

# Co-ordination in software agent systems

H S Nwana, L Lee and N R Jennings

---

*The objective of this paper is to examine the crucial area of co-ordination in multi-agent systems. It does not attempt to provide a comprehensive overview of the co-ordination literature; rather, it highlights the necessity for co-ordination in agent systems and overviews briefly various co-ordination techniques. It critiques these techniques and presents some conclusions and challenges drawn from this literature.*

---

## 1. Introduction

Co-ordination is a central issue in software agent systems in particular, and in distributed artificial intelligence (DAI) in general. However, it has also been studied by researchers in diverse disciplines in the social sciences, including organisation theory, political science, social psychology, anthropology, law and sociology. For example, organisation theorists have investigated the co-ordination of systems of human beings, from small groups to large formal organisations [1, 2]. Economists have studied co-ordination in markets of separate profit-maximising firms [3]. Even biological systems appear to be co-ordinated though individual cells or 'agents' act independently and in a seemingly non-purposeful fashion. Human brains exhibit co-ordinated behaviour from apparently 'random' behaviours of very simple neurones. Essentially, co-ordination is a process in which agents engage in order to ensure a community of individual agents acts in a coherent manner.

The scope of this paper is limited mainly to software multi-agent systems (MAS). The breakdown of the paper is as follows. Section 2 argues more cogently why co-ordination is vital, not only in DAI and agent-based systems, but also to open distributed systems. Section 3 overviews briefly various co-ordination techniques. Negotiation, overviewed in section 4, is one of the techniques used to achieve co-ordination in agent-based systems. It is overviewed separately because it has an extensive literature in its own right. Section 5 presents some lessons learned and some heuristics on co-ordination, and section 6 concludes the paper.

## 2. The necessity for co-ordination

What is co-ordination? Why is it necessary or desirable? Co-ordination has already been defined, in section 1, as a process in which agents engage in order to ensure their community acts in a coherent manner. Coherence means that the agents' actions gel well, and that they do not conflict with one another. In other words, coherence refers to how well a system of agents behaves as a unit [4]. There are several reasons why multiple agents need to be co-ordinated [5, 6].

- Preventing anarchy or chaos — co-ordination is necessary or desirable because, with the decentralisation in agent-based systems, anarchy can set in easily. No longer does any agent possess a global view of the entire agency to which it belongs. This is simply not feasible in any community of reasonable complexity. The chairman of BT, for example, cannot possibly be aware of the detailed activities of all his 130 000 employees. Consequently, agents only have local views, goals and knowledge which may conflict with others. They can enter into all sorts of arrangements with other agents or agencies. Like in any society, such haphazard arrangements are prone to anarchy; to achieve common goals, which is a *raison d'être* for having multiple agents in the first place, a group of agents need to be co-ordinated.
- Meeting global constraints — there usually exist global constraints which a group of agents must satisfy if they are to be deemed successful. For example, a system of

agents constructing a design may have to work within the constraints of a pre-specified budget. Similarly, agents doing network management may have to respond to certain failures within seconds and others within hours. Agents need to co-ordinate their behaviour if they are to meet such global constraints.

- Distributed expertise, resources or information — agents may have different capabilities and specialised knowledge in a similar manner to paediatricians, neurologists and cardiologists. Alternatively, they may have different sources of information, resources (e.g. processing power, memory), reliability levels, responsibilities, limitations, charges for services, etc. In such scenarios, agents have to be co-ordinated in just the same way that different medical specialists including anaesthetists, surgeons, ambulance personnel, nurses, etc, work together to treat someone who has been in a near-fatal accident. In this example, none of these experts working in isolation possesses all the necessary expertise, information, or the casualty and medical resources (e.g. equipment).
- Dependencies between agents' actions — agents' goals are frequently interdependent. Consider two agents solving the trivial blocks world problem shown in Fig 1.

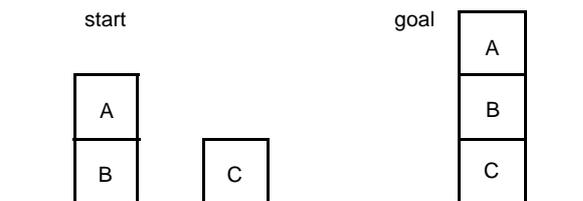


Fig 1 A blocks world problem.

The easiest way to solve it would be for the first agent to take on the sub-goal of stacking B on C while the second stacks A on top of the stack B-C in order to achieve A-B-C. Clearly, the sub-goals are interdependent; the second agent has to wait for the first agent to complete its sub-goal before it can perform its own. Where such interdependencies exist, and they invariably do in MAS, the activities of the agents must be co-ordinated.

- Efficiency — even when individuals can function independently, thereby obviating the need for co-ordination, information discovered by one agent can be of sufficient use to another agent that both agents can solve the problem twice as fast.

Co-ordination, in turn, may require co-operation; but it is important to emphasise that co-operation among a set of agents would not necessarily result in co-ordination; indeed, it may result in incoherent behaviour. This is because for agents to co-operate successfully, typically, they must

maintain models of each other as well as develop and maintain models of future interactions. If agents' beliefs about each other are wrong, for example, incoherent behaviour may well result. Co-ordination may occur without co-operation. For example, if a person is running towards you, and you get out of his way, you have co-ordinated your actions with his. However, you have not entered into co-operation with him. Likewise, non-co-operation among agents does not necessarily lead to incoherent behaviour (it may just happen to end up co-ordinated); however, it is likely to do so. Competition is a form of co-ordination involving antagonistic agents. To achieve co-ordination, agents may have to communicate with one another. However, as Huhns and Singh [7] point out, agents may achieve co-ordination without communication, provided they possess models of each others' behaviours. In such a situation, co-ordination can be achieved mainly via organisation. However, to facilitate co-ordination, where agents have to co-operate through communication, it is vital that they make known their goals, intentions, results and state to other agents. The literature on co-ordination abounds because it is used to address several DAI and distributed computing issues, including:

- network coherence — maximising how well a distributed system of agents work together,
- task and resource allocation among agents,
- recognising and resolving disparities or conflicts in goals, facts, beliefs, viewpoints and behaviour of agents,
- determining the organisational structure (i.e. architecture) of an agent set-up, i.e. defining the roles, responsibilities and chains of authority between agents.

The above list is not exhaustive. The outcome of a co-ordination process is, hopefully, a coherent set-up of agents — a set-up in which deadlock and livelock are avoided. Deadlock refers to a state of affairs in which further action between two or more agents is impossible; on the contrary, livelock refers to a scenario where agents continuously act (e.g. exchange tasks), but no progress is made. A co-ordinated set-up also maximally exploits the capabilities of individual agents and minimises conflicts and resource contentions between them. A coherent system will minimise or avoid conflicting and redundant efforts among agents. Clearly then, co-ordination is a necessary and desirable property of agent systems.

Co-ordination is not only a concern to agent researchers. With emerging technologies such as open systems, client-server computing, the World Wide Web (WWW) and the Internet, there is an unmistakable trend towards distributed computing systems, and this trend is set to continue well into the next millennium. The autonomous components or

parts of such open distributed systems, which are analogous to agents in agent-based systems, need to be co-ordinated.

### 3. An overview of co-ordination techniques

There are many approaches which have been devised to achieve co-ordination in agent systems. Co-ordination techniques are classified here in the following four broad categories:

- organisational structuring,
- contracting,
- multi-agent planning,
- negotiation.

There follows an overview and brief critique of each category separately.

#### 3.1 Organisational structuring

This is the simplest co-ordination scenario which exploits the *a priori* organisational structure. This is because the organisation defines implicitly the agent's responsibilities, capabilities, connectivity and control flow. It provides a framework for activity and interaction through the definition of roles, communication paths and authority relationships. Durfee et al define this as the pre-defined long-term relationships between agents [8]. Hierarchical organisations abound, yielding the classic master/slave or client/server co-ordination technique, used typically for task and resource allocation among slave agents by some master agent. This technique is implemented in a couple of ways.

- The master agent plans and distributes fragments of the plan to the slaves. The slaves may or may not communicate among themselves, but must ultimately report their results to the master agent.

In this case, while the master has full autonomy with respect to the slaves, the slaves have only partial autonomy with respect to their master.

- Blackboard negotiation exploits the classic blackboard architecture [9] to provide a co-ordinating base. In this scheme the blackboard's knowledge sources are replaced by agents who post to and read from the general blackboard. The scheduling agent (or master agent) schedules the agents' reads/writes to/from the blackboard. This scheme is employed by Werkman in his DFI system [10]. This approach may be used when the problem is distributed, a central scheduling agent is present or when tasks have already been assigned, *a priori*, to agents. Sharp Multi-Agent Kernel (SMAK) also adopts a blackboard strategy [11].

The latter point highlights the fact that organisational work ought not to be solely associated with hierarchies. For example, in the DVMT system [12] which also exploits a blackboard, co-ordination occurs among peer agents.

Other organisational structures exist, of course, including the centralised and decentralised market structures. The centralised market structure employs a master/slave co-ordination approach while a contracting technique, described in the next section, is more suitable for a decentralised market structure.

#### Critique

These strategies are useful where there are master/slave relationships in the MAS being modelled. Much control is exerted over the slaves' actions, and hence the problem-solving process. However, such control, in its extremes, mitigate against all the benefits of DAI — speed (due to parallelism) reliability, concurrency, robustness, graceful degradation, minimal bottlenecks, etc. In the blackboard co-ordination scheme, with no direct agent-to-agent communication, a severe bottle-neck may result if there are many agents, even in the case of multi-partitioned blackboards. Furthermore, all agents would need to have a common domain understanding (i.e. semantics). For this latter reason, most blackboard systems tend to have homogeneous and rather small-grained agents as is the case in the DVMT prototype [12]. Durfee et al [13] point out that such centralised control as in the master/slave technique is contrary to the basic assumptions of DAI. It presumes that at least one agent has a global view of the entire agency — in many domains, this is an unrealistic assumption. Even when a master/slave co-ordination technique is used, the designer should ensure that the slaves are of sufficient granularity to compensate for the overheads which result from goal distribution. Distributing trivial or small tasks can be more expensive than performing them in one location [8].

#### 3.2 Contracting

A, now classic, co-ordination technique for task and resource allocation among agents and determining organisational structure is the contract net protocol [14, 15]. In this approach, which assumes a decentralised market structure, agents can assume two roles:

- a manager who breaks a problem into sub-problems and searches for contractors to do them, as well as to monitor the problem's overall solution,
- a contractor who does a sub-task — however, contractors may recursively become managers and further decompose the sub-task and sub-contract them to other agents.

Managers locate contractors via a process of bidding which proceeds as follows:

- a manager announces a task,
- contractors evaluate the task with respect to their abilities and commitments,
- contractors table bids to the manager,
- the manager evaluates received bids, chooses a contractor and awards the contract to it,
- the manager waits for the result of the contract.

This is a completely distributed scheme where a node can both simultaneously be manager and contractor. This approach has been used in many applications, e.g. Parunak [16]. It has also been further generalised in other projects like the multistage negotiation in Conry et al [17].

#### Critique

Huhns and Singh [7] note that the contract net is a high-level co-ordination strategy which also provides a way of distributing tasks, and a means for self-organising a group of agents. It is best used when:

- the application task has a well-defined hierarchical nature,
- the problem has a coarse-grained decomposition,
- there is minimal coupling among sub-tasks.

The advantages of the contract net include the following — dynamic task allocation via self-bidding which leads to better agreements, agents can be introduced and removed dynamically, it provides natural load-balancing (as busy agents need not bid), and it is a reliable mechanism for distributed control and failure recovery [7].

Its limitations include the fact that it does not presume agents with contradictory demands; hence, the approach neither detects nor resolves conflicts, which is one key reason why co-ordination is needed in the first place. The agents in the contract net are rather passive, benevolent and non-antagonistic which for many real-world problems is unrealistic. Conry et al's generalisation of the contract net approach [17] essentially introduces an iterative mechanism for getting agents with conflicting goals to arrive at a consensus. Lastly, the contract net approach is rather communication-intensive, the costs of which may outweigh some of its advantages in real-world applications.

### 3.3 Multi-agent planning

Another approach to co-ordination in agent-based systems is to engage the agents in multi-agent planning. In order to avoid inconsistent or conflicting actions and interactions, agents build a multi-agent plan that details all the future actions and interactions required to achieve their goals, and interleave execution with more planning and re-planning. There are two types of multi-agent planning:

- centralised multi-agent planning,
- distributed multi-agent planning.

In centralised multi-agent planning, there is usually a co-ordinating agent who, on receipt of all partial or local plans from individual agents, analyses them in order to identify potential inconsistencies and conflicting interactions (e.g. conflict between agents over limited resources). The co-ordinating agent then attempts to modify these partial plans and combines them into a multi-agent plan where conflicting interactions are eliminated. Georgeff [18, 19] exemplifies such an approach where conflicting interactions are identified and grouped into critical regions. In the final multi-agent plan, communication commands are inserted to synchronise agents' interactions appropriately. Cammarata et al [20] also employ centralised multi-agent planning in a simulated air-traffic control domain. In this demonstrator, agents (i.e. aeroplanes) in a potential conflict scenario, e.g. two aeroplanes are heading for collision, decide on one of their number to act as co-ordinating agent to whom the plans of the other agent are sent. This co-ordinating agent will then attempt to modify its own flight plan in order to resolve the conflict. Jin and Koyama [21] propose their MATPEN model for co-ordinating autonomous and distributed agents based on centralised planning. Their approach exploits an 'expectation-based negotiation protocol' which draws from the roles of agents in the organisational structure. When two agents share a conflict, they form a conflict group and initiate a negotiation process. The exchange 'expectations' (which represent the expected behaviour of other agents in order to resolve the conflict) in order to decide who should play what role in the negotiation process. Eventually, they generate a multi-agent plan to resolve the conflict. They have constructed a collision avoidance system for ships using MATPEN.

In distributed multi-agent planning, the idea is to provide each agent with a model of other agents' plans [22]. Agents communicate in order to build and update their individual plans and their models of others' until all conflicts are removed. An exemplification of this approach is Lesser and Corkill's functionally accurate and co-operative (FA/C) protocol [23]. In this approach, loosely-coupled agents form high-level (but possibly incomplete) plans, results and hypotheses which they exchange with each other. Next, they refine these until they all converge on

some global complete plan. Local inconsistencies are detected and only the part of each agents' results that are consistent with local information are integrated into the local databases. This approach has been used in Lesser and Corkill's Distributed Vehicle Monitoring Testbed (DVMT) — a system for testing co-ordination strategies [12]. Another exemplification of distributed multi-agent planning is Durfee's partial global planning (PGP) approach [24]. In this technique, agents execute their local plans with each other, which, in turn, are modified continuously based on partial global plans (built by exchanging local plans). Therefore, agents are always looking out for potential improvements to group co-ordination.

### Critique

Generally, multi-agent planning of whatever form requires that agents share and process substantial amounts of information; hence, it is likely to require more computing and communication resources than other approaches. The centralised multi-agent planning technique shares many of the limitations of the master/slave co-ordination technique. Naturally, the co-ordination in distributed multi-agent planning is much more complex than in the centralised form as there may not be any agent who possesses a global view of the distributed system. Furthermore, since co-ordination in some existing multi-agent planning techniques, such as PGP, is a gradual process, the scopes of their applicability may be better for some domains than others [7].

## 4. Negotiation

A significant part of the co-ordination work which has been done or is being done world-wide goes under the heading 'negotiation'. This is because most co-ordination schemes involve some sort of negotiation. For this reason, it has merited a section on its own in this paper. The literature abounds because negotiation is a key co-ordination technique used to address several DAI issues such as those listed in section 2. However, there are probably as many similar definitions of negotiation as there are negotiation researchers. Perhaps, a more succinct and basic one is that of Bussman and Muller [25]:

“...negotiation is the communication process of a group of agents in order to reach a mutually accepted agreement on some matter.”

Sycara further points out that to negotiate effectively, agents must reason about beliefs, desires and intentions of other agents [4], and this has led to the development of techniques for the following:

- representing and maintaining belief models,
- reasoning about the other agents' beliefs,
- influencing other agents' intentions and beliefs.

These latter topics lead to the usage of all sorts of artificial intelligence (AI) and mathematical techniques including logic, case-based reasoning (CBR), belief revisions, distributed truth maintenance, multi-agent planning, model-based reasoning (MBR), optimisation and game theory. Furthermore, there is literature which proposes and reports on multi-agent test-beds, languages, protocols and interlinguas. In addition, there are also the purely social, cognitive and sociologically based papers. Thus, it is not difficult to explain the huge and varied literature on this subject of negotiation.

It is also conceded that the distinction made between negotiation and the other co-ordination approaches discussed in section 3 is quite fuzzy. Indeed, another paper, which overviews negotiation strategies [5], covers many of the aforementioned co-ordination techniques. This overview is not meant to be so exhaustive as to provide a road map through the more important literature on negotiation. Negotiation techniques are classified in the following three broad categories:

- game theory-based negotiation,
- plan-based negotiation,
- human-inspired and miscellaneous AI-based negotiation approaches.

### 4.1 Game theory-based negotiation

There is now a growing body of work on negotiation which is based on game theory [26]; perhaps the origins of such work can be traced back to Rosenschein's doctoral thesis [27]. With the help of one of his former students, Rosenschein has since refined, synthesised and collated his earlier work in one volume [28]. Using the tools of game theory, this book outlines an approach which shows how to achieve co-ordination among a set of rational and autonomous agents without an explicit co-ordination mechanism built into these agents *a priori*. In other words, it does not presume the 'benevolent agent assumption'.

The key concepts in this game theory approach to negotiation are the following — utility functions, a space of deals, strategies and negotiation protocols. Utility is defined as the difference between the worth of achieving a goal and the price paid in achieving it. A deal is an action an agent can take which has an attached utility. The negotiation protocol defines the rules which govern the negotiation, including how and when it ends (e.g. by agreement or no deal). In the book [28], several protocols and strategies for negotiation are outlined.

The actual negotiation proceeds as follows. Utility values for each outcome of some interaction for each agent are built into a pay-off matrix, which is common knowledge

to both (typically) parties involved in the negotiation. The negotiation process involves an interactive process of offers and counter-offers in which each agent chooses a deal which maximises its expected utility value. There is an implicit assumption that each agent in the negotiation is an expected utility maximiser. At each step in the negotiation, an agent evaluates the other's offer in terms of its own negotiation strategy.

Zlotkin and Rosenschein [29] have also extended work to cover agents that are not truthful, i.e. they can be deceptive; interestingly, using some simple demonstrators, they show that if an agent withholds certain information or deliberately misinforms other agents, this may result in better negotiation deals for the agent. In this work they view negotiation as a two-stage process — the actual negotiating and the execution of the joint plans. In their recent work, they are working on a general theory of automated negotiation in which they classify complex domains into three — task-oriented domains, state-oriented domains and worth-oriented domains [28].

Kraus and Wilkenfield also examine negotiation using game-theory techniques with appropriate modification [30]. In this paper, they propose a strategic model that claims to take time into consideration during the negotiation process. Time influences the outcome of the negotiations and avoids delays in reaching agreements.

#### Critique

Game theory-based negotiation fails to address some crucial issues [31]. Firstly, agents are presumed to be fully rational and acting as utility maximisers using predefined strategies. Secondly, all agents have knowledge of the pay-off matrix, and therefore full knowledge of the other agent's preferences — this is certainly unlike the real world where agents only have partial or incomplete knowledge of their own domains, let alone those of others. Therefore, this is unrealistic for truly non-benevolent and loosely-coupled agencies. Furthermore, the pay-off matrix could become very large and intractable for a negotiation involving many agents and outcomes. Thirdly, agents only consider the current state when deciding on their deal — past interactions and future implications are simply ignored. Fourthly, agents are considered to have identical internal models and capabilities. Fifthly, much of the work presumes two agents negotiating, though some later work is addressing n-agent negotiation [32]. Sixthly, and consequently, despite Zlotkin and Rosenschein's provision of mathematical proof of their ideas [32], it is unlikely that game theory-based negotiation will suffice for real-life, industrial agent-based applications for the reasons already offered. In brief, its assumptions are untenable in real applications.

#### 4.2 *Plan-based negotiation*

Alder et al [33] investigate negotiated agreements and discuss methods of conflict detection and resolution in the domain of telephone network traffic control. They imbue their agents with planning knowledge as they strongly maintain that negotiation and planning are very tightly intertwined due to the fact that agents need information from others to function effectively and efficiently.

Kreifelt and von Martial propose a negotiation strategy for autonomous agents [34]; they view negotiation as a two-stage process — firstly, agents plan their activities separately, and then, secondly, they co-ordinate their plans. The co-ordination of all the agents' plans is done by a separate co-ordination agent, though they note that this role may be played by any of the agents. They then proceed to present a negotiation protocol in terms of agents' states, message types and conversation rules between agents. This proposal also has many limitations. As Bussman and Muller [25] rightly point out, it does not really present a negotiation model but just prescribes one, and it is really left to the agents how they really achieve consensus. The protocol itself also needs some further clarification. In general, plan-based negotiation suffers from the limitations of centralised or distributed multi-agent planning (see section 3.3), depending on which of these two types is used.

#### 4.3 *Human-inspired and miscellaneous AI-based negotiation approaches*

It appears that almost every form of human interaction requires some degree of explicit or implicit negotiation [10]. Hence, it is not very surprising that many negotiation researchers draw from human negotiation strategies. As noted earlier, these often lead to the usage of miscellaneous AI techniques including logic, CBR, constraint-directed search, etc, as is evidenced by the following. However, some of the examples below just provide AI-based negotiation approaches which do not necessarily borrow from human interactions.

Sycara's work [4], which sees negotiation as an iterative activity, proposes a general negotiation model that attempts to handle multi-agent, multiple-issue, single and repeated type negotiations. Her view of negotiation leads her to exploit CBR and multi-attribute utility theory. She argues for a case-based approach based on her belief that human negotiators draw from the past negotiation experiences to guide present and future ones. In the absence of past cases, she resorts to preference analysis based on multi-attribute theory. In the latter case, issues involved in negotiation are represented by utility curves. By combining these curves in additive and multiplicative fashions, a proposal is chosen which maximises the utility. Sycara has constructed a

system called PERSUADER that resolves adversarial conflicts in labour relations with the aid of two practising negotiators. Agents can modify others' beliefs, behaviours and intentions via persuasion. Sycara strongly believes that the latter is a necessity in order to create co-operative interactions.

Sathi and Fox [35] argue that negotiation may be viewed as a constraint-directed search of a problem space using negotiation operators. They see negotiation as composed of two phases — a communications phase, where all information is communicated to participating agents, and a bargaining phase, where deals are made between individuals or in a group. In their approach, agents negotiate via relaxing various conflicts and constraints until agreement is reached. Alternatively, the solutions may be modified. Initially, preferences of negotiation are modelled as constraints. The negotiation operators are drawn from human negotiation studies [36] and include operators which simulate relaxation, reconfiguration and composition, which are used to generate new constraints. This approach has been used to build a system for resource allocation which 'performs marginally better than experts'.

The main limitation of this iterative approach stems from the fact that selecting relaxations to achieve a compromise is a major problem as no criteria are provided, and hence agents easily get caught in an infinite loop of exchanging offers (i.e. livelock).

Werkman proposes a knowledge-based model of an incremental form of negotiation [10]. Werkman's Designer Fabricator Interpreter (DFI) model, it is claimed, is based largely on various human models of negotiation. This scheme uses a shared-knowledge representation, called shareable agent perspectives, which:

"...allows agents to perform negotiation in a manner similar to co-operating (or competing) experts who share a common background of domain knowledge."

Essentially, it exploits a blackboard with partitions for requested proposals, rejected proposals, accepted proposals, a communications partition and shared knowledge. Such rich detail and knowledge of the perspectives of other agents provide invaluable information for agents to make better proposals in the future.

Werkman [10] sees negotiation as a three-phase cycle. The first phase involves some proposing agent announcing a proposal which is received and evaluated by the receiving agent. The second phase involves generating a counter proposal if the latter is not happy with the initial proposal, or it may be simply accepted. Phase three involves the submission of the counter proposal for review by other agents. An arbitrator agent assists when two agents get into deadlock by reviewing the negotiation dialogue and uses

their mutual information network to generate alternative proposals. This is done using issue relaxation techniques or some intelligent proposal generator. This may fail, in which case the arbitrator may set time limits or use other techniques.

This proposal is interesting and definitely worth studying in greater detail. The use of arbitration is relatively novel. His agents communicate via the blackboard through a speech-act based language. The centralised blackboard could be a bottle-neck, and, without an explicit scheduler, reading and posting to the blackboard seems chaotic. However, this proposal does seem to have a good understanding of the negotiating process. It has also been implemented in the DFI system.

Conry et al [37] specifically concerned themselves with negotiation strategies for distributed constraint satisfaction problems where a network of agents have a goal, but each node or agent has only limited resources. The local constraints give rise to a complex set of global and inter-dependent constraints. This investigation was done in the context of the long-haul transmission and more complex communication networks which they were researching. Their implementation involved developing algorithms for multi-agent planning while taking the inevitable conflicts into consideration. This piece of work is interesting because it begins the investigation of whether specific generic tasks may be linked with specific negotiation strategies. Researchers on negotiation will need to identify such links if the literature is to be made less *ad hoc*.

Bussmann and Muller's negotiating framework for co-operating agents [25] is, arguably, one of the most useful papers on negotiation in the literature. Drawing from socio-psychological theories of negotiation, particularly Gulliver's eight phases of the negotiation process [38], they evolve a cyclic negotiation model which is both general and simple. They attempt to address many of the limitations of other negotiation proposals/models. The cyclic nature of the model addresses the thorny issue of conflict resolution.

The general strategy is that negotiation begins with one, some or every agent making a proposal. Then agents evaluate and check the proposals against their preferences, and criticise them by listing their preferences which are violated by the proposals. The agents then update their knowledge about other agents' preferences and the negotiation cycle resumes with a new proposal or proposals in the light of the newly gleaned information. Conflicts between agents are handled in a concurrent conflict resolution cycle.

Bussmann and Muller's suggestion is untested but it does seem to draw from and address the limitations of many others. In this respect, it is an interesting proposal.

Finally, other researchers have proposed negotiation protocols including Kuwabara and Lesser's extended protocol for multistage negotiation [39] and Durfee and Montgomery's hierarchical protocol for co-ordinating multi-agent behaviour [40]. Naturally, protocols of this nature implicitly implement some negotiation strategy (e.g. Kuwabara and Lesser's implements the multistage negotiation strategy of Lesser's group at the University of Massachusetts), and, if it is suitable for some particular task, then it may as well be reused.

### 5. Lessons learned and some heuristics on co-ordination/negotiation

This section synthesises some of the key lessons culled from this review of the co-ordination literature, and presents some conclusions which have been drawn from the literature on co-ordination.

- One-off co-ordination strategies (or combination of strategies) are devised and used in one-off projects. Hence, solid conclusions as to their scope, applicability, usability, etc, have not been established.
- There is little empirical or theoretical support for any strategy or strategies. Not enough studies have been done to validate many of the proposals.
- It is not very clear when, where, how and why various negotiation and co-ordination strategies or combinations of them are used in various applications or proposals.
- The contract net and the master/slave models and variations of them appear to be the most used strategies due to their simplicity. Various AI, mathematical and operational research techniques are being used to realise co-ordination. This paper has highlighted some, e.g. including viewing negotiation as constraint-directed search [35]. Most of them seem to be either problem-dependent or were just used because the researchers had experience using those techniques; hence, their choice seem rather *ad hoc*. However, the problem-dependent issues cannot be ignored as they are crucial to the success of individual applications.
- Various techniques make fundamental assumptions about agents which must be understood. For example, the contract net approach presumes truly benevolent, trustful, non-conflicting and helpful agents.
- Most co-ordination or negotiation strategies do not involve any complex meta-reasoning required of most domains. For example, few take into consideration the time aspects, and the fact that agents' goals, beliefs, intentions, etc, change with time. Hence, in real life, something an agent is currently doing may no longer

matter because, say, the deadline is past, in which case it is best that it stops or forgets about that particular task/sub-task. Alternatively, the resources may not be available at the times required by the agents. These introduce non-trivial dimensions to co-ordination.

- There is a lack of fundamental analysis of the process of co-ordination and negotiation. Jennings [6] is a step in this direction.

### 6. Conclusions

This paper has overviewed the diverse literature on co-ordination in agent-based systems. Successful co-ordination is a key design objective for most multi-agent system builders. Without good co-ordination mechanisms, many of the benefits of the multi-agent paradigm simply disappear. For this reason, researchers have invested significant resources in devising, implementing, and evaluating a large array of co-ordination techniques. As discussed in this paper, the various approaches have their relative advantages and disadvantages. At this time, there is no universally best method. In general, the theoretical methods produce good results in well-constrained environments, but many of their underpinning assumptions are not suited to developing real-world systems. Extant implementation-oriented work, on the other hand, operates well in limited domains but suffers from a lack of grounding and rigorous evaluation.

Based on our experiences and observations, it is possible to identify four major components which must be present in any comprehensive co-ordination technique [6]. Firstly, there must be structures which enable the agents to interact in predictable ways. Secondly, there must be flexibility so that agents can operate in dynamic environments and can cope with their inherently partial and imprecise viewpoint of the community. Thirdly, there must be social structures which describe how agents should behave towards one another when engaged in the co-ordination process. Finally, the agents must have sufficient knowledge and reasoning capabilities to exploit both the available structure (individual and social) and the flexibility.

For the future, further theoretical work is needed to develop adequate, and appropriately broad, models of co-ordination which address all of the four major components identified above. This should then lead to better specified implementations of co-ordination algorithms which have clearly delimited ranges of applicability. This approach will also enable more systematic evaluation of the different algorithms to be undertaken so that system builders can base their design decisions on empirical evidence (rather than on their *ad hoc* preconceptions).

## Acknowledgements

This paper has drawn extensively from some of the work done by the first author as part of a short-term fellowship at BT Laboratories in the summer of 1994 while he was an academic. The help, in particular, of Mark Wiegand and Robin Smith, who both initiated the fellowship, is acknowledged.

## References

- 1 Galbraith J R: 'Organization design', Addison-Wesley (1977).
- 2 Thompson G, Frances J, Levacic R and Mitchell J (Eds): 'Market, Hierarchies and Networks — The Coordination of Social Life', Sage Publications (1991).
- 3 Malone T W: 'Modeling Coordination in Organizations and Markets', *Management Science*, 33, No 10, pp 1317—1332 (1987).
- 4 Sycara K: 'Multi-agent Compromise via Negotiation', in Gasser L and Huhns M (Eds): 'Distributed Artificial Intelligence 2', Morgan Kaufmann (1989).
- 5 Nwana H S: 'Negotiation Strategies: An Overview', BT Laboratories internal report (1994).
- 6 Jennings N R: 'Coordination Techniques for Distributed Artificial Intelligence', in O'Hare G M P and Jennings N R (Eds): 'Foundations of Distributed Artificial Intelligence', London, Wiley pp 187—210 (1990).
- 7 Huhns M and Singh M P: 'CKBS-94 Tutorial: Distributed Artificial Intelligence for Information Systems', Dake Centre, University of Keele (1994).
- 8 Durfee E H, Lesser V R and Corkill D D: 'Coherent Cooperation among Communicating Problem Solvers', *IEEE Trans Comput*, 11, pp 1275—1291 (1987).
- 9 Hayes-Roth B: 'A Blackboard Architecture for Control', *Artificial Intelligence*, No 25, pp 251—321 (1985).
- 10 Werkman K J: 'Knowledge-based model of negotiation using shareable perspectives', Proc of the 10th Int Workshop on DAI, Texas (1990).
- 11 Kearney P, Sehmi A and Smith R: 'Emergent behaviour in a multi-agent economics simulation', in Cohn A G (Ed): 'Proceedings of the 11th European Conference on Artificial Intelligence', London, John Wiley (1994).
- 12 Lesser V and Corkill D: 'The Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem-Solving Networks', *AI Magazine*, 4, No 3, pp 15—33 (1983).
- 13 Durfee E H, Lesser V R and Corkill D D: 'Trends in Cooperative Distributed Problem Solving', *IEEE Knowledge and Data engineering*, 1, No 1, pp 63—83 (1989).
- 14 Davis R and Smith R G: 'Negotiation as a metaphor for distributed problem solving', *Artificial Intelligence*, No 20, pp 63—109 (1983).
- 15 Smith R G: 'The Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver', *IEEE Trans on Comput*, 29, No 12 (December 1980).
- 16 Parunak H V-D: 'Manufacturing Experience with the Contract Net', in Gasser L and Huhns M (Eds): 'Distributed Artificial Intelligence 2', Morgan Kaufmann (1989).
- 17 Conry S E, Meyer R A and Lesser V R: 'Multistage Negotiation in Distributed Planning', COINS Technical Report, University of Massachusetts, Amherst, Boston (1986).
- 18 Georgeff M: 'Communication and Interaction in Multi-Agent Planning Proc 1983 National Conf Artificial Intell, pp 125—129 (August 1983).
- 19 Georgeff M: 'A Theory of Action for Multi-Agent Planning', Proc 1984 National Conf Artificial Intell, pp 121—125 (August 1984).
- 20 Cammarata S, McArthur D and Steeb R: 'Strategies of Cooperation in Distributed Problem Solving', Proc of the Eighth Int Joint Conf on Artificial Intell, pp 767—770 (1983).
- 21 Jin Y and Koyama T: 'Multi-agent planning through expectation-based negotiation', Proc of the 10th Int Workshop on DAI, Texas (1990).
- 22 Corkill D D: 'Hierarchical Planning in a Distributed Environment', Proc 6th Int Joint Conf Artificial Intell, pp 168—179 (August 1979).
- 23 Lesser V and Corkill D: 'Functionally accurate, cooperative distributed systems, *IEEE Trans on Systems, Man and Cybernetics*, 11, No 1, pp 81—96 (1981).
- 24 Durfee E H and Lesser V R: 'Using Partial Global Plans to Co-ordinate Distributed Problem Solvers', Proc of the 1987 Int Joint Conf on Artificial Intell, pp 875—883 (1987).
- 25 Bussmann S and Muller J: 'A Negotiation Framework for Co-operating Agents', in Deen S M (Ed): 'Proc CKBS-SIG', Dake Centre, University of Keele, pp 1—17 (1992).
- 26 Luce R D and Raiffa H: 'Games and Decisions', John Wiley & Sons (1957).
- 27 Rosenschein J S: 'Rational Interaction: Cooperation Among Intelligent Agents', PhD Thesis, Stanford University (1985).
- 28 Rosenschein J S and Zlotkin G: 'Rules of Encounter: Designing Conventions for Automated Negotiation among Computers', MIT Press (1994).
- 29 Zlotkin G and Rosenschein J S: 'Blocks, Lies, and Postal Freight: The Nature of Deception in Behaviour', Proc of the 10th Int Workshop on DAI (1990).
- 30 Kraus S and Wilkenfield J: 'The Function of Time in Cooperative Negotiations: Preliminary Report', Department of Computer Science, University of Maryland (1991).
- 31 Busuioc M and Winter C: 'Negotiation and Intelligent Agents', Project NOMADS-001, BT Laboratories internal report (1995).
- 32 Zlotkin G and Rosenschein J S: 'One, Two, Many: Coalitions in Multi-Agent Systems', Proc of the 5th European Workshop on Modelling Autonomous Agents in a Multi-Agent World (August 1993).
- 33 Alder M R, Davis A B, Weihmayer R and Forest R W: 'Conflict-resolution strategies for non-hierarchical distributed agents', in Gasser L and Huhns M N (Eds): 'Distributed Artificial Intelligence 2', Morgan Kaufmann (1989).
- 34 Kreifelt T and von Martial F: 'A negotiation framework for autonomous agents', in Demazeau Y and Muller J P (Eds): 'Decentralized AI2', Elsevier Science (1991).
- 35 Sathi A and Fox M: 'Constraint-directed negotiation of resource allocations', in Gasser L and Huhns M (Eds): 'Distributed Artificial Intelligence 2', Morgan Kaufmann (1989).
- 36 Pruitt D G: 'Negotiation Behaviour', Academic Press (1981).
- 37 Conry S E, Meyer R A and Lesser V R: 'Multistage Negotiation in Distributed Planning', in Bond A H and Gasser L (Eds): 'Readings in Distributed Artificial Intelligence', San Mateo, Morgan Kaufmann (1988).
- 38 Gulliver P H: 'Disputes and Negotiations — A cross-cultural perspective', Academic Press (1979).

## CO-ORDINATION IN SOFTWARE AGENT SYSTEMS

39 Kuwabara K and Lesser V R: 'Extended Protocol for Multistage Negotiation', Proc of the 9th Workshop on Distributed Artificial Intell (1989).

40 Durfee E H and Montgomery T A: 'A hierarchical protocol for coordinating multi-agent behaviour', Proc of the 8th National Conf on Artificial Intell, Boston, MA, pp 86—93 (1990).



Hyacinth S Nwana holds a BSc in Computer Science and Electronic Engineering (1985), an MSc in Computer Science (1986) and a PhD in the area of Artificial Intelligence (1989). Between 1989 and 1995, he worked in various roles (lecturer, researcher, contract and visiting scientist) for the Universities of Liverpool, Keele, Calgary (sponsored by the Royal Society), Shell Research Laboratories and BT Laboratories. He joined BT in October 1995 and leads part of BT's software agents research programme, as well as a team researching into the engineering of agent-based systems and is contactable on e-mail at:

hyacinth@info.bt.co.uk. He has a broad range of AI-related interests and has written on software agents, AI and education, distributed AI, machine learning, neural networks, knowledge acquisition, computer science education and cognitive science. In 1991, he principally won the Digital Equipment Corporation (DEC) European Artificial Intelligence prize. He is currently pursuing a part-time MEd degree at Cambridge. He is a member of the British Computer Society and a Chartered Engineer.



Lyndon Lee received an MSc in Computer Science from the University of Essex in 1989. He is currently finishing his PhD degree in distributed artificial intelligence at Essex. He joined BT's intelligent systems research group in 1996 and is currently working on software agents.

His research interests lie primarily in building formal models of multi-agent systems based on decision theory, game theory and evolutionary search algorithms.



Dr Nick Jennings is a reader in the Department of Electronic Engineering at the Queen Mary and Westfield College (University of London, e-mail: n.r.jennings@qmw.ac.uk). He is head of the Distributed AI Unit which conducts basic and applied research into a range of issues related to the development of multi-agent systems. He has been collaborating with BT's intelligent systems research group via several projects. His main research interest is in building well-grounded, real-world agent systems. He has been involved in the development of agent-based systems in the

following domains — process control, business process management, health care, and information management.

