

A Common Basis for Agent Organisation in BDI Languages^{*}

Anthony Hepple, Louise Dennis, and Michael Fisher

Department of Computer Science, University of Liverpool, Liverpool, U.K.

{A.J.Hepple,L.A.Dennis,M.Fisher}@csc.liv.ac.uk

Abstract. Programming languages based on the BDI style of agent model are now common. Within these there appears to be some, limited, agreement on the core functionality of agents. However, when we come to multi-agent organisations, not only do many BDI languages have no specific organisational structures, but those that do exist are very diverse. In this paper, we aim to provide a unifying framework for the core aspects of agent organisation, covering groups, teams and roles, as well as organisations. Thus, we describe a simple organisational mechanism, and show how several well known approaches can be embedded within it. Although the mechanism we use is derived from the METATEM programming language, we do not assume any specific BDI language. The organisational mechanism is intended to be independent of the underlying agent language and so we aim to provide a common core for future developments in agent organisation.

1 Introduction

As hardware and software platforms become more sophisticated, and as these are deployed in less predictable environments, so the level of *autonomy* built into such systems has increased. This has allowed systems to work effectively without detailed, and constant, human intervention. However, autonomous systems can be hard to understand and even harder to develop reliably. In order to help in this area, the concept of an *agent* was introduced to capture the abstraction of an autonomously acting entity. Based on this concept, new techniques were developed for analysing, designing and implementing agents. In particular, several new programming languages were developed explicitly for implementing autonomous agents.

We can simply characterise an agent as an autonomous software component having certain goals and being able to communicate with other agents in order to accomplish these goals [29]. The ability of agents to act independently, to react to unexpected situations and to cooperate with other agents has made them a popular choice for developing software in a number of areas. At one extreme there are agents that are used to search the Internet, navigating autonomously in order to retrieve information; these are relatively lightweight agents, with

^{*} Work partially supported by EPSRC under grant EP/D052548.

few goals but significant domain-specific knowledge. At the other end of the spectrum, there are agents developed for independent process control in unpredictable environments. This second form of agent is often constructed using complex software architectures, and has been applied in areas such as real-time process control [25,20]. Perhaps the most impressive use of such agents is in the real-time fault monitoring and diagnosis carried out on NASA Deep Space One [22].

The key reason why an agent-based approach is advantageous for modelling and programming autonomous systems, is that it permits the clear and concise representation, not just of *what* the autonomous components within the system do, but *why* they do it. This allows us to abstract away from low-level control aspects and to concentrate on the key feature of autonomy, namely the goals the component has and the choices it makes towards achieving its goals. Thus, in modelling a system in terms of agents, we often describe each agent's *beliefs* and *goals*, which in turn determine the agent's *intentions*. Such agents then make decisions about what action to perform, given their beliefs and goals/intentions. This kind of approach has been popularised through the influential BDI (Belief-Desire-Intention) model of agent-based systems [25]. This representation of behaviour using *mental* notions has several benefits. The first is that, ideally, it abstracts away from low-level issues: we simply present some goal that we wish to be achieved, and we expect it to act as an agent would given such a goal. Secondly, because we are used to understanding and predicting the behaviour of rational agents, the behaviour of autonomous software should be relatively easy for humans to understand and predict too. Not surprisingly, therefore, the BDI approach to agent modelling has been successful and has led to many novel programming languages based (at least in some part) upon this model; these are often termed *BDI Languages*. Although a wide variety of such languages have been developed [2] few have strong and flexible mechanisms for *organising* multiple agents, and those that do provide no agreement on their organisational mechanisms. Thus, while BDI languages have converged to a common core relating to the activity of individual agents [9], no such convergence is apparent in terms of multi-agent structuring.

Our overall aim is to provide a common logically based framework for BDI style agent programming (which incorporates organisational aspects) to facilitate agent verification [4]. As a result a clear goal is to develop a simple, intuitive and semantically consistent organisation mechanism. In this paper we show how a simple model can, in BDI languages, encompass many proposed models of multi-agent organisation and teamwork. The formal semantics of this approach is considered in detail in [10].

Paper Structure Section 2 surveys some of the leading approaches to agent organisation that have already been proposed and illustrates their diverse nature. In Section 3 we describe the structuring mechanism we propose for unifying the multi-agent concepts. Section 4 demonstrates how our framework can be used to model concepts such as joint-intentions, roles, etc., which form the basis of the

approaches surveyed in Section 2. Finally, in Section 5, we provide concluding remarks and outline future work.

2 Approaches to Agent Organisation

In this section we overview some of the key approaches to the organisation of agents that have been proposed. It is important to note that we are particularly concerned with *rational agents*, predominantly using the BDI model of computation. While we have not listed *all* approaches, the selection we give covers many of the leading attempts at teamwork, organisational structuring and role-based computation. In addition, while we are primarily interested in developing BDI languages with clear logical semantics and logic-based mechanisms, we also consider organisational approaches beyond this class.

2.1 Cohen and Levesque: *Joint Intentions*

With a respected philosophical view on agent co-operation, Cohen and Levesque produced a significant paper ‘Teamwork’ [8] extending previous work [21, 6, 7]. They persuasively argue that a team of agents should *not* be modelled as an aggregate agent and propose new (logical) concepts of *joint intentions*, *joint commitments* and *joint persistent goals* to ensure that teamwork does not break down due to any divergence of individual team members’ beliefs or intentions. The authors’ proposals oblige agents working in a team to retain team goals until it is mutually agreed amongst team members that a goal has now been achieved, is no longer relevant, or is impossible. This level of commitment is stronger than an agent’s commitment to its individual goals which are dropped the moment it (individually) believes they are satisfied. Joint intentions can be reduced to individual intentions if supplemented with mutual beliefs.

2.2 Tidhar, Cavedon and Rao: *Team-Oriented Programming*

Tidhar [27] introduced the concept of *team-oriented programming* with social structure. Essentially this is an agent-centred approach that defines joint goals and intentions for teams but stops short of forcing individual team members to adopt those goals and intentions. An attempt to clarify the definition of a ‘team’ and what team formation entails is made using concepts such as ‘mind-set synchronisation’ and ‘role assignment’. Team behaviour is defined by a temporal ordering of plans which guide (but do not constrain) agent behaviour. A social structure is proposed by the creation of *command* and *control* teams which assign roles, identify sub-teams and permit inter-team relationships. In [5], the authors formalise their ideas of social structure with concepts of commitment expressed using modal logic. This allows the formal expression of commitment between teams, such as

team A intends to achieve task B for the sake of team C.

2.3 Ferber, Gutknecht and Michel: *Roles and Organisations*

Ferber *et al.* [13] present the case for an organisational-centred approach to the design and engineering of complex multi-agent systems. They cite disadvantages of the predominant agent-centred approaches such as: lack of access rights control; inability to accommodate heterogeneous agents; and inappropriate abstraction for describing organisational scenarios. The authors propose a model for designing language independent multi-agent systems in terms of *agents*, *roles* and *groups*. Agents and groups are proposed as distinct first class entities although it is suggested that an agent ought to be able to transform itself into a group. (We will see later that this is close to our approach.)

In [14], Ferber continues to argue for an organisational-centred approach, advocating the complete omission of mental states at the organisational level, defining an organisation of agents in terms of its capabilities, constraints, roles, group tasks and interaction protocols. Clearly articulated here is a manifesto of design principles.

2.4 Pynadath and Tambe: *TEAMCORE*

Pynadath *et al.* [24] describe their interpretation of team-oriented programming that aims to organise groups of heterogeneous agents to achieve team goals. A framework for defining teams is given that provides the following concepts:

- Team*—an agent without domain abilities;
- Team-ready*—agents with domain abilities that interface with team agents;
- Sub-goal*—a goal that contributes to the team goal; and
- Task*—the allocation of a sub-goal to a team-ready agent.

An implementation of their framework, TEAMCORE, provides organisational functionality such as multicast communication between agents, assigning tasks, maintaining group beliefs and maintaining hierarchies of agents (by role). Heterogeneous agents are accommodated by wrapper agents that act as proxies for the domain agent.

2.5 Fisher, Ghidini and Hirsch: *Groups as Agents*

Beginning within the context of executable temporal logics [1], Fisher *et al.* produced a series of papers [17, 16, 18, 15] that developed the METATEM language into a generalised approach for expressing dynamic distributed computations. As we will see more about this model in Section 3, we just provide a brief outline here.

Organisational structuring within the METATEM language [15] consists of a simple nested grouping structure where groups comprise communicating elements (objects, agents, or other software components). The key aspect of this approach is that groups themselves are also agents, providing a homogeneous, simple, yet expressive, model. In [16], it is argued that systems composed of components as diverse as objects, web services and abstract features can be modelled with this general approach.

2.6 Hübner, Sichman and Boissier: *Roles and Permissions*

Hübner *et al.* believed that the agent organisational frameworks proposed prior to their 2002 paper [19] overlooked the significant relationship between structural and functional properties of an organisation. Thus, in [19], they propose a three component approach to the specification of agent organisations that combines independent structural and functional specifications with a deontic specification, the latter defining among other things the roles (structural) having permission to carry out group tasks (functional). The approach provides a proliferation of constructs for specifying multi-agent systems, including the ability to concisely express many additional aspects, such as

- the ability to specify *compatibility* of group membership, akin to the members of a government expressing a conflict of interest.
- enabling the *cardinality* of group membership to be defined and thus defining a well formed group as a group whose membership is between its specified minimum and maximum size.
- control of the organisation’s goal(s), with an ability to specify sequential, branching and parallel execution of sub-goals.
- the ability to express a variance in the agents’ permissions over time.

It is argued that such an approach improves the efficiency of multi-agent systems by focusing agents on the organisation’s goals. Indeed, we note that of all the proposals discussed in this section this approach provides the developer with the widest vocabulary with which to express agent behaviour when defining the organisation.

2.7 Dignum, Esteva, Sierra and Vázquez-Salceda: *Institutions*

These authors made formal [12] and practical [11, 28] contributions to a method of agent organisation that enjoys much current popularity [23]. An electronic institution aims to provide an *open* framework in which agents can contribute to the goals of *society* without sacrificing its own self-interest; the implication being that an autonomous agent will be motivated to participate in the institution by its desire to satisfy its own goals, but that its participation will be structured by the framework in such a way that institutional goals are achieved. A key concept is that of institutional norms.

In [12], the institution remains independent of agent-architecture by modelling agents as roles, of which there are two types—internal and external (to the institution)—with different rights. A *dialogue* defines valid locutions, a *scene* is a unit of interaction within an institution and a *performative structure* defines an objective as a network of scenes. In an attempt to allow more agent autonomy these ideas were refined and in [28] more concepts were introduced, including *landmarks* that can be used to guide agents through an interaction when a prescriptive dialogue is considered too constraining.

Perhaps the most noteworthy aspect of these proposals is the change of focus from the agents themselves onto the interactions that take place between agents.

2.8 Summary

It should be noted that none of the above organisational approaches can comprehensively model all forms of co-operative multi-agent systems. Rather they represent attempts to discover practical and beneficial ways of specifying distributed computational systems, and facilitating the focus of computation on a system's main purpose whilst not compromising the autonomy of the system's components. In achieving this aim it may be convenient to categorise groups of agents in terms of cohesion and co-operation. For instance, a *group* of agents may be individually autonomous, existing as a group solely due to their proximity to one another rather than their co-operation. In contrast, the word *team*, implies a high degree of co-operation and adhesion with an *organisation* fitting somewhere in between. As Cohen stated in [8]

“teamwork is more than co-ordinated individual behaviour”.

Thus, the more expressive proposals reviewed here enable the specification of more cohesive groups but often at significant cost to the agents involved.

3 Structuring Mechanisms

The approach we propose is based on that of METATEM described previously in [15]. However, we advocate this grouping approach, independent of the underlying language for agents. The aim of our grouping structure is to provide a simple organisational construct that enables the definition of a variety of multi-agent systems, ranging from unstructured collections of uncoordinated agents to complex systems that are often described using the high-level abstractions described in the last section.

The basic restrictions we put on any underlying language is that, as in most BDI-based languages, there are logically coherent mechanisms for explicitly describing *beliefs* and *goals*. As in the METATEM framework, the grouping approach involves very few additional constructs within the language [10]. Specifically, we require just two additional elements within each agent's state. We also, as is common, require that first-class elements such as beliefs, goals and plans, can be communicated between agents. Delivery of messages should be guaranteed, though the delay between send and receipt is not fixed. Finally, we expect asynchronously concurrent execution of agents.

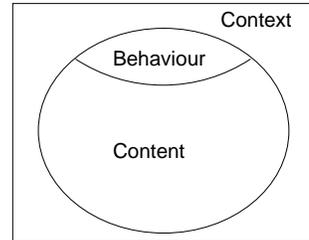
3.1 Extending Agents

Assuming that the underlying agent language can describe the behaviour of an agent, as has been shown for example in [9], we now extend the concept of agent with two sets, **Content** and **Context**. The agent's **Content** describes the set of agents it contains, while the agent's **Context** describes a set of agents it is contained within. Thus, the formal definition of an agent is as follows [17].

Agent ::= Behaviour: Specification
Content: $\mathcal{P}(\text{Agent})$
Context: $\mathcal{P}(\text{Agent})$

Here, $\mathcal{P}(\text{Agent})$ are sets of agents and **Specification** is the description of the individual agent's behaviour, given as appropriate in the target BDI language.

On the right, we provide a graphical representation of such an agent. The agent (the circle) resides within a **Context** and itself comprises its own behavioural layer and its **Content**. This **Content** can again contain further agents. Note that, for formal development purposes, the **Behaviour** may well be a logical specification.



The addition of **Content** and **Context** sets to each agent provides significant flexibility for agent organisation. Agent teams, groups or organisations, which might alternatively be seen as separate entities, are now just agents with non-empty **Content**. This allows these organisations to be hierarchical and dynamic, and so, as we will see later, provides possibilities for a multitude of other co-ordinated behaviours. Similarly, agents can have several agents within their **Context**. Not only does this allow agents to be part of several organisational structures simultaneously, but it allows the agent to benefit from **Context** representing diverse attributes/behaviours. So an agent might be in a context related to its physical locality (i.e. agents in that set are 'close' to each other), yet also might be in a context that provides certain roles or abilities. Intriguingly, agents can be within many, overlapping and diverse, contexts. This gives the ability to produce complex organisations, in a way similar to multiple inheritance in traditional object/concept systems. For example, see Fig. 1 for sample configurations.

An important aspect is that this whole structure is very dynamic. Agents can move in and out of **Content** and **Context** sets, while new agents (and, hence, organisations) can be spawned easily and discarded. This allows for a range of structures, from the *transient* to the *permanent*. From the above it is clear that there is no enforced distinction between an agent and an agent organisation. All are agents, all may be treated similarly.

While it may seem counter-intuitive for an organisation to have beliefs and goals, many of the surveyed systems required team constructs such as tasks or goals that can naturally be viewed as belonging to a team/group agent. Some also required *control agents* to manage role assignment and communication which in this framework can be handled by the containing agent itself if so desired. On the other hand it is possible to distinguish between agents (with empty **Content**) and organisations (with non-empty **Content**) and for a programmer to exclude certain constructs from organisations in order to allow an organisation-centred approach, if required.

Finally, it is essential that the agent's internal behaviour, be it a program or a specification, have direct access to both the **Content** and **Context** sets. As

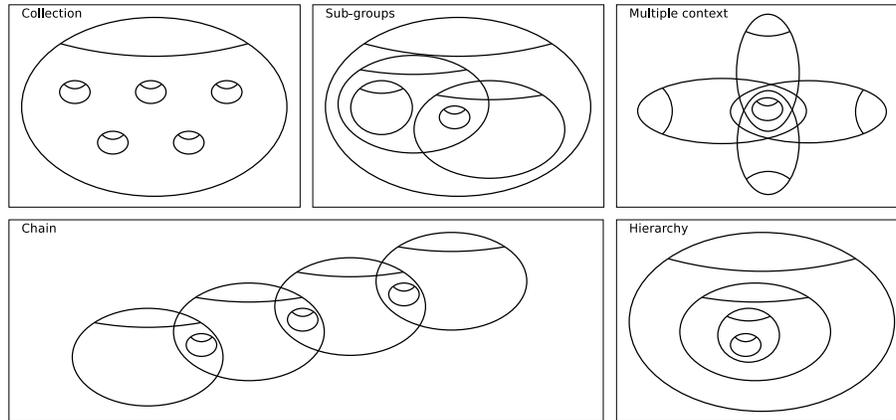


Fig. 1. A selection of possible organisation structures.

we will see below, this allows each agent to become more than just a ‘dumb’ container. It can restructure, share information and behaviour with, and control access to its **Content**. To describe fragments of the agent’s behaviour during the rest of the paper, we will use simple **IF...THEN...ELSE** statements. Again, this does not prescribe any particular style of BDI language.

3.2 Communication

The core communication mechanism between agents in our model is broadcast message-passing. The use of broadcast is very appealing, allowing agent-based systems to be developed without being concerned about addresses/names of the agents to be communicated with. The potential inefficiency of broadcast communication is avoided by the use of the agents’ **Content** and **Context** structures. By default, when an agent broadcasts a message, that message is sent to all members of the agent’s **Context** set and forwarded to agents within the same context. This, effectively, produces *multicast*, rather than full broadcast, message-passing.

This is clearly a simple, flexible and intuitive model, and the system developer is encouraged to think in this way. However, it is useful to note that multicast, or ‘broadcast within a set’, is actually implemented on top of point-to-point message passing! We will assume that the BDI language has a communication construct that can be modelled as the action $send(recipient, m)$ which means that the message m has been sent to the agent *recipient*, and a corresponding $received(sender, m)$ which becomes true when the *recipient* agent receives the message m from *sender*. Let us consider an example where an agent wishes to broadcast to all other members of one of its **Context** sets. For simplicity, let us term this context set ‘*group*’. An agent wishing to ‘broadcast’ a message, m , to members of the *group* sends a message, $send(group, broadcast(m))$, to the group agent alone, as illustrated in Fig. 2.

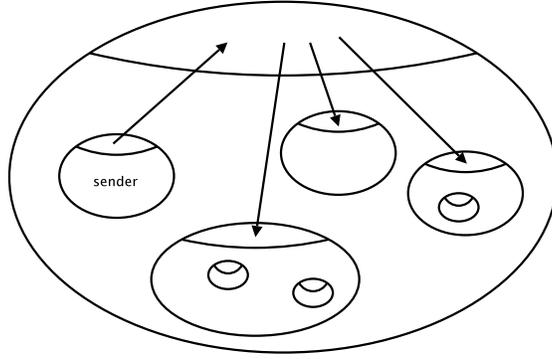


Fig. 2. Broadcast within a Group.

The effect of sending a broadcast message to the *group* agent is that the *group* acts as a proxy and forwards the message to its **Content**, modifying the message such that the message appears to have originated from the proxy. In this way agents maintain their anonymity within the group.

IF $received(from, broadcast(m))$
THEN for each x **in** $\{Content \setminus from\}$ $send(x, m)$

Being an agent-centred approach to multi-agent organisation there does not exist an (accessible) entity that references *all* agents in the agent space, thus *true* broadcast is not possible. However a number of recursive group broadcasts can be specified, allowing a message to be propagated to all agents with an organisational link to the sender.

For example, reaching all accessible agents requires the sending agent to send a message to all members of its **Content** and **Context** sets and for each first-time recipient to recursively forward that message to the union of their **Context** and **Content** (excluding the sender). Clearly this is not an efficient method of communication as it is possible for agents to receive multiple copies of the same message, and so it may not be practical in very large societies, but what it lacks in sophistication it makes up for in simplicity and clarity [16].

IF $received(from, broadcastAll(m))$ **AND not** $received(_, m)$
THEN for each x **in** $\{Content \cup Context\}$
 $send(x, m)$ **AND** $send(x, broadcastAll(m))$

Perhaps more useful than indiscriminate broadcasting would be the case of an agent who wants to reach all other members of the ‘greatest’ organisation to which it belongs. This requires a message to propagate up through the agent structure until it reaches an agent with an empty context, at which point the message is sent downwards until all members and sub-members have been reached.

To illustrate this, consider the situation of agent **E** in Fig. 3(a), who wants to send a message to its entire organisation—the organisation specified by **A**.

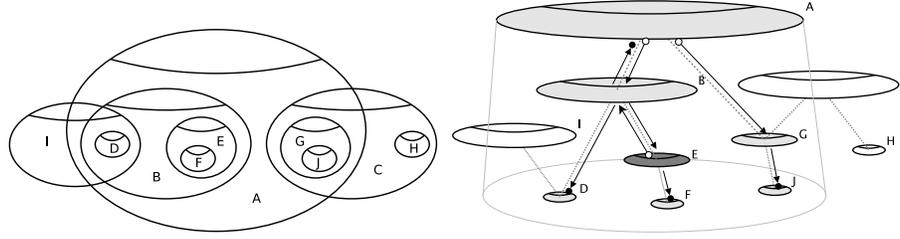


Fig. 3. (a) Nested Organisations (b). Propagation of Messages

A *propagateUp(m)* message originates from agent E who sends it to agent B. B's context is non-empty so the message continues upwards to A. Since A is the widest organisation to which E belongs (it has an empty *Context* set), it modifies the message, converting it to *propagateDown(message)* and broadcasts it along with the message to all members of its *Content*. Upon receipt of this message, agents B and G send it to their *Content* and so it continues until the message reaches an agent with an empty *Content* as illustrated by Fig. 3(b). This might be specified as follows.

IF *received*($_$, *propagateUp(m)*) **AND** *Context* $\neq \emptyset$
THEN **for each** x **in** {*Context*}
send(x , *propagateUp(m)*)

IF *received*($_$, *propagateUp(m)*) **AND** *Context* = \emptyset
THEN **for each** x **in** {*Content*}
send(x , m) **AND** *send*(x , *propagateDown(m)*)

IF *received*($_$, *propagateDown(m)*) **AND** *Content* $\neq \emptyset$
THEN **for each** x **in** {*Content*}
send(x , m) **AND** *send*(x , *propagateDown(m)*)

3.3 Refining and Restricting Communications

Further restriction of communication is possible by, for example, restricting the type of communications agents can make. Employing the concept of speech acts [26] we can use the group agent as a communication filter that restricts intra-group messaging to those messages that conform to permissible protocols or structures.

If, for example, a fact-finding agent contains a number of agents with access to information resources, it may be necessary to restrict their communication to **inform** speech acts. In such circumstances it is possible to modify the default behaviour by imposing a message filter.

IF *received*(*from*, *broadcast(m)*) **AND** *informFilter(m)*
THEN **for each** x **in** {*Content* \ *from*} *send*(x , m)

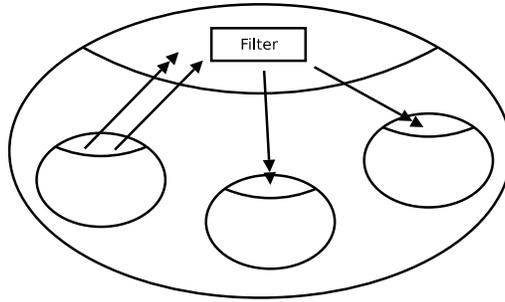


Fig. 4. Filtering communication by group.

See Fig. 4 for an example of this. In this way filters can be adapted for many purposes, enabling organisations to maintain:

relevance — ensuring communication is relevant to group goals and intentions;
fairness — allowing each member of a group an equal opportunity to speak; and
legality — assigning permissions to group members to restrict communication.

3.4 Communication Semantics

The above variations on *broadcast* define varying semantics for a message. A key feature of the grouping approach is that the semantics of communication is flexible and, potentially, in the hands of the programmer. Such semantics can also, potentially, be communicated between agents in the form of plans allowing an agent to adopt different semantics for communication as its **Context** changes.

Adherence to particular common communication protocols/semantics also allows groups to establish the extent to which a member is autonomous (e.g., a group can use a semantics for **achieve** speech acts which forces recipients to adopt the communicated goal). This is important because organisational approaches vary from those in which group behaviour is specified by the organisation and imposed on its members with little option for autonomy to those in which group behaviour emerges from an appropriate combination of individual agents without any explicit coordination at all.

4 Common Multi-Agent Structures

In this section we will examine some of the key structuring mechanisms that are either explicit or implicit within the approaches surveyed in Section 2, and show how each might be represented appropriately and simply, using the approach outlined above. Table 1 lists the mechanisms identified by our surveyed authors as being useful in the specification of agent co-operation. We believe that our approach is flexible enough to model all of these but for brevity we will only demonstrate a sample of them.

		Joint intentions	Roles	Shared beliefs	Joint commitments	Global goals	Communication	Organisation/Team abstraction	Mutual beliefs	Social beliefs	inter-[team/group] structure	Sub-teams	Social commitment	Groups	Interactions/scenes	Task structure	Organisational	Agent centred
Cohen & Levesque	Joint intentions	✓																
Esteva et al	Institutions		✓															
Ferber	Roles and groups		✓															
Fisher et al.	Groups																	
Hübner	Roles and permissions		✓															
Pynadath & Tambe	Team-oriented		✓	✓														
Tidhar, Cavedon & Rao	Team-oriented	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓

Table 1. Multi-agent organisation concepts.

4.1 Sharing Information

Shared beliefs Being a member of all but the least cohesive groups requires that some shared beliefs exist between its members. Making the contentious assumption that all agents are honest and that joining the group is both individual rational and group rational, let agent i hold a belief set BS_i . When an agent joins a group¹ j it receives beliefs BS_j from the group and adds them to its own belief base (invoking its own belief revision mechanism in case of conflicting beliefs). The agent in receipt of the new beliefs may or may not disseminate them to the agents in its content, depending on the nature and purpose of the group. Once held, beliefs are retained until contradicted.

Joint beliefs Joint beliefs are stronger than shared beliefs. To maintain the levels of cohesion found in teams each member must not only believe a joint belief but must also believe that its team members also believe it. Let us assume the agent is capable of internal actions such as $addBelief(Belief, RelevantTo)$ adding $Belief$ to its belief base, and recording the context that $Belief$ is relevant to, and $removeBeliefs(Context)$. Upon joining a group, an agent is supplied the beliefs relevant to that context, which it stores in its belief base along with the context in which they hold. This behaviour is captured in the rule below.

IF $received(from, membershipConfirm(beliefSet))$
THEN for each b **in** $\{beliefSet\}$ $addBelief(b, from)$.

The presence of such **Context** meta-information can be used to specify boundaries on agent deliberation, thus mitigating the complexity caused by introducing

¹ Let us refer to such an agent as a *group* to distinguish it from the agent within its **Content**.

another variable. When leaving a **Context** an agent might either choose to drop the beliefs relevant to that **Context** or to retain them.

4.2 Sharing Capabilities

Let agent Ag_i have a goal G , for which a plan P exists. However, Ag_i does not have plan P and therefore must find an agent that does. Two options available to Ag_i are to find an agent Ag_j , who has P , and either: request that Ag_j carries out the plan; or request that Ag_j sends P to Ag_i so that Ag_i can carry out the plan itself. The first possibility suggests a closer degree of co-operation between agents i and j , perhaps even the sub-ordination of agent j by agent i . Whereas, in the second possibility, agent i benefits from information supplied by j .

In the first scenario we might envisage a group in which a member (or the group agent itself) asks another member to execute the plan. In the second case, we can envisage agents i and j *sharing* a plan. This second scenario is typical if groups are to capture certain capabilities —agents who join the **Content** of such a group agent are sent (or at least can request) plans shared amongst the group. Either scenario can be modelled using our approach.

4.3 Joint Intentions

An agent acting in an independent self-interested way need not inform any other entity of its beliefs, or changes to them. On the other hand, an agent who is working, as part of a team, towards a goal shared by itself and all other members of the team has both an obligation and a rational interest in sharing relevant beliefs with the other team members [8]. This gives an agent a *persistent* goal with respect to a team. Such that the agent must intend the goal whilst it is the team’s mutual belief that the goal is valid (not yet achieved, achievable and relevant) — it must not give up on a goal nor assume the goal has been achieved, independently. The implications of this impact on agent’s individual behaviour when it learns, from sources external to the group, that the goal is no longer valid. In such a situation the team/group agent maintains its commitment to the invalid goal but informs its team members of the antecedent(s) that lead it to believe the goal is invalid. Only when the agent receives confirmation that the entire team share its belief does it drop its commitment.

The intuitive implementation of this joint intention is not via a team construct but as an extension of an agent’s attributes, however, increases in expressiveness of this sort are often accompanied by increased complexity. The organisational or team construct may overcome this problem but we believe that our simple group approach is sufficient to implement joint intentions, mutual beliefs and common goals. Consider the scenario given in Fig. 5.

Agent A. On joining group T , agent A accepts goal JI and confirms its adoption of the goal. Whilst T remains a member of A ’s **Context**, A informs T of all beliefs that are relevant to JI . Finally, all communications from agent T must be

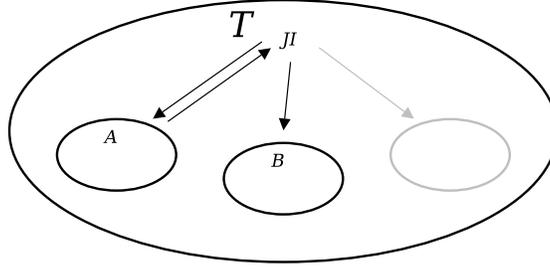


Fig. 5. Communicating Joint Intentions.

acknowledged, with an indication of the agent's acceptance (or non-acceptance) of the information.

A simple specification of this might be:

IF *received*(*from*, *jointIntention*(*JI*)) **AND** *inContext*(*from*)
THEN *achieve*(*JI*) **AND** *send*(*from*, *ack*(*JI*))

IF *belief*(φ) **AND** **there is** *x* **in** {*Context*} *relevantTo*(φ , *x*)
THEN *send*(*x*, *inform*(φ))

IF *goal*(γ) **AND** **there is** *x* **in** {*Context*} *relevantTo*(γ , *x*)
THEN *achieve*(γ)

Thus, an agent is obliged to inform its group of beliefs relevant to jointly held intentions and will maintain a goal whilst it remains relevant to its **Context**.

Agent T. Evaluates group beliefs and communicates both the adoption, and dropping, of intentions when mutual agreement is established. Since *T* has details of the agents in its **Content** and can send messages to interrogate them, it can maintain knowledge of *common* information and behaviours, and reason with this knowledge.

4.4 Roles

The concept of a role is a common abstraction used by many authors for a variety of purposes [19, 14, 28], including:

- to define the collective abilities necessary to achieve a global goal.
- to provide an agent with abilities suitable for team activity.
- to constrain or modify agent behaviour for conformance with team norms.
- to describe a hierarchy of authority in an organisation of agents and hence create a permissions structure.

Roles are most obviously integrated into our framework as further agents whose **Content** is those agents fulfilling the role and whose **Context** is the organisation to which the role belongs. However in some cases, in particular strict hierarchies, it may be possible to associate roles directly with the organisational agent. Below we examine a variety of such role types and consider in more detail how each could fit into our model.

Ability roles Let plan P be a complex plan that requires abilities x, y and z if it is to be fulfilled. An agent A is created (without any domain abilities of its own) to gather together agents that have the necessary abilities. Agent A might generate a new agent in its **Content** for each of the abilities required to fulfil plan P .

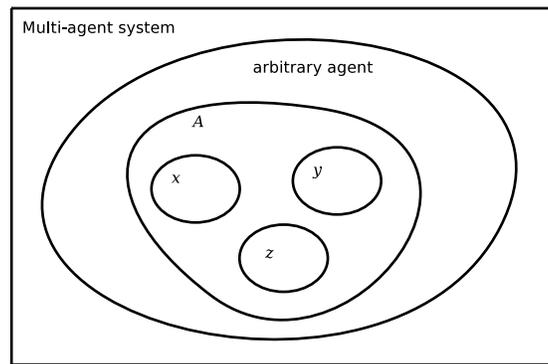


Fig. 6. Roles according to abilities.

When agents with abilities x, y or z join the **Content** of agent A , A adds them to the **Content** of the appropriate group (agent), analogous to assigning roles.

A talented agent might become a member of several ability sets. The ability set, itself an agent, may be a simple container or could exhibit complex behaviour of its own. One basic behaviour might be to periodically request (of the agents in its **Content**) the execution of its designated ability. Note that, in the case of an ability that is hard to carry out, it may be provident to include many agents with that ability. Similarly, the desired ability might be a complex ability that must be subjected to further planning, resulting in a number of nested abilities.

Roles in society Joining an institution, organisation or team of agents commonly involves the adoption of the norms of that institution, organisation or team. Whether these norms are expressed as beliefs, goals, preferences or communication protocols, our approach allows them to be transmitted between group members, particularly at the time of joining. For example, if team membership

requires that members acknowledge receipt of messages then each new member of a group might be given the new rule (behaviour)

IF $received(ag, \theta)$ **THEN** $send(ag, ack(\theta))$.

A stronger constraint might require an agent to believe all messages received from its **Context**:

IF $received(ag, \theta)$ **AND** $ag \in \mathbf{Context}$
THEN $addBelief(\theta, ag)$ **AND** $send(ag, ack(\theta))$.

Of course, agents can not be certain that another agent will keep with given constraints or comply with norms of the society, the most it can do is demand formal acknowledgement of its request and a commitment to do so. Group membership can be denied if an agent fails to satisfy the *entry criteria*.

Authority roles None of the structures discussed usefully reflect hierarchies of authority. Each allow almost arbitrary group membership, with transitive and cyclic structures possible making them unsuitable for expressing a hierarchy of authority, which by its nature must be acyclic with exactly one root.

A common use for such a hierarchy is for creating channels of communication. Our approach to grouping enables communication restrictions for free, as agents may only communicate with their immediate superiors (**context**), or their direct subordinates (**content**). Communication to peers (by multicast) can only be achieved by sending a single *broadcast* message to the agent common to the contexts of the intended recipients. The receiving [superior] agent will, if it deems it appropriate, forward the message to the other agents in its content.

5 Concluding Remarks

In this paper, we have proposed a simple but clear model for multi-agent structuring in agent languages based on varieties of the logical BDI approach. Although derived from work on METATEM, we propose this as a general approach for many languages. To support this, we first show how simple and intuitive the approach is and how the underlying structures of any appropriate language can be modified. (Note that more detailed operational semantics for our grouping approach in logic-based BDI languages is given in [10].) We then showed, in a necessarily brief way, how many of the common teamwork and organisation aspects can be modelled using our approach.

In order to evaluate the approach, we have also implemented it in AgentSpeak (actually, Jason [3]) and have developed several simple examples of dynamic organisations. This simple additional layer has so far proved to be convenient and powerful. Obviously, the **Content/Context** approach has also been extensively used in previous work on METATEM [16–18]. In addition, it has been incorporated in the semantics of AIL [9], a common semantics basis for a number of languages, including AgentSpeak and 3APL; see [10] for formal details.

5.1 Future Work

Our immediate aim with this work is to apply the model to larger applications, particularly in the areas of ubiquitous computing and social organisations. This will give a more severe test for the approach and will highlight any areas of difficulty.

As mentioned above, the approach is being integrated into the AIL semantics [9], which provides a common semantics basis for a number of BDI languages. Since translations from AgentSpeak, 3APL, etc are being produced, we also aim to translate the organisational aspects used into the above model.

Finally, since the aim of the work on AIL is to provide generic verification techniques for BDI languages (that can be translated to AIL) [4]. In extending the AIL semantics, we also aim to provide verification techniques for teams, roles and organisations developed within BDI languages.

References

1. H. Barringer, M. Fisher, D. Gabbay, R. Owens, and M. Reynolds, editors. *The Imperative Future: Principles of Executable Temporal Logic*. Research Studies Press, 1996.
2. R. H. Bordini, M. Dastani, J. Dix, and A. El Fallah Seghrouchni, editors. *Multi-Agent Programming: Languages, Platforms and Applications*. Springer-Verlag, 2005.
3. R. H. Bordini, J. F. Hübner, and M. Wooldridge. *Programming Multi-Agent Systems in AgentSpeak using Jason*. Wiley, 2007.
4. R. H. Bordini, M. Fisher, W. Visser, and M. Wooldridge. Verifying Multi-Agent Programs by Model Checking. *Journal of Autonomous Agents and Multi-Agent Systems* 12(2):239-256, 2006.
5. L. Cavedon, A. S. Rao, and G. Tidhar. Social and Individual Commitment. In *Proc. PRICAI Workshop on Intelligent Agent Systems, Theoretical and Practical Issues*, pages 152–163. Springer, 1997.
6. P. R. Cohen and H. J. Levesque. Intention is Choice with Commitment. *Artificial Intelligence*, 42(2-3):213–261, 1990.
7. P. R. Cohen and H. J. Levesque. Confirmations and Joint Action. In *Proc. International Joint Conference on Artificial Intelligence (IJCAI)*, pages 951–959, 1991.
8. P. R. Cohen and H. J. Levesque. Teamwork. Technical Report 504, SRI International, California, USA, 1991.
9. L. A. Dennis, B. Farwer, R. H. Bordini, M. Fisher, and M. Wooldridge. A Common Semantic Basis for BDI Languages. In *Proc. 7th International Workshop on Programming Multiagent Systems (ProMAS)*, 2007. (To appear in LNAI, Springer.)
10. L. A. Dennis, M. Fisher, and A. Hepple. Language Constructs for Multi-Agent Programming. In *Proc. 8th Workshop on Computational Logic in Multi-Agent Systems (CLIMA)*, 2007. (To appear in LNAI, Springer.)
11. M. Esteva, D. de la Cruz, C. Sierra. ISLANDER: and electronic institutions editor. In *Proc. 1st International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 1045–1052, ACM, 2002.
12. M. Esteva, J. A. Rodríguez-Aguilar, C. Sierra, P. Garcia, J. L. Arcos. On the Formal Specification of Electronic Institutions. Chapter in *Agent Mediated Electronic*

- Commerce. The European AgentLink Perspective.*, pages 126–147. LNAI Number 1991. Springer-Verlag, 2001.
13. J. Ferber and O. Gutknecht. A Meta-model for the Analysis and Design of Organizations in Multi-agent Systems. In *Proc. 3rd International Conference on Multi-Agent Systems (ICMAS)*, pages 128–135, 1998.
 14. J. Ferber, O. Gutknecht, and F. Michel. From Agents to Organizations: An Organizational View of Multi-agent Systems. In *Proc. Workshop on Agent-Oriented Software Engineering (AOSE)*, volume 2935 of *LNAI*, pages 214–230. Springer, 2003.
 15. M. Fisher. METATEM: The Story so Far. In *Proc. 3rd International Workshop on Programming Multiagent Systems (ProMAS)*, volume 3862 of *Lecture Notes in Artificial Intelligence*, pages 3–22. Springer, 2006.
 16. M. Fisher, C. Ghidini, and B. Hirsch. Programming Groups of Rational Agents. In *Proc. International Workshop on Computational Logic in Multi-Agent Systems (CLIMA)*, volume 3259 of *LNAI*. Springer-Verlag, 2004.
 17. M. Fisher and T. Kakoudakis. Flexible Agent Grouping in Executable Temporal Logic. In *Intensional Programming II*. World Scientific, 2000.
 18. B. Hirsch. *Programming Rational Agents*. PhD thesis, Department of Computer Science, University of Liverpool, June 2005.
 19. J. F. Hübner, J. S. Sichman, and O. Boissier. A Model for the Structural, Functional, and Deontic Specification of Organizations in Multiagent Systems. In *Proc. 16th Brazilian Symposium on Artificial Intelligence (SBIA)*, pages 118–128. Springer, 2002.
 20. N. R. Jennings and M. Wooldridge. Applications of Agent Technology. In *Agent Technology: Foundations, Applications, and Markets*. Springer, 1998.
 21. H. J. Levesque, P. R. Cohen, and J. H. T. Nunes. On Acting Together. In *Proc. 8th American National Conference on Artificial Intelligence (AAAI)*, pages 94–99, 1990.
 22. N. Muscettola, P. P. Nayak, B. Pell, and B. Williams. Remote Agent: To Boldly Go Where No AI System Has Gone Before. *Artificial Intelligence*, 103(1-2):5–48, 1998.
 23. P. Noriega, J. Vázquez-Salceda, G. Boella, O. Boissier, V. Dignum, N. Fornara, and E. Matson, editors. *Coordination, Organization, Institutions and Norms in agent systems II (COIN)*, volume 4386 of *Lecture Notes in Computer Science*. Springer-Verlag, 2007.
 24. D. V. Pynadath, M. Tambe, N. Chauvat, and L. Cavedon. Towards Team-Oriented Programming. In *Intelligent Agents VI*, volume 1757 of *LNAI*, pages 233–247. Springer, 1999.
 25. A. S. Rao and M. Georgeff. BDI Agents: from Theory to Practice. In *Proc 1st International Conference on Multi-Agent Systems (ICMAS)*, pages 312–319, San Francisco, USA, 1995.
 26. I. A. Smith and P. R. Cohen. Toward a Semantics for an Agent Communications Language Based on Speech-Acts. In *Proc. American National Conference on Artificial Intelligence (AAAI)*, pages 24–31, 1996.
 27. G. Tidhar. Team-Oriented Programming: Preliminary Report. Technical Report 1993-41, Australian Artificial Intelligence Institute, Melbourne, Australia, 1993.
 28. J. Vázquez-Salceda, V. Dignum, and F. Dignum. Organizing Multiagent Systems. *Journal of Autonomous Agents and Multi-Agent Systems* 11(3):307-360. Springer, 2005.
 29. M. Wooldridge and N. R. Jennings. Intelligent Agents: Theory and Practice. *The Knowledge Engineering Review*, 10(2):115–152, 1995.