Defence Research and Development Canada
Recherche et développement pour la défense Canada

DEFENCE DÉFENSE

Coordination in distributed intelligent systems

A. Boukhtouta
J. Berger
M. Allouche
DRDC Valcartier

A. Bedrouni
Concordia University

Defence R&D Canada – Valcartier

Technical Report
DRDC Valcartier TR 2006-780
October 2009

Canada

Coordination in distributed intelligent systems

A. Boukhtouta
J. Berger
M. Allouche
DRDC Valcartier

A. Bedrouni
Concordia University

Defence R&D Canada- Valcartier

Technical Report

DRDC Valcartier TR 2006-780

October 2009

Author

A. Boukhtouta, J. Berger, A. Allouche, A. Bedrouni

Approved by

Éloi Bossé

Decision Support Systems / Section Head

Approved for release by

Christian Carrier

Chief Scientist

Abstract

Intelligent and autonomous agents have emerged as a novel approach offering the ability to analyze, design, build, and implement complex systems. The agent-based conceptual view provides a large repertoire of tools, techniques, and metaphors that can be adopted to considerably improve the way many types of complex systems are conceptualized and implemented in various domains. Agent-oriented technologies are increasingly used in a variety of applications, ranging from comparatively small systems to large and complex, mission-critical applications such as air traffic control.

This report covers issues and comprehensively answers commonly asked questions about coordination in agent-oriented systems. It is dedicated to providing a state-of-the-art review of current coordination strategies, protocols and mechanisms. In exploring the progress achieved, this study has unveiled the lack of coherence and order that characterizes the area of research pertaining to coordination. Based on current practical deployed applications and future opportunities, the report identifies and thoroughly examines trends, challenges, and future agent-oriented research directions.

Résumé

Les approches utilisant les agents autonomes et intelligents sont une nouvelle tendance innovatrice qui offre la possibilité de faire le design, la mise en œuvre et l’analyse des systèmes complexes. Le concept de systèmes à base d’agents englobe une panoplie d’outils, de techniques et de métaphores dont l’emploi ne pourra qu’améliorer considérablement la conceptualisation et l’implémentation de systèmes complexes, et ceci dans différents domaines. Il est à signaler aussi que les technologies à base d’agents sont de plus en plus utilisées dans différentes applications variant des petits systèmes aux systèmes complexes. On n’a qu’à citer, à titre d’exemple, les missions critiques de contrôle du trafic aérien.

Ce rapport couvre des aspects et des éléments de réponses à des questions concernant la coordination dans les systèmes à base d’agents. Il présente l’état actuel des stratégies de coordination, des protocoles et des mécanismes qui lui sont associés. En explorant les travaux dans le domaine, cette étude a révélé le manque de vision globale et cohérente commune dans tous les travaux de recherche effectués dans le domaine de la coordination. Par ailleurs, en se basant sur les applications actuelles et les tendances futures, ce rapport identifie et examine les tendances, les défis et les futures directions de recherche dans le domaine de la coordination dans les systèmes à base d’agents.

This page intentionally left blank.

Executive summary

It is commonly acknowledged that the number of deployed systems is still surprisingly limited, despite intense and persistent efforts devoted to building practical commercial and industrial applications. In addition, currently deployed agent-oriented applications are typically characterized as closed systems incorporating pre-defined communication and interaction protocols. Thus, it is likely that future research in the area of agent-based systems will be directed towards bridging the gap between academic promises and reality.

Emerging application domains such as *Ambient Intelligence*, *Grid Computing*, *Electronic Business*, *Semantic Web*, *Bioinformatics*, and *Computational Biology* will certainly determine future agent-oriented research and technologies. Over the next decade, it is thus expected that real-world problems will impose the emergence of trulyopen, fullyscalable agent-oriented systems, spanning across different domains, and incorporating heterogeneous entities capable of reasoning, learning, and adopting adequate protocols of communication and interaction.

Among all emerging issues associated with building such sophisticated applications, none is more fundamental, challenging, and complex that the need to dynamically ensure adequate management of activities attributed to a large number of heterogeneous entities. Thus, coordination will remain a central issue in agent-oriented engineering research. In this context, the agent research community will continue to strive to establish appropriate coordination models and frameworks, devise novel strategies, and build adequate protocols and mechanisms.

The present report is dedicated to providing a state-of-the-art review of current coordination strategies, protocols and mechanisms. In exploring the progress achieved, this study has unveiled the lack of coherence and order that characterizes the area of research pertaining to coordination. Based on current practical deployed applications and future opportunities, the report identifies and thoroughly examines trends, challenges, and future agent-oriented research directions. Finally, it provides specific recommendations to the military staff.

A. Boukhtouta, B. Berger, M. Allouche, A. Bedrouni. 2009. Coordination in distributed intelligent systems. Defence Research and Development Canada-Valcartier.

This page intentionally left blank.

Sommaire

On sait bien que le nombre de systèmes déployés est limité en dépit des efforts consacrés au développement de systèmes qui traitent des applications commerciales et industrielles. En outre, les systèmes à base d'agents déployés actuellement sont décrits dans la littérature comme des systèmes fermés incorporant des protocoles d'interaction et de communications prédéfinis. Ainsi, les futurs travaux de recherche dans le domaine des systèmes à base d'agents représenteront vraisemblablement un pont reliant les promesses académiques à la réalité.

Les nouvelles technologies émergentes comme l'intelligence d'ambiance, le calcul en réseau, les affaires électroniques, les réseaux sémantiques, la bioinformatique et les calculs biologiques détermineront sûrement les recherches futures dans le domaine des agents. On s'attend à ce que les problèmes réels favorisent l'émergence de systèmes à base d'agents « ouverts », extensibles à différentes échelles, touchant plusieurs domaines et incorporant des entités hétérogènes capables de « raisonner », d'apprendre et de s'adapter aux protocoles d'interaction et de communication adéquats.

Parmi les enjeux associés au développement de systèmes à base d'agents, aucun n'est aussi fondamental, complexe et important que le besoin de gérer dynamiquement les activités attribuées à des entités hétérogènes. Ainsi, la coordination dans ce cas reste un élément central de la recherche liée à l'ingénierie des systèmes à base d'agents. Dans ce contexte, la communauté de chercheurs va continuer ses efforts pour développer des modèles de coordination appropriés et des protocoles et mécanismes de coordination adéquats.

Ce rapport présente l'état actuel des stratégies de coordination, des protocoles et des mécanismes qui lui sont associés. Le rapport couvre des aspects et des éléments de réponses à des questions concernant la coordination dans les systèmes à base d'agents. En explorant les travaux dans le domaine, cette étude a révélé le manque de vision globale et cohérente commune dans tous les travaux de recherche effectués dans le domaine de la coordination. Par ailleurs, en se basant sur les applications actuelles et les tendances futures, ce rapport identifie et examine les tendances, les défis et les futures directions de recherche dans le domaine de la coordination dans les systèmes à base d'agents. Le rapport formule des recommandations précises en matière de modèles de coordination aux militaires intéressés à la coordination dans les systèmes intelligents distribués.

A. Boukhtouta, B. Berger, M. Allouche, A. Bedrouni. 2009. Coordination in distributed intelligent systems. Defence Research and Development Canada-Valcartier.

This page intentionally left blank.

Table of contents

Abstract i

Executive summary iii

Sommaire v

Table of contents vii

List of figures xi

List of tables xiii

1. Introduction 1

2. Real world applications of the coordination in distributed intelligent systems 7

3. The emergence of the coordination theory in distributed artificial intelligence 11
   - 3.1 Coordination, a challenging issue 11
   - 3.2 Coordination, an ill-defined concept 12
   - 3.3 Coordination, objective or subjective? 15
     - 3.3.1 Objective coordination 16
     - 3.3.2 Subjective coordination 18

4. Main coordination frameworks 21
   - 4.1 Sikora & Shaw framework 21
   - 4.2 Mintzberg Model 22
   - 4.3 Coordination Reference Model 26
   - 4.4 Malone & Crowston theory of Coordination 27

5. Classification of coordination techniques 31
   - 5.1 Organizational structuring 31
   - 5.2 Contracting 32
     - 5.2.1 Different types of contracts 33
     - 5.2.2 Contract Net Protocol 35
   - 5.3 Multi-agent planning 37

5.4 Negotiation 39
5.4.1 Negotiation techniques 40
5.4.1.1 Game theory-based negotiation 40
5.4.1.2 Human-inspired AI-based negotiation 43
5.4.1.3 Plan-based negotiation 47
5.4.2 Negotiation models 48

6. Coordination through protocols 51
6.1 Auctions for resolving conflicts 51
6.1.1 Type of auctions 53
6.1.1.1 English auctions 53
6.1.1.2 Dutch auctions 54
6.1.1.3 First-price sealed-bid auctions 54
6.1.1.4 Vickery auction 55
6.1.2 Issues related to the type of auctions 57
6.2 Coalition formation in coordination 60
6.3 Argumentation in coordination 62
6.3.1 Logic-based argumentation 65
6.3.2 Dialogues and dialogue Systems 69

7. Coordination in distributed problem solving and planning 73
7.1 Properties of DPS systems 74
7.2 Applications, examples and motivations 76
7.3 Task sharing 78
7.3.1 Task sharing among homogeneous agents 79
7.3.2 Task sharing in heterogeneous systems 80
7.3.3 Task sharing for distributed sensor net establishment 82
7.3.4 Task sharing for interdependent tasks 83
7.4 Result sharing 84
7.4.1 Functionally accurate cooperation 84
7.4.2 Shared repositories and negotiation search 86
7.4.3 Distributed constrained heuristic search 87
7.4.4 Organizational structuring 87
7.5 Distributed planning 88

7.5.1 Centralized planning for distributed plans.. 88
7.5.2 Distributed planning for centralized plans.. 90
7.5.3 Distributed planning for distributed plans 91
7.6 Dependence theory for action coordination... 94
7.6.1 Social interaction models 94
7.6.1.1 Dependence relations between tasks.............................. 96
7.6.1.2 Dependence relations between agents............................ 96
7.6.2 Main characteristics of a social dependence..................................... 96
7.6.3 Social dependence and autonomy 97
7.6.4 Dependence relations in practice.. 98
7.6.4.1 Coalitions formation 98
7.6.4.2 Incoherencies solving.. 99
7.6.5 Dependence relations and social power.. 100
7.6.6 The individual and social characteristics of dependence relations.. 101

8. Scalability and metrics... 103
8.1 Coordination and scalability.. 103
6.2 Coordination performance metrics 110

9. Research directions and future trends.. 117
9.1 Future trends... 117
9.2 Technological challenges 121
9.2.1 Appropriate design methodologies... 122
9.2.2 Effective, Agreed-upon standards 123
9.2.3 Incorporate learning and reasoning capabilities 124

10. Conclusion and recommendations... 129

11. References 133

List of symbols/abbreviations/acronyms/initialisms 159

This page intentionally left blank.

List of figures

Figure 1. An Integrative view of DAI [184] ... 4

Figure 2. The basic communications paradigms [26] ... 17

Figure 3. The basic parts of the organization [249-250] ... 23

Figure 4. The five coordination mechanisms in organization [249-250] ... 24

Figure 5. A structure for Tolksdorf coordination reference model [14] ... 26

Figure 6. Three Basic Types of Dependencies among activities [22] ... 28

Figure 7. The Space of negotiation agreements [157] ... 63

Figure 8. The Tower of Hanoi [17] ... 77

Figure 9. Distributed delivery example [17] ... 78

Figure 10. Derivation of dependence relations ... 95

Figure 11. Agent platforms around the world registered in the AgentCities Network- February 2003 [267] ... 119

This page intentionally left blank.

List of tables

Table 1. The five structural configurations of organizations [249-250]. 25

Table 2. Examples of common dependencies between activities and alternative coordination processes for managing them [21-22]. 30

Table 3. : Walton and Krabbe's classification of dialogues [69,180,182]................................ 69

This page intentionally left blank.

1. Introduction

The research community working in the area of Distributed Artificial Intelligence (DAI) unanimously endorses the idea that coordination – a fundamental ***paradigm*** – represents a challenging ***research area***. A clear consensus has thus emerged around a forged, well-articulated, and strongly advocated common vision that coordination is a **central issue** to agent-based systems engineering research.

Commonly presented as a subfield of AI, DAI is characterized by a lack of consensus on a useful recognized definition required to achieve clarity and common understanding. It is described as a research area that "investigates the behaviour of societies of artificial agents" [12, 13]. It is also presented as a field concerned with "concurrency in AI computations, at many levels" [9]. According to Findler and Elder [15], DAI involves "collections of distributed computers (intelligent decision-making agents) sharing information and resources in order to solve a common set of tasks". On the other hand, it is defined as an area of research "concerned with how automated agents can interact and solve problems effectively" [16].

In an early attempt to clarify and decompose broader research objectives associated with DAI, Bond and Les Gasser divided this field into two primary areas [9]:

1. **Distributed Problem Solving (DPS).** – DPS is a subfield of DAI where the emphasis is on "getting agents to work together to solve problems that require collective effort" [17]. Due to the inherent distribution of resources such as knowledge, capability, information, and expertise among individual entities, an agent in a distributed problem-solving environment is thus unable to accomplish its own task alone. When working with others, it can at least accomplish its task better – more quickly, completely, precisely, or certainly. Using a DPS system to solve a particular problem amounts to dividing this problem among a number of nodes that divide and share knowledge about it and the developing solution [9,11].

   According to Ephrati [13], agents are, within the DPS paradigm, assumed to be created by the same designer or group of designers, and thus, work together to solve particular problems. Using economical terms, he notes that the agents have a common preference profile which yields an identical utility function. The common goal may indeed be decomposed into sub-goals that are allocated to different agents in the group. Coordination is thus needed to allow efficient distributed activity towards the achievement of the global goal.

2. **Multi-agent Systems (MAS).** – MAS is a subfield of DAI initially described as a research area concerned with "coordinating intelligent behaviour among a collection of (possibly pre-existing) autonomous agents" [9]. As a result, it focuses on enabling these agents to coordinate their knowledge, goals, skills, and activities and plan to take action and solve problems [9,11].

   In a multi-agent system, autonomous agents "may be working either towards a single global goal or common task or to achieve separate individual goals that interact" [9,12]. In the absence of a global control mechanism, these agents have thus to reason about the process of coordination [9,11,12].

Whereas research in the area of DPS is confined to the development of methods and approaches to allow task decomposition and solution synthesis, work on MAS is described as having a vocational emphasis on behaviour coordination. However, Durfee and Rosenschein have pointed out that the term "multi-agent system" that served to distinguish between the prevalent research activities in DAI and an emerging body of work has been generally applied to any system that is, or considered to be, composed of multiple interacting agents [10]. They observed that such distinctions have been lost except within a rather small segment of the community. The authors then argued that the overuse and abuse of the term has made it difficult for researchers who are relatively new in the field to be aware of the distinctions the terms MAS and DPS once meant. They finally added that many researchers who have been aware of these terms "might have very different views as to what the distinction really is between them".

In an attempt to tackle this issue, Durfee and Rosenschein revisited the terms DPS and MAS so as to bring clarity to what they might mean and encourage the community to consider useful decompositions of the broader research objectives of DAI. In this context, the authors exposed several views of how MAS research differs from DPS research, while attempting to provide some insight into important questions in the field and into different ways of solving problems and designing systems.

In recognizing that DPS and MAS are not mutually exclusive, and in fact built upon each other to some extent, Durfee and Rosenschein identified, formulated and presented three views of the relationship between these areas of research:

1. **DPS is a subset of MAS.**
2. **MAS provide a substrate for DPS.**
3. **MAS and DPS are complementary research agendas.**

Earlier, Ephrati has argued that the distinction between DPS and MAS should really be seen more as a distinction between research agendas, rather than between running systems. He noted that it will not be obvious to an outside observer whether a given distributed system falls into one paradigm or the other. While the DAI community still debates this issue [18], Weiss has recently made, according to contemporary usage, no explicit distinction between DPS and MAS [11].

Finally, in a contribution to the debate on DAI, Ossowski describes and relates the area of Autonomous Agents (AA) to the classical areas of DAI, namely MAS and DPS [184]. From this point of view, DAI is characterized as an aggregate of three main currents: AA, MAS, and DPS. Following this classification, an attempt is indeed made to characterize the components of DAI through ***endogenous*** properties. As a result, it is stated that:

- **AA** focuses on "*questions of autonomy of individual agents*",
- **MAS** is concerned with "*heterogeneous systems of non-benevolent agents with conflicting goals*",

- **DPS** amounts to building "*homogeneous systems*" characterized by "*the cooperation of benevolent agents pursuing common goals*".

Referring to Durfee and Rosenschein [10], Ossowski points out that using endogenous properties to characterize DAI systems would set a rather "fuzzy border" between the research areas stated above. To make this point of view more explicit, the author puts forward the example of the Persuader system [48] where union agents and other agents representing the management of a company can negotiate about labour contracts. In this respect, Ossowski explains that while agents have a goal of finding an agreement – *DPS form of solving a problem* – their interests (unions versus management) are however definitely opposed – *MAS involving conflicting goals*. According to the author, this example highlights the possibility that ***exogenous*** properties can be used in the same way as endogenous properties to classify areas relevant to DAI.

In going further into this important analysis, Ossowski suggests that AA, MAS, and DPS should be considered descriptions of the specifically defined operational objectives associated with DAI systems rather than conceived as labels for particular kinds of systems. Thus, a researcher in the area of AA is specifically concerned with providing formalisms and architectures for designing agents *endowed with autonomy* and thus capable of 'surviving' in a certain class of environments. A designer in the area of MAS is rather interested in the emergence of desired properties of agent interactions attributed to a group of agents with "*uncontrollably varying characteristics*". In the area of DPS, a designer involved in building an agent-based system to cope with a given task is finally interested in establishing a desired functionality among a group of agents displaying "*manipulable characteristics*".

On the other hand, Ossowski notes that a classification based on exogenous properties would make AA, MAS, and DPS appear as the different parts of a "*divide and conquer*" approach to DAI. In this respect, a designer of a DPS-based system is devoted to providing a means to solve a particular problem through a distributed environment incorporating methods that presuppose certain properties of agent interactions. As noted, the designer may, for example, decide to provide efficient and robust problem solving via the Contract Net mechanism, which indeed relies on the existence of benevolent agents. Since the condition of having benevolent agents cannot always be satisfied, the designer would thus, as Ossowski argues, adopt an alternative approach which consists in coercing self-interested decision-makers into emulating benevolent agents via–among other things–taxation mechanisms. In substance, the arguments put forward above provide the ability to state, as shown in Figure 1, that:

- AA research is ***agent-centered***.

- MAS research is ***interaction-centered***. This research can focus on providing internal properties – such as stability – in an environment where agents show individual properties that vary.

- DPS research is ***problem-centered***. This research can be based on the idea of building environments involving agents designed to demonstrate desired external properties – such as robustness and high-performance.

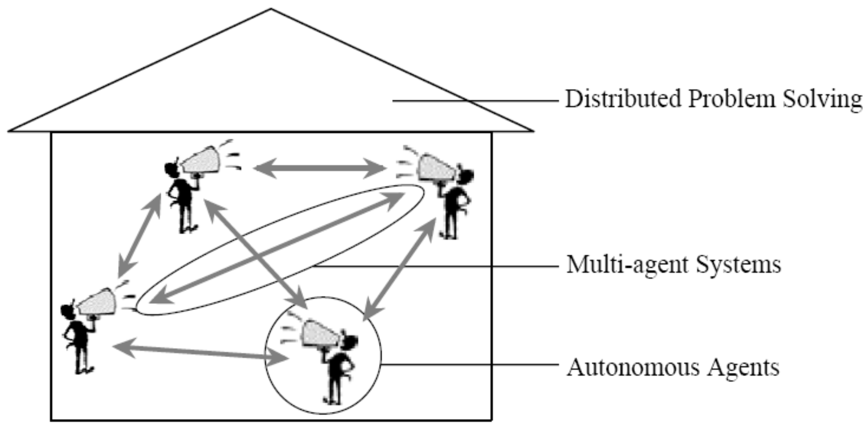

Figure 1. An Integrative view of DAI [184]

Building an agent-based system is more than just incorporating a set of individual agents into the boundaries of a common environment. The entire process of assembling and organizing autonomous entities into a smoothly functioning ensemble cannot be handled separately from the need to incorporate tools and mechanisms designed to achieve coherence and synergy between the individual members of the resulting community. Coordination is thus a fundamental **design dimension** of all aspects related to the collective behaviour in the space of agent interaction.

In addressing this issue, agent-oriented research has led to the development of a wide range of approaches: **frameworks**, **models**, **abstractions**, **strategies**, **techniques**, **protocols**, and **mechanisms**. However, depending on whether coordination addresses inter-agent dependencies or refers to intra-agent aspects, such approaches could well be categorized into two main classes: the *subjective* and *objective* coordination [25]. This classification approach stems from the idea that interaction can be observed from either the inside or outside interacting entities. In other words, agent interaction can be viewed and characterized as either subjective or objective depending on whether the observer is indeed a participating agent or simply an external entity not directly involved in the interaction. While subjective coordination focuses on the behaviour of agents as social entities, objective coordination puts the emphasis on the behaviour of an agent-based system as a whole.

This classification scheme offers the opportunity to impose conceptual order and introduce more coherence in a research area already characterized by rapid innovation, growth, and expansion in various directions simultaneously. Rather than being a clear expression of a deliberate preference over other existing taxonomies, the classification method is introduced as a strong plea in favour of a concerted approach to address and resolve the apparent confusion that prevails in the research area of multi-agent coordination.

Indeed, a common source of great confusion in this area lies in the ***lack of consistent terminology***. It highlights a basic problem of semantics in an area of research which requires some precision for uniform interpretation of fundamental concepts. The broad literature regarding coordination reveals the extent of the confusion associated with the use of inconsistent terminology to either define or describe common coordination concepts and principles. Researchers often use either the same term to define different concepts – *polysemy* – or different terms to refer to similar entities – *synonymy* . For instance, while Durfee refers to "task decomposition", "task allocation", and "results synthesis" as steps of a coordination strategy [17], Ossowski [184] identifies these concepts under the heading of coordination techniques. Without being specific to the research area of coordination, the lack of consistent use of terminology concerns newcomers to the field as well as well-established experts.

On the one hand, aspects regarding multi-agent coordination have been studied by a great number of individuals from many different fields, including psychology, biology, political sciences, social sciences, managerial and organization science, economics, and computer science. Consensus among researchers is thus more likely to loosen whenever the question is raised regarding what coordination is all about. Despite significant progress achieved in this area, the literature still highlights recurrent attempts to resolve the legitimate question of "what coordination is". In the area of computer science, the term coordination, still ***ill-defined***, is differently appreciated and used in subfields such as robotics, concurrent programming, and software engineering.

Interdisciplinary in its nature, research pertaining to multi-agent coordination transgresses traditional and established disciplinary boundaries to integrate a wide range of ideas, theories, concepts, methodologies, and approaches formulated in different areas, ranging from social sciences to DAI. It resulted in the development and deployment of a variety of applications, ranging from comparatively small systems to large and complex, mission-critical applications such as airtraffic control. However, it is commonly acknowledged that agent-oriented applications currently deployed are typically characterized as closed systems incorporating pre-defined and specifically designed communication and interaction protocols and mechanisms. Early interest in applications using benevolent agents working towards a common goal led to the emergence of ***cooperation*** strategies, protocols, and mechanisms to deal and guide agent interactions. On the other hand, managing interactions between self-interested entities has been centered on the notion of ***negotiation*** to resolve conflicts.

In examining the literature, research pertaining to multi-agent coordination appears much like the quest for the philosopher's stone: *to identify a universal approach offering unparalleled opportunities to resolve the coordination issue*. A fundamental question is thus to what extent the existing methodologies and approaches do address inter-agent dependencies and interactions to enable the realization of future real-world agent-oriented systems.

Indeed, emerging application domains such as *Ambient Intelligence*, *Grid Computing*, *Electronic Business*, *Semantic Web*, *Bioinformatics*, and *Computational Biology* will certainly determine future agent-oriented research in the area of inter-agent coordination. Over the next decade, it is thus expected that real-world problems will impose the emergence of trulyopen, fullyscalable agent-oriented systems, spanning across different domains, and incorporating heterogeneous entities capable of reasoning, learning, and adopting adequate protocols of communication and interaction.

In this respect, the primary objectives assigned to the present study include the need to:

1. Trace and examine the on-going debate over what constitutes the area of DAI.

2. Conduct a comprehensive review of the extensive literature regarding multi-agent coordination.

3. Identify open issues and highlight current trends and future research directions.

Built on various existing surveys, this review is shaped according to the way coordination is dealt with in the literature so as to remain accurate and better highlight the real mess characterizing this area of research. As a result, the remainder of this report is organized into ten Chapters directly tied to the objectives enumerated above. Chapter 2 describes a set of problems drawn from real-world domains and using the distributed artificial intelligence paradigm. Chapter 3 discuss the emergence of the coordination theory. On the other hand, Chapter 4 describes various models of coordination and outlines the theory of coordination. Furthermore, Chapter 5 is specifically devoted to the classification of the coordination techniques. Chapter 6 provides a more complete coverage of coordination strategies, protocols, and mechanisms. Chapter 7 discusses the coordination from a planning perspective. However, Chapter 8 addresses a key issue regarding open, large-scale agent-oriented systems. In this respect, it explores the notion of scalability in agent-based systems and discusses various approaches and techniques designed to allow the development of scalable coordination mechanisms. Chapter 6 also presents a few proposed studies regarding coordination evaluation criteria and performance measurement metrics. Chapter 9 identifies a range of open and emerging issues and identifies future research directions pertaining to multi-agent coordination. Finally, some concluding remarks and recommendations are given in Chapter 10.

2. Real world applications of the coordination in distributed intelligent systems

It is commonly recognized that many different disciplines can contribute to a better understanding of the coordination function as a way to build and provide appropriate tools and adequate approaches for designing complex organizations and systems. Agents incorporated into these self-regulating entities represent “communities of concurrent processes which, in their interactions, strategies, and competition for resources, behave like whole ecologies”. We describe in this chapter applications requiring coordination or distributed decision-making. A common characteristic of these applications is the highly distributed nature of them. These applications can be also seen as challenging problems for multi-agent coordination (See Panait and Luke [273]).

Military applications: The ability to coordinate military operations with the aid of information technology has become an imperative for command and control systems for the past several years. One of the most important military problems (in vogue) where coordination mechanisms should be used is in the control of swarms of UAVs (unmanned aerial vehicles). The UAVs are considered in this case as highly mobile agents that can be used for performing reconnaissance in a combat environment. In this case the communication between the UAVs is too costly because it would reveal their location to the enemy. The UAVs are sent out to collect information from certain locations and then return. The information sought by a UAV may depend on other information that it has collected by another UAV. There are many other cooperative control problems of the UAVS, we enumerate: cooperative target assignment, coordinated UAVs intercept, path planning, feasible trajectory generation, .etc (See [276]). A recent study of performance prediction of an unmanned airborne vehicle multi-agent system has been developed by Lian and Deshmukh [275]. The Markov Decision Processes (MDP) techniques and the Dynamic Programming approach have been widely used to solve the problems cited above (See Goldman and Zilberstein [277]).

Coordination techniques are also useful for the systems that address the Joint Fire Support (JFS) problem. The mandate of Joint fires is to assist air, land, maritime, amphibious, and special operations forces to move, manoeuvre, and control territory, populations, and airspace. The JFS and coalition operations require shared approaches and technologies. The challenge is large, however, with a legion of command-and-control computer systems all developed along slightly different lines of approach. Some of these systems operate jointly within their spheres of influence. The bigger challenge is tying these and other efforts together in a shared global information network, linking ground, air and sea forces.

Power management: Potential applications in this domain include control of power grids or any other distribution of a resource such as water stored in the reservoirs, gas, etc. Usually the electricity distribution management problem consists in maintaining an optimal power grid configuration that satisfies the suppliers (customers) while minimizing losses and dealing with possible damage of the network. We cite among the properties of the problem: variable and stochastic demand from customers, scheduled maintenance operations, equipment failures, and stochastic inflows (in the case of hydroelectric production). See Schneider et al. [274] for a reinforcement learning approach to managing a power grid.

Transportation and Network management (routing): Agents are deployed in this problem domain to cooperatively control and manage a distributed network. Agents can also handle failures and balance the flows in the network. Various approaches are developed in [278] to route packets in static and ad-hoc networks.

Distributed artificial intelligence (DAI) offers also suitable tools to deal with the hard transportation problems [279]. Fischer describes in [279] the modeling autonomous cooperating shipping companies system (Mars), which models cooperative order scheduling within a society of shipping companies. Three important instances for DAI techniques that proved useful in the transportation application, i.e., cooperation among the agents, task decomposition and task allocation, and decentralized planning are presented in this paper. Indeed, the problem tackled in this paper is a complex resource allocation problem. The auction mechanism is used for schedule optimization and for implementing dynamic replanning.

Coordination mechanisms are also used for distributed vehicle monitoring problems or traffic management. The objective in such problems is to maximize the throughput of cars through a grid of intersections or a network. Each intersection is equipped with a traffic light controlled by an agent. The agents associated to different intersections need to coordinate to deal with the traffic flows [280, 281].

Another application concerns the Air fleet control or airspace deconfliction. A multi-agent approach for air fleet control is reported in [282]. In this application, the airspace used by the traffic controllers is divided into three-dimensional regions to guide the airplanes to their final destination. Each region can hold a limited number of airplanes. The multi-agent solution is used to guide the planes from one region to another along minimal-length routes while handling real-time data.

Manufacturing and supply chains: Modern and big manufacturing depends heavily on computer systems. In fact, in many manufacturing applications the centralized software is not as effective as distributed networks. Indeed, to provide competitiveness in global markets, manufacturers must be able to implement, resize, design, reconfigure, respond to unanticipated changes, and maintain manufacturing facilities rapidly and inexpensively. These requirements are more easily satisfied by distributed small modules than by large monolithic systems. Moreover, small modules allow system survivability. Multi-agent systems offer a way to build production systems that are decentralized rather than centralized, emergent rather than planned, and concurrent rather than sequential.

From another side, the emergence of internet-based computing and communication to execute business processes has emerged as key in the supply chain domain. By using the internet, businesses gain more visibility across their network of trading partners, and it helps them to respond quickly to customer demands, resource shortages, etc. Min and Bjornsson present a method based on computer agents technology for building virtual construction supply chains [283].

Ambient Intelligence: The early developments in Ambient Intelligence took place at Philips in 1998. The Philips vision of Ambient Intelligence is : “people living easily in digital environments in which the electronics are sensitive to people's needs, personalized to their requirements, anticipatory of their behavior and responsive to their presence”.

The main purpose of Ambient Intelligence applications is to coordinate the services offered by small smart devices spread in a physical environment in order to have a global coherent intelligent behaviour of this environment. From a computing perspective, ambient intelligence refers to electronic environment that is sensitive and responsive to the presence of people. Ambient intelligence paradigm builds upon several computing areas and human-centric computer interaction design. The first area is **ubiquitous** or **pervasive** computing. Pervasive computing devices are very tiny (can be invisible) devices, either mobile or embedded in almost any type of object imaginable, including cars, tools, clothing and various consumer goods. All these devices communicate through increasingly interconnected networks (ad hoc networking capabilities). The second key area is **intelligent systems research.** Learning algorithms and pattern matchers, speech recognition and language translators, and gesture classification belong to this area of research. A third area is **context awareness**; research on this area lets us track and position objects of all types and represent objects’ interactions with their environments. See [294] and [295] for more details about Ambient Intelligence applications.

This page intentionally left blank.

3. The emergence of the coordination theory in distributed artificial intelligence

Creative and innovative ideas are unmasked through a "synergy of knowledge" resulting from the interaction among research areas that were once considered distinct disciplines. Indeed, Cognitive Science has recently emerged at the interface of several disciplines, in particular neurobiology, psychology, AI and epistemology. Chaos Theory has found applications in such diverse disciplines as mathematics, physics, economics, and psychology. In the same way, DAI is an exciting and fruitful domain of scientific research that "brings together and draws on results, concepts and ideas from different disciplines, computer science (including AI), sociology, economics, organization and management science, and philosophy" [11].

Confusion, discomfort or disquiet often accompanies the use of multiple methods, concepts, and approaches using theoretical frameworks from different disciplines. Much interdisciplinary and cross-disciplinary research has not yet produced a synthesis which would go beyond disciplinary boundaries, and hence produce innovative solutions to deal with and solve real-world problems. The development and implementation of feedback loops which can link researchers and experts of different disciplines would promote understanding and eliminate confusion. It is thus believed that the current practice of using multiple methods, techniques and technologies in the area of DAI should be grounded in a comprehensive and pragmatic framework, if results are to be integrated and interpreted in a meaningful way.

Indeed, the DAI area of research appears to be only concerned with the transfer of tools and methods across disciplines in response to specific problems. Based on an instrumental view of knowledge, DAI thus lacks a disciplinary approach that would introduce order and control required to ensure that knowledge produced is legitimate and reliable. In this respect, there have been relatively few or no attempts to reconstitute the past, shape the future, clarify the objectives, carefully set up rules, and establish reasonably identifiable boundaries of this emergent research and application field. We will discuss in the sequel of this chapter the challenging issue related to the coordination theory and we will explain the ill-defined concept related to the coordination.

3.1 Coordination, a challenging issue

Among all presently identified issues inherent in DAI, none is more fundamental, challenging, and complex than the need to adequately manage inter-agent activities. Thus, coordination has emerged as a central issue in agent-oriented engineering research. Hence, DAI focuses primarily on designing coordination protocols and mechanisms to manage activities of artificial, intelligent, and autonomous entities.

Indeed, there is a growing interest in questions regarding how activities of autonomous entities in complex systems are or can be adequately coordinated. Scientific research focuses on coordination in parallel and distributed systems, in human systems, in organizations, and in complex systems involving human-computer interactions. On the other hand, research regarding collective activities in insect societies has long discovered that insects, relatively simple natural agents, form a society with a very high level of coordination. Coordination is thus the subject of interdisciplinary research in a variety of different disciplines including computer science, sociology, economics, organizational theory, management science, biology, etc.

Present in every activity in human society, coordination remains an elusive concept. Almost everybody has an intuitive sense of what the term coordination means. While good coordination is often invisible, a lack of coordination is obviously highly noticeable. In attempting to investigate and characterize this area of research, it becomes necessary to have a clear and more precise definition of the term coordination. However, due to its multi-disciplinary nature, the coordination concept has generated a diversity of definitions.

Over the last three decades, the particular area of DAI has witnessed intense research activities regarding problems associated with inter-agent processes. In this respect, efforts directed towards developing practical agent-oriented application systems have brought a considerable progress, ranging from the establishment of coordination models to the design and development of coordination protocols and mechanisms.

Indeed, different application contexts exhibit different needs with respect to coordination, and the choice of a coordination model or mechanism can have a great impact on the design of multi-component applications. Research on coordination languages and models provided the ability to evaluate and compare the impact of different classes of coordination models and languages. On the other hand, research addressing coordination strategies in agent-based systems led to different techniques, protocols and mechanisms, ranging from the simple but effective Contract Net protocol to more sophisticated market mechanisms. More recently, formal dialogue games have found application as a foundation for interaction protocols between autonomous agents. Thus, dialogue game protocols have been proposed for agent team formation, negotiation, consumer purchase interactions, and joint deliberation.

Finally, research activities are currently focusing on achieving a smooth transition from existing closed agent-based applications to interconnected automated and scattered systems in large and open environments – AgentCities, GRID computing, Semantic Web, etc[1]. Building and deploying such ubiquitous, open, and scalable applications will certainly pose a major challenge in terms of providing adequate and specifically adapted coordination models and protocols.

3.2 Coordination, an ill-defined concept

Presented as an inherently interdisciplinary concept, coordination thus becomes diverse enough to accommodate a broad spectrum of definitions. As Coates et al. point out [240], this concept still remains ill-defined despite the significant amount of research devoted to developing a wide range of approaches, models, strategies, and mechanisms. Differences

[1] Agent-Cities, GRID computing and Semantic Web applications also consider some properties specific to distributed systems: the *heterogeneity* of the system entities and the *openess* of the system.

in the definitions of coordination can indeed represent a contributing factor to the reigning confusion in the area of coordination. At first sight, these differences appear to stem from the perception and conception that scientists, experts, and researchers from different disciplines have on the functional role ascribed to coordination.

A rough examination of the literature provides the ability to highlight at firsthand the absence of a single operational definition of the term coordination which is acceptable to all. In seeking to identify the differences that experts use to define coordination, it may be at least possible to help promote the idea that a consensus or a common understanding would enhance communication between researchers from different disciplines and thus pave the way for a larger cooperation in the area of coordination. Indeed, the literature provides a number of alternative definitions of the concept of coordination. The purpose is to address the question regarding the meaning of this term, further attempt to outline the differences, examine whether these differences in meanings are complementary or contradictory, and finally seek to identify a common thread that could be used by researchers and experts from various disciplines as a basis for communication and research.

First recorded in 1605, the term coordination generally means "***orderly combination***"[2]. Without changing the basic traditional meaning that prevails over centuries, contemporary English brought fresh insights into the meaning and usage of the term coordination. Thus, literally, the term coordination is generally increasingly associated with harmonious functioning of different parts. In the American Heritage Dictionary of the English Language[3], coordination is defined as "***harmonious adjustment or interaction***". The Hyperdictionary[4] describes coordination as "***the regulation of diverse elements into an integrated and harmonious operation***". Furthermore, the Oxford Advanced Learner's Dictionary proposes an entry for the word coordination which is defined as "***the act of making parts of something, groups of people, etc. work together in an efficient and organized way***". Finally, the Web-based AllWords.com[5] offers a number of other words that can be used as synonyms of the term coordinate: organize, harmonize, integrate, synchronize, correlate, match, relate, arrange, codify, grade, and graduate.

Indeed, in "Social Co-ordination among Autonomous Problem-solving Agents" [241], Ossowski et al. draw attention to the differences in how the research issue pertaining to coordination is addressed in different disciplines. While referring to research addressing coordination in Social Sciences as primarily *analytic*, the authors identify coordination-related interest in the area of Distributed Artificial Intelligence as *constructive*. In this respect, Ossowski et al. argue that a scientist in a Social Sciences discipline "observes the outside world and builds a model of how human agents mutually adapt their activities as part of societies or organizations". On the other hand, research in Distributed Artificial Intelligence is primarily concerned with the design and implementation of appropriate and adequate coordination mechanisms in accordance with known requirements and

---

[2] The Barnhart Concise Dictionary of Etymology, 1994.
[3] The American Heritage Dictionary of the English Language, 4th Edition, 2000.
[4] http://www.hyperdictionary.com/dictionary/
[5] http://adams.allwords.com/

anticipated operating conditions of groups of artificial agents. More precisely, a central designer in the sub-area of Distributed Problem-solving constructs interaction patterns among benevolent agents, so as to enable these entities to efficiently achieve a common goal. In addition, a researcher in Multi-agent Systems is concerned with how to install desired global properties into "heterogeneous groups of autonomous agents that pursue partially conflicting goals in an autonomous fashion". On the other hand, Ossowski also points out that research in the area of economics is focused on the structure and dynamics of the market as a particular coordination mechanism [242]. As noted, in organizational theory, the emphasis is, however, placed on "predicting future behaviour and performance of an organization, assuming that a certain coordination mechanism is valid".

In "Coordination Languages and their Significance" [243], Gelernter and Carriero advocate making a clear distinction between computation and coordination to provide the ability to build a complete programming model. Such separation provides the advantage of both facilitating the reuse of the components of a program and allowing the identification of coordination patterns that could be applied in similar situations [244]. Based on these considerations, coordination is specifically defined as "***the process of building programs by gluing together active pieces***" [245]. Within the framework of this research, the term ***coordination*** is thus used to refer to the process of ***gluing together*** independent active pieces to build programs whose pieces – i.e. a process, a task, etc. – can effectively communicate and synchronize with each other.

According to Van de Ven et al. [246], the organizational theory is based on the premise that all organizations need coordination. As noted, coordination of organizations is achieved either through a programming approach or a feedback approach. In this context, three predominant modes are thus identified to be frequently used to coordinate work activities within an organization. While the programming approach is said to be exercised through an *impersonal mode*, the feedback approach, also known as mutual adjustments, is described as a method of coordination that occurs either through *personal channels* or *group meetings*. Whatever approach or mode is used, however, Van de Ven et al. define the coordination of work activities within an organization as a process of "***integrating or linking together different parts of an organization to accomplish a collective set of tasks***".

In a paper entitled "A note on Hierarchy and Coordination: An Aspect of decentralization" [247], Kochen and Deutsch focus on the costs and benefits associated with coordination and attempt to define in operational terms the concept of coordination. In this context, the authors ask the question of what is logically intended by the term coordination, while recalling that in popular usage it means "***harmonious functioning of several actors***". Then, Kochen and Deutsch put forward the 'managerial-based view' that coordination is "a task involving planning and execution of how two or more jobs should be synchronized in time and space". From this point of view, the coordination concept is thus defined as "***a means of directing the operation of functional units so that their joint behaviour attains a specific goal with a higher probability and/or at a lower cost***". In this respect, the authors draw attention to the fact that in case the functional units are human, a common expectation of reward is usually also required. In tackling the question of the operational meaning of "coordination", Kochen and Deutsch observe that in case several activities are required for obtaining a joint reward, coordination then "consists in bringing about a deployment in time and space of the combined activities such as to increase the expected reward".

While recalling research efforts pertaining to coordination in diverse disciplines, Nwana et al. note that systems – even biological systems – appear to be coordinated through "individual cells or agents acting independently and in a seemingly non-purposeful fashion" [29]. In this context, the authors further argue that human brains exhibit coordinated behaviour from apparently 'random' behaviours of very simple neurons. Apparently based on these arguments, Nwana et al. state that coordination is essentially "***a process in which agents engage in order to ensure a community of individual agents acts in a coherent manner***".

3.3 Coordination, objective or subjective?

A critical challenge to providing the foundation of an organized DAI discipline is to capture, classify and evaluate knowledge regarding the concept of coordination so as to facilitate the transfer of information and foster the development of new concepts and strategies. Indeed, coordination in DAI is fundamentally an interdisciplinary topic with research efforts scattered across disparate communities. Coordination concepts, techniques and technologies developed in one area are not easily transmitted to researchers and practitioners in other disciplines.

Given the confusion that characterizes the area of DAI, it would seem impossible for experts and researchers to correlate, categorize, analyze and act on the unstructured information, concepts, results and data hidden in reports and published literature. Despite significant difficulties, many attempts gave rise to a series of empirical coordination-related studies and surveys in order to organize ideas, assess knowledge, evaluate progress, track down trends, and shape future research.

Based on the widely accepted framework introduced by Malone and Crowston, coordination in MAS is indeed viewed as the art of managing dependencies [25]. Within this framework, Schumacher recently proposed to distinguish between objective and subjective dependencies as a way "to define an original and comprehensive framework for modeling, designing, and implementing interactions in complex MAS".

Indeed, *objective dependencies* are referred to as inter-agent dependencies, or in other words, the configuration of a system in terms of the basic interaction means, agent generation/destruction and organization of the environment. On the other hand, *subjective dependencies* are described as intra-agent dependencies often involving "mentalistic" categories. Correspondingly, Schumacher classifies coordination into two categories:

1. Objective coordination to refer to the management of the objective dependencies that are external to the agents and essentially concerned with inter-agent aspects.

2. Subjective coordination to refer to the management of intra-agents dependencies towards other agents.

The author argues that a failure to differentiate between these two levels of coordination would complicate the design and subsequent implementation of MAS, since it leads to using intra-agent aspects for describing system configurations. He gives the example of a multi-agent system that is intended for modeling the hierarchy in an organization but would model this hierarchy within each agent through knowledge representation, instead of describing it by establishing communication flows that represent it. He finally notes that this confusion between subjective and objective aspects is typically present in MAS composed of mental agents that use agent communication languages (ACLs) in order to communicate.

According to Schumacher, subjective coordination is dependent on objective coordination, as the first is based on and supposes the existence of the second. Thus, adequate mechanisms specifically designed and implemented to ensure subjective coordination in MAS must have access to mechanisms associated with objective coordination. Otherwise, no subjective coordination can be possible.

3.3.1 Objective coordination

Objective coordination is, as already stated, mainly associated with the organization of "the world of a multi-agent system". This type of coordination is indeed achieved through:

1. ***A description of how the environment of a multi-agent system is organized***. – In this respect, M. Schumacher introduces the notions of implicit and explicit organizations. An organization given or imposed by the underlying logical structure on which a MAS evolves is identified as the implicit organization, since it does not explicitly model the environment. However, the explicit organization establishes a model of an environment that does not necessarily reflect the intended logical structure.

2. ***Proper handling of the agents' interactions***. – Handling agent interactions implies the need to describe the interactions between an agent and its environment, and the interactions between the agents evolving in that environment. Through its perception, an agent establishes a relation with its environment. It can also influence this environment through specific actions.

   In this respect, it is argued that the interaction with the environment can be understood by a given agent and thus used to communicate. Since all information is transmitted within the environment, communication thus becomes an action that influences the agent's milieu. To explain this view, the author cites the case of a robot spreading or diffusing information like an insect that spreads pheromone in its environment. Another robot can then sense this information and certainly notice the presence of the source in the same way an insect can detect the trace of another roaming insect.

However, interactions between agents are based on specific communication means. As shown in **Erreur ! Source du renvoi introuvable.**, Schumacher classifies these communication means into four basic paradigms [26]:

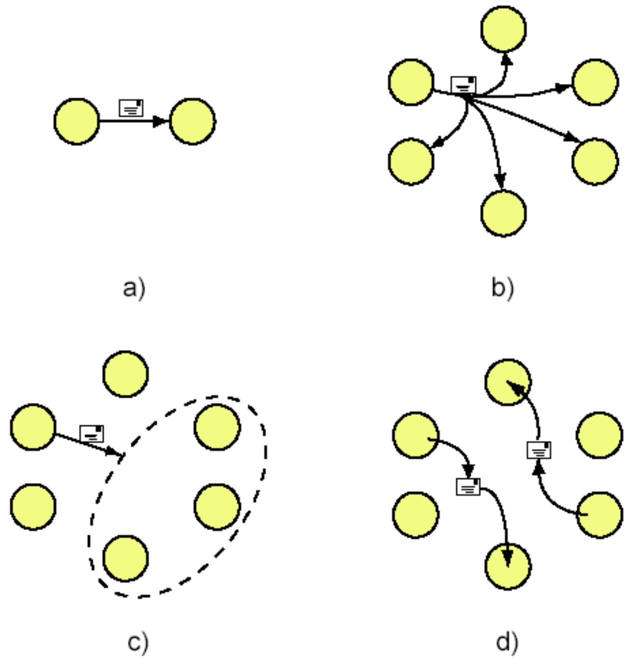

a) *Peer-to-peer communication*

b) *Broadcast communication*

c) *Group or multi-cast communication*

d) *Generative communication*

Figure 2. The basic communications paradigms [26]

a. ***Peer-to-peer communication***. – This form of communication is characterized by messages sent directly to specific agents. Agents usually identify their partners through, for instance, email-like addresses. It is however possible that an intermediate channel is used to ensure this form of communication and takes charge of the transmission of data.

b. ***Broadcast Communication***. – In this form of communication, a message is sent to all agents. In this context, agents are offered the opportunity to evaluate or ignore the transmitted and received data.

c. ***Group communication***. – A group communication is achieved when a message is sent to a specific group of agents.

d. ***Generative communication***. – Generative communication is a particular form of communication between agents realized through a blackboard, where agents generate persistent messages which can be read by other agents. Access to messages in the blackboard is achieved independently of the time a message is provided. This form of communication is thus uncoupled.

Schumacher further distinguishes between *identified* and *anonymous communication*. He identifies *identified communication* as a form of communication requiring an identity of a partner. *Anonymous communication* is, on the other hand, described as a form of communication where the agent producing the message ignores the recipient, and vice-versa.

3.3.2 Subjective coordination

Based on how the management of dependencies is handled, subjective coordination has been divided into two categories:

1 **Explicit subjective coordination**. Schumacher notes that the research in DAI has induced the development of several coordination techniques that are designed to tackle explicit subjective coordination. Quoting Jennings, the author clearly identifies these techniques as approaches that typically consider coordination as "*the process by which an agent reasons about its local action and the (anticipated) actions of others to try and ensure the community acts in a coherent manner*". According to him, these techniques are considered as subjective coordination because of a common tendency to resolve subjective dependencies by means of intra-agent structures. He notes that the latter often involves high-level mentalistic notions and appropriate protocols.

Characterized as explicit because of their explicit handling of coordination, the techniques identified above are further divided into three categories in accordance with Ossowski's classification [27]:

- *Multi-agent Planning techniques*. – Multi-agent planning techniques can indeed be centralized or decentralized. Using these techniques, agents establish and commit to behave according to a plan. The latter contains and describes all actions required for agents to appropriately achieve their respective goals.

- *Negotiation techniques*. – According to Schumacher, negotiation techniques represent the most significant part of DAI coordination-related research. Borrowed from S. Bussmann and J. Mueller [28], he proposes a definition where negotiation is described as "*the communication process of a group of agents in order to reach a mutually accepted agreement on some matter*". He finally notes that the negotiation process which starts from contradictory demands generates agreements that can be later re-negotiated.
- *Organization techniques*. – Schumacher observes that these techniques support a priori organization patterns by defining *roles* for each agent. He explains that a role determines the expectations about the agent's individual behaviour by describing its responsibilities, capabilities and authorities within a MAS. He adds that the agent then consults its organizational knowledge and the role devoted to him and acts accordingly.

  According to Schumacher, the purpose of organization is associated with objective coordination, while the notion of role refers to a subjective dependency towards other roles in a MAS. This argument is indeed used to label organization techniques as methods belonging to subjective coordination. The author finally notes that organization methods generally suffer from a lack of dynamicity in the structure of a given organization, since roles are considered as long-term relationships.

2 **Implicit subjective coordination**. – In collective robotics, different methods and approaches to using autonomous robot teams are implemented to efficiently fulfill predefined missions. The use of multiple mobile robots offers significant advantages over the use of a single robot: the possibility of distributed sensing, distributed action, task dependent "reconfigurability", and system reliability through redundancy. Multiple robots sharing a common environment organize their actions based on the result associated with the actions attributed to other robots.

Locally perceived through distributed sensors, this result provides the ability to resolve subjective dependencies, namely the necessity to sense specific information in order to act. This type of coordination, commonly known as "*stigmergic coordination*" allows indirect communication between autonomous entities through sensing and modification of the local environment. It literally means "*an incitement to work by the product of the work*". Using research on coordination in the area of collective robotics as an example, Schumacher outlines how agents may also coordinate themselves implicitly, without having explicit mechanisms of coordination.

4. Main coordination frameworks

Widely reported are attempts to develop and implement a more coherent, carefully crafted, and acutely relevant framework or model that serves as a sound basis for understanding existing principles and approaches and directing future research pertaining to the coordination function. Stemming from different disciplines, a large variety of models on how entities interact are indeed provided. These models can be used either to build middleware and languages or to express coordination strategies that lead to a coherent behaviour of interacting entities [14]. Models of coordination are thus required to provide adequate technologies and support the design of applications and interaction of computational units or entities.

When identified, commonalities amongst coordination models provide the ability to transfer coordination patterns across disciplines. In this respect, Tolksdorf outlines the need to build models that "have certain qualities such as being complete with respect to interaction forms and open to new patterns of interactions". He considers that models of coordination have to be "easy and safe to use to facilitate efficient software engineering", must be "scalable and efficient to implement" so as "to cope with the number of units to coordinate", and be "aware of the characteristics of future environments". On the other hand, Tolksdorf recognizes the lack of consensus on the relations between coordination, communication and cooperation and outlines the need to work towards a standardized terminology which contains definitions and clarifications of basic notions including the term "coordination".

Despite a growing amount of research in this area, Coates et al. have also similarly outlined the lack of a single widely accepted perception of coordination and its constituents [7]. Following a survey on research in coordination approaches and systems, these authors noted the relative absence of efforts designed to satisfy "the requirement to comprehensively understand the broader aspects of coordination". Within the same discipline, there is, according to Coates et al., "a clear difference of opinion as to what constitutes coordination". This confusion resulted in "a rather liberal use and application" of the word *coordination,* and thus produced a diversity in definitions and understanding of the concept.

We present below four of the most important (from our point of view) coordination frameworks used in distributed intelligent systems.

4.1 Sikora & Shaw framework

Sikora and Shaw presented a general multi-agent framework for understanding and analyzing the different issues involved in coordinating and integrating a collection of stand alone units, each of which can be viewed as an agent [8]. This framework is presented to form "the foundation for the development of a complete theory of coordination". Sikora and Shaw suggest that future work would look at extending the proposed formalisms so as to:

1. Provide formal proofs of the best coordination schemes associated with different scenarios, and

2. Develop formal methods in order to derive coordination mechanisms suitable for any given scenario based on the interdependencies among agents.

Applied to a manufacturing information system designed to manage a circuit boards production process, the proposed framework exhibited, according to the authors, several advantages, including the ability:

- to provide uniformity of a general framework by treating heterogeneous components in a homogeneous fashion,

- to represent distributed control by modeling agents as having local control with the capability to interact and coordinate their activities through message-passing.

4.2 Mintzberg Model

In an attempt to synthesize a large and varied literature research regarding organizational structuring, Mintzberg reveals what he calls “a curious tendency to appear in five’s” [249-250]. Interestingly, these findings do indeed suggest the possibility to logically isolate five basic parts of an organization, distinguish five basic mechanisms of coordination in the organization, and identify five fundamental types of decentralization. In exploring this possibility, Mintzberg proposes a theory on the structure of organization based on the postulate that any organization is composed of five parts.

Thus, according to Mintzberg, the organization can – as shown in Figure 3 – be described in terms of the following basic elements:

1. The operating core. – The operating core forms the basis of the organization where employees produce basic products and services, or directly support their production.

2. The strategic apex. – The strategic apex refers to the top of the organization where top-ranking managers and their personnel staff are responsible for making strategic decisions.

3. The middle line. – The middle line represents a chain of managers who “sit in a direct line of formal authority” standing between the staff of the strategic apex and the employees of the operating core. These managers are assigned the task of implementing decisions by “supervising subordinates and reporting to the supervisors”.

4. The technostructure. – Out of the formal ‘line’ structure, the technostructure includes all members of the staff (e.g., accountants, work schedulers, long-range planners) responsible for “applying analytic techniques” to the design and maintenance of the structure and to allow the organization to achieve adaptability to its environment.

5. The support staff. – The support staff includes people such as legal counsels, public relations, payroll, cafeteria employees, etc. This staff provides a number of indirect services to the rest of the organization.

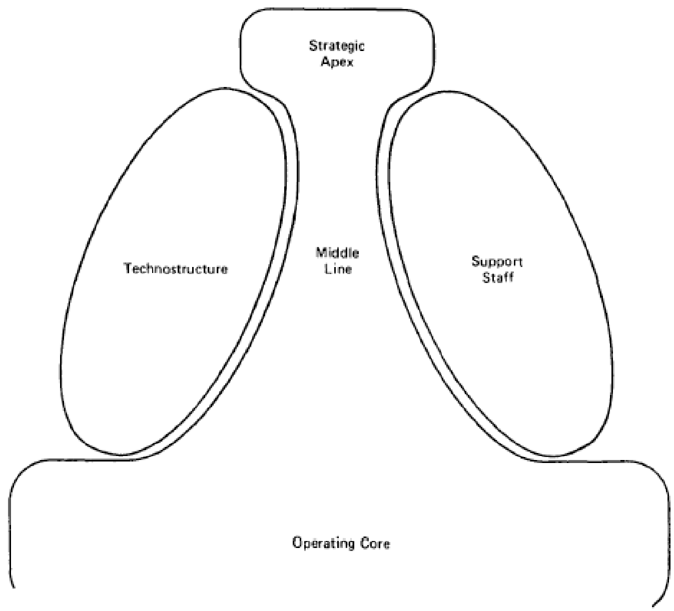

Figure 3. The basic parts of the organization [249-250]

As pointed out, organizational structuring focuses on job specialization and a clear division of labour of a given organizational mission into a set of distinct tasks. Then, it concentrates on appropriate coordination of the resulting tasks so as to allow the organization to accomplish its goals or mission in a unified way. In this respect, Mintzberg again identifies five basic ways or mechanisms through which the coordination of the organizational mission tasks can be affected. As illustrated in Figure 4, these mechanisms include:

1. *Direct supervision*. – Direct supervision refers to coordination of tasks conducted through a supervisor who gives specific orders or instructions to his or her subordinates.

2. *Standardization of work processes*. – Standardization of work processes indicates coordination of tasks that is achieved through standards – typically imposed by analysts of the technostructure in the form of work orders, rules and regulations, etc. – that govern work in the organization.

3. *Standardization of outputs*. – Standardization of outputs points to coordination of tasks conducted through standard performance measures and specifications regarding the output of the work – again this type of coordination mechanism is typically imposed by analysts of the technostructure.

4. *Standardization of skills*. – Standardization of skills refers to coordination of work achieved as individuals internalize standard skills and knowledge – usually prior to starting a work.

5. *Mutual adjustment*. – Mutual adjustment is a form of coordination of work conducted as participants communicate informally with each other.

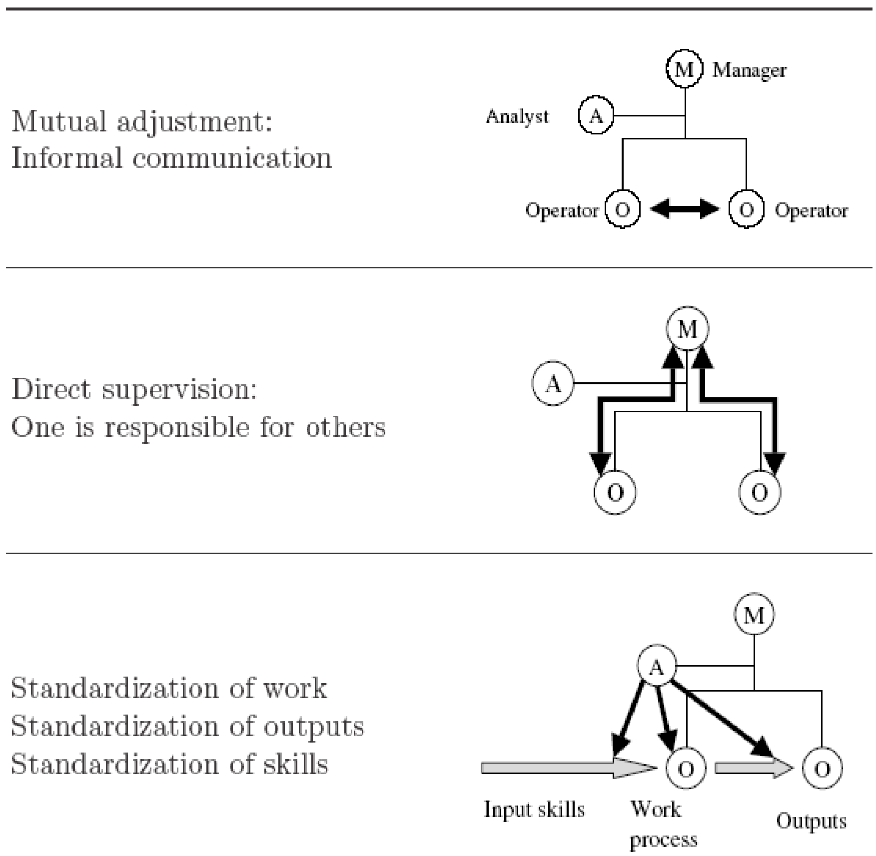

Figure 4. The five coordination mechanisms in organization [249-250]

Furthermore, Mintzberg identifies a list of nine mechanisms – design parameters – organizations can use to design their structures. Then, he presents these parameters as levers organizations can turn to effect the division of labour and coordination. In addition, the author points to research efforts in organizational structuring designed to assess the effect of various so-called *contingency factors* on the design parameters. According to Mintzberg, these research efforts are based on the *congruence* hypothesis, stating that effective structuring requires "a close fit between contingency factor and design parameter". While the congruence hypothesis relates the effectiveness of an organization to the fit between a given design parameter and a contingency factor, a second proposed hypothesis – called the *configuration* hypothesis – indicates that "effective structuring requires an internal consistency among the design parameters".

In substance, Mintzberg presents, as mentioned above, a typology in three parts:

1. a set of design factors susceptible of being used to characterize a structure of a given organization,

2. a number of contingency factors to provide the ability to characterize an organization's context,

3. five ideal types of organizational configurations described in terms of the above design parameters and contingency factors:

- Simple structure.
- Machine bureaucracy.
- Professional bureaucracy.
- Divisionalized form.
- Adhocracy.

According to this theory, each of the above configurations depends on fundamentally different mechanisms of coordination. Thus, Mintzberg selects five configurations of coordination mechanism and the associated pre-eminent part in the organization, as shown in Table 1.

Table 1. The five structural configurations of organizations [249-250].

Name	Coordination mechanism	Key part
Simple structure	Direct supervision	Strategic apex
Machine Bureaucracy	Standardization of work processes	Technostructure
Professional bureaucracy	Standardization of skills	Operating core
Divisionalized form	Standardization of outputs	Middle line
Adhocracy	Mutual adjustment	Support staff

Tolksdorf notes that the Mintzberg model assumes "a role for actors in an organization" [14]. He argues that the choice of coordination mechanisms is determined by the choice of the organizational structure. Hence, mechanisms of coordination are not easily exchangeable. On the other hand, Harold Doty reports the Mintzberg typology and the underlying theory have little support, since they received little or no systematic empirical examination in large-scale comparative studies [251].

4.3 Coordination Reference Model

Following a brief review regarding various coordination models, Tolksdorf proposes "a coordination reference model" [14] inspired by the 4-layer metamodel architecture for UML [252]. First, the author recalls the existence of situations associated with the real world where coordination is "neither present nor necessary". Then, he further specifies that the model he proposes takes into consideration only the part of the world in which "activities are managed in order to be coordinated".

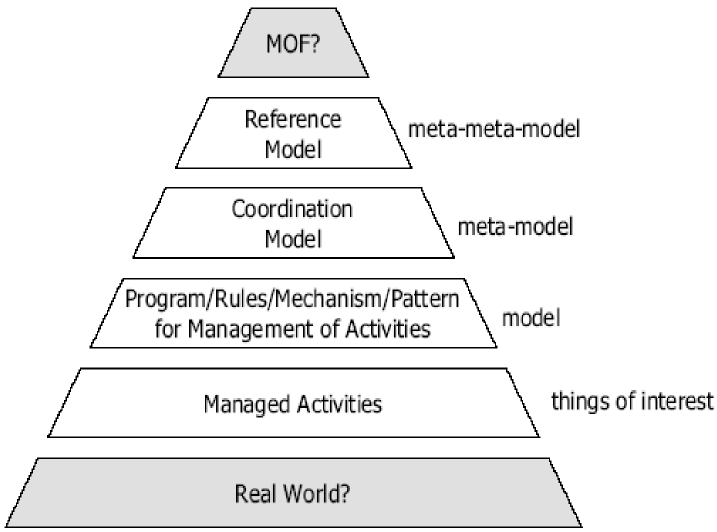

Figure 5. A structure for Tolksdorf coordination reference model[6] [14]

As shown in Figure 5, the real world coordinated activities form the object-level in the proposed model-hierarchy. Considering a specific set of objects, the concrete management of activities can, as Tolksdorf explains, be described through a set of rules, specific mechanisms, programs or a selection of coordination patterns. According to the author, these objects are indeed models of specific managed activities or, in other terms, the blueprints for actual interactions.

As proposed by Tolksdorf, the coordination reference model is a meta-model to other coordination models such as the Mintzberg model described above. As a result, it contains terminologies and other concepts required to describe these coordination models. In this respect, Tolksdorf considers a set of concepts that he believes are necessary to incorporate into the reference model:

- *Interactors*. Interactors represent those entities related to other interactors.

[6] Meta Object Facility (MOF) is a standard of Object Management Group (OMG).

- *Relations.* Relations are designed to associate in some way two or more interactors.
- *Non-interactors.* Non-interactors represent entities with no relations to interactors.
- *Operations.* Operations are performed by interactors on non-interactors.
- *Attributes.* Interactors, non-interactors or their corresponding current states can be described using attributes.
- *Meta-attributes*. Models built from the previous five concepts can be described through characteristics contained in the meta-attributes.

Finally, Tolksdorf claims that the reference model he proposed can be used to describe various coordination models, such as those he reviewed. As suggested, the reference model can thus be used to compare these coordination models through a set of meta-attributes: *distinction*, *orthogonal*, *coupling*, *autonomy*, *external*, *awareness*, *interactors stability*, relations *stability*, *reliability*, *scalability*, *programming*, *qualitative* or *quantitative measures*.

4.4 Malone & Crowston theory of Coordination

Malone and Crowston advocate the notion of a “coordination theory” to refer to principles and ideas about how coordination can occur in diverse kinds of complex systems [3]. In this context, they estimate that many of the researchers whose efforts can contribute to and benefit from this new area are not yet aware of each other’s work. The authors report having used the term “theory” with some hesitation since it connotes a degree of coherence that is not present in this field. They further described the field as a collection of intriguing analogies, scattered results, and partial frameworks. In proposing an interdisciplinary study of coordination, they thus hoped to both define a community of interest and suggest useful directions for future progress [19].

Malone and Crowston outlined the need for and introduced ***a coordination framework*** required to “transport” concepts and results back and forth between the different kinds of systems. In their view, coordination involves organized actors engaged in sequences of interdependent activities to achieve desired goals. As a result, actors face coordination problems arising from dependencies that “constrain how tasks can be performed” [20]. Coordination is viewed as a response to problems induced by dependencies. The duality between coordination and dependencies implies the need to consider that there is nothing to coordinate in the absence of interdependence. Accordingly, coordination processes are, within the theory of coordination, viewed as ways of “***managing dependencies between activities***”.

Defined in a way that emphasizes its interdisciplinary nature, coordination can thus occur, as outlined above, in many kinds of systems: human, computational, biological, and others. To facilitate the transfer of concepts and results among these different disciplines, Malone and Crowston proposed as a second step to identify and study the basic processes involved in coordination. Since coordination is defined as managing dependencies between activities, it should be possible, according to them, to characterize different kinds of dependencies and identify the coordination processes that can be used to manage them. In this respect, the authors identify several common dependencies between activities and provide an analysis of alternative coordination processes to manage them.

The coordination theory concept suggests that a process consists of three types of elements: resources, activities, and dependencies. A resource is produced and/or consumed during a process. While an activity is indeed a partitioned action that produces and/or consumes resources, a dependency is a relation among activities mediated by producing or consuming resources. Arising from resources related to multiple activities, three basic types of dependencies are identified [22], as shown in Figure 6**:**

1. **Flow dependencies.** This type of dependency occurs all the time in almost all processes. It arises whenever one activity produces a resource that is used by another activity.
2. **Sharing dependencies.** This kind of dependency arises whenever multiple activities all use (or could use) the same resource.
3. **Fit dependencies.** This type of dependency occurs when multiple activities collectively produce the same resource.

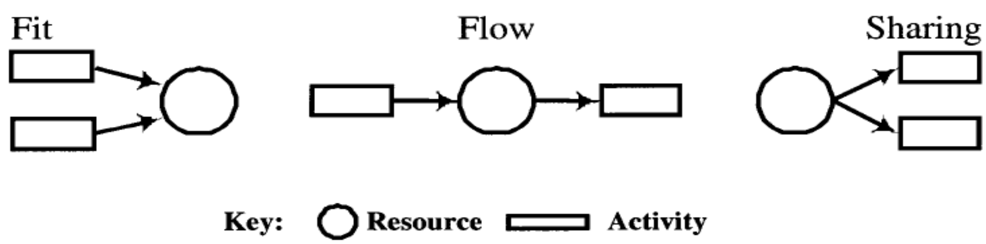

Figure 6. Three Basic Types of Dependencies among activities [22]

Later, Crowston developed taxonomy of dependency types by considering possible combinations of activities using resources [23]. This taxonomy includes task-resource dependencies and three types of task-task dependencies: shared resources, producer-consumer and common output. For each type of dependency, alternative coordination mechanisms are described.

To extend the above analysis within the theory of coordination, Table 2 is established to categorize an exhaustive list of dependencies and their associated coordination processes.

According to Malone and Crowston, many specific processes that arise in particular kinds of systems can be seen as instances of more generic processes [19]. In this respect, dependencies that are not shown can be usefully analyzed as specializations or combinations of those illustrated in Table 2. It is thus believed that it is possible to identify and systematically analyze a wide variety of dependencies and their associated coordination processes. Designing a “handbook” of coordination processes could, in the view of Malone and Crowston, facilitate interdisciplinary transfer of knowledge about coordination. It could also provide a guide for analyzing the coordination needs in particular situations and generating alternative ways of fulfilling them. The authors further noted that the power of analyzing processes in terms of dependencies and coordination mechanisms is greatly increased by access to a rich library of alternative coordination mechanisms for different kinds of dependencies.

Indeed, the Center for Coordination Science (CCS) at the Massachusetts Institute of Technology (MIT) conducted a project designed to develop a “Process Handbook” [22]. This project involved collecting examples of how different organizations perform similar processes. The handbook is a process knowledge repository intended to help users:

1. Redesign existing organizational processes,

2. Invent new organizational processes, and

3. Share ideas about organizational practices.

In this context, tools have been developed to provide read-only access and the ability to edit the Handbook repository contents [24]. The Handbook is based on powerful concepts such as dependencies and coordination in order to capture and organize process knowledge. According to Klein and Dellarocas [24], it is under active use and development by a highly distributed group of more than 40 scientists, teachers, students and sponsors for such diverse purposes as adding new process descriptions, teaching classes, and business process re-design.

Table 2. Examples of common dependencies between activities and alternative coordination processes for managing them [21-22].

Dependency	Examples of Coordination Mechanisms for Managing Dependency
Flow	
Prerequisite ("right time")	Make to order vs. make to inventory ("pull" vs. "push")
	Place orders using "economic order quantity", "Just In Time" (kanban system), or detailed advanced planning
Accessibility ("right place")	Ship by various transportation modes or make at point of use
Usability ("right thing")	Use standards or ask individual users (e.g., by having customer agree to purchase and/or by using participatory design)
Sharing	"First come/first serve", priority order, budgets, managerial decision, market-like bidding
Fit	Boeing's total simulation vs. Microsoft's daily build

5. Classification of coordination techniques

Rather than attempting to provide a comprehensive overview of the coordination literature, Nwana et al.[29] examine the crucial area of coordination in MAS [29]. In this context, they highlight the importance of coordination, not only in DAI and agent-based systems, but also in the design and implementation of open distributed systems. To effectively support this view, the authors gathered and put forward many arguments. In short, coordination is thus presented as a process in which agents engage coherently in order to:

- Prevent anarchy or chaos.
- Meet global constraints.
- Ensure distribution of expertise, resources or information.
- Manage dependencies between actions.
- Enhance efficiency.

Nwana et al. classify coordination techniques into four broad categories:

1. Organizational structuring.
2. Contracting.
3. Multi-agent planning.
4. Negotiation.

We discuss below these four coordination techniques.

5.1 Organizational structuring

Organizational structuring is described as the simplest coordination scenario, which exploits the a priori organizational structure. To support this claim, it is argued that the organization defines implicitly the agent’s responsibilities, capabilities, connectivity and control flow. In defining the roles, communication paths and authority relationship, this type of coordination provides, according to Nwana et al., a framework for activity and interaction. Citing Durfee et al. [30], the authors present the organizational structuring techniques as the pre-defined long-term relationships between agents.

Within this category, hierarchical organizations are identified as the dominant mode of coordination. This type of organization yields the classic master/slave or client/server coordination technique, which is used to enable a master agent to allocate tasks and resources among slave agents. It is explained that this technique is implemented in two ways:

1. In the first, the slaves within a master/slave configuration have limited autonomy with respect to the master, while the latter exercises full autonomy with respect to the slaves. In this context, the master is devoted the task of planning and allocating sub-plans to the slaves. Ultimately bound to report their results to the master, the slaves may or may not communicate with each other.

2. In the second, a blackboard negotiation scheme exploiting the classic blackboard architecture is used to provide a “coordinating base”. Agents replacing the blackboard’s knowledge sources are used to post to and read from the general blackboard. A master agent is, on the other hand, devoted the role of scheduling other agents’ reads/writes to/from the blackboard.

This approach based on the blackboard architecture has been, as reported by Nwana et al., implemented in various systems, such as Werkman's DFI system [31] and the Sharp Multi-Agent Kernel (SMAK) system [32].

Nwana et al. also mentioned other organizational structures, including both centralized and decentralized market structures. The centralized market structure is identified as employing a master/slave coordination approach. The decentralized market structure suitably includes contracting techniques such as those described below.

Referring to techniques belonging to the organizational structuring category, Nwana et al. note that these strategies are indeed useful in a MAS configuration based on master/slave relationships. It is argued that the approach implies much control over the slaves’ actions, and consequently over the problem-solving process. At its highest level, such control can however compromise the numerous advantages offered by DAI: speed, reliability, concurrency, robustness, graceful degradation, minimal bottleneck, etc. Without direct communication between agents, a blackboard-based scheme engaging a great number of agents may result in a severe bottleneck, even if it is a multi-partitioned blackboards-based configuration.

As a result, it is explained that most blackboard-based systems tend to incorporate homogeneous and rather small-grained agents as the Distributed Vehicle Monitoring Testbed (DVMT) prototype [33]. Citing Durfee et al., the authors point out that centralized control as required in the master/slave technique is not in accordance with the basic assumptions of DAI. They add that the centralized control presumes that at least one agent has a global view of the entire agency, an assumption that is unrealistic in many domains.

5.2 Contracting

Deployed in a multi-agent environment, an agent may be faced with certain tasks that cannot be accomplished locally or that can be performed more efficiently through appropriate assistance of other agents. In this respect, Kraus suggests that a self-interested agent may be able to take advantage of the presence of other agents and hire the services of appropriate entities so as to ensure that the assigned tasks are effectively accomplished.

Commonly known as contracting, the approach consists in allowing a self-interested agent to take the necessary steps – say through promises of rewards – in order to convince another self-interested agent to assist it in performing its tasks.

An attempt to implement this approach would thus raise the important question of how one agent can convince others to provide it with assistance in carrying out some tasks, while these agents are not only self-interested but do not share a global task. In addition, another issue worth exploring is how the manager-agent can convince the contractor-agent to perform a given task with a desired level of effort without the need for close observation, while the latter is susceptible of choosing different levels of effort when carrying out an assigned task.

5.2.1 Different types of contracts

According to Kraus, the issue of incentive contracting has indeed been the subject of numerous studies over the last three decades [131-136]. Conducted in both economics and game theory, these and other similar publications consider different contracts for different applications. Different types of contracts thus examined are between:

- A firm and an employer or employers [137-140].
- A government and taxpayers [141].
- A landlord and a tenant [131].
- An insurance company and a policy holder [142-145].
- A buyer and seller [146-147].
- A government and firms [148].
- Stockholders and managements [131].
- A professional and a client [149]. Etc.

As noted by Kraus, two parties usually exist in situations reflected through such types of contracts:

Commonly known as *the agent* in the economics literature, the first party is required to choose “an action or a level of effort from a number of possibilities.

Called *the principal*, the second party has the additional function of prescribing payoff rules.

Prior to choosing the action, a task devoted to the agent or first party, the principal determines a rule, i.e. a contract, designed to specify “the fee to be paid to the other party as a function of the principal's observations". Though similar, the applications mentioned above reveal differences in several aspects, such as the number of agents, the amount of information available to the parties involved, and the observations peculiar to the principal. As Kraus points out, several concepts and techniques are applied to the principal-agent paradigm in the relevant economics and game theory literature.

Indeed, the Contract Net protocol is a well-known framework for achieving automated contracting. A contract initiated through the Contract Net protocol represents an explicit agreement between, on the one hand, an agent that generates a task (the manager) and, on the other hand, an agent willing to execute the task (the contractor). While the manager has the responsibility to both monitor the execution of a given task and process the results, the contractor has the duty of actually executing the task.

The Contract Net protocol is indeed designed to provide a scheme that enables the manager to issue an announcement of the existence of a task. In response, available agents are thus required to evaluate the task announcements made by managers. These agents acting as potential contractors then submit bids or proposals on tasks which they are well suited to undertake. As the protocol was initially developed for and implemented in DPS environments where benevolent agents are organized to achieve common goals, there is thus no need to stimulate and motivate an agent so as to bid for tasks and deploy all its capabilities to better execute the contracted tasks.

Smith and Davis originally developed the Contract Net protocol and addressed problems related to task distribution, sub-tasks distribution, and synthesis of the overall solution [34-35,150]. On the other hand, the protocol has also been implemented into numerous applications relevant to various domains [36,151-153]. Initially, the Contract Net protocol was applied to a simulated distributed acoustic sensor network to provide a framework of a negotiation process that involves a mutual selection by both managers and contractors. In this application, the agents are totally cooperative, and the selection of a contractor is based on suitability, such as adjacency, processing capability, and current agent load. However, Sandholm highlights the absence of a formalized decision-making process regarding tasks announcement, bidding, and awarding [99]. Absent in the original Contract Net protocol, a formal model is thus proposed to allow agents to locally calculate their marginal costs for performing desired tasks. In this context, self-interested agents with different local criteria can interact to distribute tasks within a more effective system. The proposed pricing mechanism thus offers a new generalized version of the Contract Net protocol that can effectively handle both competitive and cooperative agents.

On the other hand, a backtracking method, known as *leveled commitment contract*, has been developed and proposed by Sandholm et al. to allow each party involved in a contract to unilaterally de-commit by paying a predetermined penalty [155-156]. This method has demonstrated the ability to improve expected social welfare even if agents de-commit strategically in Nash equilibrium.

In a paper entitled "an Overview of Incentive Contracting" [154], Kraus proposes techniques designed to enable agents to efficiently make incentive contracts in various situations of automated agent environments. Different situations and contexts are thus explored:

- Certainty versus uncertainty.
- Full information versus partial information.
- Symmetric information versus asymmetric information.
- Bilateral situation versus a situation where more than two automated agents are involved in an environment.

Drawn from the game theory and economics literature, appropriate economics-based mechanisms and techniques that can be used for contracting in automated agents environments are, according to Kraus, fitted for each of the above mentioned situations and contexts. Given the constraints of the other entities, the agent that designs the contract is, in all the cases, provided with techniques to maximize its personal expected utilities.

5.2.2 Contract Net Protocol

In this category, Nwana et al. identify the *Contract Net Protocol* [34-35] as a classic coordination technique for decentralized allocation of tasks and resources among agents. In this approach, where the description of activities is specified in terms of roles, agents can be assigned two functions:

1. *A manager* to assume the role of dividing a problem into sub-problems and searching for contractors to tackle them. The manager is also assigned the task of monitoring the overall solution of the global problem.

2. *A contractor* to accomplish a sub-task. The contractors may however recursively assume the role of managers to further decompose the sub-tasks assigned to them and find sub-contractors to undertake them.

Nwana et al. further specify clearly and explicitly how managers locate possible contractors. In this context, it is stated that agents endowed with managerial responsibilities locate contractors via the following bidding process:

- A manager issues a task announcement,
- Contractors evaluate the announced task with respect to their abilities and commitments,
- Contractors send bids to the manager to announce their willingness and ability to perform a given task,
- The manager then evaluates received bids, selects the most appropriate contractor to which it awards a contract,
- The manager finally waits for results emanating from the contractor.

Furthermore, it is pointed out that the Contract Net Protocol is a completely distributed scheme, where a node can simultaneously assume the role of manager and contractor. Finally, it is reported that the approach has been implemented and generalized in various applications, such as the discrete manufacturing environment that partitions tasks [36] or the multistage negotiation paradigm involved in the monitoring and control of a complex communication system [37].

Citing Huhns and Singh [38], Nwana et al. note that the Contract Net Protocol is a high-level coordination strategy which also provides the ability to distribute tasks and self-organize a group of agents. They further state certain conditions in which this protocol can be best used:

- A well-defined hierarchical nature of the application task,
- A coarse-grained decomposition of the problem,
- A minimal coupling among sub-tasks.

According to the authors, the Contract Net Protocol is a reliable mechanism for distributed control and failure recovery that indeed provides various advantages, such as:

- Better agreements can be made through self-bidding in dynamic task allocation,
- Dynamic introduction and removal of agents is ensured,
- Natural load-balancing is achieved as busy agents need not bid,

However, a basic drawback associated with the Control Net Protocol lies, as noted by Nwana et al., in the inability to anticipate the existence of agents with contradictory demands. Hence, the approach can neither detect nor resolve possible conflicts, whereas coordination is, as noted, required to detect and resolve conflicts. In their view, agents in the Contract Net are rather passive, benevolent, and non-antagonistic. It is believed that having agents with such attributes is unrealistic in many real-world problems. It is reported that Conry et al. [37] tackled this limitation by essentially introducing an iterative mechanism to enable agents with conflicting goals to reach a consensus. On the other hand, it is finally argued that the Contract Net approach is rather communication-intensive. Consequently, the associated costs may indeed outweigh its advantages in real-world applications.

5.3 Multi-agent planning

Described as a coordination approach designed to avoid inconsistent and conflicting actions and interactions, multi-agent planning is indeed classified under a unique category. Nwana et al. thus explain that this approach is implemented to allow agents to build a multi-agent plan containing details of all future actions and interactions so as to achieve their goals, and interleave execution with more planning and re-planning. In this respect, the authors identify two types of multi-agent planning, namely:

1. Centralized multi-agent planning.

2. Distributed problem solving.

The centralized multi-agent planning is further described as an approach usually characterized by the need to incorporate a coordinating agent. On receipt of all partial or local plans from individual agents, the latter analyzes these plans in order to identify potential inconsistencies and conflicting interactions, such as those associated with limited resources. It is also added that the coordinating agent then attempts to modify the partial plans so as to eliminate conflicting interactions before they are combined into a multi-agent plan.

In this respect, Nwana et al. report research related to the implementation of the centralized multi-agent planning approach conducted separately in the 1980’s by Georgeff [38-39] and Cammarata et al. [40]. Georgeff introduced a method designed to synthesize multi-agent plans from single-agent plans. The idea was to insert communication acts into the single-agent plans so that agents can synchronize activities and avoid harmful interactions. He acknowledged that the approach does not guarantee solutions to some classes of problem involving complex interactions between single-agent plans. However, he added that the method has a wide applicability in many real-world domains, such as in automated factories and cooperative robot assembly tasks. Later, Georgeff proposed a more general model of action and showed how it can be used in the synthesis or verification of multi-agent plans and concurrent programs.

On the other hand, Cammarata developed and simulated in the domain of airtraffic control various problem solvers based on an approach characterized as *task centralization*. Within this scheme, agents associated with aircrafts involved in any conflict will have to select and decide which one of them can appropriately assume the role of the coordinating agent. In this context, the latter is required to modify only his flight plan to resolve the conflict. Adopting passive information-gathering roles, the remaining agents will thus perform no planning or actions, but merely send in their intentions or plans to the selected coordinating agent.

Another application reported by Nwana et al. concerns the MATPEN model proposed by Jin and Koyama [41] to provide the ability to coordinate autonomous and distributed agents based on centralized planning. Using MATPEN, Y. Jin and T. Koyama developed and implemented a ship collision avoidance system. In this approach, two agents sharing a conflict form a conflict group and initiate a negotiation process through an expectation-based negotiation protocol. “Expectations” are thus exchanged in order to decide which agent should assume what role in the negotiation process. As a result, agents generate a multi-agent plan required to resolve the conflict.

Referring Corkill [42], Nwana et al. note that the idea behind the distributed multi-agent planning lies in the need to provide each agent with a model of other agents’ plans. In a system based on this idea, agents thus communicate until all conflicts are resolved. In this context, each agent is devoted to building its individual plan and a model of other agents’ plans.

Furthermore, Nwana et al. report research efforts dedicated to performing this style of distributed processing. In this respect, the authors briefly mentioned the *Functionally Accurate and Cooperative* (FA/C) protocol proposed by Lesser and Corkill [43]. As reported, this approach was indeed used in the Distributed Vehicle Monitoring Testbed, a system developed in early 1980’s by Lesser and Corkill for testing coordination strategies [44]. FA/C is described by Nwana et al. as an approach where loosely coupled agents form high-level, but possibly incomplete, plans, results and hypotheses. The objective is to bring agents sharing the same environment to exchange and refine incomplete input data and possibly incomplete, incorrect and inconsistent tentative results received from other nodes until they all reach a convergence on some global complete plan.

Another application of distributed multi-agent planning is also briefly mentioned. It concerns the Partial Global Planning (PGP) approach proposed by Durfee and Lesser [45]. As described by Nwana et al., this approach is designed to allow agents to execute their local plans with each other. According to them, these plans are, in turn, continuously modified based on partial global plans, built by exchanging local plans. It is further explained that agents within this framework are always searching for potential improvements to group coordination.

Finally, Nwana et al. briefly note that the multi-agent planning, whether it is centralized or distributed, implies the need for agents to share and process substantial amounts of information. It is thus argued that such an approach is likely to require more computing and communication resources than other approaches. Though without much explanation, the authors subsequently note that the centralized multi-agent planning approach and the master/slave coordination technique described above share many common limitations. However, it is added that the coordination in the distributed multi-agent planning is much more complex than in the centralized form. As explained, there may not be any agent which possesses a global view of the distributed system. Referring to the work of Huhns and Singh [46], Nwana et al. finally argue that the scope of applicability of some existing multi-agent planning techniques, such as PGP, may be better in some domains than in others, as coordination resulting from these techniques is a gradual process.

5.4 Negotiation

Indeed, negotiation constitutes the most significant part of DAI research in coordination [26]. It has been the subject of considerable interest in multi-agent systems (MASs), as it has been in economics, political science, and social studies. In multi-agent domains, negotiation is used to resolve conflicting situations arising over the usage of joint resources or task assignment, problematic issues concerning document allocation in multi-server environments, and conflicts between a buyer and a seller in electronic commerce.

In a comprehensive survey on negotiation [68], Kraus describes various techniques on reaching agreements in multi-agent environments. She mainly discusses both game theory and economics-based techniques. She also presents logical based mechanisms for argumentations. In this review, Kraus particularly focuses on negotiation of self-interested agents. She also mentions several mechanisms designed for cooperative agents trying to resolve conflicts arising from conflicting beliefs about different aspects of their environment.

On the other hand, Wooldridge dedicates an entire chapter of his book on “MultiAgent Systems” [69] to aspects related to the issue of cooperation and conflict resolution in a society of self-interested agents via negotiation and argumentation. In this respect, the author first examines in more detail the question associated with the design of mechanisms or protocols that govern negotiation scenarios. In other words, Wooldridge provides a comprehensive description of desirable properties that should characterize negotiation protocols to govern multi-agent interactions. In the remainder of the chapter, the author presents auctions, negotiation protocols and strategies, and finally discusses argumentation.

Indeed, the choice of a protocol - i.e. mechanism - will thus depend on what properties the designer wants the overall system to have. Since the agents are self-interested, the evaluation of the results of multi-agent protocols becomes a difficult task, as Kraus points out. According to her, each agent engaged in a negotiation process is concerned only about its own benefits and losses. Hence, the success of any negotiation outcome is not easy to assess, as the appropriate question of which agent was successful remains to be answered.

Both Kraus and Wooldridge provide various criteria or properties designed to provide the ability to evaluate negotiation protocols or mechanisms. These criteria include: negotiation time, guaranteed success, maximizing social welfare, efficiency and Pareto efficiency, individual rationality, stability, simplicity, distribution, money transfer.

Nwana et al. claim that a significant part of research in the coordination area has been or is being conducted world-wide under this category. To support this claim, they argue that negotiation mechanisms are indeed involved in most coordination schemes. While recognizing the abundance of literature in this specific area of research, the authors devote a large section to coordination through negotiation mechanisms.

Entirely borrowed from Bussmann and Muller [47], negotiation is succinctly defined as "*the communication process of a group of agents in order to reach a mutually accepted agreement on some matter*". Referring to Sycara [48], it is suggested that agents must reason about beliefs, desires and intentions of other agents in order to negotiate effectively. This approach to negotiation led, according to Nwana et al., to the development of numerous techniques designed to:

- Represent and maintain belief models,
- Reason about other agents' beliefs,
- Influence intentions and beliefs of other agents.

While outlining the existence of a huge and varied literature on this topic, the authors proposed a classification of negotiation techniques into three broad categories:

1. Game theory-based negotiation.
2. Human-inspired and miscellaneous AI-based negotiation approaches.
3. Plan-based negotiation.

5.4.1 Negotiation techniques

We discuss below the three negotiation techniques : game theory-based negotiation, Human-inspired AI-based negotiation and plan-based negotiation

5.4.1.1 Game theory-based negotiation

As outlined by Nwana et al., game theory-based negotiation has become a growing area of research. The authors roughly suggest that the origin of such work can be traced to the doctoral thesis of Rosenschein [49]. According to them, the content of this thesis and results of earlier research conducted by the same author have been later synthesized, refined and compiled [50].

The resulting published collection outlines several strategies and protocols for negotiation. It proposes an approach based on game theory, which, according to Nwana et al., shows how to achieve coordination among a set of rational and autonomous agents with no a priori and explicit mechanism. Stemming from the compiled research of Rosenschein, the approach does not thus presume “the benevolent agent assumption”.

The key elements associated with this approach are, in this context, first identified and enumerated and then succinctly defined. They concern the concepts of utility functions, space of deals, and strategies and negotiation protocols. Utility is first defined as “*the difference between the worth of achieving a goal and the price paid in achieving it”*. A deal is, on the other hand, presented as “*an action one can take which has an attached utility”*. The negotiation protocol is finally viewed as a concept designed to “*define the rules which govern the negotiation, including how and when it ends*”.

Next, Nwana et al. describe how the actual negotiation proceeds. They explain that utility values for each outcome associated with a given interaction are built for each agent into a payoff matrix. The latter is, as noted, a common knowledge to (typically) both parties involved in the negotiation. It is further added that the negotiation process involves an interactive process of offers and counter-offers in which each agent chooses a deal which maximizes its expected utility value. The approach implicitly assumes that each agent involved in the negotiation is an expected utility maximizer. As noted by Nwana et al., an agent engaged in this process is devoted the task of evaluating at each step of the negotiation the offer emanating from another party.

On the other hand, Nwana et al. report that Zlotkin and Rosenschein [51] have extended the work described above to cover agents that are not truthful or, in other words, agents that can be deceptive. In this respect, they used simple demonstrators to show that an agent may strike better negotiation deals if it withholds certain information or deliberately misinforms other agents. As part of this work, the authors viewed negotiation as a two-stage process: the actual negotiation and the execution of the resulting plans. In addition, it is reported that Zlotkin and Rosenschein later worked on a general theory of automated negotiation in which complex domains are classified into: task oriented domains, state-oriented domains, and worth-oriented [52].

Finally, Nwana et al. mentioned a research work on negotiation conducted by Kraus & Wilkenfield [53]. Using game theory techniques with appropriate modifications, these authors propose a strategic model that takes time into consideration during the negotiation process. It is explained that time is incorporated into the model so as to influence the outcome of negotiations and avoid delays in reaching agreements.

Referring to a report on project NOMADS-001 conducted at British Telecom Labs in the U.K. [54], Nwana et al. claim that game theory-based negotiation fails to address some crucial problems. In this respect, it is explained that the theory presumes that agents are fully rational entities that act to maximize utility using pre-defined strategies. It is added that agents are, unlike the real world, also presumed to have, through the payoff matrix, full knowledge of other agents' preferences. In this context, it is further noted that for truly non-benevolent and loosely-coupled agencies, this assumption is unrealistic as an entity which only has partial or incomplete knowledge of its own domain cannot determine other agents' preferences. As also mentioned, another problem associated with this approach arises from the fact that a payoff matrix generated through a negotiation that involves too many agents and outcomes could well become both very large and intractable.

According to Nwana et al., agents using the game theory-based negotiation approach simply ignore past interactions and future implications to only consider the current state when deciding on their deal. In this approach, it is noted that agents are also considered to have identical internal models and capabilities. According to them, much of the research related to the game theory-based negotiation presumes two agents negotiating "though some later work is addressing n-agent negotiation" [55]. Considering that these assumptions are untenable in real-world problems, Nwana et al. raise doubt and uncertainty about the effectiveness of the game theory-based negotiation in dealing with real-life industrial agent-based applications despite mathematical proofs provided by Rosenschein and Zlotkin.

Sub-game perfect equilibria –S.Kraus tackles the problem of how a rational agent can choose a suitable negotiation strategy. In this respect, the author puts forward notions and concepts for choosing a strategy or analyzing a negotiation process. Described as useful, the first notion considered is commonly known as ***the Nash Equilibrium*** [92-93].

Given a set of agents $\{A_1, ..., A_N\}$, a strategy profile $F = \{f_1, ..., f_N\}$ is a Nash equilibrium of a model of alternating offers, if, as defined, "each agent Ai does not have a different strategy yielding an outcome that it prefers to that generated when it chooses f_i, given that every other agent A_j chooses fj. Briefly, no agent can profitably deviate, given the actions of the other agents". In an attempt to comprehensively clarify the above definition through more fundamental and necessary details, Kraus explains that having all agents use the strategies designed and specified for them in the strategy profile of the Nash equilibrium induces a state where all of them lose the motivation to deviate and implement another strategy. However, the use of the Nash equilibrium in a model of alternating-offers leads, as she points out, to absurd Nash equilibria. Referring to the Theory of Industrial Organization by Tirole [94], Kraus thus states that an agent may use a threat that would not be carried out if it were put in the position to do so, since the threat move would give this entity lower payoff than it would get by not taking the threatened action. According to her, the reason behind this lies in the fact that Nash equilibrium strategies may be in equilibrium only at the beginning of the negotiation process, but may be unstable in intermediate stages.

In addition, Kraus further presents another stronger concept for negotiation analysis known as the ***sub-game perfect equilibrium*** [91]. In her attempt to define this concept, she states that "a strategy profile is a *sub-game perfect equilibrium* of a model of alternating offers if the strategy profile induced in every sub-game is a Nash equilibrium of that sub-game". In other terms, Kraus attempts to better clarify the above definition, explaining briefly that at any step of the negotiation process, no matter what the history is, all agents lack the motivation to deviate and use a different strategy that is not defined in the strategy profile. According to her, the sequential equilibrium [95], which takes the beliefs of the agents into consideration, can be used in incomplete information situations where there is no proper sub-game.

Summarizing the above discussion, Kraus states that the strategic negotiation model provides a unified solution to a wide range of problems. This model is, according to her, appropriate for dynamic real-world domains. She reports that the strategic negotiation model was, in addition to the above application, used to tackle resource allocation and task distribution problems and other problems related to pollution allocation [87]. In this respect, Kraus explains that the strategic negotiation implemented in these domains provides the negotiators with ways to reach mutually beneficial agreements without delay. On the other hand, she finally reports that the strategic negotiation model was also applied to domains characterized by human high-pressure crisis negotiations [96-97].

5.4.1.2 Human-inspired AI-based negotiation

Referring to the work of Werkman [59], Nwana et al. first note that almost every form of human interaction requires some degree of explicit or implicit negotiation. With this in mind, they notice, with no surprise, that many researchers engaged in this area draw from human negotiation strategies, often leading to the usage of miscellaneous AI techniques including logic, case-based reasoning, constraint-directed search, etc. They finally conclude that different approaches described below provide examples of AI-based techniques that do not necessarily borrow from human interactions.

In this context, Sycara is reported to have proposed a general negotiation model that attempts to handle multi-agent, multiple-issue, single and repeated type negotiations [60]. Considering negotiation as an interactive activity, the author has accordingly exploited both case-based reasoning and multi-attribute utility theory. As far as the case-based approach is concerned, Nwana et al. point out that Sycara argues that human negotiators draw from past experience to guide present and future negotiations. However, in the absence of past cases, it is reported that the author resorts to preference analysis based on multi-attribute theory, where issues involved in negotiation are represented by utility curves. In this respect, it is further explained that the curves are combined in additive and multiplicative fashions so as to select a proposal that maximizes the utility.

As reported by Nwana et al., Sycara has indeed developed and implemented PERSUADER, a framework based on the integration of case-based reasoning and multi-attribute utility theory. As stated, this system is designed to operate in the domain of labour management dispute to provide negotiation support in conflicts resolution with the aid of two practicing negotiators. While further recognizing that Sycara's use of case-based reasoning is quite novel to the negotiation problem, Nwana et al. explain that PERSUADER creates and handles cooperative interactions where agents can modify beliefs, behaviours and intentions of other entities through persuasion.

As described by Nwana et al., another approach to negotiation proposed by Sathi and Fox [61] amounts to a *constraint-directed search* of a problem space using negotiation parameters. As stated, the authors view negotiation as a two-phase process including a communication phase and bargaining phase. While the communication phase is designed to enable the transfer of information to all participating agents, the bargaining phase offers a framework where deals are made between individuals or in a group.

In an attempt to provide a more detailed description of this approach, Nwana et al. explain that agents using the technique proposed by Sathi and Fox negotiate through a relaxation of various conflicts and constraints until agreement is reached. While noting that the solutions may be modified, they further point out that negotiation preferences are initially modeled as constraints. As also mentioned, negotiation operators including operators which simulate relaxation, reconfiguration and composition are drawn from studies regarding human negotiation [62] and indeed used to generate new constraints. Based on this approach, a system that "performs marginally better than experts" has been built to handle, as reported, resource allocation. The method described above is thus an iterative approach which, according to Nwana et al., presents a major problem as agents can easily get caught in an infinite loop of exchanging offers. As further explained, this limitation stems from the fact that the process of selecting relaxations is conducted with no criteria provided.

Also reported by Nwana et al. is the Designer Fabricator Interpreter (DFI) model proposed by Werkman [31]. It is a knowledge-based model of an incremental form of negotiation which is, according to them, largely inspired from various human models of negotiation. Furthermore, Nwana et al. add that the proposed scheme uses a shared knowledge representation called *shareable agent perspectives* which offers agents the ability to perform negotiation in a manner similar to cooperating or competing experts sharing a common background of domain knowledge.

Detailing this description, Nwana et al. explain that Werkman's scheme is designed to basically exploit a blackboard with partitions for requested proposals, rejected proposals, and accepted proposals, a partition for communications, and shared knowledge. They further argue that the scheme proposed by Werkman provides agents with rich detail and knowledge of the perspectives of other entities, thus giving them invaluable information to make better proposals in the future.

As reported by Nwana et al., Werkman considers the negotiation process as a three-phase cycle:

Phase I – A proposing agent announces a proposal that is received and evaluated by a receiving agent.

Phase II – The initial proposal is then simply accepted. However, in case the receiving agent is not satisfied with the initial proposal, a counter proposal is thus generated.

Phase III – The counter proposal is finally submitted for review by other agents.

Using their mutual information network, two agents may, as noted, engage in an incessant process of reviewing the negotiation dialogue to generate alternative proposals. An arbitrator agent is thus expected, as explained, to handle and break a deadlock in negotiations using issue relaxation techniques or some intelligent proposal generator. However, it is explained that in case this procedure fails, the arbitrator may, on the other hand, set time limits or use other techniques.

Considering the use of arbitration as a relatively novel approach, Nwana et al. describe the Werkman proposal as both interesting and definitely worth examining and studying in more detail. In this respect, they first remind the reader of the bottleneck problem associated with the centralized blackboard, where reading and posting to the blackboard seem to generate chaos in the absence of an explicit scheduler. Then, they maintain that the approach proposed by Werkman provides agents with the ability to communicate via the blackboard through a speech-act based language. According to them, Werkman's proposed approach, which has been implemented in the DFI system, seems to have a good understanding of the negotiation process.

In addition, Nwana et al. attribute to Conry et al. research aiming to develop negotiation strategies designed to tackle *distributed* *constraint satisfaction* problems where local constraints give rise to a complex set of global and inter-dependent constraints [37]. As clearly mentioned, this research was specifically designed to investigate problems where agents in a network have a goal, but each node or agent has only limited resources. Conducted as part of a study regarding long-haul transmission systems and more complex communication networks, it involved developing algorithms for multi-agent planning while taking the inevitable conflicts into consideration. Presenting this investigation as an interesting research issue, Nwana et al. note that it starts with looking into whether specific generic tasks may be linked with specific negotiation strategies.

Research work by Bussmann and Muller centered on developing a negotiation framework for cooperating agents [28] is further referred to as one of the most useful papers on negotiation in the published literature. In this work, Bussmann and Muller attempt, as Nwana et al. point out, to address various limitations associated with other negotiation proposals and models. With a particular emphasis on Gulliver's eight phases of the negotiation process [64], the authors draw their inspiration from socio-psychological theories of negotiation to gradually develop a general and simple cyclic negotiation model as a way to address the thorny issue of conflict resolution.

In an attempt to provide more details about the model proposed by Bussmann and Muller, Nwana et al. explain that the general strategy adopted consists in starting the negotiation process with one, some or every agent making a proposal. In response, activities are thus engaged where agents evaluate and check out the resulting proposals against their own preferences. The agents can then present a criticism of these proposals through a list of violated preferences.

According to Nwana et al., the evaluation procedure offers agents the opportunity to update their knowledge regarding the preferences of other entities. It is then followed by a resumption of the negotiation cycle through which new proposals are thus introduced in the light of newly gleaned information. Described as an interesting approach, this proposed negotiation model provides the ability to handle conflicts between agents in a concurrent resolution cycle.

Finally, Nwana et al. briefly mention various other negotiation protocols, including Kuwabara and Lesser's extended protocol for multistage negotiation [65] and Durfee and Montgomery's hierarchical protocol for coordinating multi-agent behaviour [66].

In conclusion, Nwana et al. synthesize key lessons learned and put forward conclusions drawn from conducting this review. In this respect, they question the scope, applicability, and usability of one-off coordination strategies. They further notice the lack of enough studies designed to validate many proposed strategies. In this context, Nwana et al. highlight the absence of information regarding when, where, how, and why various negotiation and coordination strategies or a combination of them are used in various applications.

On the other hand, Nwana et al. state that the contract net and master/slave models, including variations of them, are simple and the most used strategies. According to them, the contract net approach, like various other techniques, presumes truly benevolent, trustful, non-conflicting, and helpful agents. They add that most coordination or negotiation strategies do not involve any complex meta-reasoning required of most domains. Finally, Nwana et al. refer to the lack of fundamental analysis of the process of coordination to explain that the work of Jennings [67] is a first step in this direction.

5.4.1.3 Plan-based negotiation

Plan-based negotiation is the second classified area of research where many experts have been involved. Adler et al. [56] are first cited to have investigated negotiated agreements and discussed methods of conflict detection and resolution in the area of telephone network traffic control. According to Nwana et al., the authors strongly maintain that negotiation and planning are very tightly intertwined. While arguing that agents need information from others to function effectively and efficiently, Adler et al. imbue these entities with planning knowledge.

Within the same area of research, Nwana et al. explain that Kreifelt and von Martial view negotiation as a two-stage process: agents plan their activities separately and then coordinate the resulting plans. Based on this approach, it is reported that the authors propose a negotiation strategy for autonomous agents [57]. According to Nwana et al., Kreifelt and von Martial devote a separate agent solely to activities related to the coordination of all the other agents. As further noted, the authors present a negotiation protocol in terms of agents' states, message types, and conversation rules between agents.

In an attempt to highlight the limitations associated with this protocol, Nwana et al. referred to Bussman & Muller [58] to point out that the proposal made by Kreifelt and von Martial prescribes rather than really presents a negotiation model. According to them, the proposed approach leaves to the agents how to really achieve consensus. Nwana et al. finally explain that the protocol itself also needs some further clarification, stating that, in general, plan-based negotiation suffers from limitations related to centralized or distributed multi-agent planning.

5.4.2 Negotiation models

In "Automated Negotiation and Decision Making in Multiagent Environments", Kraus presents a short survey of various negotiation approaches used in DAI and social sciences [68]. She then attempts to demonstrate the applicability of a social sciences-related approach to MASs and describes how Moehlman et al. use negotiation as a tool for distributed planning [71]. Applied in the Phoenix fireman array, this approach is, according to her, designed to offer each agent, subjected to certain important constraints, the opportunity to find a feasible solution using a negotiation process. On the other hand, Kraus also mentions a method based on negotiation search proposed by Lander and Lesser [72]. The approach is described as a multi-stage negotiation used as a means of cooperation while searching and solving conflicts among the agents. In the area of multi-agent environments, Kraus subsequently provides a relatively more detailed description of negotiation-related research work attributed to Rosenschein and Zlotkin [50]. Sycara is reported to have developed a model of negotiation that combines case-based reasoning and optimization of multi-attribute utilities (see [73,74,76] for more details). In this approach, the buyer and the seller can update their beliefs about the opponent's reservation price using the Bayesian rule.

Without giving much information, Kraus further provides a description of the work she conducted with Lehmann in the area of negotiation in multi-agent environments [75,77]. Sandholm and Lesser are finally reported to have discussed issues, such as levels of commitment that arise in automated negotiation among self-interested agents whose rationality is bounded by computational complexity [78].

In an attempt to provide an idea of how issues pertaining to negotiation are really tackled in areas other than DAI, Kraus gives a concise overview of negotiation approaches in social studies. According to her, there are two main approaches to the development of theories related to negotiation in this area. Described as the formal theory of bargaining [79-80], the first approach constitutes a formal, game-theoretic approach. Referred to as the negotiation guides approach, the second approach includes informal theories that attempt to identify possible general beneficial strategies for a negotiator [81-86].

Based on Rubinstein's model of alternating offers [88], the strategic-negotiation model proposed by Kraus consists of a protocol designed to govern interactions between agents, utility functions associated with these entities, and negotiation strategies that each agent is required to adopt [87,89].

According to Kraus, a specific example of a distributed information system lies in the Data and Information System component of the Earth Observing System[7] (EOSDIS) attributed to NASA. She describes it as a distributed system supporting archival data and distribution of data at multiple and independent data centers, called Distributed Active Archive Center (DAAC) .

[7] Consult EOSDIS home page at http://www-v0ims.gsfc.nasa.gov/v0ims/index.html

Arranging agents randomly in a specific order before the negotiation begins represents, according to Kraus, a fair and reasonable method pertaining to the decision on the order in which these entities will make offers. In this respect, the author notes that a distributed algorithm for randomly ordering the agents can be based on the methods presented in [90]. Referring to the work of M. J. Osborne and A. Rubinstein [91], Kraus briefly notes that a strategy profile represents a collection of strategies, one for each agent.

TAEMS, MOISE+, OPERA, ODML, AGR, etc. are models that can be used to reach coordination by an external framework that constrain the agents' interactions (e.g. See [296]).

This page intentionally left blank.

6. Coordination through protocols

We discuss in this chapter, from one side, how auctions can be used to tackle the coordination issues. We also describe, from another side, the different type of auctions. The coalition formation and argumentation will also be tackled from a coordination perspective.

6.1 Auctions for resolving conflicts

In [68,87], Kraus briefly tackles the question of using auctions, another game-theory based technique, as a way of reaching agreements in multi-agent environments. While recalling the example relating to the information server environment described above, Kraus emphasizes that most of the conflicts occurring in such domains where agreements concerning the distribution of a set of items should be reached can be resolved efficiently by providing agents with a monetary system. Referring to Ronen's work on algorithms for rational agents [98], she explains that agents in this system are modeled as buyers and sellers to provide the ability to resolve occurring conflicts using money transfer.

On the other hand, Wooldridge devotes a lengthy discussion to address rules, terms and conditions pertaining to using auctions for resolving conflicts in agent-based environments. Both Kraus and Wooldridge put forward the emergence of the Internet to highlight how auctions have been transformed from a comparatively rare practice to an area of increased interest with a large, international audience. In this respect, they quote the example of eBay to explain how large virtual businesses and houses have recently sprung up around the idea of online auctions. According to Wooldridge, the current popularity of online auctions lies in the fact that auctions are extremely simple interaction scenarios. Easy to automate, auctions have thus become, as further noted, a good first choice to implement as a way for agents to reach agreements. While considered as a rich collection of problems for the research community despite their simplicity, auctions represent, as Wooldridge notes, a powerful tool that automated agents can indeed use for allocating goods, tasks, and resources.

As described by Wooldridge, an auction takes place between an agent known as the *auctioneer* and a collection of agents known as the *bidders*. The author notes that the auction offers the auctioneer the opportunity to allocate the good to one of the bidders. In his terms, Wooldridge adds that in most settings – and certainly most traditional auction settings – the auctioneer desires to maximize the price at which the good is allocated, while bidders desire to minimize price. In this respect, he further explains that the auctioneer attempts to achieve his desire through the design of an appropriate auction mechanism – the rules of encounters – while the bidders attempt to achieve their desires by using a strategy that will conform to the rules of encounter, but that will also deliver an optimal result.

On the other hand, Kraus identifies two patterns of interactions in auctions. According to her, the most common are one-to-many [99,100,101] and many-to-many [102] auction protocols. She explains that the first pattern of interactions is the one-to-many auction protocols designed to enable one agent to initiate an auction so that a number of other agents can bid in the auction. As also stated, the second pattern represents the many-to-many protocols where several agents initiate an auction and several other agents can bid in the auction. Kraus then points out that given the pattern of interaction, the first issue is to determine the type of protocol to use in the auction [103]. Similarly, she explains that given the protocol, the agents then need to opt for their bidding strategy.

For Wooldridge, several factors are indeed susceptible of affecting the protocol and the strategy that agents can use. He further states that the most important factor is related to whether the good for auction has a *private* or a *public*/*common* value. In this respect, he adds that a good for auction has a common value if it is worth exactly the same to all bidders in the auction. According to him, a good for auction is said to have a private value, however, if each agent values it differently. Another type of valuation is, on the other hand, identified by Wooldridge as *correlated value*. He then defines it as a setting at which an agent's valuation of the good depends partly on private factors, and partly on other agents' valuation of it. In an attempt to provide a clearer understanding of this type of valuation, Wooldridge gives the example of an agent that was bidding for a painting it liked, but wanted to keep open the option of later selling the painting. He further explains that, in this case, the amount it would be willing to pay would depend partly on how much it liked it, but also partly on how much it believed other agents might be willing to pay for it if it put it up for auction later.

Furthermore, Wooldridge considers and identifies three dimensions along which auction protocols may indeed vary:

1. The first dimension is identified as *Winner determination*. It is about who gets the good that the bidders are bidding for. It is observed that the answer to this question is obvious since in familiar auctions, the agent that bids the most is commonly allocated the good. As stated, such protocols are commonly known as *first-price* auctions. However, it is pointed out that another possibility could be that the good is allocated to the agent that bid the highest, but pays only the amount of the second highest bid. As finally noted, this type of protocol is known as *second-price* auctions.

2. As identified, the second dimension that characterizes an auction protocol lies in the answer to whether or not the bids made by the agents are known to each other. In this respect, it is stated that an *open-cry* auction imposes common knowledge bids. In other terms, an *open-cry* auction is defined as an auction where every agent can see what every other agent is bidding. It is also specified that if the agents are not able to determine the bids made by other agents, the auction is thus said to be a *sealed-bid* auction.

3. The third dimension referred to is about the mechanism through which bidding proceeds. In this respect, it is pointed out that auctions defined as *one shot* auctions represent the simplest event characterized by a single round of bidding followed by the allocation of the good to the winner. It is also mentioned that a second possibility consists in starting with a low

price – often at a *reservation price* – and then successive bids will be conducted for increasingly large amounts. As stated, such protocols are known as *ascending* auctions. The alternative possibility finally referred to – *descending* auctions – imposes upon the auctioneer that he start off with a high value, and that he decrease the price in successive rounds.

Indeed, both Kraus and Wooldridge attempt to examine and classify various types of auctions according to the set of dimensions extensively described above. While Kraus provides a brief discussion, Wooldridge proposes a classification together with an extensive and informative description of the most familiar auctions. Covered in both studies, the most important ones are discussed below.

6.1.1 Type of auctions

6.1.1.1 English auctions

Made famous by the houses of Sothebys, English auctions are, according to Wooldridge, the most commonly known type of auction. Kraus describes the English auction as an ascending auction in which “the price is successively raised until only one bidder remains, and that bidder wins the item at the final price”. She reports that in one variant of the English auction “the auctioneer calls higher prices successively until only one willing bidder remains, and the number of active bidders is publicly known at all times”. According to her, the bidders in other variants of the English auction call out prices themselves, or have the bids submitted electronically and the best current bid is posted.

On the other hand, Wooldridge describes the English auctions as *first-price*, *open cry*, *ascending* auctions where:

- The auctioneer starts off by suggesting a *reservation price* for the good which may be 0. However, the good is allocated to the auctioneer if none of the agents is willing to bid more than the reservation price.

- Otherwise, the auction process continues and agents are invited to submit their bids that must be more that the current highest bid.

- In case no other agent is willing to raise the bid, the good is then allocated to the agent that has made the current highest bid. The price paid for the good is the amount of that bid.

Wooldridge then brings up the question of what strategy should an agent use to bid in English auctions. According to him, the dominant strategy for an agent is to successively bid a small amount more than the current highest bid until the bid price reaches its current valuation, and then to withdraw.

While considering English auctions as simple bidding process, Wooldridge attributes some interesting properties to this type of auctions. He thus highlights a common interesting feature of English auctions arising from the uncertainty about the true value of the good being auctioned. In this respect, he gives the example of a sale where he imagines an auctioneer who has limited geological information to provide about some land he is selling to agents desiring to exploit the mineral resources. As Wooldridge further points out, none of the agents knows exactly the value of the land being auctioned in an environment where scarce information is deployed. The author then assumes that the agents, each using the dominant strategy mentioned above, engage in an English auction to obtain the land. In this context, he raises the question about the attitude and feelings the winner should have at the end of the auction. In other terms, Wooldridge wonders whether the winner should feel happy that it has obtained the land for less than or equal to its private valuation or worried that no other agent valued the land so highly. As he finally points out, the latter situation, known as the *winner's curse*, occurs most frequently in English auctions.

6.1.1.2 Dutch auctions

Dutch auctions are examples of *open-cry descending* auctions. Both Kraus and Wooldridge describe this type of auctions as an event where:

- The auctioneer begins by offering the good at artificially high value.
- He then continuously lowers the offer price of the good by some value, until an agent makes a bid for the good which is equal to the current offer price.
- The good is finally allocated to the agent that made the offer.

According to Wooldridge, there is no dominant strategy for Dutch auctions in general.

6.1.1.3 First-price sealed-bid auctions

Considered as the simplest of all the auction types, first-price sealed-bid auctions are, according to Wooldridge, examples of *one-shot* auctions. It is thus separately described by both Kraus and Wooldridge as a single round event in which bidders submit a bid for the good to the auctioneer. Then, without subsequent rounds, the good is, as explained, awarded to the agent that made the highest bid. It is finally stated that the winner pays the price of the highest bid. In this respect, Wooldridge highlights the absence of any opportunity for the participating agents to offer larger amounts for the good.

He then raises the question of how an agent should act in first-price sealed-bid auctions. To answer this question Wooldridge considers an event where every agent bids its true valuation. Since there is only a single round, the author notes that the good is allocated to the agent that bid the highest amount. Considering the second highest bid, he then observes that the winner still has been awarded the good while it could have offered an amount that is just a tiny fraction more than the second highest price. As a result, most of the difference between the highest and the second highest price is, in effect, as Wooldridge points out, money wasted as far as the winner is concerned. As a result, the author finally highlights the need for a participating agent to bid less than the good's true valuation, depending on what the other agents bid.

Described as a sealed-bid auction, the second-price auction represents, according to Kraus, the event where the buyer making the highest bid claims the object, but pays only the amount of the second highest bid.

6.1.1.4 Vickery auction

Wooldridge considers the Vickery auction [104] as "the most unusual and perhaps the most counterintuitive" of all the auction protocols described in [69]. On the other hand, Kraus reports that this type of auction is widely used in DAI [49,105-108] and in research on electronic commerce [107-108] for the case of one-to-many auctions. She further claims that, under various assumptions, the Vickery auction-based protocol is incentive compatible. In other terms, she further explains that each participating bidder has incentives to bid truthfully.

Indeed, the Vickery auction falls, according to Wooldridge, within the category of the *second-price sealed-bid* auctions. In an attempt to give a more comprehensive description, he explains that the Vickery auction, which is based on a single negotiation round, is designed to provide each bidder with the opportunity to submit a single bid in the absence of appropriate information about the bids made by other agents. He further points out that the good is finally awarded to the agent that made the highest bid. In this respect, he notes that the amount the winner pays for that good represents the price of the second highest bid.

Kraus states that, under various assumptions, the Vickery auction-based protocol is incentive compatible. In other terms, each participating bidder has incentives to bid truthfully. In response to the question of why one would even consider using Vickery auctions, Wooldridge similarly argues that this type of auctions "make truth telling a dominant strategy". In an attempt to explain why a bidder's dominant strategy in a private value Vickery auction is to bid his true valuation, he supposes that an agent bid more than the good's true valuation. In this case, Wooldridge observes that the agent may indeed be awarded the good, but it runs the risk of paying more than the amount of its private valuation. In such circumstances, the agent will, as stated, make a loss, since it paid more than it believed the good was worth. On the other hand, Wooldridge imagines that the same agent bid less than its true valuation. In this case, he notes that the bidder stands less chance of winning than had it bid its true valuation. He further explains that even if the agent wins, the amount it pays will not have been affected by the fact that it bid less than its true valuation. The reason lies in the fact that the agent will, as noted, pay the price of the second highest bid. In conclusion, Wooldridge states that the best conduct for an agent engaged in a Vickery auction is to bid truthfully, no more and no less than its private valuation.

Referring to a research work on distributed rational decision-making in multi-agent systems [70], Wooldridge notes that Vickery auctions have received a lot of attention in the literature. According to him, this wide interest of the research community in the Vickery auctions lies in the fact that "truth telling" is indeed imposed as a dominant strategy. He further reports that this type of auctions have not, however, been widely adopted in auctions among humans. The reason mentioned as perhaps the most important is that humans frequently find the Vickery mechanism hard to understand, since, according to him, it seems so counterintuitive at first sight.

Furthermore, Wooldridge highlights the fact that Vickery auctions can lead to antisocial behaviour. In this respect, he describes a situation where an agent wants to acquire a good through a Vickery auction. While its private valuation is $90, he knows that another agent desires the same good and values it at $100. He explains that in case truth telling is the dominant strategy, the first agent can do no better than bid $90, while the second agent will eventually bids $100. In this context, the latter will be awarded the good, but will pay only $90. However, Wooldridge further argues that the first agent may not be too happy about the whole procedure and may thus decide to punish its successful opponent. The author then suggests that the agent can bid $99 instead of $90. As a result, the second agent will still be awarded the good, but it will pay $9 more than it would do if the first agent had been truthful.

According to Kraus, situations are likely to occur where the value of some items to a bidder depends upon which items are won. As explained, such situations may indeed result in the need for bidders to submit bids for combinations of items. Often referred to as *combinatorial auctions*, such auctions highlight the need to determine, as Kraus points out, the revenue of maximizing sets of non-conflicting bids. In this respect, the author mentions various approaches where researchers attempt to develop polynomial algorithms, either to respond to specific cases [109] or to provide sub-optimal solutions [110-112]. Also mentioned is another interesting research work pertaining to combinatorial auctions [113] where the author considers two aspects related to this market mechanism, namely the bidding language and the allocation algorithm.

Although not classified as one of the major four auction types, the *double auction* is another type of auction finally referred to as the most known auction protocol for many-to-many auctions. In double auction, buyers submitting bids and sellers submitting minimal prices are, according to Kraus, treated symmetrically [114]. In addition, the author mentions the existence of several algorithms that are implemented to match buyers and sellers and determine the transaction price. Referring to the work of Wurman et al. on flexible double auctions for electronic commerce [115], Kraus finally underlines the need for the double auction-based protocol to be *incentive compatible*, *individual rational* and *Pareto optimal*.

On the other hand, Wooldridge tackles several issues relating to the type of auctions discussed above.

6.1.2 Issues related to the type of auctions

We discuss below the most important issues related to the auctions and that can influence the coordination mechanisms. We discuss the expected revenue, lies and collusion, and the Counter-speculation.

- ***Expected revenue*** – It is observed that in all likelihood the overriding consideration of an auctioneer will be to maximize his revenue. In other terms, the auctioneer is much more likely to prefer an auction protocol that will get him the highest possible price for the good on offer. In addition, the question of whether or not agents are truthful may well not be of interest to the auctioneer. As a result, it is pointed out that some protocols – Vickery’s mechanism in particular – may seem not to address these considerations. With reference to the work performed by Sandholm [70], it is stated that the answer to the question of which protocol the auctioneer should choose depends partly on the attitude to risk of both auctioneers and bidders. The attitude to risk of these actors is summed up as follows:

    - *Risk-neutral bidders* – In all four main types of auctions discussed above, the expected revenue to the auctioneer is probably identical in case the attitude of bidders is risk-neutral. Using all of these types of auction, the auctioneer can, in other terms, expect on average to get the same revenue for the good.

- *Risk-averse bidders* – These bidders would prefer to get the good even when they have to pay slightly more than their private valuation. In this context, higher expected revenue for the auctioneer can result from Dutch and first-price sealed-bid protocols, since a risk-averse agent using these protocols can 'insure' itself by bidding slightly more for the good than would normally do a risk-neutral bidder.

- *Risk-averse auctioneers* – In this case, an auctioneer can do better using Vickery or English auctions.

Finally, Wooldridge recommends guidelines for treating very carefully the results pertaining to which protocol the auctioneer should choose. Given risk-neutral bidders, the author notes, for example, that the first result relating to the revenue equivalence of auctions depends critically on the fact that bidders really do have private valuations. He further adds that it becomes critical to ensure that the properties of the auction scenario – and the bidders – are correctly understood when choosing an appropriate protocol.

- ***Lies and collusion*** – Another interesting issue brought forward and discussed concerns the extent to which the auction-based protocols presented above are susceptible to antisocial behaviour likely to come from both bidders and auctioneer. Indeed, an auctioneer would ideally prefer a protocol that is immune to collusion by bidders. In other terms, the protocol can be designed to make it in the best interest of bidders not to engage in collusion with other bidders. Similarly, a potential bidder is also likely to prefer a protocol that imposes honesty as the dominant strategy on the part of the auctioneer.

  The four main auction types are believed to be vulnerable to collusion. In this respect, it is claimed that all agents involved in any of these types of auction can form a 'grand coalition' so as to agree beforehand to put forward artificially low bids for the good on offer. As noted, the bidders can thus obtain the true value of the acquired good and split the profits. Then, it is argued that the most obvious approach to preventing collusion is to modify the protocol to prevent bidders from identifying each other. As stated, such an approach is of course unpopular with bidders engaged in open-cry auctions, since the concerned bidders prefer to be sure the information about the bids placed by other agents is accurate.

  As further noted, the auctioneer engaged in a Vickery auction can overestimate the price of the second highest bid so as to force the winner to pay more than the true bid. To overcome this problem, bids can in some way – through a digital signature – be signed to offer the winner the opportunity to independently check the value of the second highest bid. Another possibility is to appoint a trusted third party to handle bids. However, open-cry and first-price sealed-bid auctions are reported to be lies-resistant protocols that are not susceptible to allow the auctioneer to make false statements. While open-cry auctions provide all agents with the ability to see bids formulated by others, the winner in a first-price sealed-bid auction is, on the other hand, allowed to examine the offers made by the other agents. In other cases, the auctioneer may place bogus bids to artificially inflate the current bidding price. Commonly known as shills, these bids represent however a potential problem only in English auctions.

- ***Counter-speculation*** – Counter-speculation is widely recognized as an issue worth tackling. It is characterized by "the process of a bidder engaging in an activity in order to obtain information either about the true value of the good on offer, or about the valuations of other bidders". Indeed, every agent would be tempted to engage in such activity if it were free and accurate. However, in most types of auction, counter-speculation may result in time and monetary costs. Investing in this activity would be wasteful and counterproductive, since the time cost is an important factor in auctions – e.g. English or Dutch auctions – that depend heavily on the time at which a bid is made. On the other hand, engaging in counter-speculation can be advantageous and worthwhile only in the case where the bidder can, as a result, expect to be no worse off than if it behaved otherwise. Hence, it appears that prior to engaging in counter-speculation, a clear trade-off analysis has to be made between the potential gains and the associated costs.

Market-oriented programming: Kraus provides a brief overview of market-oriented programming largely based on the research and experience of Wellman [116-118]. This methodology is presented as a programming paradigm and a new approach to distributed computation based on market price mechanisms.

In exploiting the institution of markets and the models associated to them, the approach is implemented to build, as Kraus stated, computational economies to address and solve particular problems of distributed resource allocation. Based on Wellman's view, the author wrote that the idea of market-oriented programming is inspired "in part by economists' metaphors of market systems 'computing' the activities of the agents involved, and also by AI researchers' view of modules in a DAI system as autonomous agents". Referring to the agents involved in such market systems, Kraus explains that the modules interact in a very restricted manner by offering to buy and sell quantities of commodities at fixed unit prices. Indeed, as the system reaches equilibrium, the computational market has, as noted, computed the allocation of resources throughout, and dictates the activities and consumptions of the various modules.

While it is applicable in systems characterized by incomplete information, the market-oriented programming approach does not necessarily, as Kraus points out, require money transfer. The author adds that the method is however applicable only in the presence of several units of each kind of goods and in the case the number of agents is large. It is thus argued that otherwise it is not rational for the agents to ignore the effect of their behaviour on the prices when they actually have an influence. Referring to the work of Wellman [117], Kraus reports that reaching equilibrium may be time-consuming and that the systems based on the market-oriented programming approach may not even converge. Finally, the author notes that the approach also requires the implementation of mechanisms – possibly distributed mechanisms – to manage auctions.

6.2 Coalition formation in coordination

Kraus outlines the usefulness of the formation of coalitions for executing tasks in both MASs and DPS environments. Referring to the literature pertaining to this method, the author presents the act of creating coalitions as an important way for agents to cooperate.

DPS is, according to Kraus, usually characterized by the absence of the need to motivate the individual agent to join a coalition. In this context, agents involved in a distributed problem-solving environment can be built to try to maximize the overall performance of the system. It is thus required, as stated, to only consider the problem pertaining to the nature of the coalitions that should be formed to maximize the overall expected utility of these entities. Note that the problem of finding the coalition structure that maximizes the overall utility of the system is NP-complete.

On the other hand, it is reported that a self-interested agent in MASs will join a coalition only if it could derive a clearer advantage from the coalition formation than it could gain otherwise. In this context, the problem associated with the division of the coalition's joint utility represents indeed another very important issue. As Kraus suggests, this issue can be tackled using game theory techniques for coalition formation. Indeed, numerous research papers and reports devoted to game theory [119-122] are reported to provide a description of which coalitions will form in NP-person games under different settings and how the players will distribute the benefits resulting from the inter-players cooperation. As noted, the profit-sharing process is conducted through the application of several stability-based notions such as the core, Shapley value and the kernel [123]. In this respect, Kraus argues that each of the stability notions just mentioned "is motivated by a different method of measuring the relative strengths of the participating agents". She further explains that the game-theory solutions to the coalition formation problem "do not take into consideration the constraints of a multi-agent environment – such as communication costs and limited computation time – and do not present algorithms for coalition formation".

In "Feasible Formation of Stable Coalitions among Autonomous Agents in Non-super-additive Environments" [124], Shehory and Kraus present a multi-agent approach to the coalition formation. As stated, the main contribution of this work lies in the provision of algorithms susceptible to dedicated self-interested agents. The approach is indeed designed to handle both the coalition structure and the division of the utility problems. It rests on an anytime algorithm developed to enable, as already stated, the formation of coalitions that satisfy a certain stability based on the kernel stability criteria mentioned above. As reported, simulations were indeed conducted in order to examine the properties of the proposed algorithm. According to Shehory and Kraus, the results which were applied to the formation of coalitions among information agents [125] show that:

- The model increases the benefits of the agents within a reasonable time period, and
- More coalition formations provide more benefits to the agents.

In a study pertaining to the key topic of coalition formation in multi-agent environments, Sandholm et al. focus on establishing a worst case bound on the quality of the coalition structure while only searching a small fraction of the coalition structures [126]. This work shows that a threshold – a minimum number of structures – is, as Kraus points out, necessary to search in order to establish a bound. As a result, an anytime algorithm is devised and implemented to establish a tight bound within this minimal amount of search. In case additional time is available, the algorithm is allowed to search further and progressively establish a lower bound.

As reported by Kraus, another interesting research relative to coalition formation has been the subject of a paper entitled “Coalition Formation among Bounded Rational Agents” [127]. In this paper, Sandholm and Lesser analyze coalitions among self-interested agents that need to solve combinatorial optimization problems to operate efficiently. Then, they adopt a model of bounded rationality where computational resources are costly. Furthermore, Sandholm and Lesser introduce a domain classification for rational and bounded rational agents. In this study, the authors concentrate, as Kraus points out, on the problem of computing the value of a coalition. In the adopted model, the computed value depends on the computation time available to the agents.

Kraus also mentions a research study conducted by Zlotkin and Rosenschein in the area of coalition formation [128]. In this study, the authors consider the general case of n-agent coalition formation and present a simple mechanism that uses cryptographic techniques for sub-additive Task-Oriented Domains. Considering only the grand coalition structure where all the agents belong to the same coalition, the authors provide, as Kraus reports, a linear algorithm that guarantees each agent an expected utility that is equal to its Shapley value. On the other hand, Ketchpel has also devoted a study to the analysis of a more difficult problem where agents have different estimates of the value that a coalition will obtain [129]. As noted, the author proposes and describes a utility distribution mechanism designed to perform in situations where there is uncertainty in the utility that a coalition obtains. The “two agent auction” mechanism suggested for complementing an existing coalition formation algorithm to solve the problem presents certain properties:

First, the agent valuing the collaboration more highly is always selected manager.

Second, the agreement price is a deterministic function of the agent’s initial offers and estimates of the value of collaboration.

Finally, another research work mentioned by Kraus refers to the use of coalition formation to manage task execution in DPS environments – multi-agent environments requiring cooperation among agents – [130]. While tackling only the coalition structure problem, the study presents efficient distributed algorithms with low ratio bounds and with low computational complexities for task allocation among computational agents in a non super-additive environment. The proposed approach is suitable for both:

- Coalition formation where each agent must be a member of only one coalition,
- Overlapping coalitions.

6.3 Argumentation in coordination

Centered on the trading of proposals, the approaches to reaching agreement discussed so far - game-theoretic and heuristic approaches - have a number of advantages. Perhaps the most important advantage associated to these approaches lies in the ability to prove some desirable properties of the negotiation protocols [69]. However, various disadvantages associated with these types of negotiation approaches have been identified [69,157]:

According to Jennings et al. [157], the proposals made in the course of a negotiation process generally denote a single point in the space of negotiation agreements, as shown in Fig. 3 where X and O represent single points in space. In addition, the authors note that the only feedback that can be made to a proposal through game theory and heuristic approaches is a counter-proposal – which itself is another point in the space – or an acceptance or withdrawal.

In this respect, Wooldridge observes that positions in a negotiation process cannot be justified using game-theoretic approaches [69]. He argues that humans engaged in a negotiation process tend to justify their negotiation stances. In this context, the author gives the example of a person who attempts to sell his car to somebody else. He explains that the seller may justify the price of the car with respect to a set of features – say a particularly powerful engine. He adds that the buyer may, in response, justify his proposal for a lower price by pointing out that he needs the car for short inner-city journeys, making the use of a powerful engine less useful.

On the other hand, Wooldridge outlines another limitation associated with game-theoretic techniques [69]. He argues that using this technique may make it very hard to understand how an agreement has been reached. According to him, this issue is particularly important in case a person wants to delegate specific tasks such as buying and selling to particular agents. In this respect, Wooldridge provides an example of someone who delegates the task of buying a car to his agent. He explains that once the agent purchases the car, the person concerned would, reasonably enough, be eager to know how the agreement was reached. He would thus ask for clarification on the details of this agreement. Without such clarification, the person who delegated the task of buying the car to the agent may however find the agreement rather hard to accept. In this context, Wooldridge finally outlines the need for people to be able in such scenarios to trust and relate to the decisions made by other agents acting on their behalf.

Another limitation associated with game-theoretic and heuristic approaches lies in the difficulty of changing the set of issues under negotiation in the course of a negotiation process [157]. As pointed out, changing the issues under negotiation corresponds to changing the negotiation space of Figure 7 by adding new dimensions.

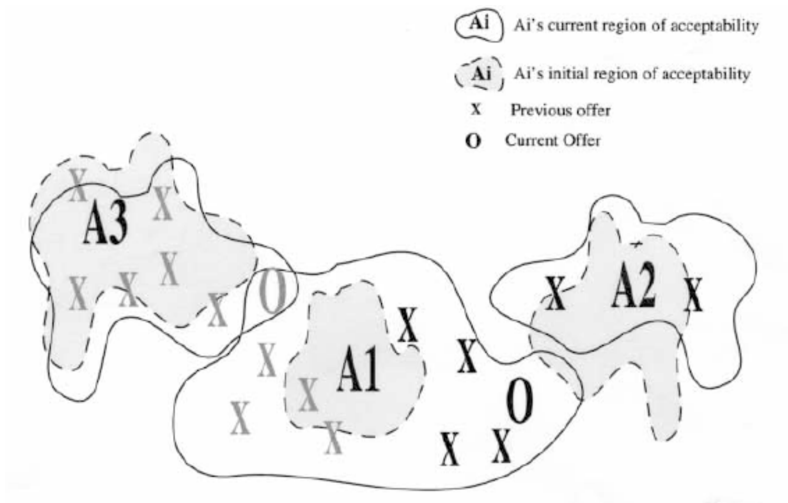

Figure 7. The Space of negotiation agreements [157]

On the other hand, Wooldridge notes that game theory tends to assume that a utility function associated with an agent is fixed and immutable [69]. He points out that an objective, external, omniscient observer can indeed consider an agent's utility function as fixed during negotiation. However, he further adds that from a subjective, personal point of view, someone's preferences certainly do change as he negotiates. Bringing back the car-buying example, Wooldridge argues that a person setting out to buy a car may initially desire a car with an electric sun roof. He then adds that this person might however well change his preferences if he learns that electric sun roofs are unreliable and tend to leak.

Argumentation-based negotiation has thus emerged as an alternative approach designed to overcome the limitations discussed above [69,48,157,159]. The basic idea behind this approach is to allow "additional information to be exchanged over and above proposals [157]. As Wooldridge points out, argumentation in a multi-agent context can be crudely assimilated to "a process by which one agent attempts to convince another of the truth or falsity of some state of affairs" [69]. As also noted, this process involves agents using arguments for and against propositions, together with justifications regarding the acceptability of these arguments.

Indeed, the information that can thus be exchanged over and above proposals conveys, in a number of different forms, arguments that explain explicitly the opinion of the agent making the argument. As Jennings et al. explain, an agent can, in addition to rejecting a proposal, offer an explanation on why the proposal is unacceptable [157]. Such explanation results in the identification of "an entire area of the negotiation space" as inappropriate for the other agent to explore. On the other hand, an agent can similarly provide an argument along with a given proposal in an attempt to convince the other agent to accept it. Through such argument, it would thus be possible for one agent to alter the preferences of the other agent, change the region of acceptability of the latter, and provide a means of modifying the negotiation space itself. In other terms, it is shown that without the ability to argue for the worth of a new element in the negotiation object, the receiving agent would not, in general, have any basis on which to determine its value.

According to Jennings et al. [157], this type of persuasive argumentation does not have to be directly tied to proposals. While recalling that it is possible to make threats or offer rewards in human argumentation, the same authors note that these kinds of arguments can be captured in this approach [160]. When generating the arguments, the agents may not, as in human argumentation, be truthful. Thus, based on its own perception of the argument's degree of credibility, the recipient has to assess the argument on its own merits and then provide the required modification so as to work out an adequate response.

Considering argumentation as it occurs between humans, Gilbert [161] suggests that arguments be classified into four different modes:

1. Logical mode – The logical mode is the kind of argument that is generally employed in "courts of law and scientific papers". Wooldridge suggests that this mode is perhaps the paradigm example of argumentation [69]. While it resembles mathematical proof, it tends to be deductive in nature.

2. Emotional mode – The emotional mode of argumentation is indeed recognized through arguments that appeal to feelings, attitudes, and the like.

3. Visceral mode – The visceral mode of argumentation allows a participant to make arguments stemming from the physical or social human aspect. For instance, this mode occurs when a participant bangs his fist down firmly on a table as a way to emphasize the strength of his feeling.

4. Kisceral mode – The kisceral mode – in Japanese the term "ki" means "energy" – covers the intuitive, mystical, religious, and non-sensory arenas.

Wooldridge reports that people might not be inclined to accept some modes of argumentation [69]. In this respect, he provides the example of courts of law in most western societies where the emotional and kisceral modes of argumentation are not permitted. He further notes that lawyers, of course, still attempt to use such arguments, but one of the roles of a judge is to declare them unacceptable. In contrast, other societies explicitly allow people to refer to religious beliefs in legal settings. While arguments based on emotions are ruled out in a court of law, people are, as noted, inclined to permit such arguments with children or spouse.

6.3.1 Logic-based argumentation

Indeed, Wooldridge suggests that the logical mode of argumentation be regarded as the “purest” or “most rational” kind of argument [69]. In “Reaching Arguments” [69], he introduces a system of argumentation inspired from research work attributed to Fox et al. [162] and Kraus et al. [163]. This system functions through the construction of a series of logical steps or arguments against or in support of propositions of interest. Already used by Parsons and Jennings [158] as a basis for negotiation, the system introduced represents, as stated, a promising option for building a framework for dialectic argumentation through which agents can negotiate.

In an extensive review of logical approaches to argumentation [68], Kraus reports that several researchers developed frameworks for negotiation through argumentation. As noted, these frameworks allow agents to exchange proposals and counter-proposals backed by arguments that summarize the reasons for accepting the proposal. In addition, the author further states that this type of argumentation is persuasive since the exchanges are able to alter the mental state of the participant – the agent involved.

Indeed, most of the frameworks referred to above are based on logical models representing the mental states – such as beliefs, desires, intentions, and goals – of the agents participating in the negotiation-based exchanges. A logic-based model of argumentation has been used in argumentation-based negotiation systems [158,159]. The approach is designed to allow, for example, a given agent to negotiate with a peer over which agent will carry out a particular task. As Wooldridge points out [68], the idea is to bring an agent to argue for the other agent intending to perform the task in an attempt to convince it of the acceptability of the argument that it should carry out the task.

In an attempt to provide an adequate treatment of the collaborative behaviour exhibited in dialogues, a formal model is developed to deal with collaboration [164]. As stated, this model is used as a specification for agent design to constrain certain planning and negotiation processes. On the other hand, the formal model can also be used either by the agents or as a way to check the agents’ behaviour.

In a complex, dynamic multi-agent setting, the coherence of collective actions from agents is often harmed in the presence of conflicting beliefs about different aspects of the environment, resource availability, and individual or collective capabilities and performance. In this context, team members are thus forced to communicate and negotiate to restore team coherence. Indeed, numerous research projects on teamwork have made progress in enabling agents to coordinate and communicate in a flexible manner without addressing the problem of inter-agent negotiation to resolve conflicts. However, Zhun Qiu et al. focus on the problem of negotiation in teamwork to resolve such conflicts [165]. As noted, the basis of such negotiation lies in inter-agent argumentation where agents assert their beliefs to others with supporting arguments. According to the authors, this work introduces a key novelty which consists in using agents' argumentation that exploits previous research on general, explicit teamwork models. Through such models, it is possible, as Zhun Qiu et al. argue, to categorize the conflicts that arise into different classes, and more importantly provide a generalized and reusable argumentation facility based on teamwork constraints. This approach to resolving agents' conflicting beliefs is effectively implemented into the Collaborative Negotiation System based on Argumentation (CONSA).

Indeed, as self-motivated agents in a multi-agent environment are inclined to pursue their own goals, cooperation is indeed organized and achieved through communication and negotiation. Often, negotiation involves argumentation in the form of an exchange of "messages or a dialogue". In "Reaching Agreement through Argumentation: a Logic Model and Implementation" [160], Kraus et al. develop a formal logic that forms a basis for the development of a formal "axiomatization" system for argumentation. In addition, the authors present a logic model of the mental states of the agents based on a representation of their beliefs, desires, intentions, and goals. Using categories identified in human multi-agent negotiation, the authors demonstrate how the logic can be used both to specify argument formulation and evaluation and describe different types of agents. In this context, Kraus et al. present argumentation as an "iterative process emerging from exchanges among agents to persuade each other and bring about a change in intentions". As pointed out, argumentation is seen as a mechanism that can be implemented to achieve cooperation and agreement.

In addition, Kraus et al. further introduce a general Automated Negotiation Agent based on the logical model mentioned above [160]. The Automated Negotiation Agent is implemented to act and negotiate in a simulated multi-agent environment. As stated, several Automated Negotiation Agents can thus be defined and created. The resulting simulation system infrastructure allows a user to analyze and explore different methods for negotiating and arguing within a non-cooperative environment that lacks a centralized coordination mechanism. An example of agents that are able to plan, act, and resolve conflicts via negotiation in a Blocks World environment is used to illustrate the development of negotiating agents in the framework of the Automated Negotiation Agent.

As proposed by Kraus et al. [160], the formal modal consists of "a set of agents, not necessarily cooperative, with the ability to exchange messages". Concepts such as beliefs, goals, desires, intentions, and local preferences are indeed incorporated in the formal models so as to characterize the mental states of these agents. Having a set of desires, each agent develops and implements activities motivated by the will to fulfill these desires. Thus, at a given time, a participating agent selects a consistent and suitable subset of desires that serve as its peculiar set of current goals. Preferring to fulfill goals of higher importance, the agent “ascribes different degrees of importance to different goals”. This exercise results in a set of properly prioritized goals that motivate and actively direct the agent’s planning process.

According to Kraus, the planning process may generate several intentions [68]. Referring to actions that are within the direct control of the agent, some intentions are classified as the “*intent-to-do*” category. On the hand, classified in the “*intent-that*” category [164,169,170-171], other intentions represent propositions that are not directly within the agent’s realm of control. In order to fulfill or satisfy such propositions, the agent must, in this case, rely on other agents. However, in the presence of “intent-that” actions in a plan, there is often room for argumentation, as Kraus reports [68].

Argumentation is thus the process through which one agent – the persuader – attempts to modify the intention structure of another agent – the “persuadee” – so as to convince it to include the actions it wants it to achieve. In the course of a negotiation process, the persuader and the persuadee are attributed or assigned a dynamic rather than a static role. Thus, each participating agent may, once it receives a message from another agent, update its own intentions and goals. Furthermore, the agent dispatching an unsuccessful argumentation is then required “to revise its arguments, its plans”, and/or seek other sources of satisfying the portion of its plan in question” [68].

A belief set peculiar to an agent includes beliefs concerning the mental states of other agents and beliefs related to the world. Associated with each agent, the mental model of other agents is thus used to determine the actions a particular agent undertakes. As Kraus reports [68], an agent may indeed be mistaken in both kinds of beliefs. As a result, it may update its beliefs following the reception of messages from other agents and through the observation of the world.

Kraus argues that arguments offer the opportunity either to add an intention to the belief set of a persuadee, change the latter’s preferences, or retract an intention [68]. Without being exhaustive, she then attempts to present a list of several argument types. However, presenting an authoritative classification proved to be difficult to achieve, since arguments are both liable to interpretation and effective within a particular context and domain [166]. Given this consideration, Kraus proposes a list of arguments loosely classified into six categories. The argument types presented are commonly thought to have persuasive force in human negotiations [62,168,172]. The author suggests that argumentations which proved to be successful in human negotiation may also be successful in automated inter-agent negotiation. In this respect, Kraus further expresses the willingness about the need to provide agents with the ability to negotiate with humans. As noted, these agents are therefore required to at least understand human argumentation.

Finally, the arguments classified in distinct categories are presented as follows [68]:

1. Arguments containing threats designed to lead a persuadee to adopt or abandon a goal.

2. Arguments accompanied by a promise of a future reward intended to entice or incite a persuadee.

3. Arguments in the form of an appeal to past reward.

4. Arguments appealing to precedents as counterexamples to convey to the persuadee a contradiction between what she/he says and past actions.

5. Arguments referring to “prevailing practice” as an appeal designed to convey to a persuadee that the proposed action will further his/her goals as it has helped others to successfully achieve their goals in the past.

6. Arguments appealing to self-interest in order to convince a persuadee that taking the proposed action is susceptible of enabling the achievement of a high-importance goal.

Indeed, the most common arguments used in human negotiation contain threats and promises as a strategy of persuasion [167]. On the other hand, it is reported that the appeal to prevailing practice represents the most common argument used in the legal system. In this respect, presenting example instances – in other terms prevailing practice cases – appears to be much more persuasive than presenting statistical summaries [173-176]. Assuming that a person will be willing to keep his/her promises as a way to maximize the internal consistency of his/her cognition, the cognitive dissonance theory supports arguments based on an “appeal to past promise” [168]. In repeated interactions, arguments appealing to past promise are indeed important since agents prefer to preserve their established credibility. Finally, bounded rational agents with limited inferential resources are likely or prefer to use the types of arguments in the form of “an appeal to self-interest” and in the form of “a counter example” [160].

6.3.2 Dialogues and dialogue Systems

Wooldridge points out that many authors are interested in studying agents “that argue with themselves, either to resolve inconsistencies or else to determine which set of assumptions to adopt” [69]. While recalling that an agent engages in a dialogue in order to convince another agent of some state of affairs, he reports that, in contrast, his interest lies in the study of agents that are involved in a dialogue with other agents.

Table 3.* : *Walton and Krabbe’s classification of dialogues [69,180,182].

Type	*Initial Situation*	*Main goal*	*Participants aim*
Persuasion	Conflict of opinion	Resolve the issue	Persuade the other
Negotiation	Conflict of interest	Make a deal	Get the best for oneself
Inquiry	General ignorance	Growth of knowledge	Find a ‘proof’
Deliberation	Need for action	Reach a decision	Influence outcome
Information seeking	Personal ignorance	Spread knowledge	Gain/pass on personal knowledge
Eristics	Conflict/antagonism	Reaching an accommodation	Strike the other party
Mixed	Various	Various	Various

In “Reaching Agreement” [69], Wooldridge defines the notion of dialogue and investigates the concept of winning an argument. On the other hand, in “Agents that Reason and Negotiate by Arguing” [159], Parsons and al. present a well-grounded framework for describing the reasoning process of negotiating agents that is based upon the use of argumentation both at the level of an agent’s internal reasoning and at the level of negotiation between agents. In addition, they further provide a concrete example of how a dialogue – a series of arguments – can be used to solve negotiation problems.

As mentioned in various studies pertaining to argumentation-based dialogue [69,159,177,178,180,181], Walton and Krabbe identify a set of dialogues classified into different categories [182]. Considered as an important contribution, the proposed typology is based exclusively on dialogue encounters between human interlocutors [181]. Although it is not exhaustive, this classification is, according to Chris Reed [181], both formal and robust to withstand application to MAS domain.

Indeed, Walton and Krabbe introduce a classification that distinguishes seven types of dialogue summarized in Table 3. As noted, ‘eristics’ represents a category of dialogue that “serves primarily as a substitute for physical fighting” [181]. As such, it is thus unlikely that it plays a significant role in current research pertaining to MAS. However, Wooldridge explains that an ‘eristic’ dialogue may result from the need to air in public a conflict between agents [69]. According to him, the objective of such dialogue may be to reach an accommodation, but need not be. On the other hand, the remaining types of dialogue are characterized through:

1. the initial situation – in particular, the absence or absence of conflicts,

2. the private objectives of the participating agents, and

3. the joint intents or purposes to which all participating agents “implicitly subscribe”.

Initiated from a position of recognized conflict – both agents are mutually aware of the situation, a ***persuasion dialogue*** is conducted through a sequence of argumentation where “the elements at stake” are primarily a particular set of beliefs. In this context, each agent adhering to a system of beliefs is motivated by the objective of offering compelling arguments in support of its thesis. Entering a persuasion dialogue implies the need to implicitly agree to the joint aim of resolving a particular conflict. Hence, a successful persuasion dialogue requires that each participating agent had beforehand prepared itself to alter its beliefs with respect to its thesis. In this respect, an agent participating in a persuasion dialogue can therefore “assume that the other agent is at least prepared to alter its beliefs”.

In contrast, a ***negotiation dialogue*** is an argumentation process that “involves a utility”. From another point of view, this type of dialogue arises from a situation characterized by a real, potential or perceived “conflict of interest”. It may thus involve, for instance, an attempt to reach agreement either over scarce resources or on a division of labour between agents. Sharing the common and strong overarching aim of reaching a deal, selfish agents adopt goals aiming at maximizing “their own share of any resource”. The goal of reaching a deal is not highly important in many multi-agent scenarios, as any agent may be able to negotiate and strike a better deal with a third party [181]. However, there are multi-agent scenarios where agents are faced with extremely limited options. In this case, sharing a common desire to reach a deal can be considered as or becomes an important factor affecting the progress and outcome of the negotiation process. In addition to the contextual differences, another important difference between negotiation and persuasion refers to the procedure or way of doing things. As Chris Reed points out, coherence between beliefs is not required and the relevant beliefs of the participating agents may very well remain at odds after negotiation [181].

In an ***inquiry dialogue***, participating agents engage in a process of argumentation pertaining to a matter of common interest, where the object of the inquiry is "a belief". Resulting from a perceived lack of knowledge, the inquiry dialogue thus provides evident differences with the negotiation and persuasion dialogues, in that the argumentation process is not based on any conflict. Rather than challenging any still existing knowledge, the participants would instead try to establish "truth or falsity". In this respect, Wooldridge provides the best-known example of an inquiry: a public inquest into some event such as a train crash [69]. He notices that the aim of an inquiry is simply to determine facts, as it takes place when a group of people have some mutual interest in determining something.

A ***deliberation dialogue*** represents the process through which a group of participants attempt to form and decide upon a course of action. This type of dialogue expresses shared similarities with negotiation: the joint goal is to reach an agreement and the individual goal associated with each participating agent is to influence this agreement to suit its benefit [181]. While the deliberation dialogue results from the need for action, negotiation arises from a situation of conflict. Rather than displaying the proposal-counterproposal associated with a negotiation process, the deliberation dialogue is thus usually oriented around means-ends discussion.

Closely related to an inquiry, an ***information-seeking dialogue*** occurs as a result of an agent taking the individual necessary step "to find out something for itself". Chris Reed argues that the information-seeking dialogue is initiated to deal with a recognized asymmetry between one agent having more information that another "in regard to some particular data" [181]. He then outlines the agreement between the joint aim and the individual goals of the agents. Finally, a ***mixed dialogue*** occurs through the combination of different types of dialogues discussed above. In this respect, Wooldridge recalls that committee meetings involve negotiation, deliberation, inquiry, and, frequently, eristic dialogue [69].

In "Dialogue Frames in Agent Communication" [181], Chris Reed presents an initial work designed at introducing rich models developed in informal logic into multi-agent settings. This work provides a formal characterization of inter-agent communication which clearly distinguishes persuasion from negotiation. It further introduces the notion of a dialogue frame that is used to explore the dialogue typology and the concept of functional embedding. As noted, the approach demonstrated that "the typology designed on the basis of empirical research into natural argumentation can be successfully formalized and applied in a multi-agent domain”.

Furthermore, in an attempt to advance knowledge in the area pertaining to the use of argumentation techniques as a basis for negotiation dialogues, Amgoud et al. present a model for inter-agent dialogues based on argumentation [180]. It is a general model incorporating a precisely defined protocol for the exchange of arguments. As explained, this model extends the general approach suggested by Chris Reed as a way to capture a range of dialogue types [181]. In contrast, S. Parsons et al. concentrate more on the interaction between beliefs and intentions in a specific negotiation/deliberation form of dialogue [159]. Introduced by Amgoud et al. [180], the general model provides an underlying argumentation system and the illocutions necessary to carry out the kinds of dialogues discussed in [181]. As suggested, this approach can thus be regarded as an attempt to bridge the gap between the low-level detail of handling beliefs and intentions described in [159] and the general approach discussed in [181].

In “An Analysis of Formal Inter-agent Dialogues” [177], Parsons et al. explore three types of argumentation-based dialogues between agents – *information seeking*, *inquiry*, and *persuasion* –, define a precise protocol for each type of dialogue, and examine various important properties of the resulting protocols. While also considering some aspects of the complexity associated with these dialogues, the authors show in particular that each protocol leads to dialogues that are guaranteed to terminate. In a follow-on work [178], Parsons et al. extend this analysis of formal inter-agent dialogues to provide a first detailed characterization of the outcome of such dialogues. Then, they investigate the extent to which outcomes are dependent “on tactical play” by the participating agents. The results of this investigation show that tactics can have an important effect on the outcome. In this respect, Parsons et al. identify how to rule out the effect of tactics. According to them, excluding the effect of tactics is desirable from the perspective of mechanism design.

7. Coordination in distributed problem solving and planning

The concepts and algorithms that establish the foundations of distributed problem solving and planning are addressed by Durfee in [17]. In this respect, the author devotes a great deal of effort to outlining through practical examples how protocols of interaction can be used to ensure that appropriate information is conveyed to the right agents to adequately accomplish problem solving and planning.

Durfee refers to distributed problem solving as a subfield of distributed AI where efforts are directed towards “getting agents to work together well to solve problems that require collective effort”. While underlying the inherent distribution of resources such as knowledge, capability, information, and expertise among agents in a distributed problem-solving system, the author argues that an agent incorporated in such an environment is unable to accomplish its own tasks alone, or at least can accomplish its tasks through collective action.

Indeed, solving distributed problems requires both **coherence** – a common property of multiple agents associated with the need to have the incentive to work together – and **competence** – an attribute of multiple agents related to the ability to work together well. He then points out that group coherence is both hard to realize among individually-motivated agents and essential for multiple agents to solve a problem in a distributed way [17]. Considering that agents are designed to work together in distributed problem solving, Durfee further notes that a fair degree of coherence is thus already present in such environments. In addition, he attributes this coherence either to the implementation of some sort of stimulating payoffs to self-interested agents that are only accrued via collective efforts or to the introduction of disincentives for agent individualism. Finally, Durfee thus concludes that distributed problem solving concentrates on competence to provide the ability to tackle problems that need to be solved so as to meet expectations about what constitutes viable solutions.

In “Trends in Cooperative Distributed Problem Solving”, Durfee et al. compare the problem of designing a DPS system to the problem of building a house [183]. This process obviously involves several expert agents – each expert in a different area: one agent may be an expert in the strength of materials, another in the required space for different types of rooms, the next in plumbing, the other in electrical wiring, etc. As this task is beyond the individual capability of agents, building the house requires that agents cooperate.

In this respect, partial results have to be exchanged – *the architect passes the blueprint to a construction company*, tasks may be delegated from one agent to another – *the construction company engages subcontractors*, information related to the state of the problem-solving process has to be exchanged – *the plumber will start his work only when the building workers notify they have finished the required task*, and, finally, unforeseen conflicts have to be solved – *the same plumber may have to negotiate with a building worker on how to solve a given conflict*. Indeed, the above house-building example highlights the boundary delimitation between DPS and ordinary distributed systems.

According to Ossowski, an ordinary distributed system is designed to solve "well-defined problems in the frame of fixed control structures and rigid problem decomposition" [184]. As stated, this type of system is mainly adapted to or concerned with issues such as "synchronization and technical aspects of communication". On the other hand, DPS tackle complex tasks where diverse aspects cannot be totally anticipated beforehand. In this type of systems, the focus lies on aspects of cooperation within a dynamic problem-solving process. In this context, Ossowski observes that DPS tackles the issue of how the problem-solving activities of agents can be coordinated to successfully deal with the overall task. This issue includes many topics and questions related to finding and designating the suitable agent to perform a given task, defining the appropriate partial results to be exchanged, and determining the right time for communication to take place.

7.1 Properties of DPS systems

In "Distributed Artificial Intelligence" [184], Ossowski points to the problem of using the vaguely defined term "***cooperation***" when tackling aspects related to DPS. In this respect, he reports that various attempts have recently been made to characterize the area of DPS without the need to refer to the ambiguous notion of cooperation. However, Durfee and Rosenschein define and propose various properties designed to enable the identification of DPS systems [10]. Defined in the form of assumptions likely to be made with respect to DPS systems, these properties include:

1. The benevolent assumption. – This assumption proposed as a standard property required to determine whether a distributed architecture is a DPS system implies that the agents in the system are benevolent – *the agents help each other whenever they are requested to do so* –. While Sascha Ossowski notes that benevolent agents slavishly "do what they are told" [184], Durfee and Rosenschein clarify the meaning of benevolence by clearly stating that benevolent agents help each other whenever possible [10]. In groups of benevolent agents, any one entity might request others to do something without having to put forward the possibility of a reward or any kind of compensation. As a result, the benevolent assumption does not imply the existence or the necessity to implement a sort of authority structure or any other kind of organization between agents.

When using the Contract Net Protocol [35], an agent first determines whether the task it should perform can be broken into subtasks, then announces the resulting tasks that could be transferred and requests bids from agents that could perform any of these subtasks. In this context, tasks are allocated based on suitability and availability of individual entities, without any sense of an agent asking whether it wants or has to perform a task for this or the other agent. Upon hearing the task announcement, an eligible agent will make an honest bid with a reference to his ability to perform this task. In the end, the agent making the best bid will be awarded the task. As already stated, in the absence of any sense of reward, agents do not need to be bribed or otherwise persuaded to undertake tasks that other agents want them to do.

However, Durfee and Rosenschein argue that the benevolent assumption is not a guarantee of cooperation and coherent coordination [10]. Even with agents willing to help each other, difficulties associated with timing and local perspectives can, as stated, lead to uncooperative and uncoordinated activities. To support this view, the authors provide the example of the Contract Net protocol where, according to them, important tasks could go unclaimed when suitable agents are busy with tasks that others could have performed. While referring to the Contract Net protocol, Durfee and Rosenschein also put forward the argument of a more general case where tasks could be improperly assigned so as to result in redundant or incompatible agents' activities.

2. The common goals assumption. – Pursuing a common goal indeed represents a motivation for benevolence among agents sharing a common environment. Engaged in a group activity, these agents will thus attempt, in whatever way it is possible, to maximize "the same global utility measure". In this context, conflicts between agents are, according to Sascha Ossowski, essentially subjective [184]. Indeed, rather than arising from objective contradictory interests, such conflicts are due to "incomplete or incorrect views of the state of the world".

On the other hand, the common goals assumption is considered to lie at the heart of the Contract Net protocol [10]. As argued, this assumption is also at the core of cooperation in inherently distributed tasks where agents value "the development of a global result". Indeed, the Distributed Vehicle Monitoring Testbed (DVMT) simulates a number of geographically distributed sensor nodes that share the system-wide goal of integrating the nodes-associated local maps into a global map of vehicle movements [33]. In working together to achieve the convergence of the distributed system towards the global solution, each node is indeed helping the others make "good local interpretations as quickly as possible".

However, as Durfee and Rosenschein point out, local views on the problem to be solved – such as the above DVMT example – may indeed lead to globally incoherent local decisions [10]. Thus, in the absence of strict guidelines about responsibilities or interests, agents can "inundate each other with superfluous information". A more negative aspect lies in the fact that agents can, through information they send, distract other entities into pursuing unimportant tasks.

Finally, Durfee and Rosenschein highlight the absence of clarity regarding the level at which goals should be common so that a given system can be considered a DPS system [10]. In this respect, the authors imagine an environment where agents are meeting to hold a competition. Within this environment, the agents might, as explained, share a high-level goal of holding the competition, while having opposing goals regarding which agent is supposed to win the competition. Indeed, the PERSUADER system provides an interesting example of a distributed environment where agents, representing opposing sides in a labor contract, share a goal of reaching an agreement while having diverse preferences in rating candidate contracts [185]. As far as the PERSUADER system is concerned, the question is raised as to whether this system falls within the DPS category [10].

3. The centralized designer assumption. – This assumption refers to the argument according to which a DPS system is a system with a centralized designer. Concerned about getting all the parts of the system to perform effectively as a whole, the central designer would likely design and implement agents having the required attribute of benevolence. With the task of building a DPS system, the central designer would thus both calibrate the preferences associated with the agents and standardize the mechanisms required to express and act on these preferences.

As in the case of the common goals assumption, Durfee and Rosenschein raise the open question regarding the detail to which the common designer must specify each agent's design so as to incorporate all the agents into a DPS system [10]. Finally, the authors argue that the approach which consists in categorizing systems will likely remain arbitrary until aspects of agents that need to be dictated by a central designer to make an environment a true DPS system are exactly defined.

However, Ossowski refers to the third property as the ***homogeneous agents assumption*** [184]. In this particular context, homogeneity is defined through the absence of "unnecessary" conflict or incongruity between agents. According to Ossowski, homogeneous agents would thus use compatible representation and communication languages, without pursuing conflicting goals.

7.2 Applications, examples and motivations

Durfee identifies various technical motivations that stimulate and induce the development of both distributed problem solving and distributed planning [17]. At the same time, he enumerates examples of the kinds of applications through which these motivating factors can be clearly recognized and understood.

According to Durfee, an obvious motivation lies in the need to use distributed resources concurrently to allow a speedup of problem solving, thanks to the advantages conferred by parallelism. As noted, any possible improvements that can be achieved through parallelism obviously depend on the degree of parallelism inherent to the problem.

In this context, the ***Tower of Hanoi*** (ToH) puzzle is, as shown in Figure 8, described by Durfee as a classical toy problem that allows a large amount of parallelism during planning. The author reports that the puzzle is well known to students, since it appears in virtually any introductory AI course. He then explains that the ToH consists of a board supporting three pegs and n disks of different sizes initially placed on one peg with the largest one on the bottom and the other disks following in order of size. Under the condition that the disks must be moved one at a time, without ever placing a larger disk on top of a smaller disk, the problem, then, is to find an appropriate sequence of moves that will achieve the goal state: move the disks from the start peg to another peg.

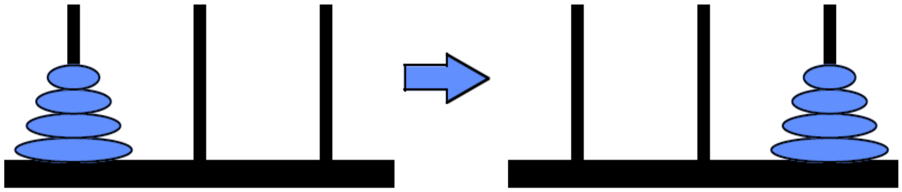

Figure 8. The Tower of Hanoi [17]

According to Durfee, expertise or other problem-solving abilities may be inherently distributed to provide a motivation for distributed problem solving and planning. In this respect, he provides a practical example of a problem pertaining to concurrent engineering where a team of specialized agents could be involved in the design and manufacture of an artifact such as a car. As stated, this problem is characterized by the need to allow the agents to individually formulate components and processes, and combine these into a collective solution.

On the other hand, Durfee mentions other examples of applications where different problem-solving abilities are inherently distributed, such as supervisory systems for air-traffic control, factory automation, or crisis management. He explains that these applications can involve an interaction between tasks pertaining to event monitoring, situation assessment, diagnosis, prioritization, and response generation. Associated with these kinds of systems is, according to him, the need to employ diverse capabilities to solve both large and multi-faceted problems. However, Durfee prefers to use an application problem commonly known as ***distributed sensor network establishment*** (DSNE) – an application designed to provide the ability to monitor a large area for vehicle movements – as an example to further address strategies employed in the area of distributed problem solving and planning. The author then argues that the overall monitoring task cannot be accomplished through a central location, since a large area cannot be sensed from a single location. Thus, the problem in DSNE consists, according to him, in decomposing the large monitoring task into subtasks that can be appropriately allocated to geographically distributed agents.

Related to the second, the third motivation introduced by Durfee is based on the fact that beliefs and other data can be distributed. In this respect, the author further provides an example of application where the problem is to actually conduct ***distributed vehicle monitoring*** (DVM) through a centralized mechanism. He points out that the procedure could be designed to allow each of the distributed sensor agents to transmit raw data to a central site to be interpreted into a global view. However, this centralized mechanism could, according to him, involve unnecessary communication compared to a strategy where separate sensor agents could formulate local interpretations that could then be transmitted selectively.

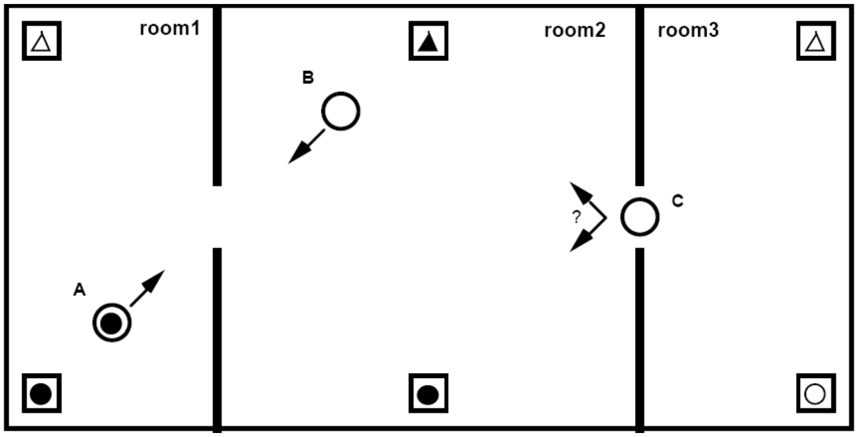

Figure 9. Distributed delivery example [17]

A fourth motivation is finally considered and presented by Durfee as a possible case where the results of problem solving and planning might be distributed to be acted on by multiple agents. This possibility is then detailed through an example of a ***distributed delivery*** (DD) system – see Figure 9 – where agents can act in parallel to conduct the delivery of objects between locations. In this context, Durfee suggests that plans development could either be undertaken at a centralized site through a dispatcher or involve distributed problem-solving among agents. According to the author, features of the environment that unexpectedly change or were unknown earlier during planning can, during the execution phase, trigger changes in what the agents are supposed to do. In this respect, Durfee observes that all these decisions could be routed through a central coordinator. He further explains that it could be preferable for a variety of reasons that agents modify their plans unilaterally or with limited inter-agent communication.

Following the identification and brief description of the motivations underlying the development of applications in distributed problem solving and planning, various classes of distributed problem-solving strategies are then addressed in a detailed and very structured and helpful way.

7.3 Task sharing

Task sharing, also known in the literature as ***task passing***, represents the first class of distributed problem-solving strategies that Durfee considers. As reported, the idea behind this strategy consists in allowing an agent that has many tasks to enlist the help of other entities that have few or no tasks. The task-sharing strategy is indeed presented as a four-step process that includes:

1. ***Task decomposition*** – This is the initial step characterized by the need to generate a set of tasks susceptible of being attributed to other agents. In general, it consists in decomposing large and complex tasks into subtasks that could be undertaken by different agents.

2. ***Task allocation*** – This crucial step in the process consists in assigning the subtasks to appropriate agents.
3. ***Task accomplishment*** – The purpose of this step is to allow the agents to adequately accomplish the allocated subtasks. This process could eventually lead to further decomposition of the subtasks into smaller subtasks so as to allow each agent to accomplish alone the tasks it is handed.
4. ***Result synthesis*** – As a final step in this process, each agent that accomplishes its subtask is instructed to pass the result to the appropriate agent. The latter is usually the original agent which is behind the decomposition process and is thus likely to effectively and adequately compose the results into a complete overall solution.

In conclusion, Durfee notes that different steps might, under different circumstances, be more or less difficult. In this respect, he gives the example of an overburdened agent that begins with a bundle of separate tasks, making the decomposition of tasks into subtasks an unnecessary step. He also mentions the effect on the allocation process resulting from the case of the agent that passes tasks off to any of a number of identical agents. He finally points out that accomplishing the tasks does not always yield results that need to be synthesized in any complex way.

7.3.1 Task sharing among homogeneous agents

Durfee uses the example of the ToH problem described above in order to highlight the possibilities of task sharing. Concomitantly, he provides a thorough understanding of how this strategy can be implemented among homogeneous agents. Considering the task-sharing four-step process outlined above, the author describes how this very simple puzzle can be solved:

1. Task decomposition – Tackling the first step of the process through means-ends analysis leads to a recursive decomposition. It first amounts to solving the problem of getting to the state where the largest disk can be moved, and once it is moved, the next approach is to eventually determine how to get from this intermediary state to the goal state. Resulting from the ToH puzzle are thus a collection of sub-problems that can indeed be further decomposed into problems corresponding to the need to move the second largest disk to the middle peg so as to get it out of the way, etc.

2. Task allocation – Assuming an indefinite number of identical idle agents having enough skills and knowledge to solve the ToH problem, the allocation is thus reduced to a simple process of randomly assigning a task to one of these idle agents. On the other hand, the recursive decomposition terminates if the decomposed problems are such that the start and goal states are the same.

3. Task accomplishment – In general, using means-ends analysis, an agent can find the most significant difference between the start and goal states for which it is responsible. It can further decompose the problem based on this analysis.

4. Results synthesis – Once an agent has solved a problem, it then proceeds to pass the solutions back on up. In addition, the agent that has received solutions to all of the sub-problems it passed down can subsequently compose these into a single solution corresponding to a more comprehensive sequence of moves. This sequence is then passed up by the agent as its own solution.

Due to its hierarchical nature, the ToH problem represents, as stated, an ideal case for demonstrating the possibilities of distributed problem solving. However, it is worthwhile noting that most frequently encountered problems are more complicated than the ToH problem. In "Distributed Problem Solving and Planning" [186], Durfee enumerates the reasons for which most real-world problems are more complicated:

1. As the solution of one problem can affect the solution of others, problems will often "require backtracking back upward in the abstraction hierarchy".

2. Rather than scaling up the number of agents with the problem size, the number of available agents is generally maintained at a fixed level.

3. In various domains problems are decomposed into qualitatively different sub-problems that require different expertise. As a result, problems to be solved cannot usually be decomposed into equal-sized sub-problems.

4. Substantial time is often required to conduct the process of decomposing problems, distributing sub-problems, and collecting results.

7.3.2 Task sharing in heterogeneous systems

In the ToH problem tackled above, the resulting sub-problems requires agents with identical skills and knowledge – competencies. As a result, the decisions regarding the assignment of tasks and sub-problems are "extremely simple". However, the distribution of sub-problems becomes a complicated process if the global problem is decomposed into sub-tasks that require different capabilities from agents having different expertise.

Indeed, it is difficult to train humans or build artificial entities that can handle every possible task. In this respect, Durfee points out [186] that building or training an omni-expert agent would result in an obvious waste of most of its capabilities. As in human-based systems, the strategy adopted in many distributed problem-solving systems consists in bringing together on demand combinations of experts from different areas [17,186]. These agents can thus combine their skills and knowledge in order to collectively solve problems that are beyond their individual capabilities.

Having a “table” designed to allow the identification of the agents’ skills and knowledge – i.e. capabilities, an agent can indeed simply select a suitable agent to which it can also assign and send an appropriate sub-problem. However, Durfee notes that the decisions regarding the distribution of tasks are usually based on more “dynamic information” [17,186]. He further provides a typical example of a problem that may arise when an agent attempts to assign and distribute sub-problems to several other candidate agents. In this respect, Edmund H. Durfee raises the question of how this agent can discover which entities among the candidate agents that are capable of solving the sub-problem are already committed to other tasks.

According to Durfee, one way of discovering which candidate agents are already committed to other sub-problems is to use the **Contract Net protocol** with directed contracts or focused addressing [17,186]. Based on the table of capabilities mentioned above, the agent – *called the manager in the Contract Net protocol* – announces a sub-problem to a specific agent – *in directed contracts* – or to a subset of other agents – *in focused addressing* – and requests that bids highlight acceptance and/or availability. Based on the collected returned bids, the manager can then award the sub-problem to the most favourable bid – either to the directed contractor or to one of the focused available contractors. However, it happens that none of the agents is available to respond to the announcement as requested. In this context, the manager has various available options [17,186]:

Broadcasting contracting. – In an open environment, a manager using the Contract Net protocol is unlikely to be familiar with all of the possible contractors. In case first attempts to find an available agent through directed contracts or focused addressing fail, the manager may desire to extend its area of activity and update its knowledge regarding eligible contractors. In this respect, broadcasting the announcement remains the most commonly considered mode of operation in the Contract Net available to the manager in order to reach agents of whose existence and capabilities it is unaware.

Retry. – Another simple strategy available to the manager is to retry the announcement. Assuming that a contractor can eventually free up, the manager can retry the announcement periodically at appropriate intervals. While quick retries may cause the network to bog down with messages, slow retries prevent agents from efficiently using each other.

However, the Contract Net protocol can be reversed in order to overcome this problem. In other words, potential contractors can use the protocol to announce availability and managers can, instead of announcing sub-problems and collecting bids, respond to the announcements through bids of pending tasks.

Announcement revision. – Indeed, the eligibility specifications that can be examined by every potential contractor are contained in the announcement message. In case no contractors respond to an announcement, the manager may have been “too exclusive in whom it would entertain bids from”. In this respect, another option offered to the manager is to relax the eligibility requirements until it begins to receive bids. Hence, the revision of the announcement can indeed be an iterative process.

However, the eligibility specifications could well reflect certain preferences over the quality of services that different contractors can provide. In this case, the manager will otherwise handle the lack of bids by taking a decision concerning “the relative importance of having a preferred contractor eventually pursue the sub-problem compared to finding a sub-optimal contractor sooner”. According to Durfee, these preferences and trade-offs between them can, in many cases, be captured using economic representations [17,186].

Alternative decomposition. – In response to the lack of available contractors, another procedure consists in decomposing the overall problem so as to obtain a set of alternative sub-problems. As pointed out, the relationship between the problem decomposition process and the allocation of tasks is, in general, extremely complex [17,186]. While sometimes the space of decompositions can be very restrictive, the manager should, in other times, first determine the space of alternative contractors so that it can later focus on problem decomposition. In addition, the manager has to make appropriate decisions regarding the number of problems to decompose into sub-problems and the granularity of the resulting tasks. According to Durfee, these decisions depend on other features relating to the application environment [17,186].

7.3.3 Task sharing for distributed sensor net establishment

In “*Negotiation as a Metaphor for Distributed Problem Solving*” [34], Davis and Smith examine the concept of distributed problem solving and explore the use of the Contract Net framework of communication and control to demonstrate the solution of a simulated problem in area surveillance – *ship or air traffic control*. In this context, the authors discussed the mode of operation of a distributed sensing system – the Distributed Sensor Net Establishment (DSNE), a network of nodes scattered throughout a relatively large geographic area. Although the distributed sensor system may have several functions, ranging from passive analysis to active control over vehicle courses and speeds, R. Davis and R. G. Smith focus on the analysis function. While the task involves detection, classification, and tracking of vehicles, the solution of the problem is, as stated, to provide a dynamic map of traffic in the area.

7.3.4 Task sharing for interdependent tasks

Durfee has also looked into the question regarding task-sharing in problems characterized by the presence of interdependent tasks [17,186]. Indeed, problems similar to the ToH are widely acknowledged to generate tasks that can be accomplished independently. In other terms, the sequence of actions required to get from the initial state to an intermediate state can be determined completely separately from the sequence required to move from that intermediate state to the goal state. As noted by Durfee, the sub-problems resulting from the task decomposition process can thus be successfully accomplished in any order or concurrently. Furthermore, the synthesis – the act of passing the result to the appropriate agent – need “only wait to complete until all the subtasks are carried out”.

However, it is obvious that often tasks or subtasks contracted by eligible agents are interdependent. In this respect, Durfee puts forward the example of a concurrent engineering application where “process planning subtasks usually need to wait until product design tasks have progressed beyond a certain point” [17,186]. Consequently, a manager in charge of coordinating the execution of the subtasks can initiate a given subtask based on the progress of another, or, as noted, by communicating to contractors interim results that are required to accomplish other interdependent subtasks.

On the other hand, Durfee recalls the complexity associated with the process of solving assigned interdependent sub-problems [17,186]. Rather than being a priori defined through the process of problem decomposition, aspects related to subtasks interdependencies might, more generally, only become apparent during the course of problem solving. Indeed, the task-sharing strategy can, as discussed above, be implemented through the Contract Net framework to establish, for example, a distributed sensor network. However, an important and challenging question arises on how to discover and exploit exiting interrelationships between different classes of tasks allocated through the task-sharing strategy. While it can effectively be used to establish a distributed sensor network, this strategy does not provide sufficient basis for using the network, since “the runtime relationships between what is being monitored in different areas is as variable as the possible movements of vehicles through the areas”. In this respect, it is interesting to note that the most appropriate approach is to generate and share tentative results so as to be able to accomplish interdependent tasks.

7.4 Result sharing

Given the same task, different problem solvers sharing a distributed environment will separately and independently (likely) derive different solutions – like students that are often given the same homework problem. Derived within the context of each problem solver, the results of the same problem-solving task performed by different problem solvers could thus well differ. Incentives that effectively motivate problem solvers to share results in order to enhance group performance include:

1. Confidence. – Problem solvers can independently accomplish the same task from which results are separately derived. These results obtained independently can indeed help corroborate and/or clarify individual approaches and be used to achieve and provide a collective solution with a higher level of confidence.

2. Completeness. – Results emanating from each problem solver are assembled so that altogether they can form a more complete portion of the overall task.

3. Precision. – Results formulated independently by a set of problem solvers can be used by any agent in order to refine its own solution.

4. Timeliness. – Solving subtasks in parallel offers the ability to optimize the time required to achieve an overall solution, though in principle an agent could solve a large sub-problem alone.

In a distributed problem-solving environment, agents obviously need, in some cases, to share results in order to accomplish subtasks. However, the process of sharing results is not as simple as it may at first sight appear. Emerging from this discussion is an important question on how a particular agent should assimilate shared results into its own solution. In response to this question, Durfee highlights the need for agents "to first know what to do with shared results" [17,186]. On the other hand, the author points to the need for agents to also practice a selective exchange of results. Indeed, communicating large volumes of results can obviously be costly, and managing a large volume of assimilated results "incurs overhead".

7.4.1 Functionally accurate cooperation

In various conventional applications – such as the ToH –, problem-solving agents operate separately and independently without needing to exchange results. In such applications, the problem-solving process amounts to the decomposition of the global task in such a way that each subtask can be performed completely by a single agent, without the need for this particular agent to see the intermediate results associated with the activities of other agents. Indeed, having all information and a complete specification regarding the subtasks assigned to them, these agents are "completely accurate in their computations".

Naturally well-suited to distributed implementation, a number of other applications - such as sensor networks, tasks involving mobile robots, automotive and air-traffic control, etc. – lack the task decomposition characteristics of conventional distributed applications. In these applications, the lack of information makes it difficult to partition the data required to achieve a solution in such a way that an agent can locally complete a subtask without having to examine "the intermediate state of processing" elsewhere in the network. Apparently ill-suited to conventional approaches, such applications are indeed characterized by the lack of information and, hence, uncertainty and errors.

As suggested, an alternative approach is for the processing agents to engage in a different style of problem-solving characterized by highly cooperative interactions. Agents may thus formulate and exchange tentative results with one another. Initially proposed by Lesser and Corkill [43], this approach is commonly known as **functionally-accurate cooperation**. In the presence of uncertainty, the functionally-accurate cooperation paradigm provides indeed an alternative model for task decomposition and agent interaction in distributed problem-solving architectures.

In "A Retrospective view of FA/C Distributed Problem Solving" [200], Lesser examines again the paradigm of the functionally-accurate cooperation, traces the development of a series of increasingly sophisticated cooperative control mechanisms for coordinating agents, and presents ongoing and new research directions in functionally-accurate cooperative (FA/C) systems. In this context, the author raises the key question of how to structure cooperative interactions among agents in order to limit costs associated with communication and still generate an acceptable answer with a reasonable amount of time.

In response, Lesser notes that the FA/C paradigm is designed to provide the answer to this question when an application can be structured in such a way that agents produce more complete partial results, resolve solution uncertainty due to competing, alternative partial solutions, detect inconsistencies in previously generated results, and speed up local problem-solving as the space of possible solutions that needs to be examined is constrained. According to the author, errors introduced as a result of incomplete, inconsistent and out-of-date local information are, in this way, resolved as an integral part of the asynchronous, co-routine exchange of tentative, high-level partial results among agents. In general, the objective behind FA/C distributed problem-solving is to allow agents to cooperate effectively in the presence of limited and inconsistent information related to the activities of other agents, different criteria regarding the most appropriate activities to perform, contradictory raw information and "conflicting long-term problem-solving knowledge".

On the other hand, Durfee reports that the functionally accurate cooperation approach has been extensively used in distributed problem-solving applications where agents "only discover the details of how their sub-problems results interrelate through tentative formulation and iterative exchange" [17,186]. This author notes that participating agents need to consider the partial results that are formulated and received as tentative.

Furthermore, it is interesting to note that exchanging tentative partial solutions can impact completeness, precision and confidence. However, the iterative exchange of partial results is, eventually, expected to lead to agents accumulating enough information "to keep moving the overall problem-solving forward". On the other hand, FA/C based problem solving could obviously incur dramatic communication overhead and wasted time, a problem that could be tackled through some control decisions. Another phenomenon known as **distraction** may well arise as a result of agents sharing too many results.

7.4.2 Shared repositories and negotiation search

The adverse side effects of increased communication associated with the functionally-accurate cooperation approach has already been outlined above. A common strategy designed to reduce "potential flurry of multicast messages" is to instead concentrate and store tentative partial results from agents sharing the same environment in a single, shared repository. Indeed, the blackboard architecture is a commonly known strategy designed to allow "cooperating knowledge sources to exchange and use results by communicating through a common, structured blackboard".

Using a shared repository in a design application would allow agents with different design criteria to search through alternative designs and critique these alternatives [202]. While it differs from traditional formulations, the search process amounts, in many ways, to a distributed constraint satisfaction problem [17,186]. In attempting to highlight the important differences, Durfee noted that agents are first not "assumed to know whose constraints might be affected by their design choices". Since agents would not know whom to notify of their decisions, a shared repository is thus highly recommended. On the other hand, Durfee adds that these entities can also "relax constraints in a pinch". Agents might, at any given time, need to choose between improving solutions, rejecting solutions, or relaxing expectations. In this context, heuristics to control the distributed search are also highly needed.

Negotiated search is applicable to diverse application areas and problem-solving environments [201]. As stated, the approach requires only basic search operators and allows maximum flexibility in the distribution of these operators. *Initiate-solution* represents the basic operator that proposes an initial state for searching a solution. *Critique-solution* is the operator that provides an evaluation of the solution in the form of feedback analysis together (when necessary) with information regarding detected conflicts. *Extend-solution* is applied "to extend and evaluate a partially specified composite solution". This operator is required in domains where solutions include interacting components, each developed by an agent. *Relax-solution-requirement* represents the operator that allows agents to change local requirements to reach mutual acceptability. At any given time, an operator can be invoked by an agent. According to Durfee, the problem domains for negotiated search are complex enough to require heuristic guidance [17,186]. Indeed, heuristic measures designed to guide an agent in its approach to invoke an operator – such as the *relax-solution-requirement* – are generally application-specific.

7.4.3 Distributed constrained heuristic search

Problems associated with contention for resources can frequently arise in distributed environments. The idea that these problems, commonly known as constraint satisfaction problems, can be solved by searching a space of states has been extensively explored. Instead of using a shared repository, a search strategy can thus be implemented as a way of associating an "agent" with each resource and then bringing that agent to process "the contending demands for the resource".

Indeed, **market-oriented programming** provides the ability to associate auctions with resources to support the search for "equilibria" in which resources are allocated efficiently [116]. Another strategy consists in allowing resources to compute their associated demands that each competing agent can take into account as it attempts to solve its constraint-satisfaction problem.

Sycara et al. propose a decentralized problem solving model known as distributed constrained heuristic search that allows the use of aggregate demand to guide asynchronous distributed search [203]. In this respect, Durfee observes that more informed search decisions are susceptible of decreasing wasted backtracking effort. He adds that constraint satisfaction heuristics such as variable and value ordering can be beneficially employed in a distributed environment.

7.4.4 Organizational structuring

In distributed problem solving, shared repository cannot be systematically supported. In this case, exploiting the structure of task decomposition can be a suitable alternative strategy for reducing communication. In an attempt to explain this strategy, Durfee puts forward an example of a distributed design problem [17,186]. In this respect, he argues that it is naturally reasonable that designers working on components that must "connect" communicate more frequently with each other rather than with designers working on more remote parts. The idea behind this example is that agents can make appropriate decisions based on knowledge about general roles each agent has to play in the collective effort.

This idea can be explicitly implemented in an organizational structure that defines roles, responsibilities, and preferences that in turn determine the control and communication patterns between agents. In addition, the organizational structure indicates capabilities associated with each agent and usually provides a prioritization of such capabilities. Through the prioritization of capabilities, the structure permits overlapping responsibilities that can increase probability of success.

In the organizational structure, it is important to have mechanisms through which an agent can learn about partial results that affect its assigned task. However, it is even more important not to communicate to agents results that are not susceptible of affecting their respective actions. As a result, the organizational structure provides adequate mechanisms designed to allow the identification of agents or the agent that may be potentially interested in a partial result.

7.5 Distributed planning

A distributed planning architecture consists of a group of computational agents that interact, share resources, and coordinate activities in order to achieve an overall common goal: design and integrate individual plans to yield a final global plan. Assuming that the task of designing and accomplishing the global plan is a problem assigned to these computational agents, distributed planning can thus, in many respects, be considered as a branch of distributed problem solving. Specifically characterized by particular features, planning problems are however addressed through particularly suitable techniques.

In this respect, Durfee highlights the ambiguity that lies behind the term "distributed planning" [17,186]. In his view, it is unclear exactly what is "distributed" in distributed planning. However, the author admits the existence of various issues that explain the use of the term "distributed planning". Indeed, as a consequence of planning, the formulated plan could be distributed among a variety of execution systems. On the other hand, whether the resulting plan(s) is/are distributed or not, the planning process could alternatively be distributed. Finally, another possible issue regarding distributed planning is that both the planning process and the resulting plan(s) could be distributed.

7.5.1 Centralized planning for distributed plans

Consider a partial order planner that is used to generate desired plans where there is no need for strict ordering between some actions, except the fact that these actions can be executed in parallel. One can thus imagine an agent acting as a centralized coordinator that takes the resulting plan and breaks it into separate threads. Using task-passing techniques, the resulting separate pieces of the plan are then passed to other agents and executed locally. Based on the assumption that the world is predictable and the knowledge is correct, the agents operating in parallel will, in the end, achieve “a state of the world consistent with the goals of the plan”.

While the actions of the plan are, as required, locally executed in a distributed manner, the formulation of the plan is nonetheless achieved centrally. Durfee examines algorithmically the process regarding the issue of centralized planning for distributed plans [17,186]. As presented, this process involves the following actions:

1. **Generate** a partial order plan based on a set of operators and a description of both the goal and the initial state. It is better to have a plan characterized by few ordering constraints among its steps.

2. **Decompose** the plan into sub-plans. The decomposition task should be conducted so as to minimize ordering relationships between steps contained in the same sub-plan, as suggested by Lansky [187].

3. **Define** and **introduce** synchronization actions into the resulting sub-plans.

4. **Allocate** these sub-plans to the appropriate operators or agents through adequate task-passing mechanisms. Following a successful allocation, an additional action is to insert remaining bindings into sub-plans – such as binding names of agents to send synchronization messages. However, in case of failure, the previous steps should be reviewed in order to provide a different decomposition of the plan, or generate a different partial order plan, etc.

5. **Initiate** the execution of the plan, and conduct monitoring activities (optional) in order to ensure, for example, a complete plan execution.

A close examination of the algorithm described above offers the opportunity to understand the nature and impact of the specific issues of decomposition and allocation. In planning, the decomposition and allocation steps essentially amount to the objective of finding, among all the possible plans susceptible of accomplishing a given goal, the most appropriate plan that can be effectively decomposed and distributed. This statement raises the question of the availability of appropriate agents to execute the sub-plans resulting from the decomposition process. It is obviously not totally certain that the plan considered “the most decomposable and distributable” can be allocated in any context [17,186]. Indeed, the availability of agents to locally execute sub-tasks is naturally difficult to determine without first having decomposed the plan.

Another question worth pondering concerns the impact of the communication infrastructure on the degree to which plans should be decomposed and distributed. While the distributed sub-plans may require synchronization, the communication channels could at the same time be slow or undependable. In this extreme case, it is suggested to form “a more efficient centralized plan” and take into account the monetary and/or time costs associated with distributing and synchronizing plans [17,186].

7.5.2 Distributed planning for centralized plans

In various areas, such as manufacturing, the planning process is too complex to be conducted through a single agent. As a result, the tedious planning function could well be distributed among a group of agents. In this context, each agent is assigned the task of contributing pieces to the plan, until a more acceptable or overarching plan is formulated. Just like generating a solution to a complex problem would require, creating a complex plan might indeed impose collaboration among a number of cooperative planning experts.

Durfee appropriately highlights existing parallels between task-sharing and result-sharing strategies used in DPS, on the one hand, and distributed planning for centralized plans, on the other hand [17,186]. According to this author, the overall problem-formulation task can be thought of as being decomposed and distributed among a group of planning experts, which individually proceed to generate their contribution to the plan. These distributed planning-related activities raise the crucial issue regarding the interactions among the planning experts.

Indeed, many real-world planning problems, such as planning in a manufacturing [188], unmanned vehicles control [189], or logistics [190] domain, involve interactions between planning experts through the exchange of a partially-specified plan. In an attempt to highlight the use of this model in various domains, Durfee [17,186] refers to an investigation regarding the implementation of a hybrid model that utilizes “a set of specialists to complement the overall expressiveness and reasoning power of a traditional hierarchical planner”

[188]. Used in the manufacturing domain, the approach consists in coupling a general-purpose planner with specialist planners for geometric reasoning and fixturing. As mentioned, the planner is a hierarchical non-linear planner similar to NONLIN [191]. While the geometric specialist augments the specification of the problem as seen by the planner and detect interactions that he planner itself cannot detect, the fixturing specialist utilizes the generated plan to make its own further commitments. In this respect, the geometric specialist uses solid models of the part and features to detect a variety of geometric interactions that

may affect the machining or fixturing of parts. Finally, the geometric specialist generates an abstract plan as an ordering over the geometric feature to put into the part and constructs or updates the interaction graph. Specific ordering constraints are then conveyed to the planner via the interaction graph. The hierarchical general-purpose planner uses these ordering constraints to plan the required machining operations. The resulting augmented plan is then passed on to the fixturing specialist. The task of the fixturing specialist is finally achieved in two phases: the first phase consists in proposing adequate set-ups and the second phase amounts to testing these set-ups.

In contrast, a result-sharing approach would impose on each planning agent the necessity to individually generate a partial plan in parallel and then share and merge these plans to "converge on a complete plan in a negotiated search mode". Indeed, mechanisms based on this approach have been developed and implemented to solve problems in the domain of communication networks [192]. In this class of problems, known as distributed constraints satisfaction problems, the satisfaction of each goal requires a coordinated set of actions distributed over a subset of nodes for completion. The mechanisms are thus designed to allow agents, each of which has incomplete knowledge about system resources and awareness of only partial solutions of system problems, to cooperate in solving complex distributed constraint satisfaction problems. Using a result-sharing strategy amounting to a distributed constraint satisfaction problem, localized agents can, in sum, "tentatively allocate network connections to particular circuits and share these tentative allocations with neighbours". However, in the case inconsistent allocations are detected, other allocations are tried without the interruption of the process until a consistent set of allocations are determined.

7.5.3 Distributed planning for distributed plans

Indeed, the distributed planning for distributed plans represents the most challenging issue of distributed planning. While both the planning process and the results are distributed, it is not however necessary and required to have a complete global plan available locally at each node. To avoid conflicts, the distributed pieces of the global plan should be compatible to allow agents to effectively execute the plans or at least rationally help each other achieve their respective planning task.

Research regarding distributed planning for distributed plans is relatively "rich and varied". As a result, various techniques for managing agents-related planning activities have been proposed in the literature. In "Distributed Problem Solving and Planning" [17,186], Durfee covers extensively some useful techniques:

Plan Merging. – A question that may quite normally arise in this area concerns the challenge related to the need to identify and resolve potential conflicts. Having to individually formulate plans for themselves, a group of agents are considered to be consequently faced with the problem of ensuring that the resulting separate plans can be executed without any conflict. In this context, it is thus assumed that the assignment of goals is either systematic – the application domain is inherently distributive – or has been conducted through task-sharing-based techniques.

A potential approach that may allow agents to avoid conflicts is to use a centralized plan coordination technique. Thus, the solution consists in appointing an agent to collect together the individual plans, analyze these plans to identify the sequences of actions that may lead to conflicts, and consequently remove these conflicts through a plan modification procedure. Given a set of possible initial states and a set of action sequences that can be executed asynchronously, the problem amounts to the task of enumerating "all possible states of the world that can be reached". In conducting this task, the agent can then determine the subset of worlds that should be avoided so as to be able to eliminate them through insertion of adequate constraints on the sequences.

However, the task which consists in enumerating the reachable state space is a difficult process. Appropriate strategies are thus required to deal with the complexity of this search. In "Distributed Problem Solving and Planning" [17,186], Durfee provides an extensive discussion on a search-easing method adapted from the work of Georgeff [38].

Given the plans of several agents, the merging method is first implemented to provide the ability to analyze interaction between pairs of actions to be taken by different agents. From this analysis, unsafe situations are thus identified. Actions that commute with all others can thus be dropped to "reduce the agents' plans to relatively short sequences". Given the simplified sequences, the merging process can enable the identification of the space of unsafe interactions by "considering the number of interleavings". Since the unsafe interactions are discovered, the synchronization of actions can finally be added to the plans to allow agents to avoid conflicts with ongoing actions.

Other alternative methods are also available to adequately deal with more complex forms of the search problem. A look-ahead coordination-based mechanism to maximize expected performance and to make forecasts regarding future activities has indeed been introduced to tackle problems related to uncertainties about the time required to conduct tasks or the possibility of new emerging tasks [193]. On the other hand, the issue about how agents should decide which plans to combine to maximize their global performance in the absence of schedules has also been addressed [194]. Finally, Petri-net based mechanisms and other techniques based on model-checking have also been explored and proposed as a way to coordinate more complex representations of reactive plans [195-196].

Iterative Plan Formation. – The plan-merging approach presented above is considered to be a powerful technique to achieve increased parallelism in both the planning process and execution. However, locally made decisions are sometimes tightly dependent on decisions attributed to other agents. Such a scenario offers a great challenge to researchers faced with the problem of how to manage the "degree to which local plans should be formulated" without ignoring the presence of other planning agents that impose the necessity to look into all aspects related to coordination.

Rather than allowing each agent to propose a single specific plan, agents should, in this respect, search through larger spaces of plans as a way of "tempering proposed local plans based on global constraints". Using this approach, each agent would thus establish a set of all feasible plans that allow it to accomplish its own goal. Hence, the distributed planning process would then amount to a search through generated plans to determine the subsets that can adequately fit together.

In "Divide and Conquer in Multi-agent Planning" [197], Ephrati and Rosenschein propose a ***plan combination search*** approach to deal with this type of search problems. This approach emphasizes starting with the inclusion of sets of possible plans and the refinement of these plans to achieve a desired convergence on a "nearly optimal subset".

Designed to perform ***distributed hierarchical planning***, an alternative approach consists in exploiting the hierarchical structure of a plan space. According to Edmund E. Durfee [17,186], this approach presents substantial advantages – exemplified in the ToH problem – in that "some interactions can be worked out in more abstract plan spaces". It results in pruning away "large portions of the more detailed spaces".

Known as the **hierarchical behaviour-space search** approach, a variation of the distributed hierarchical planning approach is to allow each agent to represent its associated local behaviours at multiple levels of abstraction. Durfee characterizes this approach as a search through hierarchical behaviour space [198], since, he explains, the plans at various levels dictate the behaviours of agents to a particular degree.

Negotiation in Distributed Planning. – In various cases, determining which agent should wait for another is, as Durfee explains, both fairly random and arbitrary [17,186]. This issue has been examined through a large amount of work in the area of negotiation.

In a work conducted in the air traffic control domain, Steeb and Cammarata focus on determining which of the various aircraft should alter direction in order to reduce “potentially dangerous congestion”. While the agents exchange descriptions regarding their flexibility, the system is designed such that the agent that had the most other options is asked to alter its plan. In this case, the selection of the agent susceptible of revising its local plan is based on “models of the possibilities open to agents”.

Designed to resolve goals, these and other negotiation mechanisms presume, as Durfee points out, that agents are honest about the importance characterizing their goals and the options of how to achieve them [17,186]. The literature pertaining to this area shows that issues regarding how to incite self-interested agents to be honest have been covered. While recognizing that the space of possible conflicts between agents is large, Durfee argues that the space of possible cooperative activities can be even larger, and introduces a variety of utility assessments [17,186]. According to this author, cooperation is “better”, but the degree to which agents benefit might not outweigh the efforts required to find cooperative opportunities. Finally, in “Coordinating Plans of Autonomous Agents”, von Martial focuses on strategies that agents can exploit to achieve particular goals based on “favor relations” [199]. While it can pursue its specific goal, an agent may, as a result, be able to achieve or accomplish a goal for another agent.

7.6 Dependence theory for action coordination

Dependence relations between agents constitute a motivation for cooperation in a group of agents, and more generally, a guide for coordination of actions and interactions [284].

7.6.1 Social interaction models

The definition of a social interaction model determines how agents should use organizational information [284]. In regard to organization, interaction models are generally considered from two points of view:

- Top-down models. In these models, in order to achieve a common goal for a group of agents, interactions are constrained by the organizational relationships that are fixed (static) by design. The organization could be represented explicitly [285], or simply by a structure at the implementation level. It is important to notice that fixed rigid organizations prevent the evolution of interactions.
- Bottom-Up models. In these models, agents have non-constrained interactions. Usually, they don’t have a common goal to achieve. Instead, each agent attempts to achieve his/her own goals. Bottom-up models form two classes:
    - Utility models consider agents to be relational, that is, they will always attempt to maximize their utility in the group they belong to. These agents must then coordinate their actions in order for the whole group to work properly and coherently.

- Complementarity models suppose that agents have complementary capacities, and that to achieve his/her own goal, an agent can ask for the help of other agents. In order to determine the agents that are likely to offer help, agents may use their dependence relations.

A dependence relation *r* exists between two agents *a* and *b*, if agent *a* needs agent *b*'s services in order to achieve a local goal. The relation *r* should contain all necessary information that allows agent *a* to know that he/she depends on agent *b* and how this agent should be able to help.

Dependence relations are a key enabler for cooperation and coordination: using these relations an agent may ask other agents for help, and also may help other agents without an explicit request. Due to the dynamic nature of interactions between agents, the dependence relations also become dynamic. For this reason, an agent must continuously determine dependencies on other agents within his/her interaction network.

Generally, dependence relations between agents are only a subclass of larger dependence relations that motivate interactions between agents. These relations are defined between three entities, the agent, the task and the resource (*Agent-Agent, Agent-Task, Agent-Resource, Task-Agent, Task-Task, Task-Resource, Resource-Agent, Resource-Task, and Resource-Resource*). These dependencies are the base of coordination [286]. Figure 10 shows how dependence relations may be derived from each other.

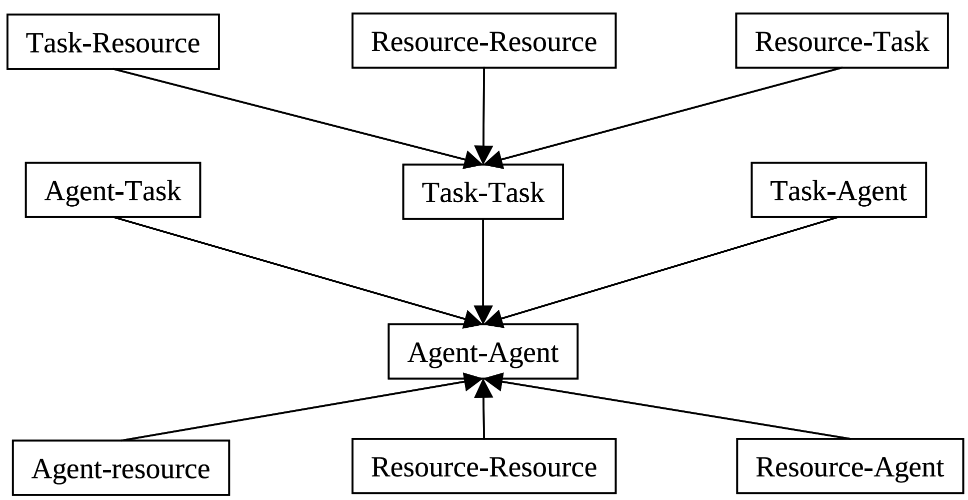

Figure 10. Derivation of dependence relations

7.6.1.1 Dependence relations between tasks

The dependence relation Task-Task is a combination of relations that a task has with its sub-tasks and the other tasks.

This relation is *negative* if the achievement of a task prevents or hinders the achievement of another task. The relation is *positive* if the execution of a task may help the execution of another task. Taking into account this kind of dependence may considerably increase the system performance.

7.6.1.2 Dependence relations between agents

Dependence relations between agents help them manage and control their interactions. In [284, 287], a formalization of two types of dependence relations is proposed. The first is resource-based (R-DEP) and is defined as the Agent-Resource relation. The second is a social dependence (S-DEP). It is defined between two agents: unable to perform an action, the first agent, to achieve a local goal, needs the performance of this action, which can be performed by the second agent. This type of relation is an Agent-Agent relation and may be defined between several agents.

7.6.2 Main characteristics of a social dependence

A social dependence may be decomposed in a conjunctive way (and-dependence) as in the situation where several actions necessary for the achievement of a goal may be performed by different agents, or in the situation where an agent is dependent on another agent to achieve several local goals.

A social dependence may be decomposed in a disjunctive way, as in the situation where several agents may perform the same action, necessary for the achievement of a goal. Another situation is that where several agents may perform several actions, necessary for the achievement of a single goal.

A dependence relation is *unilateral* if an agent is dependent on another agent regarding a goal. It is *bilateral* if it is *mutual* or *reciprocal*:

- Two agents have a *mutual dependence* relation if they are able to carry out a part of the necessary actions to achieve a common goal. These agents play complementary roles regarding this goal;
- Two agents have a *reciprocal dependence* relation if they are dependent on each other for two different goals.

Four dependence relations derived from the unilateral and reciprocal dependence relations can be defined:

- *Locally recognized mutual dependence*. An agent deduces a mutual dependence with another agent by using his/her own plans and without using the plans of the other agent;
- *Locally recognized reciprocal dependence*. An agent deduces a reciprocal dependence with another agent by using his/her plans but not those of the other agent;
- *Mutually recognized mutual dependence*. An agent deduces a mutual dependence with another agent by using indifferently his/her plans or the plans of the other agent;
- *Mutually recognized reciprocal dependence*. An agent deduces a reciprocal dependence with another agent by using indifferently his/her own plans or the plans of the other agent.

Two other types of dependence relations that imply three agents are the *external dependence* and the *transitive dependence*. The first relation portrays that an agent depends on another agent because he/she has adopted the goal of another agent. The second describes that an agent becomes dependent on another agent because this latest has adopted the goal of another agent. It is clear that before the goal adoption, the first agent was already dependent on the third agent.

7.6.3 Social dependence and autonomy

Dependence relations give agents the possibility to determine their level of *autonomy*, and hence the possibility to identify the tasks that should be part of the coordination process. An agent may be autonomous by design, that is, he/she has its own existence independently of that of the other agents. Autonomy can be expressed regarding the environment, the agent's goals or motivations (whether he/she wants to cooperate with others or not) [289].

The notion of autonomy is related to that of social interference. It exists between two agents if the achievement of an agent's goal has an impact on the achievement of at least one goal of the other agent. The impact is qualified as positive if the achievement of the goal helps the other agent achieve at least one of his/her goals. This impact is negative if the achievement of the goal prevents or hinders the other agent from achieving at least one of his/her goals (competition over the same resource or conflicting goals). The existence of *social interference* within a group of agents gives rise to different *social situations*:

- ***Exploitation.*** It is a situation where an agent, having the same goal as another agent, has just to wait for the other agent to achieve his/her goal;

- ***Influence.*** When an agent is unable to achieve a goal but believes that another agent may help him/her to do so, he/she will influence the other agent to achieve the goal. The degree of influence depends on the capacity of the other agent to help achieve the goal;
- ***Aggression.*** It is a negative interference where conflicts exist between the goals of two agents. An agent may achieve a goal in order to prevent the other agent from achieving his/her goal;
- ***Adoption.*** It is used in a cooperation context. An agent may adopt the goal of another agent if he/she has the goal that "this agent should achieve his/her goal".

By taking into account these different social situations, it is possible to make a distinction between two classes of social interactions. The first is positive, where the agents will help each other in achieving their goals. This kind of interaction takes place in a context of cooperation and social exchange. The second is negative where agents will try to prevent others from achieving their conflicting goals. This kind of interaction also exists in a context of competition where two agents, having the same goal, may perform the same actions in the environment.

7.6.4 Dependence relations in practice

Dependence relations are the means for agents to reason on each other, that is, to have social reasoning. *Social reasoning* covers all classes of reasoning where an agent makes use of information on another agent [288]. By using his/her dependence relations, an agent is able to measure his/her degree of autonomy regarding goals and other agents. When a dependence relation is detected, the concerned agents will elaborate a strategy to coordinate their actions and execute their plans.

Social reasoning requires an agent to have a representation of the other agents in the system. This representation is called *external representation* [290]. *The internal description* corresponds to the description the agent has of himself/herself. There are no standards for external description representation. It may contain goals, actions, resources controlled by an agent, and finally the plans an agent is able to execute by using a set of actions and resources in order to achieve a goal. External description may also contain dependence relations between an agent and others.

7.6.4.1 Coalitions formation

Generally, the strategy followed by an agent to achieve a certain goal (local or global), is to execute a plan of actions. Dependence relations are not sufficient for the choice of a strategy. An agent must consider the feasibility of the execution of a plan and the achievability of a goal. To do so, the author in [291] defines the notions of executable plans and achievable goals. He supposes that the agents use pre-established plans (no dynamic planning). Once a goal is chosen, agents must evaluate its importance by the means of a special function to that end. The choice of a goal requires the choice of a plan to be executed. If the agent is not able to execute the plan, he/she tries to form a coalition with other agents that are likely to help execute the plan. The choice of partners within a coalition is based on dependence relations.

7.6.4.2 Incoherencies solving

To benefit from the social reasoning, an agent must have a set of coherent beliefs. In practice, the external description that an agent J has of an agent I, must be compatible and coherent with the internal description of agent I even if the mechanism to update those descriptions is not the same. To update the external description of agent I, agent J makes use of social reasoning. This reasoning is based on information perceived in the environment, and information exchanged with others through interactions.

To contribute to problem solving, in [292], the authors have identified three sources of information: *reasoning*, *perception* and *communication*. Then, they identified a relation to rank these sources according to their credibility by using the following rule: an agent I will give maximum credibility to information coming from agent J describing himself/herself. If such information is not communicated by the agent J, agent I will first try to use his internal mechanism of reasoning to infer it. Second, agent I will try to infer it by using the perception of the environment. Finally, agent I will use any information communicated by the other agents (other than agent J) regarding agent J.

Incoherencies in the external descriptions will cause incoherencies in the dependence relations deduced by the agents [293]: dependence relations between agents (Agent-Agent) are usually deduced from other relations such as Task-Task, Task-Agent and Agent-Task. These relations are generally part of the internal description of an agent. Hence, any corresponding external description must contain the same relations. For example, let’s consider a situation where two agents I and J have the same goal G, and to achieve it, need to execute a plan P composed of two actions A1 and A2. Suppose that agent I (resp. agent J) can execute only action A1 (resp. A2), and suppose that the external description that agent I has of agent J is correct. Let’s suppose that the external description that agent J has of agent I is not correct, in that, it doesn’t contain the information that agent I is capable of executing action A1. In this example, agent I will deduce a mutually recognized mutual dependence with agent J, whereas agent J will deduce an independence.

7.6.5 Dependence relations and social power

Dependence relations between agents are objective relations, that is, they exist even if the agents don't take them into account. When taken into account, these relations may be composed with the notion of *power*. This notion is very often neglected when it comes to studying social situations such as autonomy and goal adoption. In [289], the author describes two types of power:

- An agent has the *power of a goal*, if he/she is able to achieve this goal. This notion is very important in the social reasoning, in that it is important for an agent I to know that agent J has the power to achieve a goal G, especially if agent I has plans that are dependent on goal G (G could be part of agent I's plans or cause the achievement of a Goal G that is part of agent I's plans. This type of power corresponds to the relation Task-Agent if a task is associated to a goal;

- From the above definition of power, the author defines the notion of *social power* between agents:
    - An agent I has the power on agent J regarding a goal G, if I can help J achieve G or can prevent J from achieving G. The relation Agent-Agent corresponds to this type of power in a cooperative context;
    - The notion of possession is linked to the notion of resource. It is also linked to the relations Agent-Resource and Resource-Agent. An agent I possesses X if it has the power to access X, can use or prevent another agent J from using X. Agent I has the power over agent J.
    - The power of influence that an agent I has regarding a goal G is defined as the power agent I has to increase or reduce the probability that another agent achieves the goal G. This type of power is interesting in the case where agent I is dependent on agent J regarding goal G1 and J has the power of G1. Knowing that J is dependent on I regarding goal G2, I may influence J to achieve G1 and promises to achieve G2 in return. If the achievement of goal G2 may prevent the achievement of one of the goals of J, then I may threaten J unless J promises to achieve G1. This type of social power is applicable in a situation where I wants to cooperate with J. If J accepts, he/she will adopt a

cooperative attitude with I (influenced by the proposition of I).

7.6.6 The individual and social characteristics of dependence relations

Dependence relations (Agent-Agent) are social relations that are very useful to guide interactions in a cooperation process. The relations Task-Resource, Resource-Task, Task-Task and Resource-Resource are application-oriented relations because they may be considered independently from agents. The relations Agent-Task, Task-Agent, Agent-Resource and Resource-Agent are rather *individual relations* because they include only one agent. Agent-Agent relations are very often deduced from the other relations.

The passage from individual to *social relations* is of primary importance: the individual relations indicate whether an agent is capable of executing a task or not. From the individual relations, an agent will determine the social relations that will allow him/her to ask for help or help others.

The social power may be composed with the dependence relations or considered as an extension. Indeed, dependence relations are mostly used in cooperative contexts where an agent needs the help of others. Social power may also be used in competition contexts. Whatever the context used, dependence relations composed with social power allow agents to manage their interactions and coordinate their actions.

This page intentionally left blank.

8. Scalability and metrics

Scalability is a problem for many intelligent distributed systems. The dimensionality of the search space grows significantly with the number of agents involved, the complexity of agent behaviours, and the size of the network on interactions between agents.

8.1 Coordination and scalability

Research dealing specifically with scalability in multi-agent systems is relatively scarce. Increasing demand for large-scale agent-based environments brought a limited number of researchers and experts to look into different aspects related to this issue. A review of the literature regarding scalability of multi-agent systems provides the ability to broadly categorize research in this area into several categories:

1. Definition, delimitation and further clarification of issues regarding scalability in agent-based environments.

2. Development of metrics and other adequate methods and procedures that offer the ability to conduct formal, theoretical and experimental analysis of scalability of existing mechanisms and protocols designed to allow interaction and resource allocation and distribution.

3. Development and implementation of tools to allow the prediction of agent-based systems performance and scalability.

4. Development of appropriate, scalable mechanisms and protocols.

5. Construction of scalable applications.

6. Development of algorithms and methods designed to increase the scalability of a multi-agent system through a change of the environment.

In “Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture” [204], Lesser examines key challenges facing the multi-agent community in addressing the issue regarding the development of next generation multi-agent systems consisting of a large number of agents operating in open environments. In this respect, the author raises the fundamental question of what are the basic functions and interaction patterns required to support the realization of open, large, complex, and robust systems and offers specific strategies and recommendations. In focusing on major problems associated with the design of an agent architecture that has to operate in an open and large-scale multi-agent environment, Victor R. Lesser outlines the importance of effective coordination strategies and the corresponding need for more sophisticated coordination mechanisms.

Roughly at the same time, Lee et al. highlight the lack of sufficient attention devoted to issues regarding non-functional properties of multi-agent systems such as scalability, stability and performance [205]. From this point of view, the authors explore the meaning of the term "scalability" of multi-agent systems and then attempt to demonstrate how this property can be properly analyzed. Thus, L. C. Lee et al. introduce a model of a multi-agent system, and describe a procedure that was used to analyze its performance and scalability. The objective behind this approach is to show that different designs of multi-agent systems need to be investigated in order to optimize these important non-functional requirements that underpin the quality of service. In a quite similar connection, Burness et al. recognize the significance of issues pertaining to the non-functional aspects of a complex multi-agent system [206]. According to these authors, it is quite hard to detect and correct many scalability problems through testing procedures, particularly if these problems are intimately related to the design of the system. From the analysis of a high-level design, Burness et al. demonstrate the ability to identify potential scalability problems embedded in a multi-agent system.

In an interesting paper entitled "What is Scalability in Multi-agent Systems", Rana and Stout first report that most agent-based systems built so far involve interactions between a relatively small number of agents [207]. Having in mind the growing trend towards building large-scale systems, the authors note that the contribution of this work is to propose a performance modeling scheme that can be combined with an agent-oriented design methodology in order to build large agent-based applications. While recalling that scalability, in its most general form, is defined as "*the ability of a solution to a problem to work when the size of the problem increases*", Rana and Kate Stout state that multi-agent systems will, in the context of multi agent communities, need to scale in a number of different dimensions. These dimensions include an increase in:

1. The total number of agents involved on a given platform,
2. The total number of agents involved across multiple systems or platforms,
3. The size of rules-based data on/with which agents operate,
4. The diversity among agents.

In case the total number of agents increases – scenarios 1 and 2 – the authors suggest using metrics associated with a particular platform or operating environment "to determine the total agent density and the resulting effect on system performance". As noted, memory usage, scheduling/swapping overheads for active agents, cloning or dispatching an agent to a remote site are the only metrics reported in the literature pertaining to agent performance and scalability [208-209].

On the other hand, Rana and Stout [207] provide an extensive discussion regarding the potential case where the agent density is increased as more agents will necessarily have to be added in response to an increase in the size of a particular problem. As noted, scalability in this scenario measures "the overall effect of parallel processing overheads on performance when the system and problem size scale up".

In addition, the authors point to the other case regarding the increase in the diversity of agents and the corresponding effects on agent density. In this particular case, scalability management is, according them, related to methodologies pertaining to agent analysis and design. In this respect, Rana and Stout highlight the need for "software engineering approaches for developing agent communities which extend beyond existing approaches based on object-oriented modeling techniques, or knowledge engineering". They also outline the need for other methods that have been implemented to build large complex systems based on interacting entities. Motivated by A-Life research, these methods are, according to these authors, based on the premise that societies of simple agents are capable of complex problem-solving behaviour, while processing limited individual capabilities. Referring to A-Life research [210-212], Rana and Stout highlight the absence of a central coordination of activity in such societies. It is some kind of "social coherence" that is, as reported, used to guide agents to work together towards a global objective. In this context, inter-agent communication is generally local and conducted through "time dependent signals" in such a way that no single agent is aware of the global state.

On the other hand, scalability can, according to Rana and Stout, also be stated in terms of "coordination policies in agent communities". It can thus be stated in terms of the total number of message exchanges to converge on a solution. Hence, a first approach is to group agents in order to limit message exchanges. Another approach consists in specifying a global utility function. Implemented in the COIN system [214-215], this approach is inspired by game theoretic methods applied to economic markets. In these methods, the emphasized necessity to converge to a global optimum has indeed an impact on performance, and correspondingly on scalability

Finally, Rana and Stout point out to another set of approaches that provide the ability to manage scalability. As noted, these approaches are related to modeling agent systems to predict performance. Based on the concept of a *messenger* paradigm, various approaches are mainly aimed at mobile agents. It is explained that messenger systems view agents as systems which involve mobile threads of execution without a central control. Modeling interactions in a messenger system can, as stated, be divided into: process algebras, actors and actor-related formalisms, Petri nets, coordination languages, temporal logic, category theory and others. According to Rana and Stout, Petri nets extensions – *Mobile Petri nets* [216] and *Communicative and Cooperative Nets* [217] – are proposed to handle mobility and dynamicity.

Aware of the fact that multi-agent systems have already moved out from research laboratories to commercial application environments, Phil Buckle et al. have, in a more recent paper, investigated the scalability issue [218]. In this paper, the authors propose a revision or rather a further clarification regarding the meaning of the term scalability. In addition, they describe a new model of scalability, which is also investigated through a discussion of qualitative as well as quantitative issues. According to these authors, scalability in a multi-agent system is not just a factor determined through counting the number of agents resident on a platform within any particular environment. Phil Buckle et al. claim that other qualitative issues are also crucially important. They further suggest the development or extension of a methodology in order to identify benefits and losses associated with agent behaviour against deploying larger numbers of agents within a multi-agent system.

Extensively described above, the Contract Net protocol is indeed a coordination mechanism widely used in multi-agent systems. While commonly considered as a simple and dynamic scheme, this protocol is however communication intensive due to "broadcast of task announcements". In "Learning in Multi-agent Systems" [220], Sen and Gerhard

Weiss address the problem of reducing communication in multi-agent environments. In this context, the authors consider the Contract Net protocol and highlight various disadvantages associated with this approach. While noting that the Contract Net-based approach operates quite well in small problem environments, Sen and Weiss assert that the protocol runs into problems as the number of communicating agents and the number of tasks announced – problem size – increases. In this respect, the authors advocate the development of adequate mechanisms designed to reduce the communication load resulting from broadcasting in more complex environments. Sen and Weiss then refer to various solutions – *focused addressing* and *direct contracting* – proposed by Smith [221] in the form of mechanisms designed to substitute point-to-point communication for broadcasting. However, these mechanisms present the disadvantage of imposing that a system designer must in advance know direct communication paths. Therefore, the resulting communication patterns may be too inflexible in dynamic environments. Based on this analysis, Sen and Weiss propose and describe an alternative and more flexible learning-based mechanism known as *addressee learning*. As explained, the primary idea behind this approach consists in reducing the communication efforts associated with task announcement through a procedure that enables individual agents to acquire and refine knowledge regarding task abilities of other agents. Hence, using this knowledge, task can thus be directly assigned without the need to broadcast announcements to all agents. Referring to various research papers [222,223], Sen and Weiss report that case-based reasoning is indeed used as an experience-based mechanism to acquire and refine knowledge.

Similarly, Deshpande et al. have examined and extensively investigated the scalability problems related to the Contract Net protocol [219,224]. In “Adaptive Fault Tolerant Hospital Resource Scheduling” [224], the Contract Net protocol is used as a coordination mechanism in a distributed hospital to provide the ability to share resources across nodes. In using this protocol, Deshpande et al realize that at high task arrival rate, many tasks cannot be completed within deadlines and the average task waiting time of the multi-agent system increases. In response to these scalability-related problems, an *Instance Based Learning* mechanism is used to improve the performance of the Contract Net protocol [219]. While similar to the *addressee learning* method described above, this mechanism is designed to use “the history of subtask migrations in order to choose a target node for a new subtask”. In this respect, Deshpande et al. explain that the history of subtask migrations contains instances, which consist of both the system state and the target node selected by the coordination mechanism at that state. Thus, using the Instance Based Learning mechanism, the appropriate target node is decided whenever a task is close to the deadline.

Inserted into the Instance-Based Learning mechanism is the k-Nearest Neighbour algorithm used as a technique to determine the k instances that are close to the current system state. As noted, using this technique provides the ability to both save valuable time - the bidding process is avoided through local computations – and reduce the communication load on the channels. Thus, incorporated into the coordination mechanism to allow resource sharing, the Instance-Based Learning approach contributes appreciably to improving the performance of the simulated distributed hospital system. On the other hand, results are provided to outline the ability of this approach to improve the overall scalability with respect to the number of tasks. Indeed, Deshpande et al. report that, even at high loads, a greater number of tasks complete within the required deadlines, and the average task waiting time improves considerably.

In “Scalability Analysis of the Contract Net Protocol” [225], Juhasz and Prasenjit also present early results of a study designed to investigate the performance and scalability of agent coordination protocols. Using the JADE environment, the authors implemented a test-bed system in order to measure the length of task execution under varying experimental conditions. The results provided show that the performance of the Contract Net mechanism depends largely on the number of agents and the load attributed to each entity. Under heavy load, the protocol often has to be performed repeatedly, while under light load, the delay is linear in the number of agents. Finally, Juhasz and Prasenjit report that this investigation regarding the scalability of coordination mechanisms will be extended to other protocols and approaches.

On the other hand, Turner and Jennings [226] examined various aspects related to agent-based systems deployment in distributed large-scale, open and dynamic environments. In [226], the authors raise the common question regarding the significant increase in the number of agents and address the corresponding challenging problem of scalability. In this respect, Turner and Jennings suggest that multi-agent systems are required to be *self-building* and *adaptive* in order to cope with the prominent adverse effects resulting in large-scale environments. While a self-building agent-based system could, at runtime, independently determine its most suitable organizational structure, an adaptive system would, on the other hand, modify this structure as the environment undergoes any change. According to the authors, agents that are able to build and maintain their own organizational structure are required a priori to be able to independently determine the importance of acquaintances and decide on tasks to be shared, delegated, or individually pursued. Since changing the organizational structure is a collective process, agents complying with the above requirements must be both aware and able to:

- Meta-reason about their individual internal efficiency and goals,
- Infer or question the goals of the system as a whole and the abilities and goals of their acquaintances,
- Create and annihilate other agents,
- Delegate and surrender tasks and information,
- Modify and influence their own operation,
- Influence the operation and activities of other agents.

In order to evaluate these hypotheses, Turner and Jennings implemented a self-building and adaptive multi-agent system. Applied to the domain of automated trading, the results show that adaptation and self-organization enhance the ability of a multi-agent system to cope with large numbers of agents. Despite the confirmation of the self-building and adaptive hypotheses, the authors believe that the experimental tests should be extended to several domains of application.

Similar approaches have also been proposed to address the issue of scalability through algorithms and techniques designed to deal with dynamic changes to the organizational structure of agent-based environments. Christian Gerber dedicated an extensive work – a PhD thesis – to methods designed to allow agents to adapt themselves to any application scale and nature in order to achieve and maintain efficiency and scalability in multi-agent societies [227]. In this context, the author presents GRAIL – ***G**eneric **R**esource **A**llocation & **I**ntegration a**L**gorithm* – a self-adaptive scheme designed for a society of benevolent agents. In this work, the concept of a *bounded-rational agent society* is introduced as an extension of a bounded-rational agent. According to the author, agents in such a society "optimize their behaviour to their individual resource". Implemented through special monitor agents, control instances optimize the allocation of societal resources.

The resource concept is thus extended to the multi-agent case where an abstract resource describes "an entity of the environment of an agent society, which expresses interdependencies" among members of this society. As noted, the task which consists in organizing such a society of agents is conceived as an optimization problem through a procedure designed to characterize a multi-dimensional *search space* and an *objective function*. The objective function denotes the performance of the system, while the search space describes the set of possible configurations of this system.

According to Christian Gerber, the theoretical foundations discussed above are implemented and used in the SIF – Social Interaction Framework – system which supports rapid prototyping of multi-agent scenarios. As an extension of the SIF system, another system called the SIFIRA – **S**ocial **I**nteraction **F**ramework for Integrated **R**esource **A**daptation – is conceived to provide the ability to conveniently design self-adapting agent societies. Christian Gerber reports that the GRAIL approach makes only a few assumptions on the nature of agents in a society. He explains that the autonomy of agents is in particular modified as little as possible. However, the author notes that if the autonomy of members of a group, or even all society members, can be changed more significantly, the GRAIL approach can be refined to a *holonic* agent society. In a holonic society, agents give up part of their autonomy and unite to a *holon*, which is seen by the outside world as one single entity.

Furthermore, Brooks et al. explored the area of large-scale multi-agent systems and attempt to tackle the complexity resulting from interactions of a significantly large number of agents [228-230]. The authors recall that agents in a large-scale multi-agent system are faced with "a combinatorially explosive number of potential interactions" [228]. Christopher H. Brooks et al. argue that one approach for agents to deal with this complexity is to form *congregations*. The idea behind this approach lies in the fact that most agents in a multi-agent system would rather interact with a subset of agents, instead of interacting with every other agent. As noted, this subset would ideally form a group of agents having complementary needs, goals, or preferences. Christopher H. Brooks et al. point to the idea that agents would tend to group together with agents that share some important features and obviously avoid interacting with incompatible agents.

Referred to as congregation, this grouping together represents a common phenomenon in human societies. Indeed, congregations that exist in a society – such as clubs – make it much easier for humans to find people with complementary interest or capabilities. From this behaviour, the idea thus put forward is how to emulate such congregations within a computational framework. In [228], Brooks et al. implemented the idea regarding the formation of congregations in a particular domain peculiar to information economy. As stated, an economy is an example of a multi-agent system in which congregations naturally occur, typically in the form of markets. Moreover, artificial agents may in this context represent producers and consumers in an information economy. Hence, Brooks et al. [230] addressed a particular type of congregation formation regarding the need to group together consumers and producers of information goods in such a way as to produce a desirable global state. In this respect, the authors propose two classes of strategies pertaining to the formation of desirable congregations. While the first class consists in using external mechanisms, such as taxes, the second class is about the introduction of learning to members of the agent population. The results are finally analyzed to provide a better understanding of the problems associated with the formation of desirable congregations.

Later, Brooks et al. conducted further research on congregating in multi-agent systems to explain and predict the behaviour of self-interested agents that search for other agents in an attempt to interact with them [229]. Experimental and analytical results are provided to highlight the difficulty of the congregating problem. Indeed, this problem becomes exponentially more complex as the number of agents increases. In this respect, Christopher H. Brooks et al. report how the complexity of the problem regarding the formation of congregations is reduced through the introduction of basic coordination mechanisms such as labellers. As noted, labellers are agents that “assign a description to a congregation in order to reduce agents’ search problem”. While highlighting the similarity between labellers and focal points [231], the authors explain that labellers allow agents to coordinate by “providing an external mechanism for synchronizing behaviour”.

In a more recent follow-on research work [230], Brooks et al. show how a structured label space can be exploited to simplify the labeller-related decision problem and reduce the congregation problem to linear in the number of labellers. On the other hand, experimental evidence is also presented to demonstrate that the formation of congregations among agents can reduce search costs, thereby allowing the system to scale up.

In focusing attention entirely on agent coordination strategies, Durfee dedicates a whole research paper to the issue of scalability of agent-based systems and applications [232]. While noting that advances in agent-oriented software engineering provide the ability to develop complex, distributed systems, the author highlights the importance of building agents that are able to act and interact flexibly. In this context, Durfee argues that coordination has indeed emerged as a central concern of intelligent agency. Based on this analysis, the author puts forward the idea that an effective coordination strategy should scale and respond well to being stressed along various dimensions. According to him, any attempt to understand the capabilities and limitations of coordination strategies that are designed to support flexible component agent interaction would thus require the characterization of agents’ properties as well as their task environment and collective behaviour.

In this respect, Durfee devotes an extensive discussion to various dimensions that are identified to allow mapping of the space of potential coordination approaches. In tackling properties associated with agent population, the author emphasizes the need to handle more agents as a challenge in scaling any coordination strategy. Thus, ***Quantity*** is identified as an important dimension of coordination stress. As pointed out, this property infers the absolute necessity to avoid using a centralized coordinator in order to direct the efforts of other agents. On the other hand, ***heterogeneity*** and ***complexity*** are mentioned as the two other agent population-related dimensions that impact the scalability of proposed coordination strategies and mechanisms. While complexity refers to the difficulty associated with predicting the activity and reaction of a given agent, heterogeneity reveals the presence of different kinds of agents having different goals, beliefs, or expertise and using various communication languages, onthologies, or internal architectures.

Furthermore, Durfee provides a set of task-environment properties that can also be considered in order to anticipate the performance, effectiveness, and suitability of a given coordination strategy. These properties include the ***degree of interaction***, ***dynamics***, and ***distributivity***. As stated, the degree of interaction grows as “the number of agents concerned with the same issues” increases and as “more issues become a concern to each agent”. On the other hand, the dynamics property may indeed characterize a multiagent setting where each member of a group of agents monitoring only portions of the environment can change its own perception about what goals to pursue and how to pursue them. Finally, a central feature of current multi-agent environments lies in the way both agents and tasks are distributed. In various task environments, inherently distributed tasks are often allocated to highly distributed agents. As a property of such environments, distributivity obviously impacts coordination, since it contributes to increasing uncertainty about the presence - which agents are currently sharing the task environment – and activities of individual agents.

Durfee notes that coordination strategies are still required to yield satisfactory solutions. In this respect, the author suggests the idea of making solution criteria more stringent along dimensions that include ***quality***, ***robustness***, and ***overhead limitations***. As proposed, the quality of a solution can be expressed in terms of some standards regarding the coordination of agent interactions, the use of agent resources, and the ability to settle issues. While lower quality may express a mere satisfactory level of coordination, a high quality standard may, as noted, correspond to “near optimal” coordination. Given the uncertainty and dynamic behaviour characterizing a task environment, robustness of a solution may express the ability of the coordination strategy in anticipating “the range of conditions under which the solution it provides will be followed”. Durfee then observes that costs associated with a given coordination strategy could include computation requirements, communication overhead, time spent, etc. He finally raises the question on “whether a coordination strategy can scale well to environments that impose stringent limits to the costs the strategy incurs”. However, scaling along combinations of the above dimensions would, as Durfee points out, induce even greater challenges.

6.2 Coordination performance metrics

Traditionally, a designer of a system involving multiple agents analyzes the task domain of interest, and based on this analysis, imposes upon the agents rules that constrain them into interacting and communicating according to patterns deemed desirable. The idea is to provide agents with ready-to-use knowledge in order to allow these entities to share a common environment, effectively interact, and achieve a set of desirable properties: load balancing, conflict avoidance, stability, and fairness. This approach has guided research into coordination techniques and often led to prescriptions for a variety of methods: task-sharing protocols [35], negotiation conventions [50], and rules for interaction such as social laws [233], etc.

On the other hand, another alternative approach is based on the argument that agents should, in the absence of common pre-established protocols or conventions, be able to make decisions regarding interactions within a shared environment. The idea does not imply any obligation to restrict agents from interacting through protocols, so long as these entities are not however restricted from deliberating about what to do. Without relying on protocols or conventions, an agent should be able to use knowledge regarding the environment and the capabilities, desires, and beliefs of other agents in order to rationally act, interact, and coordinate in a multi-agent setting. This conception of agents' coordination has, on the other hand, given rise to a number of methods such as the techniques based on the normative decision-theoretic paradigm of rational decision-making under uncertainty [234].

As already described in the previous sections, a wide variety of mechanisms and protocols have been proposed to address the challenging issue of coordination in DAI. Having different properties and characteristics, these mechanisms are obviously suited to different types of tasks and environments. As coordination techniques are not equally effective – each protocol or mechanism has its benefits and limitations – another proposed strategy consists in providing agents with a range of tools with varying properties so that any appropriate mechanism or protocol can be selected to handle a coordination episode [235,236].

Whatever approach described above is deemed the most effective, a fundamental and legitimate question can then be raised on how to evaluate this entire panoply of coordination protocols and mechanisms. In the absence of a standard methodology, an analysis of the literature reveals a substantial inadequacy of research focused on the development of methods and tools to evaluate coordination mechanisms and protocols in comparison to the efforts dedicated to the coordination issue.

Over the last two decades, research in the area of DAI has, among other things, focused on the design of agent-based architectures providing the ability to dynamically modify agent organizational structures. The approach is based on the idea that a single specific organizational structure cannot handle many different situations pertaining to a dynamic environment. Just as a single organizational structure is not suited to every situation, there is not either a single coordination technique that can handle all organizations well. In response, agents are thus offered, as mentioned above, the ability to operate using a variety of coordination mechanisms. In "Coordinating Distributed Decision Making Using Reusable Interaction Specifications" [237], Barber et al. investigate the issues involved in enhancing the flexibility of agents in terms of coordination capabilities. In this paper, the authors describe an approach – a representation of coordination strategies based on concepts designed to assist agents in evaluating and dynamically selecting alternative coordination strategies. Based on object-oriented concepts – *encapsulation* and *polymorphism*, this approach amounts to a representation of coordination strategies which also offers the ability to compare coordination techniques through a constant framework as well as other techniques. As coordination strategies are treated in an abstract manner, agents have the ability to change strategies and incorporate newly developed strategies.

According to Barber et al., strategic decision-making, whether performed on-line or off-line by a designer, must be made with regard to which strategies are appropriate for the given problem. In this respect, the authors recommend the analysis of the coordination strategies in terms of domain dependent and independent characteristics [237]:

- ***Strategy-related requirements.*** A given strategy may impose constraints on an agent's reasoning capabilities. This strategy may (or may not) require inter-agent communication and only partially make use of agent's abilities.

- ***Cost of strategy execution.*** When executed, each strategy uses part of the resources attributed to an agent. On the other hand, a strategy may, for instance, require a larger number of messages or a longer time. In this case, it is important to consider these factors to better deal with all aspects related to deadlines or limited agent resources.

- ***Solution quality.*** Indeed, using different strategies may lead to solutions with differing quality. Furthermore, agents may need longer deliberation to generate a better solution. Depending on the time available, an agent may perform trade-off analysis between the expected quality of a solution and the cost associated with the execution of a strategy.

- ***Domain requirements.*** Strategies may or may not be able to satisfy requirements associated with an application domain.

Furthermore, Barber et al. mentioned the existence of multiple approaches that are "applicable to the manner in which agents consider the above characteristics and select the most appropriate strategy". While it is recognized as difficult to manage, trade-off reasoning between multiple objectives is roughly classified into two categories: utility-based dynamic decision-making and ranking relations.

Furthermore, Bourne et al. have also tackled the problem resulting from the design of complex applications where autonomous agents are required to behave rationally in response to uncertain and unpredictable events [238]. In a similar way, the authors thus explored the idea of building systems where agents are fitted with the ability to coordinate their activities in accordance with demands and needs of prevailing circumstances. Instead of implementing a detailed coordination plan in which actions associated with each participant are rigidly prescribed, a more appropriate approach implies, according to Rachel A. Bourne et al., the necessity to adopt much looser coordination policies enabling agents to reason and select, at run-time, the method best suited to their current situation. As reported, such flexibility can thus be achieved through an adequate framework where agents, given a set of varied coordination mechanisms with different properties and characteristics, are offered the ability to assess the likely benefit of adopting an available mechanism deemed appropriate in the prevailing circumstances.

Based on this analysis, Bourne et al. propose a decision-theoretic model offering the ability to evaluate and select between competing coordination mechanisms. In this context, the authors identify a number of potentially differentiating features that are specifically common to a wide range of coordination mechanisms. Then, Bourne et al. show that agents can effectively evaluate and decide which mechanism to use, depending on the prevailing operating conditions. Referring to contingency theory [238], the authors observe that coordination mechanisms that are guaranteed to succeed typically have high set-up and maintenance costs, whereas mechanisms that have lower set-up costs are more likely to fail [236]. In this respect, Bourne et al. propose a framework to provide agents with the ability to select appropriate coordination mechanisms through set-up costs and probability of success.

In “Criteria for the Analysis of Coordination in Multi-agent Applications” [239], Frozza and Alvares highlight the importance of having tools designed to evaluate coordination in a society of agents. Such tools would help design more efficient and robust multi-agent systems and offer designers the opportunity to discuss the quality of coordinated actions executed by agents. In an attempt to provide a conceptual framework that can be used to conduct performance analysis of coordination mechanisms, Frozza and Alvares propose a set of criteria offering the ability to characterize various elements related to agents’ coordination:

- ***Predictivity.*** The predictivity of a coordination mechanism represents its ability to determine the future status of both agents and the environment.

- ***Adaptability.*** The adaptability represents the ability of a coordination approach to deal successfully with new situations or unexpected events.

- ***Action control.*** The action control can be either centralized or distributed. While a centralized action control is characterized by the use of a single agent that holds the knowledge of a given problem and assigns tasks to other agents, a distributed control supposes that any agent can establish rules based on its knowledge.

- ***Communication mode.*** Communication between agents can be achieved through interaction, perception, or direct communication. Coordination can also be conducted without communication.

- ***Conflicts.*** Depending on the nature of conflicts, some coordination mechanisms are unable to avoid conflicts, while others demonstrate a great ability to approach and resolve conflicts.

- ***Information exchange.*** Coordination is also achieved through handling and exchange of information.

- ***Agents.*** Agents involved in a coordination scheme demonstrate both characteristics – homogeneous or heterogeneous – and capabilities.

- ***Applications.*** Applications are either suited to specific domains or adaptable to any domain.

- ***Advantages.*** Qualities that make a coordination method more successful.

- ***Disadvantages.*** Negative aspects associated with a coordination approach.

Based on these criteria, the authors provide a table of comparison of different coordination approaches. The objective behind this comparison is to demonstrate the effectiveness and applicability of such criteria.

Finally, in the area of negotiation, Sandholm suggests the use of various criteria or properties for evaluating negotiation protocols or mechanisms [70]. Listed below, these criteria include:

Negotiation time. – Assuming that a delay in reaching an agreement induces an increase in the cost of communication and computational time, negotiations conducted without delay are thus preferred over time-consuming negotiations.

Guaranteed success. – In case a protocol ensures that an agreement is, eventually, certain to be reached, it guarantees a successful negotiation outcome.

Maximizing social welfare. – In a given solution, social welfare is the sum of all agents' payoffs or utilities. It can be used as a property or criteria for comparing alternative mechanisms. Hence, a protocol maximizes social welfare if the resulting solution maximizes the sum of the utilities of negotiation participants.

Efficiency and Pareto efficiency. – An efficient outcome is preferred since it increases the number of agents that will be satisfied by the negotiation results.

Pareto efficiency represents another property or criteria for evaluating a solution that a negotiation mechanism can lead to. A solution or negotiation outcome is said to be Pareto efficient – i.e. Pareto optimal – "if there is no other outcome that will make at least one agent better off without making at least one other agent worse off". Similarly, a negotiation solution which is not Pareto efficient is characterized by the existence of another outcome that "will make at least one agent happier while keeping everyone at least as happy".

Individual rationality. – An individually rational mechanism supposes that participation in negotiations is individually rational for all agents. In other words, it is in the best interests of all agents to participate in negotiations through a mechanism that is individually rational. Thus, in the absence of individually rational protocols, agents lack incentive to engage in negotiations. Hence, an agent participating in an individually rational negotiation will get a payoff in a solution that is no less than the payoff it would have by not taking part in the negotiation process.

Stability. – Where agents are self-interested, mechanisms or protocols should be designed to be stable so as to motivate each entity to behave in a desirable way. Indeed, a self-interested agent tends to behave in some other way than desired if it is better off. A stable protocol is defined as a protocol that "provides all agents with an incentive to behave in a particular way". According to him, the best-known kind of stability is *Nash equilibrium*.

Simplicity. – A simple protocol makes the appropriate strategy “obvious” for a negotiation participating agent. He adds that a participant using a simple negotiation mechanism can easily determine the optimal strategy. Indeed, simple and efficient protocols are better than complex mechanisms.

Distribution. – A protocol should be designed so as to avoid the presence of a single point of failure, such as a single arbitrator, and minimize communication between agents.

Money transfer. – Money transfer may be used to resolve conflicts. In this respect, a server may, for instance, “sell” a data item to another server when relocating this item. Such procedure requires the necessity to provide agents with a monetary system and a mechanism designed to secure payments. Finally, negotiation protocols that do not require money transfers are indeed preferred, since maintaining such systems requires both resources and efforts.

This page intentionally left blank.

9. Research directions and future trends

Academic research has made substantive applied research contributions in varying degree to industry. Major technological opportunities and breakthroughs have resulted from bi-directional flow of ideas across boundaries between academia and industry. Indeed, industrial innovations and technological change in both economy and society have always been significantly influenced by academic research and development. With currently deployed agent-based applications featuring characteristics of prototypes developed through academic research, it is obviously clear that the area of research related to agent technologies is not an exception.

9.1 Future trends

Insightful exploration of current industrial applications and future commercial opportunities offers a useful perspective and provides the ability to identify key issues and project trends for future multi-agent research. As already stated, despite intense and persistent efforts consented to building practical agent-oriented applications, the actual number of deployed systems is still surprisingly limited. Attempts to explore the obstacles preventing the development and consistent widespread use of such applications may offer the opportunity to better understand current progress, grasp current preoccupations of researchers and practitioners, dissect their priorities, and identify research issues in this area.

Indeed, current agent-oriented applications deployed in industry and other sectors can be described as classical systems driven by the need to solve practical problems, rather than test the possibility of existing technologies. Focusing on practical results, a project aiming to build and deploy an agent-based system is evaluated based on how well it addresses and solves a problem, instead of how sophisticated is the technology. Deployed real-world agent-oriented applications highlight the existing links between theoretical research and practice.

As noted, it is commonly acknowledged that building practical industrial-strength systems bringing together a large number of interacting agents is an inherently complex construction endeavour. However, practical applications in current deployment may indeed be characterized as ***close*** agent-based systems featuring a set of systemic properties. Indeed, existing applications are each designed and developed by a ***single team*** to allow a group of ***homogeneous*** agents to interact within a ***single corporate environment***. Only suited for a ***single problem-domain***, each system incorporates an ensemble of participating agents that share common high-level goals. Hence, the design of such systems obeys to a traditional approach which consists in analyzing the task domain so as to impose upon participating entities pre-defined rules, languages, and typically in-house interaction protocols. Provided with ready-to-use knowledge, participating agents can thus interact effectively, share a common environment, and exhibit a set of desirable properties. On the other hand, despite considerable efforts to address scalability problems, current applications are scalable under “controlled or simulated conditions”. Rather than implementing appropriate standard design methodologies, practical agent-oriented systems are thus currently conceived and developed using ad hoc design rationale inspired by the agent paradigm.

Indeed, a range of practical industrial and commercial problems can only be addressed through an open, robust, secure, scalable, and multi-corporate agent-based approach. However, growing security concerns arising in open environments may obviously slow down the development and delay the introduction of open and large-scale agent-oriented applications. Thus, it is unlikely that the acknowledged commercial need for well-secured and closed agent-based applications will decline and suddenly cease as progress in the area of multi-agent research will promote and impose novel technologies.

In addition to this broad view on actual progress related to deploying commercial and industrial agent-oriented applications, current academic research interests and activities in this area have also been monitored so as to extract appropriate information for identifying trends for future research. Without any reference to a timescale, future developments in the area of multi-agent systems can thus be distinguished along the following dimensions:

1. **Inter-corporate agent-based applications.**
2. **Heterogeneity of agents.**
3. **Scalability of next-generation systems.**
4. **Openness of emerging agent-oriented systems.**

1. Inter-corporate agent-based applications

AgentCities [267], *Grid computing* [266], and the *Universal Information Ecosystem* [268] are a few examples of many on-going strategic research initiatives designed to create worldwide networks of agents (See Figure 11). Though the end objective is to produce truly open, robust, adaptive, self-organizing, and scalable environments, it is highly likely that the next phase will, the in the ***short term***, lead to the development of systems that address the same domain ***across corporate boundaries***.

Thus, while inter-agent interactions will obviously concern the same domain, emerging systems will, in contrast with actual applications, bring together a large diversity of agents having fewer goals in common. Despite this growing diversity, agents incorporated into a given system will, in the short term, still be designed by the same team in charge of designing a whole system. Based on current research to develop Grid applications and AgentCities, emerging agent-oriented technologies will thus provide the ability to integrate, handle, and federate a large number of distributed agents within *pre-determined environments*. Thus, ***interaction protocols*** required to implement such complex applications will, in the short term, still ***remain non-standard***. However, agent-oriented systems designers will, in the foreseeable future, increasingly rely on ***standard communications languages***, such as FIPA ACL.

Figure 11. Agent platforms around the world registered in the AgentCities Network- February 2003 [267]

2. Heterogeneity of Agents

Substantial efforts to address real-world problems in areas such as e-Commerce and Bioinformatics will, in the ***medium term***, lead to agent-oriented systems that allow flexible integration of diverse ***heterogeneous agents*** conceived and developed by different designers. This step will give birth to open commercial and industrial agent-based applications. However, such systems will be specifically designed for *particular application domains*.

Available agent-oriented technologies will be such that participating agents will be designed to behave in accordance with "*publicly stated requirements and standards*". In this respect, systems' designers from different design team will, in the medium term, use either ***pre-defined*** and ***agreed upon*** or ***standardized interaction protocols***. Based on current research, it is likely that agents will, in the medium term, be allowed to evaluate and use *alternative coordination protocols* and *mechanisms* drawn from accessible *public libraries*.

On the other hand, while standardized communications languages will still be used, it is commonly acknowledged that language and terminological heterogeneity could greatly affect inter-agent communication. In this respect, problems associated with *semantic heterogeneity* will certainly be addressed particularly through the use of different ***ontologies*** about various topics. Since heterogeneous agents will use private ontologies, systems designers will have to devise an approach designed to allow *merging* and *mapping* of individual ontologies.

3. Scalability of emerging systems

Indeed, advances in networking technologies are paving the way towards building ***large-scale*** agent-oriented application systems. In the ***medium term***, it is expected that *large numbers* of autonomous and *heterogeneous* entities will be deployed in networks and the Web to pursue specific goals related to ***different domains***. In this context, specific abstractions, methodologies and tools are greatly required to enable an engineered approach to building such complex applications.

Thus, a fundamental issue underlying the design of high-quality, large-scale, and complex agent-based systems lies in the need to adequately and effectively handle interactions between continuously active and heterogeneous agents involved in different domains. As mentioned in the literature, this decisive phase in the emergence of truly scalable agent-based systems is likely to witness the development of ***bridge agents*** designed to translate between separate domains. In a multi-agent application built to automate the design of a given product, bridge agents can be used to interact with other e-Commerce systems and engage in effective commercial negotiations for ensuring access to patent-protected information. On the other hand, standard agent-specific design methodologies, including design patterns and architectural patterns, are important vehicles for constructing reliable and high-quality large-scale systems.

4. Openness of next-generation systems

In the ***long term***, agent-oriented research is expected to result in unprecedented implementation of ***open*** and loosely coupled multi-agent systems. Incorporating ***heterogeneous entities*** developed by ***different design teams***, an open system will be deployed across ***multiple application domains***. Using agent-specific design methodologies, such systems will be developed to exhibit openness and thus a genuine scalability, in that no limit or restriction will be imposed on complexity and the number of autonomous agents and users.

In this respect, handling agent *behaviour* in future open agent-based systems will present a special issue. It requires methodologies and tools to predict, control, and prevent undesirable behaviour attributed to agents in open and large-scale applications. Consequently, agents may be designed to *learn appropriate behaviour*, instead of having to adhere to a given code of conduct or behaviour prior to joining the system.

As already predicted, standard communication languages and interaction protocols will have been developed before the emergence of open agent-oriented systems. Rather than defining and imposing communication languages and interaction protocols, inter-agent coordination will be tailored to particular contexts. Thus, in future open systems, agent communication and coordination will evolve from actual interactions of participating entities.

9.2 Technological challenges

In summarizing present and future developments in the area of agent-oriented systems, the previous section shed light on a variety of new challenges in R&D over the next decade. Indeed, the area of agent-oriented systems is a rapidly expanding field of R&D. It represents an amalgam of disparate ideas, theories, and practices originating from such areas as distributed computing, software engineering, artificial engineering, economics, sociology, organizational science, and biology.

Agent-oriented R&D will have a considerable impact in various application domains such as those described in the previous section. Agent concepts and paradigms will be used as an abstraction tool and a metaphor to design and build complex, distributed computational systems. While offering an appropriate way of considering complex systems as multiple, distinct, and independent entities, agent-based approaches will remain a source of technologies for building a number of challenging real-world applications.

As stated above, current deployed applications can be described as typically closed systems, designed through ad hoc methodologies, scalable in simulations, and incorporating pre-defined communication languages and interaction protocols. However, current technology development such as Web Services, Grid computing, and peer-to-peer toolkits are rapidly changing the way systems deployed in public networks are integrated to interact. In this context, it is likely that real-world problems will, in the long term, impose the emergence of truly open, fully scalable agent-oriented systems, spanning across different domains, and incorporating heterogeneous entities capable of learning and adopting adequate protocols of communication and interaction. With the perspective of deploying such open, dynamic, and unpredictable environments, the traditional critical issue regarding the coordination of participating agents will remain a fundamental challenge in multi-agent R&D.

Indeed, located at the heart of the area of agent-based systems is a central concept of intelligent and autonomous agents interacting with one another to achieve individual or collective goals. However, agent-oriented technology is breaking with current practices to undergo a transition from monolithic systems based on a single overall design philosophy to open architectures involving conglomerates of heterogeneous, and independently designed agents and agent-based systems. Among all the identified problems inherent in such dynamic environments, none is more complex than the fundamental and imperative need to appropriately control and coordinate the activities attributed to a large number of disparate entities. Based on the practical issues discussed above, an attempt is thus made to objectively identify future academic and industrial research themes directly related to inter-agent interaction. Since short-term issues are already being addressed by the research community, this section is consequently dedicated to what is considered as strategic medium and long-term research objectives.

9.2.1 Appropriate design methodologies

As already noted, despite widely acknowledged advantages associated with agent-oriented technologies, the number of deployed agent-based applications is still limited. Thus, the discrepancy between the current state of the art in agent-oriented research and the actual deployment of commercial and industrial agent-based systems should obviously deserve greater attention. In this respect, researchers and practitioners often highlight the current use of ad-hoc design methodologies with limited specification of requirements in order to explain the gap between academic promises and reality.

The lack of widely accepted development tools and methodologies is described as one of the major roadblocks to commercial and industrial deployment of such applications. In this context, the implementation of comprehensive ***analysis*** and ***design methodologies*** dealing with both the macro-level – *societal* - and micro-level – *agent* – aspects is commonly seen as appropriate to enhancing the quality of agent-based applications to industrial standard. These methodologies would provide the ability to model and specify agent-based systems and their corresponding requirements. However, attempts to develop appropriate analysis and design methodologies imply the need to identify characteristics of real-world applications, in relation to specific domains. In describing the semantics of agent-based systems – without any concern for implementation details, specifications of requirements will, on the other hand, provide a basis for verification, testing and validation of systems' functional and non-functional properties.

According to the **F**oundation for **I**ntelligent **P**hysical **A**gents[8] (**FIPA**), existing development methodologies **-** *AOR*, *Cassiopeia*, *Gaia*, *Mase*, *Message*, *PASSI*, *Tropos* – provide different advantages when applied to specific problems. It is thus argued that a developer of an agent-based system would prefer to use phases or models coming from different methodologies so as to generate a personalized approach suited to his own problem. Focusing on the identification of an appropriate analysis and design methodology, the FIPA Technical Committee proposes the adoption of the method engineering as the referring paradigm to allow developers to reuse contributions from existing methodologies. In other words, FIPA suggests that a developer can use coherent fragments from existing and future contributions to determine the best analysis and design process for his or her needs and problems.

Anyway, a fundamental requirement associated with the process of analyzing, designing, and constructing large-scale, open, and loosely coupled agent-based systems lies in the ability to ensure certain properties and handle features, such as goals, mobility, adaptation, learning, autonomy, planning, coordination, etc. Thus, appropriate techniques will be developed and implemented to model and structure properly functional as well as non-functional properties in the suitable development stage. In this context, basic to the analysis of future open environments is the use of "*dedicated concepts and languages*":

Concepts representing dynamic – e.g., *time* and *action* – and locality aspects – e.g., *position in a space* – as well as mental state – e.g., *belief* and *desire*.

---

[8] http://www.fipa.org/activities/methodology.html

Concepts related to such aspects as coordination, interaction, organization, and society – e.g., *organization forms*, *society norms*, *interaction protocols*, *joint goals*, and *joint plans*.

9.2.2 Effective, Agreed-upon standards

Emerging properties regarding ubiquity, openness and scalability of agent-oriented systems provide a framework to support future research in this area. A key to a successful transition from current applications to building complex systems exhibiting such properties lies in the need to formulate, establish, and implement effective and largely agreed-upon standards to support interoperability between heterogeneous agents.

Agent-oriented standardization efforts are currently conducted through different bodies and communities: the Knowledge Query Meta Language9 (KQML), the Object Management Group10 (OMG), the Mobile Agent System Interoperability Facility11 (MASIF), and FIPA. However, FIPA has emerged as the leading and significant active organization. Relying on input from its members as well as the agent research community in general, this organization was formed to produce and promote standards for heterogeneous and interacting agents and agent-based systems.

Among other important activities, FIPA is currently addressing various issues directly related to agent communication and interaction:

- Develop a new semantic framework to reflect the need for verifiability and conformance. The objective is to particularly adopt or define a semantic framework that can give an account of FIPA's existing communication acts and interaction protocols as well as a number of additional constructs such as contracts, agreements, policies, trust, agent descriptions, etc.

- Establish a *roadmap* on how to develop a second generation of new interaction protocols. Indeed, the literature regarding agent-oriented systems has repeatedly outlined the importance of developing new interaction protocols in accordance with parallel efforts directed towards building open and scalable applications. In the medium term, it is likely that future communication and interaction protocols will be largely agreed upon and standardized.

- Standardize methods for knowledge sharing and filtering through ontological representation. Such an approach will allow interoperating systems to automate message processing with respect to cross-referenced

[9] http://www.cs.umbc.edu/kqml/
[10] http://www.objs.com/survey/omg.htm
[11] http://www.fokus.gmd.de/research/cc/ecco/masif/

semantic classification. These efforts will enable a structured, standardized approach designed to support multiple ontologies.

However, rather than specifying how agents process and reason about received information, FIPA focuses on specifying communication and interoperability between heterogeneous entities. In other words, this organization concentrates on agent interfaces or on specifying external communication rather than the internal processing of the communication at the receiver. In addition, a drawback associated with FIPA standards lies in the absence of tools offering the ability to certify FIPA-compliant platforms. Obviously, this drawback comes from the difficulty to validate platform behaviour through an interface designed to allow agents to exchange messages.

Furthermore, a crucial requirement of future ubiquitous, open, and scalable agent-based systems is to address the issue of having to use a large number of alternative protocols designed for specific interactions. To deal with more sophisticated interactions, future agent-oriented applications are likely to incorporate public libraries of specifically dedicated interaction protocols. Some protocols will likely use existing communication languages such as the Contract Net and other protocols based on the Dutch auction and the English auction which use FIPA ACL. On the other hand, other protocols may be implemented using in-house or ad hoc communication languages like many dialogue game protocols designed for agent argumentation.

Based on current debate over future agent-oriented applications, it is likely that activities parallel to on-going standardization efforts will be conducted to develop approaches and tools that provide heterogeneous agents in open systems with the ability to collectively evolve communication languages and interaction protocols suited for an application domain and participating agents. Research addressing this particular issue is both complex and challenging, in that it is designed to allow a group of agents with no prior experience of each other to evolve a sophisticated communication language or interaction protocol. Such research is expected to draw from linguistics, social anthropology, biology, information theory, etc**.**

9.2.3 Incorporate learning and reasoning capabilities

Coordinating heterogeneous interactions in open and scalable systems spanning across different real-world domains raises the challenging problem of uncertainty in terms of dynamism, observability, and non-determinism. In the absence of common pre-established protocols to deal with large numbers of agents, each with its own beliefs, goals and intentions, inter-agent coordination in open and scalable systems will, as already noted, rely on libraries of alternative standardized interaction protocols. In this context, designers will incorporate into future open systems learning and reasoning ability to allow heterogeneous agents to make rational decisions under uncertainty, update beliefs from information, determine new intentions, act on the basis of these intentions, and dynamically select an appropriate interaction protocol or mechanism.

Agent organizations

In "Agent-based computing: Promise and Perils" [269], Jennings argues that the development of robust and scalable systems requires agents that can *complete* their *objectives* while situated in a dynamic and uncertain environment, engage in rich *high-level interactions*, and operate within *flexible organizational structures*. Moreover, he observes that the problem regarding scalability of agent-oriented systems is increasingly complex as the number of agents in a system will, at a given time, fluctuate significantly in open and dynamic environments [226].

Advocated by Foster et al. [272], the concept of *virtual organization* can be used to design adequate flexible organizational structures in future open and scalable agent-based applications. In the area of Grid computing, a virtual organization refers to dynamic collections of individuals, institutions, and resources. Thus, the specific problem that underlies the Grid concept is *coordinated resource sharing and problem solving in dynamic, multi-institutional virtual organizations*.

On the other hand, in addressing challenging organizational problems, researchers introduced various notions such as *self-organization* and *adaptation* to support the design and deployment of complex systems suitable for real-world application domains. In other words, open agent-oriented systems should be both *self-building* – agents are able to determine the most appropriate organizational structure for the system – and *adaptive* – agents change this structure as the environment changes. Thus, self-organization represents "the process of generating social structure, adapting and changing organizational structure" [270]. It is the result of individual choices by a set of agents to engage in interaction in certain organizational patterns, depending on their own resources and the environmental context, enabling agents to reason and have such kinds of behaviours.

On the other hand, existing agent organizations cannot handle adequately issues inherent in open agent-oriented systems, such as agent heterogeneity, trust and accountability, failure and recovery, and societal change. Future research efforts may thus partly focus on developing appropriate representations of analogous computational concepts analogous to norms, rights, legislation, authorities, enforcement, etc. In this context, researchers may need to draw on political science and sociology to develop sophisticated agent societies.

Coalition formation

Agent organizations in future open and scalable systems are likely to involve dynamic coalitions of small groups so as to provide more and better services than a single group. Through coalition formation, agents in large and open systems faced with a set of goals can thus partition themselves to maximize system performance [271]. Hence, self-organization can be achieved through coalition formation. While it provides better saving of time and labour, automating coalition formation may also lead to "better coalitions than humans in complex settings".

In the past, coalition formation has been addressed in game theory to typically deal with centralized situations. Thus, computationally infeasible, the approach suffers from a number of serious drawbacks. Though it generally favours big coalitions, coalition formation is applicable for a small number of agents, thus limiting the scope of the application. However, research regarding the use of a dialect of modal propositional dynamic logic (PDL) to model games and interactions leads to appropriate representation of coalitions. This method may indeed prove efficient in introducing and formalizing reasoning about coalitions of agents.

Negotiation and argumentation

Negotiation and argumentation strategies will certainly play an important role in designing and deploying future commercial and industrial agent-oriented applications. However, research regarding negotiation in agent-based systems – see chapter 4 – can be described as disparate efforts and examples rather than a coherent negotiation science. Efforts devoted in this area have not yet provided a computational agent capable of effective negotiation in any arbitrary context. Strategies identified by economic or theoretic reasoning tend, for example, to be specific to the auction or game mechanism involved.

In addition, considered as more complex than auctions and game theoretic mechanisms, negotiation and deliberation mechanisms are not fully investigated. Parallel to enhancing research in negotiation mechanisms, appropriate argumentation strategies and mechanisms should also be investigated, developed and deployed.

Finally, future deployment of existing negotiation and argumentation strategies and mechanisms requires the need to:

- Adequately evaluate existing algorithms to determine their strengths and weaknesses in more realistic environments.

- Identify the circumstances in which well-evaluated and adequate algorithms can be implemented.

- Design and develop negotiation algorithms adapted to more open, complex environments – e.g. argumentation.

- Develop adequate argumentation engines including more adapted argumentation strategies.

- Develop appropriate strategies and techniques to allow agents to identify, create, and dissolve coalitions in inter-agent negotiation and argumentation.

- Develop ***domain-specific models of reasoning and argumentation***.

- Develop agent **ability to adapt to changes in environment**, etc.

This page intentionally left blank.

10. Conclusion and recommendations

Inherently cross-disciplinary in nature, the coordination concept is located right at the intersection of various disciplines, including computer science, sociology, psychology, biology, political science, managerial and organization sciences, and economics. While there is a broad consensus that the issue of coordination is highly interdisciplinary, the lack of a true trans-borders flow of information and knowledge between disciplines is however commonly acknowledged.

In the absence of a desirable systematic integration of ideas, concepts, and paradigms from different disciplines, trans-disciplinary exchange of knowledge pertaining to coordination, rather that being well-established, can only be described as sporadic. A strongly advocated solution thus lies in the need to bring the concerned research communities driven by common research goals around concerted actions designed to organize heterogeneous views, ideas, concepts, methods, and approaches from different disciplines within a common conceptual framework.

Common to this perspective is the idea that such actions would offer the prospect of gradually inducing cross-fertilization in the interdisciplinary research area of coordination. In this respect, designing a set of classification schemes to group a large diversity of concepts, models, and approaches is identified in this report as an attractive solution for making order out of the chaos and mess characterizing research pertaining to coordination. It is thus highly recommended to focus on defining and establishing alternative sub-themes taxonomies that are stable, formal, and recognized across disciplines in order to bridge gaps, facilitate the integration of fragmented results, bring clusters of coordination experts together, establish a true collaboration across traditional disciplinary boundaries, and provide an effective transfer of coordination knowledge and technology to industry.

From a lack of a unitary view across disciplines, a global and unanimous consensus is forged on the common idea that the concept of coordination, though ill-defined, is inherently vital to building complex systems and organizations. However, the disparate area of research pertaining to coordination is still characterized by very serious and frustrating inconsistencies in terminology related to vital corresponding objects, attributes, or features. Indeed, a key recommendation of the present report is to bring together a panel of experts from various disciplines in order to establish and incorporate a specific standard terminology relative to coordination. Such standard vocabulary would be used as a foundation for defining concepts, expressing models and frameworks, and specifying protocols and mechanisms.

Characterizing the state-of-the-art research in the area of coordination is a particularly complex endeavour, as the space of available approaches is indeed considerable, and research is independently conducted in a great number of domains. While existing surveys deal with specific aspects of coordination, the major contribution of this study lies in the attempt to provide an in-depth review covering a wide range of issues regarding multi-agent coordination in DAI. In addition to reporting various sources of confusion, this report outlines the existence of a plethora of protocols and mechanisms adapted to different problems, agent-oriented systems, environments, and domains. In short, the current study identifies the absence of a single unified approach to address multi-agent coordination problems arising in any system or organization.

In the presence of multiple coordination strategies and a corresponding wide range of protocols and mechanisms, developing adequate performance-based measurement procedures represents an important cornerstone of the area of DAI for building robust agent-oriented systems and applications. While reporting about only few studies dedicated to this issue, the present report outlines the absence of a commonly accepted standard and highlights the need for additional efforts to address the specific need for developing coordination performance measurement criteria and metrics.

Based on the examination of current status and future opportunities, the present report identifies major challenges in the area of multi-agent coordination over the next decade. Indeed, currently deployed agent-oriented applications can be typically characterized as closed systems incorporating predefined communication and interaction protocols. In the long term, current research efforts will however lead to the emergence of ubiquitous, truly open, fully scalable agent-based systems, spanning across multiple domains, and integrating arbitrary numbers of heterogeneous agents capable of reasoning and learning appropriate communication and interaction protocols.

Over the next decade, research in the area of multi-agent systems is likely to be directed towards bridging the gap between academic promises and reality. In this respect, particular efforts will focus on enhancing the quality of agent-based systems to industrial standards. In the mid-term, the research community will thus concentrate on providing widely accepted analysis and design tools and methodologies in order to remove a major obstacle to future deployment of highly sophisticated, open and scalable systems. On the other hand, similar collective efforts will trigger the use of libraries of standardized communication and interaction protocols. In addition, more complex approaches are likely to be developed to allow future open and scalable agent-oriented systems to evolve communication languages and interaction protocols specific to application domains and agents involved.

Finally, the emergence of various concepts such as virtual organizations and self-organization will result in the use of dynamic coalitions – coalitions that automatically form and disband – of small groups of agents equipped with learning and reasoning abilities. Related to this challenge, various issues that will be similarly addressed to devise novel concepts, approaches, and techniques include domain-specific models of reasoning, negotiation and argumentation. Future research associated with open and scalable systems will also be devoted to other issues related to inter-agent coordination: reusability, mobility, emergent behaviour, and fault-tolerance.

Two mains recommendations are concluded in this study:

1. Encourage increased deployment of agent-oriented technologies in the Canadian military sector.

- Identify domains where agent technology can be effectively and successfully deployed.
- Identify characteristics of agent-based real-world military applications in relation to specific military domains.
- Investigate opportunities for smooth migration from legacy systems to agent-oriented solutions, products, services, etc.
- Stimulate the need for agent-oriented technologies in all areas of the military sector.
- Encourage military take-up of agent-based technologies through provision of reasons for success.
- Provide a catalogue of practical military cases together with potential adopters of agent-based technologies.
- Evaluate past agent-oriented research conducted within the DRDC-Valcartier in order to identify success and failure.
- Design and build prototypes spanning different military domains.

2. Promote agent-oriented research activities within different participating structures of the DRDC.

- Propagate and popularize agent technologies in different departments and research groups within the DRDC.
- Seek, build, and coordinate a formal network grouping different research academic establishments and agencies in Canada and North America working in the area of agent-oriented systems.
- Seek and build links with researchers in other areas such as sociology, psychology, political science, organizational science, and economics in order to acquire relevant knowledge that can be implemented in agent-based systems.
- Bring researchers from the established formal network as well as from other disciplines to debate important issues regarding multi-agent research.
- Design and launch a newsletter to periodically report important results stemming from agent-oriented research activities.

This page intentionally left blank.

11. References

[1] Harold Seidman and Robert Gilmou, "*Politics, position, and power-from the positive to the regulatory state*", 4th ed., Oxford University Press, New York, 1986.

[2] Thomas W. Malone and Kevin Crowston, "*What is Coordination Theory and How Can it Help Design Cooperative Work Systems*", In Readings in Groupware and Computer-Supported Cooperative Work, Assisting Human-Human Collaboration, Ed. Ronald M. Baecker, Morgan Kaufmann Publishers, Inc., San Fransisco, California, 1993.

[3] B. A. Huberman, "*The Ecology of Computation*", In the Ecology of Computation, Elsevier Science Publishers, Amsterdam, Holland, 1988.

[4] Eric Bonabeau, Andrej Sobkowski, Guy Theraulaz, and Jean-Louis Deneubourg, "*Adaptive Task Allocation Inspired by a Model of Division of Labor in Social Insects*", In Biocomputing and Emergent Computation, publisher World Scientific Publisher, Dan Lundh, Bjorn Olsson, and Ajit Narayanan Editors., pp. 36-45, 1997.

[5] Deborah Gordon, "Ants at Work: How an Insect Society is Organized", Free Press, October 1999.

[6] Carl Anderson, Nigel R. Franks and Daniel W. McShea, "*The Complexity and Hierarchical Structure of Tasks in Insect Societies*", In Animal Behavior, No 62, pp. 643-651, 2001.

[7] Graham Coates, Robert Ian Whitfield, Alex H. B. Duffy and Bill Hills, "*Coordination Approaches and Systems – Part II: An Operational Perspective*" In Research in Engineering Design, No 12, pp. 73-79, 2000.

[8] Riyaz Sikora and Michael J. Shaw, "*A Multi-Agent Framework for the Coordination and Integration of Information Systems*", In Management Science, Vol. 11, Issue 11, Part 2 of 2, November 1998.

[9] Alan H. Bond and Les Gasser, "An Analysis of Problems and Research in DAI", In Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, editors, Morgan Kaufmann, Santa Mateo, CA, 1988.

[10] Edmund H. Durfee and Jeffrey S. Rosenschein. "*Distributed Problem Solving and Multi-Agent Systems: Comparisons and Examples*" In Proceedings of the Thirteenth International Distributed Artificial Intelligence Workshop, pp. 94-104, July 1994.

[11] Gerhard Weiss, “Multiagent Systems and Distributed Artificial Intelligence”, In Multiagent Systems – A Modern Approach to Distributed Artificial Intelligence”, Gerhard Weiss, Editor, the MIT Press, Cambridge, Massachusetts, 1999.

[12] Claudia V. Goldman, “Emergent Coordination Through the Use of Cooperative State-Changing Rules”, MSc. Thesis, Computer Science Department, the Hebrew University of Jerusalem, Jerusalem, August 1993.

[13] Eithan Ephrati, “Planning and Consensus among Autonomous Agents”, PhD Thesis, Computer Science Department, the Hebrew University of Jerusalem, Jerusalem, October 1993.

[14] Robert Tolksdorf, “*Models of Coordination*”, In Engineering Societies in the Agents’ World, A. Omicini, R. Tolksdorf, and F. Zambonelli, Editors, ESAW 2000, LNAI 1972, Springer-Verlag, Berlin Heidelberg, 2000.

[15] Nicholas V. Findler and Gregory D. Elder, “*Multi-Agent Coordination and Cooperation in a Dynamic Environment with Limited Resources*”, In Artificial Intelligence in Engineering, Vol. 9, pp. 229-238, 1995.

[16] Rina Azoulay-Schwartz and Arit Kraus, “*Negotiation on Data Allocation in Multi-agent Environments*”, In Autonomous Agents and Multi-agent Systems, Vol. 5, No 2, pp. 123-172, 2002.

[17] Edmund H. Durfee, “*Distributed Problem Solving and Planning*”, In Multiagent Systems – A Modern Approach to Distributed Artificial Intelligence, the MIT Press, Cambridge, Massachusetts, 1999.

[18] Brahim Chaib-draa, “Industrial Application of Distributed AI”, In Communications of the ACM, Vol. 38, No 11, November 1995.

[19] Thomas W. Malone and Kevin Crowston, “the Interdisciplinary Study of Coordination”, In ACM Computing Surveys, Vol. 26, No 1, March 1994.

[20] Kevin Crowston and Charles S. Osborn, “A coordination theory approach to process description and redesign”, Report submitted to the Center for Coordination Science, Sloan School of Management, Massachusetts Institute of Technology – MIT, July 1998.

[21] Thomas W. Malone, Kevin Crowston, J. Lee, and B. Pentland, “Tools for Inventing Organizations: Towards a Handbook of Organizational Processes”, In Proceedings of the 2nd IEEE Workshop on Enabling Technologies Infrastructure for Collaborative Enterprises, IEEE, New York, 1993.

[22] Malone, T.W., K. Crowston, Lintae Lee, Brian Pentland, Chrysanthos Dellarocas, George Wyner, John Quimby, Charles S. Osborn, Abraham Bernstein, George Herman, Mark Klein, and Elissa O'Donnell, "Tools for Inventing Organizations: Toward a Handbook of Organizational Processes", In Management Science, Vol. 45, No 3, pp. 425-443, 1999.

[23] Kevin Crowston, "A Taxonomy of Organizational Dependencies and Coordination Mechanisms", In Tools for Organizing Business Knowledge: the MIT Process Handbook, Malone, T. W., Crowston, K. and Herman, G. (Eds.), Cambridge, MA, MIT Press.

[24] M. Klein and C. Dellarocas, "Towards a Systematic Repository of Knowledge about Managing Multi-Agent System Exceptions", Massachusetts Institute of Technology, Cambridge, MA, USA, ASES Working Paper ASES-WP-2000-01, February 2000.

[25] M. Schumacher, "Objective Coordination in Multi-agent Systems Engineering: Design and Implementation", Lecture Notes in Artificial Intelligence, Vol. 2039, Springer Verlag, Heidelberg, Germany, April 2001.

[26] M. Schumacher, "Multi-agent Systems", In Objective Coordination in Multi-agent Systems Engineering: Design and Implementation, M. Schumacher, Editor, Lecture Notes in Artificial Intelligence, Vol. 2039, Springer Verlag, Heidelberg, Germany, April 2001.

[27] S. Ossowski, "*Co-ordination in Artificial Agent Societies*" In Social Structure and Its Implications for Autonomous Problem-Solving Agents, Vol. 1202, LNAI, Springer Verlag, 1999.

[28] S. Bussmann and J. Mueller, "A Negotiation Framework for Co-operating Agents", In *Proceedings of CKBS-SIG,* S. M. Deen, editor, pages 1–17, Dake Centre, University of Keele, 1992.

[29] H. S. Nwana and L. C. Lee and N. R. Jennings, "Coordination in Software Agent Systems", In the British Telecom Technical Journal, Vol. 14, No 4, pp. 79-88, 1996.

[30] E. H. Durfee, V. R. Lesser and D. D. Corkill, "Coherent Cooperation among Communicating Problem Solvers", IEEE Transactions on Computers, 36 (11), pp. 1275—1291, 1987.

[31] K. Werkman, *"Knowledge-based model of negotiation using shareable perspectives"*, In the Proceedings of the 10th International Workshop on DAI, Texas, 1990.

[32] P. Kearney, A. Sehmi, and R. Smith, "*Emergent behavior in a multi-agent economics simulation*", In the Proceedings of the 11th European

Conference on Artificial Intelligence, A. G. Cohn, Editor, London, John Wiley, 1994.

[33] V. Lesser and D. Corkill, “the Distributed Vehicle Monitoring Testbed: a Tool for Investigating Distributed Problem-Solving Networks”, In Artificial Intelligence Magazine, Vol. 4, No 3, pp 15—33, 1983.

[34] R. Davis and R. G. Smith, “*Negotiation as a Metaphor for Distributed Problem Solving*”, In Artificial Intelligence, No 20, pp. 63-109, 1983. – Reprinted in Communications in Multi-agent Systems, LNAI 2650, M.-P. Huget, Editor, pp. 51-97, 2003.

[35] R. G. Smith, “the Contract Net Protocol: High-Level Communication and Control in a Distributed Problem Solver”, In IEEE Transactions on Computers, Vol. 29, No 12, December 1980.

[36] H. Van Dyke Parunak, “Manufacturing Experience with the Contract Net”, In Distributed Artificial Intelligence, M. N. Huhns, Editor, pp. 285-310, 1987.

[37] Suzan E. Conry, Robert A. Meyer and Victor R. Lesser, “Multistage Negotiation in Distributed Planning”, In Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, Editors, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[38] M. P. Georgeff, “Communication and Interaction in Multi-agent Planning”, In Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, Editors, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[39] M. P. Georgeff, “A Theory of Action for Multi-agent Planning”, In Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, Editors, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[40] S. Cammarata, D. McArthur, R. Steeb, “Strategies of Cooperation in Distributed Problem Solving”, In Readings in Distributed Artificial Intelligence, Alan H. Bond and Les Gasser, Editors, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1988.

[41] Y. Jin and T. Koyama, “Multi-agent Planning Through Expectation-based Negotiation” In Proceedings of the 10th International Workshop on Distributed Artificial Intelligence, M. N. Huhns, Editor, Texas University, 1990.

[42] D. D. Corkill, “Hierarchical Planning in a Distributed Environment”, In Proceedings of the 6th International Joint Conference on Artificial Intelligence, pp. 168-179, August 1979.

[43] Victor R. Lesser and Daniel D. Corkill, "Functionally Accurate, Cooperative Distributed Systems", In IEEE Transactions on Systems, Man, and Cybernetics, Vol. 11, No 1, pp. 81-96, January 1981.

[44] Victor R. Lesser and Daniel D. Corkill, "the Distributed Vehicle Monitoring Testbed: A Tool for Investigating Distributed Problem-Solving Networks", In AI Magazine, Vol. 4, No 3, pp 15-33, 1983.

[45] Edmund H. Durfee and Victor R. Lesser, "Using Partial Global Plans to Coordinate Distributed Problem Solvers", In Proceedings of the 1987 International Joint Conference on Artificial Intelligence, pp. 875-883, 1987.

[46] M.N. Huhns and M. P. Singh, "CKBS-94 Tutorial: Distributed Artificial Intelligence for Information Systems", Dake Centre, University of Keele, 1994.

[47] S. Bussmann and J. Muller, "A Negotiation Framework for Co-operating Agents", In Proceedings of the Special Interest Group on Cooperating Knowledge Based Systems (CKBS-SIG), S. M. Deen, Editor, pp. 1-17, Dake Centre, University of Keele, 1992.

[48] Katia P. Sycara, "Multi-agent Compromise via Negotiation", In Distributed Artificial Intelligence 2, Les Gasser and M. Huhns, Editors, Morgan Kaufmann Publishers, Inc., San Mateo, California, 1989.

[49] Jeffrey S. Rosenschein, "Rational Interaction: Cooperation among Intelligent Agents", PhD Thesis, Stanford University, 1985.

[50] Jeffrey S. Rosenschein and Gilad Zlotkin, "Rules of Encounter: Designing Conventions for Automated Negotiation among Computers", MIT Press, 1994.

[51] Gilad Zlotkin and Jeffrey S. Rosenschein, "Blocks, Lies, and Postal Freight: the Nature of Deception in Behavior", In Proceedings of the 10th International Workshop on DAI, 1990.

[52] Rosenschein, J.S. and Zlotkin, G., "Rules of Encounter: Designing Conventions for Automated Negotiation among Computers", MIT Press, 1994,

[53] Kraus, S. and Wilkenfield, J., "the Function of Time in Cooperative Negotiations: Preliminary Report", Department of Computer Science, University of Maryland, 1991.

[54] M. Busuioc and C. Winter, "Negotiation and Intelligent Agents", Project NOMADS-001, BT Labs, Martlesham Heath, U.K., 1995.

[55] G. Zlotkin, G. and J. S. Rosenschein, J.S., “One, Two, Many: Coalitions in Multi-Agent Systems”, In Proceedings Of the 5th European Workshop on Modeling Autonomous Agents in Multi-Agent World, August, 1993.

[56] M. R. Alder, A. B. Davis, R. Weihmayer, and R. W. Forest, “Conflict-resolution strategies for non-hierarchical distributed agents”, In Distributed Artificial Intelligence II, L. Gasser and M. N. Huhns, Editors, Morgan Kaufmann, 1989.

[57] T. Kreifelt and von F. Martial, “A Negotiation Framework for Autonomous Agents”, In Decentralized Artificial Intelligence, Y. Demazeau and J. P. Muller, Editors, Vol. 2, Elsevier Science, 1991.

[58] S. Bussmann and J. Muller, “A Negotiation Framework for Cooperating Agents”, In Proceedings CKBS-SIG, S. M. Deen, Editor, Dake Centre, University of Keele, pp. 1-17, 1992.

[59] K. J. Werkman, K.J., “Knowledge Based Model of Negotiation using Shareable Perspectives”, In the proceedings of the 10th International Workshop on DAI, Texas, 1990.

[60] K. Sycara, “Multi-agent Compromise via Negotiation”, In Distributed Artificial Intelligence II, L. Gasser and M. Huhns, Editors, Morgan Kaufmann, 1989.

[61] Sathi and M. S. Fox, “Constraint-directed Negotiation of Resource Allocations”, In Distributed Artificial Intelligence II, L. Gasser and M. Huhns, Editors, Morgan Kaufmann, 1989.

[62] D. G. Pruitt, “Negotiation Behavior”, Academic Press, New York, N.Y., 1981.

[63] S. E. Conry, K. Kuwabara, V. R. Lesser, and R. A. Meyer, “Multistage negotiation for distributed satisfaction”, In IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence, Vol. 21, No 6, pp. 1462–1477, December 1991.

[64] P. H. Gulliver, “Disputes and Negotiations - a Cross-cultural Perspective”, Academic Press, 1979.

[65] K. Kuwabara and V. R. Lesser, “Extended Protocol for Multistage Negotiation”, In Proceedings of the 9th Workshop on Distributed Artificial Intelligence, 1989.

[66] E. H. Durfee and T. A. Montgomery, “a Hierarchical Protocol for Coordinating Multi-agent Behavior”, In Proceedings of the 8th National Conference on Artificial Intelligence, Boston, MA, pp. 86-93, 1990.

[67] N. R. Jennings, “Coordination Techniques for Distributed Artificial Intelligence”, In Foundations of Distributed Artificial Intelligence, G. M. P. O’Hare and N. R. Jennings, Editors, John Wiley & Sons, Inc., New York, p. 187-210, 1996.

[68] Sarit Kraus, “Automated Negotiation and Decision Making in Multiagent Environments”, Lecture Notes in Artificial Intelligence 2086, pp. 150-172, 2001.

[69] M. Wooldridge, “Reaching Agreements”, In an Introduction to Multi-agent Systems, John Wiley & Sons, Ltd, 2002.

[70] Tuomas W. Sandholm, “Distributed Rational Decision Making”, In Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, Gerhard Weiss, Editor, The MIT Press, Cambridge, MA, USA, pp. 201—258, 1999.

[71] T. Moehlman, V. Lesser, and B. Buteau, “Decentralized negotiation: An approach to the distributed planning problem”, In Group Decision and Negotiation, Vol. 2, No 2, pp. 161–191, January 1992.

[72] Susan E. Lander and Victor R. Lesser, “Customizing Distributed Search among Agents with Heterogeneous Knowledge”, In Proceedings of the First International conference on Information Knowledge Management, pp. 335–344, Baltimore, 1992.

[73] K. P. Sycara, “Resolving Adversarial Conflicts: An Approach to Integrating Case-Based and Analytic Methods”, PhD thesis, School of Information and Computer Science, Georgia Institute of Technology, 1987.

[74] K. P. Sycara, “Persuasive argumentation in negotiation”, In Theory and Decision, Vol. 28, pp. 203–242, 1990.

[75] S. Kraus and D. Lehmann, “Designing and building a negotiating automated agent”, In Computational Intelligence, Vol. 11, No 1, pp.132–171, 1995.

[76] Dajun Zeng and Katia Sycara, “Bayesian learning in negotiation”, In International Journal of Human-Computer Studies, Vol. 48, pp. 125–141, 1998.

[77] C. Sierra, P. Faratin, and N. Jennings, “A Service-oriented Negotiation Model between Autonomous Agents”, In Proceedings of the 8th European Workshop on Modeling Autonomous Agents in a Multi-Agent World (MAAMAW-97), pp. 17–35, Ronneby, Sweden, 1997.

[78] T. W. Sandholm and V. R. Lesser, “Issues in Automated Negotiation and Electronic Commerce: Extending the contract net framework, In *First* International Conference on Multiagent Systems *(ICMAS-95)*, pp. 328–335, San Francisco, 1995.

[79] M. J. Osborne and A. Rubinstein, “Bargaining and Markets”, Academic Press Inc., San Diego, California, 1990.

[80] E. Roth, “Axiomatic Models of Bargaining”, Lecture Notes in Economics and Mathematical Systems No. 170, Springer-Verlag, Berlin, 1979.

[81] H. Raiffa “The Art and Science of Negotiation”, Harvard University Press, Cambridge, MA, 1982.

[82] R. Fisher and W. Ury, “Getting to Yes: Negotiating Agreement Without Giving In”, Houghton Mifflin, Boston, 1981.

[83] D. Druckman, “Negotiations”, Sage Publications, London, 1977.

[84] L Karrass, “The Negotiating Game: How to Get What You Want”, Thomas Crowell Company, NY, 1970.

[85] R. Johnson, “Negotiation basics”, Sage Publications, Newbury Park, 1993.

[86] Lavinia Hall, editor, “Negotiation: Strategies for Mutual Gain”, Sage Publications, Newbury Park, CA, 1993.

[87] Sarit Kraus, “Strategic Negotiation in Multiagent Environments”, the MIT Press, Cambridge Massachusetts, 2001.

[88] Rubinstein, “Perfect equilibrium in a bargaining model”, In *Econometrica*, Vol. 50, No 1, pp. 97–109, 1982.

[89] R. Schwartz and S. Kraus, “Negotiation on data allocation in multi-agent environments”, In Proceedings of AAAI-97, pp. 29–35, Providence, Rhode Island, 1997.

[90] M. Ben-Or and N. Linial, “Collective Coin Flipping, Robust Voting Games and Minima of Banzhaf Values”, In Proceedings of the 26th IEEE Symposium on the Foundations of Computer Science, pp. 408–416, Portland, 1985.

[91] M. J. Osborne and A. Rubinstein, “A Course in Game Theory”, MIT Press, Cambridge, Massachusetts, 1994.

[92] R. D. Luce and H. Raiffa, “Games and Decisions”, John Wiley and Sons, New York, 1957.

[93] J. F. Nash, “Two-person Cooperative Games”, Econometrica, Vol. 21, pp. 128–140, 1953.

[94] Jean Tirole, “the Theory of Industrial Organization”, the MIT Press, Cambridge, MA, 1988.

[95] Kreps and R. Wilson, “Sequential equilibria”, In Econometrica, vol. 50, pp. 863–894, 1982.

[96] J. Wilkenfeld, S. Kraus, K. Holley, and M. Harris, “Genie: A decision support system for crisis negotiations”, In Decision Support Systems, Vol. 14, pp. 369–391, 1995.

[97] S. Kraus and J. Wilkenfeld, “A Strategic Negotiations Model with Applications to an International Crisis”, In IEEE Transaction on Systems Man and Cybernetics, Vol. 23, No 1, pp. 313—323, 1993.

[98] Amir Ronen, “Algorithms for rational agents”, In Conference on Current Trends in Theory and Practice of Informatics, pp. 56–70, 2000.

[99] T. Sandholm, “An Implementation of the Contract Net Protocol based on Marginal Cost Calculations”, In Proceedings of AAAI-93, pp. 256–262, Washington D.C., July 1993.

[100] M. R. Andersson and T. W. Sandholm, “Contract Types for Optimal Task Allocation: II Experimental Results”, In AAAI 1998 Spring Symposium: Satisficing Models, Stanford University, California, 1998.

[101] Gimenez-Funes, L. Godo, and J. A. Rodriguez-Aguilar, “Designing Bidding Strategies for Trading Agents in Electronic Commerce, In ICMAS98, pp. 136–143, Paris, 1998.

[102] R. Wurman, W. E. Walsh, and M. P. Wellman, “Flexible Double Auctions for Electronic Commerce: Theory and Implementation”, In Decision Support Systems, Vol. 24, pp. 17–27, 1998.

[103] P. Klemperer, “Auction Theory: a Guide to Literature”, In Journal of Economic Surveys, Vol. 13, No 3, 1999.

[104] William Vickrey, “Counterspeculation, auctions, and competitive sealed tenders”, In Journal of Finance, Vol. 16, pp. 8-37, March 1961.

[105] Huberman and S. H. Clearwater, “A multi-agent system for controlling building environments” In Proceedings of ICMAS-95, pp. 171–176, San Francisco, 1995.

[106] R. Schwartz and S. Kraus, “Bidding mechanisms for data allocation in multiagent environments”, In Intelligent Agents IV: Agent Theories,

Architectures, and Languages, Munindar P. Singh, Anand S. Rao, and M.J. Wooldridge, Editors, pp. 61–75, Springer-Verlag, 1998.

[107] M. B. Tsvetovatyy and M. Gini, “Toward a virtual marketplace: Architectures and strategies”, In the first International Conference on the Practical Application of Intelligent Agents and Multi Agents Technology, pp. 597–614, London, 1996.

[108] M. Tsvetovatyy, M. Gini, B. Mobasher, and Z. Wieckowski, “Magma: An agent-based virtual market for electronic commerce”, In Applied Artificial Intelligence, special issue on Intelligent Agents, Vol. 6, pp. 501–523, 1997.

[109] M. H. Rothkopf, A. Pekec, and R. M. Harstad, “Computationally manageable combinatorial auctions”, Technical Report 95-09, DIMACS, April 19 1995.

[110] Daniel Lehmann, Liadan Ita O’Callaghan, and Yoav Shoham, “Truth Revelation in Rapid, Approximately Efficient Combinatorial Auctions”, In Proceedings of the ACM Conference on Electronic Commerce (EC‘99*)*, pages 96–102, N.Y., November 3–5 1999. ACM Press.

[111] Y. Fujishima, K. Leyton-Brown, and Y. Shoham, “Taming the Computational Complexity of combinatorial Auctions: Optimal and Approximate Approaches”, In Dean Thomas, editor, Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol1), pages 548–553, S.F., July 31–August 6 1999, Morgan Kaufmann Publishers.

[112] L. Hunsberger and B. Grosz, “a Combinatorial Auction for Collaborative Planning”, In ICMAS-2000, Boston, USA, 2000.

[113] Noam Nisan, “Bidding and Allocation in Combinatorial Auctions”, In ACM Conference on Electronic Commerce”, pp. 1–12, 2000.

[114] R. Wilson, “Incentive Efficiency of Double Auctions”, Econometrica, Vol. 53, No 5, pp. 1101–1116, 1985.

[115] P. R. Wurman, W. E. Walsh, and M. P. Wellman, “Flexible Double Auctions for Electronic Commerce: Theory and Implementation”, Decision Support Systems, Vol. 24, pp. 17–27, 1998.

[116] M. Wellman, “A Market-oriented Programming Environment and its Application to Distributed Multicommodity Flow Problems”, Journal of Artificial Intelligence Research, Vol. 1, pp. 1–23, 1993.

[117] M. P. Wellman and P. R. Wurman, “Market-aware agents for a multi-agent world”, In Robotics and Autonomous Systems, Vol. 24, pp. 115–125, 1998.

[118] Gerber, C. Russ, and G. Vierke, “On the Suitability of Market-based Mechanisms for Telematics Applications”, In *Proceedings of Agents99*, pp. 408–409, 1999.

[119] Rapoport, “N-Person Game Theory”, University of Michigan, Michigan, 1970.

[120] K. I. Shimomura, “the Bargaining Set and Coalition Formation”, Technical Report 95-11, Brown University, Department of Economics, 1995.

[121] R. Vohra, “Coalitional Non-cooperative Approaches to Cooperation”, Technical Report 95-6, Brown University, Department of Economics, 1995.

[122] L. Zhou, “a New Bargaining Set of an N-person Game and Endogenous Coalition Formation”, Games and Economic Behavior, Vol. 6, pp. 512–526, 1994.

[123] J. P. Kahan and A. Rapoport, “Theories of coalition formation” Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1984.

[124] O. Shehory and S. Kraus, “Feasible Formation of Stable Coalitions among Autonomous Agents in Non-super-additive Environments”, In Computational Intelligence, Vol. 15, No 3, pp. 218–251, 1999.

[125] M. Klusch and O. Shehory, “a Polynomial Kernel-oriented Coalition Formation Algorithm for Rational Information Agents”, In Proceeding of ICMAS-96, pp. 157–164, Kyoto, Japan, 1996.

[126] T. Sandholm, K. Larson, M. R. Andersson, O. Shehory, and F. Tohm, “Coalition Structure Generation with Worst Case Guarantees”, In Artificial Intelligence Journal, Vol. 111, pp. 209–238, 1999.

[127] T. W. Sandholm and V. R. Lesser, “Coalition Formation among bounded rational agents”, In Proceeding of IJCAI-95, pp. 662–669, Montreal, 1995.

[128] G. Zlotkin and J. S. Rosenschein, “Coalition, Cryptography, and Stability: Mechanisms for Coalition Formation in Task Oriented Domains”, In Proceeding of AAAI94, pp. 432–437, Seattle, Washington, 1994.

[129] S. P. Ketchpel, “Forming Coalitions in the Face of Uncertain Rewards”, In Proceedings of AAAI94, pp. 414–419, Seattle, Washington, 1994.

[130] O. Shehory and S. Kraus, “Methods for task allocation via agent coalition formation”, In Artificial Intelligence, Vol. 15, No 3, pp. 218–251, 1998.

[131] K. J. Arrow, “the Economics of Agency”, In Principals and Agents*:* the Structure of Business, J. Pratt and R. Zeckhauser, Editors, pp. 37–51, Harvard Business School Press, Cambridge, MA, 1985.

[132] S. Ross, “the Economic Theory of Agency: the Principal’s Problem”, In the American Economic Review, Vol. 63, No. 2, pp. 134–139, 1973.

[133] Rasmusen, “Games and Information*”*, Basil Blackwell Ltd., Cambridge, MA, 1989.

[134] S. Grossman and O. Hart, “an analysis of the principal-agent problem”, In *Econometrica*, Vol. 51, No. 1, pp. 7–45, 1983.

[135] J. Hirshleifer and J. Riley, “the Analytics of Uncertainty and Information”, Cambridge University Press, Cambridge, 1992.

[136] J. Laffont and J. Tirole, “a Theory of Incentives in Procurement and `Regulation”, the MIT Press, Cambridge, Massachusetts, 1993.

[137] Nalebuff and J. Stiglitz, “Information, Competition, and Markets”, In American Economic Review, Vol. 73, No. 2, pp. 278–283, 1983.

[138] S. Baiman and J. Demski, “Economically Optimal Performance Evaluation and Control Systems”, In Journal of Accounting Research, Vol. 18, pp. 184–220, 1980.

[139] Banerjee and A. Beggs, “Efficiency in Hierarchies: Implementing the First-best Solution by Sequential Actions”, In the Rand Journal of Economics, Vol. 20, No. 4, pp. 637–645, 1989.

[140] Macho-Stadler and J. P´erez-Castrillo, “Moral hazard and cooperation”, In Economics Letters, Vol. 35, pp. 17–20, 1991.

[141] Caillaud, R. Guesnerie, P. Rey, and J. Tirole, “Government Intervention in Production and Incentives Theory: a Review of Recent Contributions”, In the Rand Journal of Economics, Vol. 19, No. 1, pp. 1–26, 1988.

[142] Rubinstein and M. Yaari, “Repeated Insurance Contracts and Moral Hazard”, In Journal of Economic Theory, Vol. 30, pp. 74–97, 1983.

[143] M. Harris and A. Raviv, “Some Results on Incentive Contracts with Applications to Education and Employment, Health Insurance, and Law Enforcement”, In the American Economic Review, Vol. 68, No. 1, pp. 20–30, 1978.

[144] M. Spence and R. Zeckhauser, “Insurance, Information and Individual Action”, In the American Economic Review, Vol. 61, No. 1, pp. 380–391, 1971.

[145] M. Landsberger and I. Meilijson, “Monopoly Insurance under Adverse Selection when Agents Differ in Risk Aversion”, In Journal of Economic Theory, Vol. 63, pp. 392–407, 1994.

[146] S. Matthews, “Selling to Risk Averse Buyers with Unobservable Tastes”, In Journal of Economy Theory, Vol. 30, pp. 370–400, 1983.

[147] R. Myerson, “Mechanism Design by an Informed Principal”, In *Econometrica*, Vol. 51, pp. 1767—1798, 1983.

[148] R. P. McAfee and J. McMillan, “Bidding for Contracts: a Principal-agent Analysis”, In the Rand Journal of Economics, Vol. 17, No. 3, pp. 326–338, 1986.

[149] S. Shavell, “Risk Sharing and Incentives in the Principal and Agent Relationship”, In Bell Journal of Economics, Vol. 10, pp. 55–79, 1979.

[150] R. Smith and R. Davis, “Framework for Cooperation in Distributed Problem Solvers”, IEEE Transactions on Systems, Man and Cybernetic, Vol. C-29, No. 12, pp. 61–70, 1981.

[151] T. Ohko, K. Hiraki, and Y. Anzai, “Reducing Communication Load on Contract Net by Case-based Reasoning: Extension with Directed Contract and Forgetting”, In Proceedings of the First International Conference on Multi–Agent Systems, Cambridge, MA, MIT Press, 1995.

[152] T. W. Malone, R. E. Fikes, K. R. Grant, and M. T. Howard, “Enterprise: a Market-like Task Schedule for Distributed Computing Environments”, In The Ecology of Computation, B. A. Huberman, Editor, pp. 177–205, North Holland, 1988.

[153] Sandip Sen and Edmund Durfee, “a Contracting Model for Flexible Distributed Scheduling”, In Annals of Operations Research, Vol. 65, pp. 195–222, 1996.

[154] S. Kraus, “an Overview of Incentive Contracting”, In Artificial Intelligence Journal, Vol. 83, No 2, pp. 297–346, 1996.

[155] T. Sandholm, S. Sikka, and S. Norden, “Algorithms for Optimizing Leveled Commitment Contracts”, In Proceedings of IJCAI99, pp. 535–540, Stockholm, Sweden, 1999.

[156] T. Sandholm and Y. Zhou, “Surplus Equivalence of Leveled Commitment Contracts”, In Proceedings of ICMAS-2000, pp. 247–254, IEEE Computer Society, Boston, MA, 2000.

[157] N. R. Jennings, P. Faratin, A. R. Lomuscio, S. Parsons, C. Sierra, and M. Wooldridge, “Automated Negotiation: Prospects, Methods, and Challenges”, In International Journal of Group Decision and Negotiation, Vol. 10, No. 2, pp. 199-215, 2001.

[158] S. Parsons and N. R. Jennings, “Negotiation through Argumentation – a Preliminary Report”, In Proceedings of the 2nd International Conference on Multi-agent Systems (ACMAS-96), Kyoto, Japan, pp. 267-274, 1996.

[159] S. Parsons, C. A. Sierra, and N. R. Jennings, “Agents that Reason and Negotiate by Arguing”, Journal of Logic and Computation, Vol. 8, No. 3, pp. 261-292, 1998.

[160] S. Kraus, K. Sycara, and A. Evenchik, “Reaching Agreements through Argumentation: a Logical Model and Implementation”, In Artificial Intelligence, Vol. 104, pp. 1–69, 1998.

[161] M. Gilbert, “Multi-modal Argumentation”, Philosophy of the Social Sciences, Vol. 24, No. 2, pp. 159-177, 1994.

[162] J. Fox, P. Krause, and S. Ambler, “Arguments, Contradictions, and Practical Reasoning”, In Proceedings of the 10th European Conference on Artificial Intelligence (ICAI-92), Vienna, Austria, pp. 623-627, 1992.

[163] Paul Krause, Simon Ambler, Morten Elvang-Gøransson, and John Fox, “a Logic of Argumentation for Reasoning under Uncertainty”, In Computational Intelligence, Vol. 11, pp. 113-131, 1995.

[164] J. Grosz and S. Kraus, “Collaborative Plans for Complex Group Activities”, In Artificial Intelligence Journal, Vol. 86, No. 2, pp. 269–357, 1996.

[165] Z. Qiu and M. Tambe, “Flexible Negotiations in Teamwork: Extended Abstract”, In Proceedings of the AAAI Fall Symposium on Distributed Continual Planning, 1998.

[166] S. Toulmin, R. Rieke, and A. Janik, “An Introduction to Reasoning”, Macmillan Publishing Co., Inc., 1979.

[167] J. Boster and P. Mongeau, “Fear-arousing Persuasive Messages”, In *Communication yearbook 8*, R. N. Bostrom, Editor, SAGE Publications, 1984.

[168] Daniel J. O’Keefe, “Persuasion: Theory and Research”, SAGE Publications, 1990.

[169] M. E. Bratman, “Shared cooperative activity”, In the Philosophical Review, Vol. 101, pp. 327–341, 1992.

[170] J. Grosz and S. Kraus, “the Evolution of Sharedplans”, In Foundations and Theories of Rational Agency, A. Rao and M. Wooldridge, Editors, pp. 227–262, Kluwer Academic Publishers, 1999.

[171] Vermazen, “Objects of Intention”, Philosophical Studies, Vol. 71, pp. 223–265, 1993.

[172] M. Karlins and H. I. Abelson, “Persuasion: How Opinions and Attitudes are Changed”, Springer Publishing Company, Inc., second edition, 1970.

[173] R. R. Jr. Koballa, “Persuading Teachers to Reexamine the Innovative Elementary Science Programs of yesterday: The effect of Anecdotal versus Data-summary Communications”, In Journal of Research in Science Teaching, Vol. 23, pp. 437–449, 1986.

[174] S. E. Taylor and S. C. Thompsons, “Staking the Elusive Vividness Effect”, In Psychological Review, Vol. 89, pp. 155–181, 1982.

[175] R. E. Nisbett, E. Borgida, R. Crandall, and H. Reed, “Popular Induction: Information is not Necessarily Informative”, In Cognition and social behavior, Lawrence Erlbaum, J. S. Carroll and J. W. Payne, Editors, 1976.

[176] R. Hamill, T. D. Wilson, and R. E. Nisbett, “Insensitivity to sample bias: Generalizing from atypical cases”, In Journal of Personality and Social Psychology, Vol. 39, pp. 579–589, 1980.

[177] S. Parsons, M. Wooldridge, and L. Amgoud, “An Analysis of Formal Inter-agent Dialogues”, In 1st International Conference on Autonomous Agents and Multi-agent Systems, ACM Press, 2002

[178] S. Parsons, M. Wooldridge, and L. Amgoud, “On the Outcome of Formal Inter-agent Dialogues”, Proceedings of the Second International Joint Conference on Autonomous Agents and Multi-agent Systems (AAMAS 2003), Melbourne, Victoria, Australia, ACM 2003, July 14-18, 2003.

[179] S. Parsons and P. McBurney, "Argumentation-Based Dialogues for Agent Coordination", In Group Decision and Negotiation, Vol. 12, pp. 415–439, 2003

[180] Leila Amgoud, Nicolas Maudet, and Simon Parsons, "Modelling dialogues using argumentation", In Proceedings of the Fourth International Conference on Multi-Agent Systems, E. Durfee, Editor, pp. 31-38, Boston, MA, USA, 2000. IEEE Press.

[181] Chris Reed, "Dialogue Frames in Agent Communication" In Proceedings of the Third International Conference on Multi-agent Systems, Y. Demazeau, Editor, pp. 246-253, IEEE Press, 1998.

[182] N. Walton and E. C. W. Krabbe, "Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning", State University of New York Press, Albany, NY, 1995.

[183] Edmund H. Durfee, Victor R. Lesser, Daniel D. Corkill, "Trends in Co-operative Distributed Problem Solving", IEEE Transactions on Knowledge and Data Engineering, 1995.

[184] Sascha Ossowski, "Distributed Artificial Intelligence", In Coordination in Artificial Agent Societies, LNAI 1535, pp. 31-63, Springer-Verlag Berlin Heidelberg, 1999.

[185] Katia Sycara-Cyranski, "Arguments of persuasion in labor mediation", In Proceedings of the Ninth International Joint Conference on Artificial Intelligence, pp. 294--296, Los Angeles, California, August 1985.

[186] Edmund H. Durfee, "Distributed Problem Solving and Planning", In LNAI 2086, ACAI 2001, M. Luck, Editor, pp. 118-149, Springer-Verlag Berlin Heidelberg, 2001.

[187] Amy L. Lansky, "Localized Search for Controlling Automated Reasoning", In Proceedings of the DARPA Workshop on Innovative Approaches to Planning, Scheduling, and Control, pp. 115–125, November 1990.

[188] Subbarao Kambhampati, Mark Cutkosky, Marty Tenenbaum, and Soo Hong Lee", "Combining Specialized Reasoners and General Purpose Planners: a Case Study", In Proceedings of the Ninth National Conference on Artificial Intelligence, pp. 199–205, July 1991.

[189] Edmund H. Durfee, Victor R. Lesser, and Daniel D. Corkill, "Cooperation Through Communication in a Distributed Problem-Solving Network", In Cognition, Computing, and Cooperation, pp. 159-186, 1990.

[190] D.E. Wilkins and K.L. Myers, “A Common Knowledge Representation for Plan Generation and Reactive Execution”, Journal of Logic and Computation, Vol. 5, No. 6, pp. 731–761, 1995.

[191] Austin Tate, “Generating Project Networks”, In Proceedings of the 5th IJCAI, 1977.

[192] Susan E. Conry, Kazuhiro Kuwabara, Victor R. Lesser, and Robert A. Meyer, “Multistage Negotiation for Distributed Constraint Satisfaction”, IEEE Transactions on Systems, Man, and Cybernetics, SMC-21(6), pp. 1462–1477, Nov. 1991.

[193] Jyi-Shane Liu and Katia P. Sycara, “Multiagent coordination in tightly coupled task scheduling”, In Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96), pp. 181–188, December 1996.

[194] Eithan Ephrati, Martha E. Pollack, and Jeffrey S. Rosenschein, “A Tractable Heuristic that Maximizes Global Utility through Local Plan Combination”, In Proceedings of the First International Conference on Multi-Agent Systems (ICMAS-95), pp. 94–101, June 1995.

[195] Jaeho Lee, “An Explicit Semantics for Coordinated Multiagent Plan Execution, PhD dissertation, University of Michigan, 1997.

[196] Amal El Fallah Seghrouchni and Serge Haddad, “A Recursive Model for Distributed Planning”, In Proceedings of the Second International Conference on Multi-Agent Systems (ICMAS-96), pp. 307–314, December 1996.

[197] Eithan Ephrati and Jeffrey S. Rosenschein, “Divide and Conquer in Multi-agent Planning”, In Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94), pp. 375–380, July 1994.

[198] Edmund H. Durfee and Thomas A. Montgomery, “Coordination as Distributed Search in a Hierarchical Behavior Space”, IEEE Transactions on Systems, Man, and Cybernetics, Special Issue on Distributed Artificial Intelligence, SMC-21(6):1363–1378, November 1991.

[199] Frank von Martial, “Coordinating Plans of Autonomous Agents”, Lecture Notes in Artificial Intelligence, Springer-Verlag, 1992.

[200] Victor R. Lesser, “A Retrospective view of FA/C Distributed Problem Solving”, In IEEE Transactions on Systems, Man, and Cybernetics, Vol. 21, No. 6, pp. 1347–1362, Nov/Dec 1991.

[201] Susan E. Lander and Victor R. Lesser, “Understanding the role of negotiation in distributed search among heterogeneous agents”, In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93), pp. 438–444, August 1993.

[202] Keith J. Werkman, “Multiple Agent Cooperative Design Evaluation using Negotiation”, In Proceedings of the Second International Conference on Artificial Intelligence in Design, Pittsburgh PA, June 1992.

[203] Katia Sycara, Steven Roth, Norman Sadeh, and Mark Fox, “Distributed constrained heuristic search”, IEEE Transactions on Systems, Man, and Cybernetics, SMC-21(6):1446–1461.

[204] Victor R. Lesser, “Reflections on the Nature of Multi-Agent Coordination and Its Implications for an Agent Architecture”, In Autonomous Agents and Multi-Agent Systems, Vol. 1, pp. 89-111, 1998.

[205] L. C. Lee, H. S. Nwana, D. T. Ndumu and P. De Wilde, “The Stability, Scalability and Performance of Multi-agent Systems”, BT Technology Journal, Vol. 16, No 3, July 1998

[206] Anne-Louise Burness, Richard Titmuss, Caroline Lebre, Katie Brown and Alan Brookland, “Scalability Evaluation of a Distributed Agent System”, In Distributed Systems Engineering, Vol. 6, pp. 129-134, December 1999.

[207] Omer F. Rana and Kate Stout, “What is Scalability in Multi-Agent Systems?”, In Proceedings of the 4th International Conference on Autonomous Agents, pp. 56-63, June 2000.

[208] M. Stasser, J. Baumann, and F. Hohl, “Mole - A Java Based Mobile Agent System”, 2nd ECOOP Workshop: Mobile Object Systems, 1996.

[209] Yamamoto and Y. Nakamura, “Architecture and Performance Evaluation of a Massive Multi-Agent System”, In Proceedings of the Third Annual Conference on Autonomous Agents, May 1999.

[210] N. R. Franks, “Army Ants”, In Scientific American, vol. 77, 1989.

[211] Holldobler and E. Wilson, “Journey to the Ants”, Bellknap Press/Harvard University Press, 1994.

[212] M. M. Millonas, “Swarms, Phase Transitions and Collective Intelligence”, Artificial Life III, Santa Fe Institute Studies in the Sciences of Complexity, Vol. 17, 1994.

[213] O. Shehory, S. Kraus, and O. Yadgar, “Emergent Cooperative Goal-satisfaction in Large-scale Automated-agent Systems”, In Journal of Artificial Intelligence (Elsevier Science), Vol. 110, pp. 1-55, 1999.

[214] Wolpert and K. Tumer, “An Introduction to Collective Intelligence”, In Handbook of Agent Technology, J. M. Bradshaw, Editor, AAAI Press/MIT Press, 1999.

[215] Wolpert, K. Wheeler, and K. Tumer, “General Principles of Learning-Based Multi-Agent Systems”, In Proceedings of the Third Annual Conference on Autonomous Agents, May 1999.

[216] Asperti and N. Busi, “Mobile Petri Nets”, Technical Report UBLCS-96-10, Department of Computer Science, University of Bologna, Italy, 1996.

[217] Sibertin-Blanc, “Cooperative Nets”, In Application and Theory of Petri Nets, R. Valette, Editor, LNCS 815, pp. 471-490, Springer-Verlag, 1994.

[218] Phil Buckle, Tom Moore, Steve Robertshaw, Alan Treadway, Sasu Tarkoma, and Stefan Poslad, “Scalability in Multi-agent Systems: The FIPA-OS Perspective”, LNAI 2403, M. d’ Inverno et al., Editors, pp. 110–130, 2002.

[219] Umesh Deshpande, Arobinda Gupta, and Anupam Basu, “Performance Improvement of the Contract Net Protocol Using Instance Based Learning”, IWDC 2003, LNCS 2918, S. R. Das, S. K. Das, Editors, pp. 290–299, Springer-Verlag Berlin Heidelberg, 2003.

[220] S. Sen and G. Weiss, “Learning in Multiagent Systems”, in Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence, G. Weiss, Editor, MIT Press, Cambridge, Mass., 1999.

[221] R.G. Smith, “A Framework for Problem Solving in a Distributed Processing Environment”, Stanford Memo STAN-CS-78-700, Department of Computer Science, Stanford University, 1978.

[222] J. L. Kolonder, “Case-based reasoning”, Morgan Kaufmann, San Francisco, 1993.

[223] Watson and F. Marir, “Case-based Reasoning: A review”, In the Knowledge Engineering Review, Vol. 9, No. 4, pp. 327-354, 1994.

[224] U. Deshpande, A. Gupta, and A. Basu, “Adaptive Fault Tolerant Hospital Resource Scheduling”, In Proceedings of the 10th International Conference on Cooperative Information Systems (CooPIS - 2002),

Irvine, USA, LNCS-2519, pp. 503–520, Springer-Verlag, Berlin, Nov. 2002.

[225] Zoltan Juhasz and Prasenjit Paul, “Scalability Analysis of the Contract Net Protocol”, In Proceedings of the 2nd IEEE/ACM International Symposium on Cluster Computing and the Grid, Berlin, Germany, May 21-24, 2002.

[226] P. J. Turner and N. R. Jennings, “Improving the Scalability of Multi-Agent Systems”, Proceedings of the 1st International Workshop on Infrastructure for Scalable Multi-Agent Systems, Barcelona, 2000.

[227] Christian Gerber, “Self-adaptation and Scalability in Multi-agent Systems”, PhD Thesis, Universität des Saarlandes, Dec. 1999.

[228] H. Brooks and E. H. Durfee, ‘‘Congregation formation in information economies’’, In Proceedings of the AAAI-99 Workshop on AI in Electronic Commerce, pp. 62– 68, 1999.

[229] H. Brooks, E. H. Durfee, and A. Armstrong, ‘‘An introduction to congregating in multiagent systems’’, In Proceedings of the Fourth International Conference on Multiagent Systems, pp. 79– 86, 2000.

[230] Christopher H. Brooks and Edmund H. Durfee, “Congregation Formation in Multiagent Systems”, In Autonomous Agents and Multi-Agent Systems, Vol. 7, pp. 145– 170, 2003.

[231] M. Fenster, S. Kraus, and J. Rosenschein, “Coordination without Communication: An Experimental Validation of Focal Point Techniques”, In Proceedings of the First International Conference on Multi-Agent Systems, pp. 102–116, San Francisco, California, 1995.

[232] Edmund E. Durfee, “Scaling up Agent Coordination Strategies”, In Computer, Vol. 34, No. 7, pp. 39-46, July 2001.

[233] Y. Shoham, “Agent-oriented Programming”, In Artificial Intelligence, Vol. 60, No. 1, pp. 51.92, 1993.

[234] Piotr J. Gmytrasiewicz and Edmund H. Durfee, “Rational Coordination in Multi-Agent Environments”, In Autonomous Agents and Multi-Agent Systems, Vol. 3, pp. 319-350, 2000.

[235] Cora Beatriz Excelente-Toledo and Nicholas R. Jennings, “The Dynamic Selection of Coordination Mechanisms”, In Journal of Autonomous Agents and Multi-Agent Systems (to appear), 2004.

[236] Rachel A. Bourne, Cora B. Excelente-Toledo and Nicholas R. Jennings, “Run-Time Selection of Coordination Mechanisms in Multi-Agent

Systems”, In Proceedings of the 14th European Conference on AI, pp. 348--352, Berlin, Germany, 2000.

[237] K.S. Barber, D.C. Han, and T.H. Liu, “Coordinating Distributed Decision Making Using Reusable Interaction Specifications”, PRIMA 2000, LNAI 1881, C. Zhang and V.-W. Soo, Editors, pp. 1-15, Springer-Verlag Berlin Heidelberg, 2000.

[238] Rachel A. Bourne, Karen Shoop, and Nicholas R. Jennings, “Dynamic Evaluation of Coordination Mechanisms for Autonomous Agents”, In Lecture Notes in Artificial Intelligence (LNAI) 2258, pp. 155-168, Springer-Verlag, 2001.

[239] Rejane Frozza and Luis Otávio Alvares, “Criteria for the Analysis of Coordination in Multi-agent Applications”, COORDINATION 2002, LNCS 2315, F. Arbab and C. Talcott, Editors, pp. 158-165, Springer-Verlag Berlin Heidelberg, 2002.

[240] Coates, R. I. Whitfield, A. H. B. Duffy, and B. Hills, “Coordination Approaches and Systems – Part II: an Operational Perspective”, Research in Engineering Design, Vol. 12, pp. 73-89, 2000.

[241] Sascha Ossowski and Ana García-Serrano, “Social Co-ordination among Autonomous Problem-solving Agents”, In Agents and Multi-Agent Systems: Formalisms, Methodologies, and Applications, W. Wobcke, M. Pagnucco, C. Zhang, Editors, Vol. 1441, pp. 134-148, January 1998.

[242] Sascha Ossowski, “Coordination”, In Coordination in Artificial Agent Societies, LNAI 1535, pp. 15-30, Springer-Verlag Berlin Heidelberg, 1999.

[243] N. Carriero and D. Gelernter, “Coordination Languages and Their Significance”, Communications of the ACM, Vol. 35, No 2, pp. 97–107, February 1992.

[244] B. Freisleben and T. Kielmann, “Coordination Patterns for Parallel Computing”, In Proceedings of the Second International Conference on Coordination Models, Languages and Applications, LNCS, Springer Verlag, September 1997.

[245] N. Carriero and D. Gelernter, “How to Write Parallel Programs”, A First Course, MIT Press, 1992.

[246] Andrew H. Van de Ven, André L. Delbecq, and Richard Koenig, JR, “Determinants of Coordination Modes within Organization”, In American Sociological Review, Vol. 14, pp. 322-338, April 1976.

[247] Manfred Kochen and Karl W. Deutsch, "A Note on Hierarchy and Coordination: An Aspect of Decentralization", In Management Science, Theory Series, Vol. 21, Issue 1, pp. 106-114, Sept. 1974.

[248] M. Gerson, "On Quality of Life", In American Sociological Review, Vol. 41, pp. 793-806, 1976.

[249] Henry Mintzberg, "The Structuring of Organizations: A Synthesis of the Research, Prentice Hall, Englewood Cliffs, N. J., 1979.

[250] Henry Mintzberg, "Structure in 5'S: A Synthesis of the Research on Organization Design", In Management Science, Vol. 26, No. 3, March 1980.

[251] Harold Doty, William H. Click, and George P. Huber, "Fit, Equifinality, and Organizational Effectiveness: A Test of Two Configurational Theories", In Academy of Management Journal, Vol. 36, No. 6, pp. 1196-1250, 1993.

[252] Cris Kobryn, "UML 2001: a Standardization Odyssey", In Communications of the ACM, Vol. 42, No. 10, pp. 29–37, October 1999.

[253] V. D. Parunak, "Practical and Industrial Applications of Agent-Based Systems", Report, Environmental Research Institute of Michigan (ERIM), 1998. http://agents.umbc.edu/papers/apps98.pdf.

[254] V. D. Parunak, "A Practitioners' Review of Industrial Agent Applications", In Autonomous Agents and Multi-Agent Systems, Vol. 3, pp. 389-407, 2000.

[255] Nicholas R. Jennings, "An agent-based approach for building complex software systems", In Communications of the ACM, Vol. 44, No. 4, April 2001.

[256] Nicholas R. Jennings and M. Wooldridge, "Applications of Intelligent Agents", In Agent Technology: Foundations, Applications, and Markets, Nicholas R. Jennings and M. J. Wooldridge, Editors, pp. 3-28, Springer-Verlag: Heidelberg, Germany, 1998.

[257] B. Chaib-Draa, "Industrial Applications of Distributed AI", In Communications of the ACM, Vol. 38, No. 11, 47-53, 1995.

[258] M. Wooldridge, "Multiagent Interactions", In an Introduction to MultiAgent Systems, John Weley & Sons, Ltd, 2002.

[259] M. Wooldridge, "Working Together", In an Introduction to MultiAgent Systems, John Weley & Sons, Ltd, 2002.

[260] Nicholas R. Jennings, Katia Sycara, and M. Wooldridge, “A Roadmap of Agent Research and Development”, In Autonomous Agents and Multi-Agent Systems, Vol. 1, pp. 275-306, 1998.

[261] L. Gasser, “DAI Approaches to Coordination” in Distributed Artificial Intelligence: Theory and Praxis, N.M. Avouris and L. Gasser, Eds. Dordrecht, The Netherlands: Kluwer Academic Publishers, 1992, pp. 31-51.

[262] Omicini, F. Zambonelli, M. Klusch, R. Tolksdorf, Editors, “Coordination of Internet Agents: Models, Technologies, and Applications”, Springer-Verlag, March 2001.

[263] Sascha Ossowski, “Co-ordination in Artificial Agent Societies”, LNAI 1535, Springer-Verlag, 1999.

[264] Papadopoulos, “Models and Technologies for the Coordination of Internet Agents: A Survey”, In Coordination of Internet Agents, Omicini et al., Editors, Springer Verlag, 2001.

[265] Giovanni Aloisio and Domenico Talia, “Grid Computing: Towards a New Computing Infrastructure”, In Future Generation Computer Systems, Vol. 18, 2002.

[266] Soumen Mukherjee, Joy Mustafi, and Abhik Chaudhuri, “Grid Computing: The Future of Distributed Computing for High Performance Scientific and Business Applications”, In IWDC 2002, LNCS 2571, S.K. Das and S. Bhattacharya, Editors, pp. 339–342, Springer-Verlag, 2002.

[267] Jonathan Dale, Bernard Burg, and Steven Willmott, “The Agentcities Initiative: Connecting Agents Across the World”, In WRAC 2002, LNAI 2564, W. Truszkowski, C. Rouff, M. Hinchey, Editors, pp. 453-457, Springer-Verlag, 2003.

[268] P Marrow, E Bonsma, F Wang and C Hoile, “DIET — a Scalable, Robust and Adaptable Multi-Agent Platform for Information Management”, In BT Technology Journal, Vol. 21, No 4, October 2003.

[269] Jennings, N, “Agent-Based Computing: Promise and Perils”, In Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99), pp. 1429–1436, 1999.

[270] M. Schillo, Klaus Fischer, Bettina Fley, Michael Florian, Frank Hillebrandt, and Daniela Spresny, “FORM – A Sociologically Founded Framework for Designing Self-Organization of Multiagent Systems”, In RASTA 2002, LNAI 2934, G. Lindemann et al., Editors, pp. 156–175, Springer-Verlag, 2004.

[271] Mark Sims, Claudia V. Goldman, and Victor Lesser, “Self-Organization through Bottom-up Coalition Formation”, In Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 867-874, 2003.

[272] Ian Foster, Carl Kesselman, and Steven Tuecke, “The Anatomy of the Grid: Enabling Scalable Virtual Organizations”, In International Journal of Supercomputer Applications, 2001.

[273] Liviu Panait and Sean Luke, “Cooperative Multi-Agent Learning: The State of the Art”, In Autonomous Agents and Multi-Agent Systems, 11, 387-434, 2005.

[274] Schneider, W.K. Wong, A. Moore, and M. Riedmiller, “Distributed value functions”, In Proceedings of the Sixteenth International Conference on Machine Learning, pp. 371-378, 1999.

[275] Zhaotong Lian and M. Abhijit Deshmukh, “Performance prediction of an unmanned airborne vehicle multi-agent system”, European Journal of Operational Research, No 172, pp. 680-695, 2006.

[276] R. W. Beard, T.W. McLain, M.A. Goodrich, and E.P. Anderson, “Coordinated Target Assignment and Intercept for Unmanned Air Vehicles”, In IEEE Transactions on Robotics and Automation, Vol. 18, No 6, pp. 911-922, December 2002.

[277] C.V. Goldman and S. Zilberstein, “Decentralized Control of Cooperative Systems: Categorization and Complexity Analysis”. Journal of Artificial Intelligence Research, pp. 143-174, 22 (2004).

[278] Y.H. Chang, T. Ho, and L. Kaelbling, “Multi-agent learning in mobilized ad-hoc networks”. Papers from the 2004 AAAI Fall Symposium, Technical Report FS-04-02-2004.

[279] Fischer, “Cooperative Transportation Scheduling: An application Domain for DAI”, Applied Artificial Intelligence, Volume 10, N1, pp. 1-34. 1996.

[280] Nunes and E. Oliveira, “Learning from multiple sources”, in AAMAS-Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems, 2004.

[281] Dresner and P. Stone, “Multiagent traffic management; A reservation-based intersection control mechanism”, In AAMAS- *Proceedings of the Third International Joint Conference on Autonomous Agents and Multiagent Systems*, 2004.

[282] Steeb R., Cammarata S., Hayes-Roth F., Thorndyke P., and Wesson R., “Distributed intelligence for air traffic control”, In A. Bond and L. Gasser (eds.). Readings in Distributed Artificial Intelligence, Morgan Kaufmann Publishers, pp. 90-101, 1988.

[283] Min J.U and Bjornsson H., Construction Supply Chain Visualization Through Web Services Integration, Technical Report, Civil and Environmental Engineering Dept., Stanford University, 2004.

[284] Sichman J.. Du raisonnement social chez les agents: une approche fondée sur la théorie de la dépendance. PhD thesis, Institut National Polytechnique de Grenoble – France, 1995.

[285] Hannoun M., Sichman J., Boissier O., and Sayettat C., Dependence relations between roles in a multi-agent system. In Multi-Agent Systems and Agent-Based Simulation, MABS’98, 1998.

[286] Lechilli K. and Chaib-Draa B., Coordination en situations non familières. In Troisièmes Journées Francophones sur l’Intelligence Artificielle Distribuée et les Systèmes Multi-Agents, April 1996.

[287] Casterlfranchi C, Micelli M., and Cesta A.. Dependence relations among autonomous agents. In E. Werner and Y. Demazeau, editors, Decentralized A. I., Vol. 3 pp. 215-227. Elsevier Science Publishers B. V., Amsterdam, 1992.

[288] Sichamn J.S., Conte R., Castelfranchi C., and Demazeau Y.. A social reasoning mechanism based on dependence networks. In Proceedings of the 11th European Conference on Artificial Intelligence (ECAI-94), pages 188-192, Amsterdam, 1994.

[289] Castelfranchi C., Social power. A point missed in Multi Agent, DAI and HCI. In Y. Demazeau and J. P. Muller, editors, Distributed Artificial Intelligence, pages 49-62. Elsevier Science Publishers B. V. (North Holland), 1990.

[290] Demazeau Y. and Muller J.P., From reactive to intentional agents. In Y. Demazeau and J.P. Muller, editors, Distributed Artificial Intelligence, Volume II, pages 3-13. Elsevier Science Publishers B. V. (North Holland), 1991.

[291] Sichman J.. On achievable goals and feasible plans in open multi-agent systems. In the 1st Ibero-American Workshop on DAI/MAS, Xalapa, Mexico, 1996.

[292] Sichman J. and Demazeau Y.. A model for the decision phase of autonomous belief revision in open multi-agent systems. Journal of the Brazilian Computer Society, 3(1):40-50, 1996.

[293] Sichamn J. and Demazeau Y.. Exploiting social reasoning to deal with agency level inconsistency. In Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS-95), pages 352-359, San Francisco, CA, June 1995.

[294] Malcolm McCullough, Digital Ground: Architecture, Pervasive Computing, and Environmental Knowing, MIT Press, 2005

[295] Wang X., Dong J.S, Chin C., Hettiarachchi S.R. and Zhang D.. Semantic Space: A Semantic Web Infrastructure for Smart Spaces. IEEE Pervasive Computing, 3(3):32-39, July-September 2004

[296] Hubner J.F, Sichman J.S., O. Boissier, Developing organised multiagent systems using the MOISE+ model: programming issues at the system and agent levels. International Journal of Agent-Oriented Software Enginnering, Vol. 1, No3/4, pp. 370-395.

List of symbols/abbreviations/acronyms/initialisms

AA	Autonomous Agents
AI	Artificial Intelligence
ACL	Agent Communication Language
CONSA	Collaborative Negotiation System based on Argumentation
DAAC	Distributed Active Archive Center
DAI	Distributed Artificial Intelligence
DND	Department of National Defence
DPS	Distributed Problem Solving
DSNE	Distributed Sensor Network Establishment
DVM	Distributed Vehicle Monitoring
FA/C	Functionally Accurate and Cooperative
FIFA	Foundation for Intelligent Physical Agents
GRAIL	Generic Resource Allocation & Integration Algorithm
MAS	Multi-Agent Systems
PGP	Partial Global Planning (PGP)
ToH	Tower of Hanoi

DOCUMENT CONTROL DATA

1. ORIGINATOR (name and address) **Defence R&D Canada Valcartier** **2459 Pie-XI Blvd. North** **Québec, QC** **G3J 1X5**	2. SECURITY CLASSIFICATION (Including special warning terms if applicable) **Unclassified**

3. TITLE (Its classification should be indicated by the appropriate abbreviation (S, C, R or U)
Coordination in distributed intelligent systems (U)

4. AUTHORS (Last name, first name, middle initial. If military, show rank, e.g. Doe, Maj. John E.)
Boukhtouta A., Berger J., Allouche M. and Bedrouni A.

5. DATE OF PUBLICATION (month and year) **October 2009**	6a. NO. OF PAGES **182**	6b .NO. OF REFERENCES **296**

7. DESCRIPTIVE NOTES (the category of the document, e.g. technical report, technical note or memorandum. Give the inclusive dates when a specific reporting period is covered.)
Technical Report

8. SPONSORING ACTIVITY (name and address)
Defence R&D Canada – Valcartier
2459 Pie-XI Blvd. North, Quebec City, Quebec, Canada, G3J 1X5

9a. PROJECT OR GRANT NO. (Please specify whether project or grant) **13du**	9b. CONTRACT NO.
10a. ORIGINATOR'S DOCUMENT NUMBER TR 2006-780	10b. OTHER DOCUMENT NOS N/A

11. DOCUMENT AVAILABILITY (any limitations on further dissemination of the document, other than those imposed by security classification)

- [x] Unlimited distribution
- [ ] Restricted to contractors in approved countries (specify)
- [ ] Restricted to Canadian contractors (with need-to-know)
- [ ] Restricted to Government (with need-to-know)
- [ ] Restricted to Defense departments
- [ ] Others

12. DOCUMENT ANNOUNCEMENT (any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in 11) is possible, a wider announcement audience may be selected.)
Unlimited

13. ABSTRACT (a brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), (R), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual).

Intelligent and autonomous agents have emerged as a novel approach offering the ability to analyze, design, build, and implement complex systems. The agent-based conceptual view provides a large repertoire of tools, techniques, and metaphors that can be adopted to considerably improve the way many types of complex systems are conceptualized and implemented in various domains. Agent-oriented technologies are increasingly used in a variety of applications, ranging from comparatively small systems to large and complex, mission-critical applications such as air traffic control.

This report covers issues and comprehensively answers commonly asked questions about coordination in agent-oriented systems. It is dedicated to providing a state-of-the-art review of current coordination strategies, protocols and mechanisms. In exploring the progress achieved, this study has unveiled the lack of coherence and order that characterizes the area of research pertaining to coordination. Based on current practical deployed applications and future opportunities, the report identifies and thoroughly examines trends, challenges, and future agent-oriented research directions.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus, e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus-identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Coordination, Distributed Artificial Intelligence, Distributed Problem Solving, Autonomous Agents, Distributed planning.

Defence R&D Canada

Canada's Leader in Defence
and National Security
Science and Technology

R & D pour la défense Canada

Chef de file au Canada en matière
de science et de technologie pour
la défense et la sécurité nationale

DEFENCE R&D DÉFENSE

www.drdc-rddc.gc.ca