither of these measures captures well the fact that there may be different agents in $\mu(A)^+$ knows different things from $\mathcal{K}_s(A)$. Under an assumption that an agent sends a single piece of information in one message, we may measure the amount of the communicated inform-type of messages. In the case of unintentional knowledge disclosure we will be measuring the amount of propose and request-type of messages that were sent outside of the agent's social neighborhood.

2.5. Alliance, coalition, team action plan

In the subject domain, we will understand as the multi-agent community Θ the whole collection of agents participating in the above-described OOTW (quasi-volunteered,

vaguely organized groups of people, non-governmental organizations, institutions providing humanitarian aid, army troops or official governmental initiatives). We will introduce the concept of an alliance as a collection of agents that share information about their resources and all agree to form possible coalitions. The alliance is regarded as a long-term cooperation agreement among the agents. Members of an alliance will all belong to one others' social neighborhood. We define the alliance as follows:

Definition 1. θ is divided into alliances $\lambda = \{\lambda_j\}$, so that:

- (i) $\forall A \in \Theta \exists \lambda_j \in \lambda : A \in \lambda_j$,
- (ii) $\forall \lambda_j \in \lambda \forall A \in \lambda_j : \lambda_j = \mu(A)$.

For an alliance to be well formed, it is necessary that: $\forall A \in \Theta A_i \in \mu(A) : \mu(A) = \mu(A_j)$. In the rest of the document we will refer to this property as *reciprocal knowledge sharing*.

A singleton agent is regarded as an alliance with just one member. From the requirements for the reciprocal knowledge sharing within an alliance follows that

$$\forall \lambda_j \in \lambda \forall A \in \lambda_j : \lambda_j = \mu(A). \quad (8)$$

Therefore, an important property – *alliance exclusivity* – of an alliance is that it cannot overlap with another alliance:

$$\forall \lambda_1, \lambda_2 \subseteq \Theta : (\exists A : A \in \lambda_1 \wedge A \in \lambda_2) \Rightarrow \lambda_1 \equiv \lambda_2. \quad (9)$$

Let us define a *coalition* as a set of agents, which agreed to fulfill a single, well-specified goal. Coalition members committed themselves to collaborate on the within-coalition-shared goal. Under the assumption $\forall A \in \Theta : A \in \varepsilon(A)$ we define coalition as follows:

Definition 2. A *coalition* achieving jointly a goal τ is a set of agents

$$\chi(\tau) \subseteq \Theta, \text{ so that: } \forall A \in \chi(\tau) : \chi(\tau) \subseteq \varepsilon(A).$$

Let us introduce a set $\varepsilon(A, \tau)$ that is an agent collaboration neighborhood with respect to a single shared goal τ . Then

$$\varepsilon(A) = \bigcup_{\tau} \varepsilon(A, \tau), \text{ and } \forall \chi(\tau) \subseteq \Theta : \forall A \in \chi(\tau) : \chi(\tau) = \varepsilon(A, \tau). \quad (10)$$

A coalition, unlike an alliance, is usually regarded as a short-term agreement between collaborative agents. As we will see in Section 6, it is better for a coalition to be a subset of one alliance, but it is not an inevitable condition. A coalition can consist of agents who are members of different alliances.

Another term that we have to introduce is a *team action plan*. In planning humanitarian relief operations, similarly as in the case of any other collaborative action planning, the agents must agree on how they will achieve the goal τ . The team action plan is thus a decomposition of a goal τ into a set of tasks $\{\tau_i\}$. The tasks will be delegated within

the coalition members. Apart from the responsible agent, each task shall be denoted by its due time, start time, and price. Provided that an agent A_j is responsible for implementing the task τ_i (where $\tau = \{\tau_i\}$) in time $\text{due}(\tau_i)$, starting at $\text{start}(\tau_i)$ for the price $\text{price}(\tau_i)$, we define the team action plan as follows:

Definition 3. A team action plan $\pi(\tau)$ is a set $\pi(\tau) = \{\tau_i, A_j, \text{start}(\tau_i), \text{due}(\tau_i), \text{price}(\tau_i)\}$.

We say that the team action plan $\pi(\tau)$ is correct if all the collaborators A_j are able to implement the task τ_i in the given time and for the given price. The team action plan $\pi(\tau)$ is *accepted* if all agents A_j get committed to implementing the task τ_i in the given time and for the given price. Similarly, we say about the goal τ to be *achievable*, if there exists such $\pi(\tau)$ that is correct. The goal τ is said to be *planned*, if there exists $\pi(\tau)$ that is accepted. Obviously, there is an important relation between the team action plan and the coalition. We say that a coalition $\chi(\tau)$ achieves a goal τ by implementing a team action plan $\pi(\tau)$ if and only if $\chi(\tau) = \{A_j\}$ and $\pi(\tau)$ is correct and accepted.

3. Agents' acquaintance model

Let us very briefly introduce the concept of agent's social intelligence and acquaintance models. Apart from its *problem-solving knowledge* that guides agent's autonomous local decision making processes (such as coalition formation, or team action planning), the agents usually exploit *social knowledge* that expresses the other agent's behavioral patterns, their capabilities, load, experiences, resources, commitments, knowledge describing conversations or negotiation scenarios [8]. This knowledge is usually stored separately from the agents' computational core – in an agent's *acquaintance model*. There have been investigated several acquaintance models previously. Based on the *tri-base acquaintance model* [10], where the social knowledge was administrated in three bases (*task-base, cooperator-base, state-base*) the social knowledge in CPlanT is organized in four separate knowledge structures (see Fig. 3):

- *community-base* (Com-BB) – which is a collection of the community members' public knowledge

$$\text{Com-BB}(A_0) = \bigcup_{A_i=z(A_i)} K_p(A_i), \quad (11)$$




Fig. 3. Structure of the CPlanT acquaintance model.

- *self-belief-base* (Self-BB) – where the agent's reflective knowledge about itself is located; here the agent stores its public knowledge that is accessible to anyone, its semi-private knowledge that is shared within the alliance and its private knowledge that is not shared by anyone

$$\text{Self-BB}(A_0) = K_p(A_0) \cup K_s(A_0) \cup K_{pr}(A_0), \quad (12)$$

- *social-belief-base* (Soc-BB) – where the agent stores the semi-private knowledge of its peer alliance members

$$\text{Soc-BB}(A_0) = \bigcup_{A_i=\mu(A_i)} K_s(A_i). \quad (13)$$

- *coalition-base* (Coal-BB) – which is a dynamic collection of the peer coalition members, the past and possible future coalitions as much as permanent coalition-formation rules.¹

See Fig. 4 for an example of the content of an acquaintance model in the OOTW scenario.

Exploitation of the acquaintance model reduces communication traffic required for collaborative activity planning. In principle, the social knowledge substantially reduces the set of agents (ideally to one) that will be requested by the coordinating agent in the CNP process [17]. An important flaw of this approach is rooted in high requirements for the social model maintenance. The social knowledge maintenance may be driven either by the owner of the acquaintance model (the coordinator) or by those which are represented in the model – hence service providers (collaborators). We refer to the former strategy as the *requestor-driven* knowledge maintenance and to the latter strategy as the *provider-driven* knowledge maintenance. As an example of a requestor-driven strategy there is the concept of *periodical revisions* [7] where the knowledge owner periodically checks consistency of the model with the potential collaborators. In other systems, there has been a *cooperation trader* [2] type of agent, which was in charge of maintaining the agents social knowledge. As explained in Section 8 we have adopted the provider-driven knowledge maintenance in CPlanT.

4. Inter-agent communication

Before explaining the lifecycle of the system let us comment the main communication techniques that have been used in CPlanT: central communication agent, contract net protocol, and acquaintance model based contraction. We have tried to minimize the role of the central communication component, as it is an important communication bottleneck of the system operation and a center where

¹ The coalition-formation rules are instances of the agent's problem-solving knowledge, while the information about the coalition members, past and future coalitions are instances of the social knowledge. Therefore the coalition base belongs in part to the acquaintance model and to the agent's body.

Self-Belief Base		
public knowledge:	Semi-private knowledge:	Private knowledge
Port: 1500 ip_address: "147.32.86.167" Country: suffer terra City: north port Type: Religious Ontologies: fipa-am, cplant-ontology	Food: 3000 Nurses: 50 Trucks: 13	Alliance restrictions: ("country","Suffer Terra") Leader restrictions: ("type","Military"). City restrictions: ("muslim",50) Cooperates with: ("type","government")
Social belief base		
	Agent: ST Police	Armed-people:30
	Agent: Christian STHO	Food: 3500 Clothing: 280 Nurses: 22 Medical-people: 12
Community belief base		
	Agent: Suffer Terra Government	Suffer Terra Government@iioip://147.32.84.131:2188/Suffer Terra Government Type: Government Services: Food, Civil-material, Medical-material, Clothing Ontologies: FIPA-AGENT-MANAGEMENT, MAP-ONTOLOGY, PORT-ONTOLOGY, CPLANT, ALLIANCE Languages: SL1, KIF, State: ACTIVE Country: Suffer Terra, City: Suffer Town
	Agent: Christian STHO	Christian Suffer Terra Humanitarian Organization@iioip://147.32.84.131:2210/Chr ST Humanitarian Organization Type: Religious Services: Food, Clothing, Medical-people, Nurses, Medical-material Ontologies: FIPA-AGENT-MANAGEMENT, MAP-ONTOLOGY, PORT-ONTOLOGY, CPLANT, ALLIANCE Languages: SL1, KIF, State: ACTIVE Country: Suffer Terra, City: North Port
Coalition Base		
Rules	(VOLCANIC-AVERAGE-SMALL-TOWN → Time: 220 (Requirements: Medical-material 60, Food 1500, Civil-material 30000, Medical-people 16, Civil-people 27, Nurses 19) ...	
Coalitions	(coalition (Members: Suffer Terra Government, Suffer Terra Police, Christian Suffer Terra Humanitarian Organization) (Services: Food, Civil-material, Medical-material, Clothing, Military-people, Food, Clothing, Medical-people, Nurses) (Price-for-coordination: 5)) (planned-coalition (Task name: Suffer-Town-24-1-2002/17-49-53.1 (Coalition members: Suffer Terra Government, Suffer Terra Police, Christian Suffer Terra Humanitarian Organization) (Coalition leader: Christian Suffer Terra Humanitarian Organization (Disaster: Volcanic, Degree: 1, (Allocations: Civil-material, 80000, Allocation Time: 350 Food, 80000, Allocation Time: 350 Medical-material, 80000, Allocation Time: 350)) ...	

Fig. 4. Instance of an agent's acquaintance model.

the agents' private knowledge may be sniffed (see Section 5).

4.1. Contract net protocol

The CPlanT implementation relied heavily on the *contract net protocol* (CNP) negotiation scenario [17]. Any agent can initiate the coalition formation process (hereafter we refer to this agent as a coalition *coordinator*) by requesting some agents in the community (*collaborators*) for specific services (see Fig. 5). Upon receiving proposals for collaboration, the coordinator carries out a computational process by which it selects the best possible collaborator(s). The coalition planning process can also be multi-staged. Such an approach requires substantial computational resources and fails in complex communities. For each single-staged CNP within a community of n agents, it is needed to send $2(n + 1)$ messages in the worst case.

At the same time many agents may not want to enter the CNP negotiation, as they would not wish to undertake the risk of disclosing their private knowledge.

4.2. Acquaintance model based contraction

The alternative communication strategy to CNP is based on exploitation of the agents' social knowledge. A coalition coordinator subscribes (by sending a *subscribe*-type of

message) the potential collaborators for specific services they may want to exploit in the future (see Fig. 6). Upon a change in the collaborators' capabilities, they provide the coordinator with an update in the form of an *inform*-type of message. When the coordinator triggers the coalition formation phase, it parses the prepared service offers and selects the best collaborator(s) without any further negotiation. The coordinator sends a request, the collaborator updates its resources and confirms the contract. Any change in collaborator resources is advertised to all coordinators which subscribed the collaborator (see Fig. 4).

If there is a single event in the community Θ that affects all the agents ($n = |\Theta|$) and all the agents are mutually subscribed, then in the worst case there is $(n(n - 1))$ messages required for the social knowledge maintenance on this event. However, this is rarely the case. Agents never subscribe all each other (we could easily use a central communication component instead).

5. CPlanT operation lifecycle

The CPlanT multi-agent system operates in four separate phases: *registration* for agents' login/logout to/from the community, *alliance formation* for forming of alliances, *coalition formation* for finding a group of agents which can fulfill a well-specified task and *team action planning* for




Fig. 5. Contraction based on a single-staged contract net protocol.

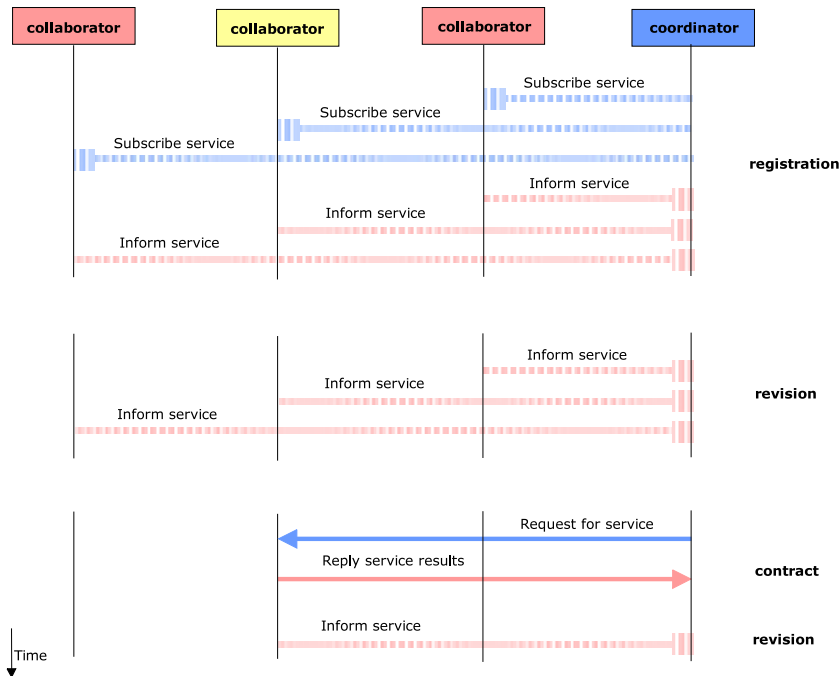


Fig. 6. Contraction based on acquaintance model exploitation.

resource allocation within the specific coalition. In the following, we will comment each of the phases.

5.1. Registration

Throughout the registration phase, a new-coming agent registers within the multi-agent community. This agent registers its public knowledge with the special central registration agent – the *facilitator*. Subsequently, the facilitator informs all the already existing agents about the new agent, and it also informs the new agent about all existing agents. Similarly, the agents can deregister with the facilitator. Any registered agent stores the public knowledge about all members of its total neighborhood $\alpha(A)$ in the Com-BB(A) base of its acquaintance model. We have used the central communication unit – facilitator in the registration phase only. As the agents register just their public knowledge, we do not breach the requirements for confidentiality of the private information.

5.2. Alliance formation

In this phase, which follows the registration process, the agents analyze the information they have about the members of the multi-agent system and make an attempt to form alliances. In principle, each agent is expected to compare its own private knowledge (i.e., alliance formation restrictions) with the public knowledge about the possible alliance members (i.e., type of the organization, its objectives, country of origin, etc.). Had the agent detected a possible future collaborator, the agent would propose joining the alliance. Throughout the negotiation process, the agent either chooses the best alliance according its collaboration preferences of agents into already existing alliances. Failing to do so, an agent may start a new alliance by itself.

According to their preferences in Self-BB and community public knowledge in Com-BB, the agents carry out a selective contract net protocol process during this phase. The *quality of an alliance* is understood in terms of maxi-

mizing the individual agent's contribution to the alliance (i.e., covering the biggest amount of services that the other members of the alliance cannot implement).

It is important to note that this process does not give us any guarantee for optimality of the alliance allocation. Each agent will join the most profitable alliance with respect to existing alliance configuration. With changing the order of agents' registration with the alliance, the formation algorithm will come up with different alliances.

Each agent undertakes intentional knowledge disclosure when *forming an alliance*. At this moment the agent's semi-private knowledge gets disclosed within the alliance members. For an action $\text{iKD}^\xi(A_0, \varphi)$, the resulting environment $\delta' = \lambda$, provided $A_0 = \lambda$ in the situation s_1 . In the case that the agent A_0 is joining an alliance λ , it sends an `inform`-type of a message to all the agents who are already members of the respective alliance. In the situation s_0 it is true that $\delta = \{A_0\}$, and $\xi = \lambda$. If another agent A_k joins the alliance λ in a later situation s_k (which is when $\delta = \lambda$ and $\xi = A_k$) the agent A_0 discloses its semi-private knowledge to the agent A_k . This results in the knowledge sharing neighborhood $\delta' = \lambda \cup \{A_k\}$.

5.3. Coalition leader selection

Unlike the previous phases that got are implemented in the initialization phase of the community, starting from the *coalition leader selection* phase we will be talking about an activity of the community that is a reaction to a single, well-specified task that needs to be accomplished. Both, the CNP technique and the acquaintance model have been used in the coalition leader selection and coalition formation process.

In the first place, there is a problem of selecting the best coalition coordinator C who will be in charge of forming the respective coalition. As already mentioned we wanted to avoid centrality in this phase. Any agent who will be notified about the request for help can participate in a specific process that we call the *coalition coordinator competition*. In the CPlanT system the agents have subscribed the geographical map for update. Such an update represents a particular call for help. Agents use their private knowledge in order to decide whether they will try to challenge others. Those who will, they parse their acquaintance models in order to evaluate coverage (in percents, and in delivery time) of the request by all their alliance members. This bid is broadcasted to all other competitors. Those who can provide better coverage or provide the delivery faster, counter-propose with a copy to all others. Provided that each agent proposed a different proposal, there is always one who did not get any counter-proposal and the others got at least one counter-proposal. The one who did not get any counter-proposal is then autonomously selected to be a coalition leader. The same functionality (with substantial less communication and computation overheads) can be achieved by a contract-net-protocol organized by an agent who takes care about the map.

5.4. Coalition formation

First, let us talk about the coalition formation process within a single alliance. The coordinator, parses its social neighborhood $\mu(A)$ and detects the set of the most suitable collaborators (cooperation neighborhood) – $\varepsilon(A, \tau)$. Upon an approval from each of the suggested agents, the respective coalition $\chi(\tau) = \varepsilon(A, \tau)$ is to be formed. Maintaining the agents' social neighborhood will save an important part of agent's interaction in the period of coalition formation. Agents will not need to broadcast a call for collaboration each time they will be required to accomplish a task. Instead, they will consult this pre-prepared knowledge and will contract the agent of which they knew it is the best to work with. The coordinator optimizes the *quality of a coalition* by seeking the coalitions that would contribute the most and in the shortest possible time. When the coalition is formed within a single alliance only neither private nor semi-private knowledge is disclosed.

However sometimes an alliance fails to form an appropriate coalition. The coordinator, who failed to form a coalition within one alliance, negotiates the proposal for collaboration by classical CNP with the agents from its total neighborhood $\alpha(A_0)$. This is the moment when the coordinator may initiate unintentional disclosure of private/semi-private knowledge. Once the agent C from an alliance λ initiates a contract-net-protocol by sending a request for proposals for delivering a service τ to the agents $\{A_1, A_2, \dots, A_n\} = \xi$, the agent C reveals the information about the *intent* (e.g., C has an intention to do something that requires τ) and information about agent's C capabilities (e.g., C cannot do τ). The environment δ is expected to be empty before starting the CNP. Generally, the agent C discloses its private knowledge, therefore it should be irrelevant whether this interaction takes place within one alliance or across the alliances. In the CPlanT system however, the contract-net-protocol is executed only across the alliances, as the alliance members have got the social knowledge about other the peer agents' capabilities. Instead of the CNP, the `subscribe-inform` social knowledge maintenance mechanism can be used. A proposal for collaboration from an agent A_n from an alliance λ_n reveals the information about agent's A_n capabilities (such as A_n can implement τ in time t_1). As any agent A_n acts on behalf of the alliance λ_n , this type of knowledge disclosure representing the loss of information about capabilities of the entire alliance has been reduced.

5.5. Team action planning

Once a coalition is formed, the agents share a joint commitment to achieve the goal τ . Within this phase, a team of collaborative agents jointly creates a team action plan $\pi(\tau)$. The team action plan, that is a result of the coalition planning activity, is a joint commitment structure that defines exactly how each team member will contribute to achieving the shared goal (amount of resources, deadlines, etc.). The

coordinator is supposed to (i) decompose a goal τ into subtasks $\{\tau_i\}$ and (ii) allocate the subtasks within the already formed coalition $\chi(\tau)$. There may be many achievable team action plans $\pi(\tau)$. The coordinator seeks for the cheapest or the fastest possible plan.

As there is no semi-private knowledge shared in between the alliances, the agents from different alliances coordinate their activities by means of the contract net protocol. The intra-alliance team action planning mechanism is not the pure acquaintance model contraction, where the team action plan would result from the coalition leader deliberation process followed by a contract. All coalition members construct the precise team action plan collaboratively.

The collaborators advertise their services in the most informative while efficient form. Therefore the coordinator's acquaintance model stores the social knowledge that is imprecise, but very compact and efficient to parse. According to this social knowledge, the coordinator suggests the most optimal request decomposition and resource allocation $\pi(\tau)$ and transforms it into a contract proposal. This proposal is sent to the other coalition members, which reply with a specific collaboration proposal. However, the coordinator may find this proposal to be different than expected owing to the fact that the approximate information provided by the collaborator was too imprecise. Instead of agreeing upon a joint commitment for this sub-optimal team action plan, the coordinator adapts the conflicting social knowledge and fires another round of negotiation. With each further negotiation stage, the team action plan should be closer to the optimal team action plan. This process is to be iterated until there is no conflict in the expected capacity of the collaborators and the proposed delivery.

6. Implementation and testing

6.1. Implementation

Verifying applicability of the CPlanT required a well-defined, formal, but realistic enough scenario that can represent, model and initiate all aspects of agents' non-trivial behavior. The above specified principles and ideas have been tested and implemented on a subset of the OOTW types of operations – humanitarian relief operations. For this purpose we designed and implemented a hypothetical humanitarian scenario Sufferterra representing a suffering island and several imaginary countries ready to help. The Sufferterra scenario was inspired by [20,12,13]. The scenario knowledge has been encoded in XML and the computational model of the scenario has been implemented in Allegro Common Lisp.

The *geographical-agents* (*g-agents*) specify the physical arrangements of the geographical objects and the resources they provide. The problem specification does not distinguish the level of modeling granularity, i.e., each physical object may be implemented as a *g-agent* or several physical objects can make together a *g-agent*. For the test-

ing purposes we have implemented a single *g-agent* that represents the entire map of the area. The “call-for-help” functionality that specifies a particular disaster (“volcanic”, “hurricane”, “flood”, “earthquake”), the degree of disaster (1, . . . , 9) and location has been integrated within the *g-agent*. The *humanitarian-agents* (*h-agents*) subscribe the *g-agent* for specific information.

CPlanT has been successfully tested on the Sufferterra humanitarian relief scenario [11]. The implementation is complemented by a visualizing meta-agent, which is implemented in Java. This meta-agent views the logical structure of the system e.g., alliances, coalitions, team action plans, and other properties of the community. There is a separate visualization for communication traffic monitoring. This component, that is not an agent, but rather a part of the multi-agent platform, serves mainly to debugging purposes. The community can be viewed and the requests can be sent from the web server via classical Internet browsers and from the WAP phones interface as well.

6.2. Experiments, testing

Several different objectives were followed within the frame of the experiments: to evaluate the communication and computation requirements, quality of the solution provided and the degree of disclosure of private and semi-private knowledge.

6.2.1. Communication traffic

An important part of the agent deliberation process has been decomposed into the inter-agent negotiation process. This is why we have concentrated our attention primarily to savings of the communication traffic in the entire system. The communication traffic has been observed in different architecture arrangements of the community (e.g., different number of alliances) and for different complexity of the tasks sent to the community (e.g., different number of contracts). Having 20 agents we have experimented with the sample of all agents in one alliance, with agents clustered into 2, 4, 7, and 20 alliances. All the experiments have been carried out on the set of 19 experiments for each of the community arrangement. From the definition of the community lifecycle (see Section 5) follows that the latter case ($\forall A : \mu(A) = \emptyset$) does not exploit any advantages of the acquaintance model contraction and the community behaves such as no social knowledge is administered and used. An important part of the communication traffic is carried out in the critical time – i.e., in the moment when the system is requested to provide a plan. By exploiting social knowledge that has been prepared in advance, we aimed at minimizing communication traffic in this moment. The cost we have paid for this was the increased communication traffic in the idle times of the community operation. In the idle times, the agents are busy with maintaining the social knowledge stored in their acquaintance models. The communication traffic grows with the increasing number of alliances as each alliance member stores a more volumi-

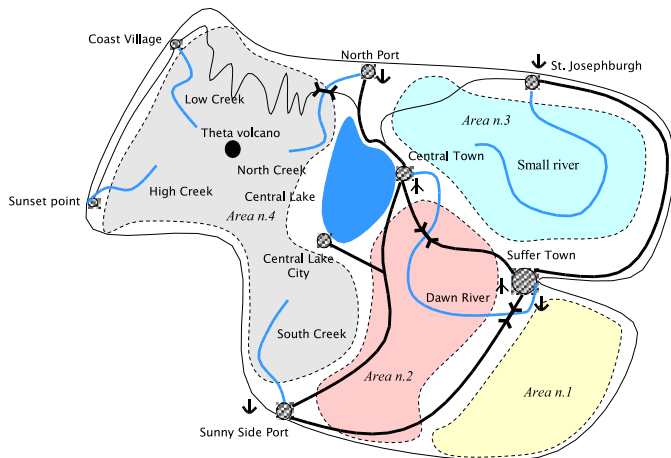


Fig. 7. Sufferterra – subject of humanitarian operations.

nous acquaintance model and it searches for a coalition by parsing the acquaintance model only.

From the graph in Fig. 7 we can see that with an increasing number of alliances (and a decreasing average number of alliance members) we reduce the communication requirements for maintenance of the model. The most of the communication in the critical time (the difference between dark and light bars in the graph) we save in the case of just one huge alliance. The optimal arrangement of the community was identified in the case of four alliances. However, it is not possible to define an optimal system structure because the agents cannot predict future tasks and the number of agents required for implementing these tasks. It is clear that for tasks requiring low number of agents, we will prefer small alliances while for the task requiring many agents, larger alliances will be preferred. Thus, the optimal size of a coalition is given by the nature of the tasks/goals under consideration.

6.2.2. Evaluation of quality of the coalition

The evaluation of quality of the formed coalition is an important aspect in any coalition formation research. In the Sufferterra scenario, there are two key attributes that influence the coalition value: (i) *success rate* – how many of the requested resources the coalition provides and (ii) *delivery time* – by when the coalition delivered the resources to the requestor. Experiments did not give any pieces of evidence to conclude any dependency between the success rate of the coalition and the used communication mechanism. However, with an increasing number of alliances, the overall delivery time is kept increasing due to additional costs of coordination among the alliances.

6.2.3. Knowledge disclosure

The key challenge has been minimization of both the private and semi-private knowledge disclosures. Once the **private** information has been identified by another agent, this agent finds about the intent of the respective agent. As already noted, this very often happens when an alliance

fails to plan all the requests and starts a contract net protocol process within members of the other alliances. Those, who will not be awarded the contract, know that the coordinator intends to operate in a mission and that it needs the resources requested.

The *semi-private* information is disclosed in the same situation, when the possible collaborator proposes a service (as a reaction to a coordinator call for collaboration) that will not be accepted by the coordinator. In such a case, the coordinator finds out about the services the suggested collaborator can provide. Both the above mentioned cases are classified as unintentional knowledge disclosures (see Section 2.4). The intentional knowledge disclosure happens in the registration phase within a single alliance and represents the amount of information that has become shared within the alliance.

As expected, the biggest disclosure of intents appears in the case of 20 alliances, as there is the highest CNP-based communication traffic among the alliances (see Fig. 8). There is no weak disclosure once the agents are utterly independent (20 alliances). On the other hand, there is no strong semi-private information disclosure in one alliance while the independent agents are starting to lose their semi-private information in the strong sense. It makes no implication to put together the strong and weak knowledge disclosures because of their different nature.

An interesting fact is that neither of the two extreme cases is the best for concealing the agents' private and semi-private knowledge. With one alliance, the semi-private knowledge becomes public while with no alliance each contract net protocol will reveal information about the contractors' intentions. It is rather difficult to find a good compromise in a number of alliances. What matters, is the probability that a request will not be fulfilled within one alliance and the coalition leader will have to subcontract other agents. The amount and structures of alliances in our domain emerge naturally according to the agents' private knowledge and other collaboration restrictions. Therefore it makes no sense to suggest an optimal number of alliances for a given community (see Figs. 9 and 10).




Fig. 8. Communication traffic in communities with different number of alliances. The light bar depicts the maintenance messages, while the dark bar illustrates the overall communication in the system.




Fig. 9. The relation between private information disclosure and number of alliances.




Fig. 10. The graph illustrating disclosure of the semi-private knowledge. The light bar depicts the weak and the dark bar strong knowledge disclosure.

7. Relation to coalition planning research

In order to position the research contribution well, let us comment on relation between the CPlanT coalition planning strategy and other coalition planning research achievements. It has been shown that finding the optimal coalition is an NP complete problem [14]. Researchers mainly suggest different negotiation strategies and analyze complexities of the coalition formation process [16]. When a subject of optimization is the quality of the formed coalition, the agents usually act *collaboratively*. There have been published many of centralized planning mechanisms for coalition formation [15]. On the other hand, the *self-interested* agents maximize their own profit when participating in a coalition, no matter how well they will perform as a group. Many researchers analyzed properties of communities of self-interested agents such as their stability, worst-case profit, or payoff division among the agents [6]. The domain we have investigated is partially of cooperative and self-interested type at the same time. The humanitarian aid providing agents tend to cooperate in the time of a crisis while they are self-interested and compete each other in a long-term horizon. Therefore, there was suggested a concept of alliances – collectives of agents that agreed to collaborate (to potentially form a coalition).

More importantly, the profit is very often the key optimization criterion when the agents optimize a coalition formation process (either collaboratively or competing each other). Besides the quality of the coalition, in the OOTW domain there are two (may be more important) aspects to be taken into account. As forming an optimal coalition is a very complex problem, the *response time* becomes an important issue. Agents are limited in resources and a reasonably good answer, that is quickly provided, is very often much better than an optimal coalition found later [18,14]. Practitioners would add that implementing a multi-agent system with a large number of agents, that are supposed to interact heavily, results in a *communication traffic overload* [3]. In our research we have tried to decompose the reasoning process and distribute it among the agents. While keeping the agents' deliberation process simple, we have concentrated our efforts on minimizing the communication interaction among the agents in order to suggest community structuring in a reasonable time. As the OOTW agents are also self-interested in certain way, they want to stay hidden in front of someone and advertise its collaborative capabilities to others. This is why we have to respect also the amount of *private information* to be disclosed. Therefore, we have also studied leaks of private information while forming the coalitions.

Research of the teamwork in a similar domain (evacuation scenarios) was reported in [19]. It was suggested to integrate the already existing software applications in the TEAMCORE wrapper agents. Unlike our acquaintance model that contains just social knowledge, the TEAMCORE wrapper agents also maintain domain specific team plans and the hierarchy of goals. Teams of agents share a team-oriented program, which is a joint knowledge structure that coordinates their activities. In CPlanT, there is no explicit team action plan distributed in agents' acquaintance models. The structure of the coalitions and the team action plan is a result of the inter-agent negotiation process. However, combination of both approaches where the agents' behavior is coordinated by a team action plan that results from the agents' negotiation seems to be an interesting topic for further research.

Investigators approaching the problem from the game-theoretic point of view solve the problem of a higher complexity. Whereas in our case, there is a hierarchy structure for each task that is sent to the community and each task is coordinated by a single agent (the coordinator), in [5] all agents are equal. The agents autonomously analyze their own value. Through negotiations, they try to find out which coalition is the most profitable for them to join. This problem is inherently more complex and causes communication problems in complex communities. There will be several stages of negotiations needed as in many cases optimality of cooperation between two agents may not be reciprocal. In our case, we did not need to solve such a complex problem. On the other hand, in CPlanT we must optimize not only which coalition to join but also which services to provide to the coalition (e.g., team action plan-

ning). One may suggest that the game-theoretic approach could be used in the alliance formation phase of our algorithm (see Section 5.2). However, the agents join the system continuously, which makes it rather difficult to maintain the overall optimality of the distribution of alliances.

Besides the contract-net-protocol, there are other negotiation strategies based on classical auctioning mechanisms. While in combinatorial actions, the motivation of an agent is usually to make the biggest profit (or to contribute to a coalition in the best way), in our case, all the auctioneers and the bidding agents collaborate. The bidding agent tries to provide the auctioneer with such a bid that approximates in the best way the resources it can provide, and will help it to suggest the best possible resource allocation. In CPlanT, the agents also do not speculate about whom to work with. As we optimize the private information loss, collaboration within one alliance is always preferred. There is a potential of using the optimization for multiple auctioning mechanisms for the team action planning within several overlapping coalitions [1].

8. Conclusions

The research described in this paper contributes to the coalition formation community by suggesting an alternative, knowledge based approach to the problem. Our research has been driven by the very specific domain of the OOTW. Apart from the classical contract net protocol techniques, we have used the communication strategy based on combination of three techniques: the centralized registration, the acquaintance models, and the contract net protocol negotiations.

The agents in the community are organized into smaller, disjunctive groups called alliances. Each agent in the alliance is able to start the negotiation process to form a coalition and to develop a team action plan for a specific task either within the alliance or in collaboration with other alliances. Inside-alliance negotiations explore mainly the social knowledge stored in the acquaintance models, but the CNP technique is used as well (especially in the phase of the team action planning). The inter-alliance negotiations are based just on the CNP principles.

The general complexity of negotiations when forming a coalition in a MAS is of an exponentially explosive nature [4]. It has been shown that finding and optimal coalition is an NP complete problem when no specific constraints are imposed. In our case, the negotiation complexity of the coalition formation problem has been significantly reduced because:

- agents are organized into several disjunctive sets (alliances) and the most of coalitions are created just inside an alliance (reduced space of negotiations),
- the coalition leader within an alliance is set randomly (each coalition member has got the same coordination capacity and can manage the negotiation process), they do not compete for the role,

- within an alliance, the negotiation process explores the acquaintance models (social knowledge) in combination with the CNP technique and the pure CNP negotiations are used just in the case of the inter-alliance negotiations.

While the contract net protocol runs rather inefficiently, it keeps the agents from different alliances independent (they do not have to disclose their semi-private knowledge across alliances). This is why, the acquaintance-model based planning has been used exclusively within the alliances.

In our approach, we have not prioritized the requirement for the global coalition optimality, as this is not the main challenge in the OOTW planning. The main issue has been to develop an acceptable plan without forcing the agencies (agents) to make their private knowledge (namely intents and resources) public. This quite specific OOTW requirement enabled to reduce the complexity of the negotiation problem significantly. It has been measured that optimality of the coalition value slightly increases with the number of alliances (the role of the acquaintance model is getting smaller), while the problem complexity with a smaller number of socially knowledgeable alliances is significantly reduced.

References

- [1] P. Anthony, W. Hall, V. Dang, N.R. Jennings, Autonomous agents for participating in multiple on-line auctions, in: Proceedings of IJCAI Workshop on E-Business and the Intelligent Web, Seattle, WA, 2001, pp. 54–64.
- [2] W. Cao, C.-G. Bian, G. Hartvigsen, Achieving efficient cooperation in a multi-agent system: the twin-base modelling, in: P. Kandzia, M. Klusch (Eds.), Co-operative Information Agents, Lecture Notes in Artificial Intelligence No. 1202, Springer, Heidelberg, 1997, pp. 210–221.
- [3] G.A. Kaminka, D.V. Pynadath, M. Tambe, Monitoring deployed agent teams, in: Proceedings of the Fifth International Conference on Autonomous Agents (Agents-2001), Montreal, Canada, 2001.
- [4] S. Ketchpel, Coalition formation among autonomous agents, in: Proceedings of the Fifth European Workshop on Modeling Autonomous Agents in a Multi-Agent World, Neuchatel, Switzerland, August 25–27, 1993.
- [5] M. Klusch, J. Contreras, F. Wu, O. Shehory, Coalition formation in a power transmission planning environment, in: Proceedings of the Second International Conference on Practical Applications of Multi-Agent Systems PAAM-97, London, 1997, pp. 335–352.
- [6] C. Li, K. Sycara, Algorithms for combinatorial coalition formation and payoff division in an electronic marketplace, Technical report CMU-RI-TR-01-33, Carnegie Mellon University, 2001.
- [7] V. Marik, M. Pechoucek, O. Stepankova, J. Lazansky, ProPlanT: multi-agent system for production planning, International Journal of Applied Artificial Intelligence 14 (7) (2000) 727–760.
- [8] V. Marik, M. Pechoucek, O. Stepankova, Social knowledge in multi-agent systems, in: M. Luck et al. (Eds.), Multi-Agent Systems and Applications, Lecture Notes in Artificial Intelligence No. 2086, Springer, Heidelberg, 2001, pp. 211–245.
- [9] M. Pechoucek, D. Norrie, Knowledge structures for reflective multi-agent systems: on reasoning about other agents registered as report number 538, Department of Mechanical and Manufacturing Engineering, University of Calgary, Alberta, Canada, 2000.
- [10] M. Pechoucek, V. Marik, O. Stepankova, Towards reducing communication traffic in multi-agent systems, Journal of Applied System

- Studies, Cambridge International Science Publishing, ISSN 1466-7738, Cambridge, UK, Spring, 2001 (Special issue on virtual organizations and e-commerce applications).
- [11] M. Pěchouček, V. Mařík, J. Bárta, A knowledge-based approach to coalition formation, *IEEE Intelligent Systems* 17 (3) (2002) 17–25.
- [12] R.A. Rathmell, A coalition force scenario ‘binni-gateway to the golden bowl of Africa’, Defense Evaluation Research Agency (1999).
- [13] G.A. Reece, A. Tate, The Pacifica NEO scenario. Technical report ARPA-RL/O-Plan2/TR/3. Artificial Intelligence Applications Institute, University of Edinburgh, Scotland, 1993.
- [14] T. Sandholm, V. Lesser, Coalitions among computationally bounded agents, *Artificial Intelligence* 94 (1) (1997) 99–137 (Special issue on Economic Principles of Multi-Agent Systems).
- [15] T. Sandholm, K. Larson, M. Andersson, O. Shehory, F. Tohme, Coalition structure generation with worst case guarantees, *Artificial Intelligence* 111 (1–2) (1999) 209–238.
- [16] O. Shehory, S. Kraus, Coalition formation among autonomous agents: strategies and complexity, in: Castelfranchi, Muller (Eds.), *Lecture Notes in Artificial Intelligence* No. 957, From Reaction to Cognition, 1995, pp. 57–72.
- [17] R. Smith, The contract net protocol: high-level communication and control in distributed problem solver, *IEEE Transactions on Computers* 29 (12) (1980) 1104–1113.
- [18] E. Steinmetz, J. Collins, M. Gini, B. Mobasher, An Efficient Anytime Algorithm For Multiple-Component Bid Selection in Automated Contracting, in: *Agent Mediated Electronic Trading* Lecture Notes in Artificial Intelligence No. 1571, Springer, 1998, pp. 105–125.
- [19] M. Tambe, Towards flexible teamwork, *Journal of Artificial Intelligence Research* 7 (1997) 83–124.
- [20] E.C.T. Walker, Panel report: coalition planning for operations other than war, Workshop at AIAI, Edinburgh, 1999.
- [21] M. Wooldridge, *Reasoning about Rational Agents*, The MIT Press, Cambridge, MA, 2000.