Chapter 6: Coordination Approaches

This chapter provides an overview of the most commonly used coordination mechanisms, classified according to the paradigm of their origin. In section 6.1, definitions of cooperation and coordination are given, as well as clarification of the scope of this thesis with respect to coordination. Coordination approaches inspired by biology are presented in section 6.2, followed by approaches inspired by organisational sciences, described in section 6.3. Basic concepts of social networks are introduced in section 6.4, together with a brief discussion on the applicability of social networks to MASs. An overview of social networks-related research in the field of MASs is given in section 6.5. Lastly, the new coordination approach, based on social networks, is presented in section 6.6.

6.1 Introduction

Agents in a Multi-Agent System (MAS) can exhibit cooperative or competitive behaviours. While competitive behaviour can be encouraged in some computational intelligence approaches, such as evolutionary computing, in robotic applications it is not often desired. In robotics the cost of building a robot is relatively high, thus the evolutionary approach where undesired specimens are discarded, is often not desirable (however, the evolutionary approach can be used in simulations and only the final optimal solution can be implemented physically). It is important to note that cooperative behaviour does not exclude market-based competitive approaches, such as the auctioning coordination technique [175].

The key to a successful MAS is to prevent negative interaction (conflict) and to promote positive interaction (cooperation). In order to promote these two goals, it is necessary to implement a coordination mechanism.

Coordination mechanisms can be broadly divided onto two distinctive groups:

- Emergent - where each agent pursues its own goals, but a coordination-like behaviour emerges through interaction within an environment. This is often the case in swarm robotics [189][177].

- Intentional - where agents actively and intentionally communicate in order to avoid conflict [144].

For the purpose of this thesis, the intentional approach [77] to coordination is followed.

Furthermore, a specific intentional coordination method, task allocation, is the focus of this thesis and emphasis is on task allocation methods as a coordination technique for multi-robot teams.

6.2 Biology-Inspired Approaches – Coordination Perspective

The main advantage of investigating a biology-inspired approach for coordination mechanisms is the existence of coordination in biological systems. In fact, there is an abundance of insects and animals that successfully coordinate, for example ants in an ant colony and wolves in a wolf pack. Biology-inspired coordination mechanisms range from simplistic mechanisms that rely on very limited communication channels (as seen in insect societies) to sophisticated mechanisms that utilise multi-channels of communication (i.e. gestures, sounds and "body language") as observed in mammalian societies.

The diversity of coordination mechanisms requires a clear separation of biology-inspired approaches into two main categories: insect society-inspired and higher mammalian society-inspired. An overview of the main differences between the two approaches follows next.

6.2.1 Overview of Differences Between Insect and Mammalian Societies (Coordination Perspective)

Before considering multi-robot systems inspired by insects (commonly referred to as swarm robotics systems) or mammalian societies, it is useful to consider the fundamental differences between real, biological agents in insect colonies and those in higher mammalian societies. The differences are discussed from cognitive and social perspectives. A full comparison is outside the scope of this thesis and only the characteristics related to cooperation and coordination are considered:

- ***Ability to learn.*** Insects have a much lower level of self-awareness and their individual learning ability is often non-existent. Insects typically do not develop memories and do not learn from past experience, whereas mammals do learn from past experience.

- ***Communication methods***. Insects have much simpler communication mechanisms that prohibit the exchange of complex messages. Mammal societies usually employ complex means of communication. Another communication-related issue is the localised nature of insect communications. Insects usually communicate by touch and/or chemical reactions [19]. Mammals often use sound. Using sound as a communication method, mammals can communicate over greater distances.

- ***Individualism.*** Most agents are homogenous in insect colonies. In other words, insect colonies are anonymous societies where agents of the same type are indistinguishable. In contrast, kinship and other social relationships are of extreme importance to mammalian societies.

Cooperation is a form of positive interaction between agents. It is a process of working together to achieve a common goal. The requirement for cooperation is the existence of a coordination and/or negotiation mechanism. A view expressed by Matarić [114] is that cooperative behaviours (such as task allocation) based on negotiations require direct communication between the agents.

Direct communication, i.e. when a specific agent is identified and addressed, is not possible in insect societies due to the lack of individualism and insects' limited communication mechanisms. It is important to note that insects do communicate (for example bees have a dance based communication mechanism), but the communication is limited in range and it is limited in the number of messages that can be communicated. However, it would be wrong to say that insects do not cooperate; they do, but through much simpler mechanisms based on interaction with their environment, using the principle of stigmergy [19]. The stigmergy principle is derived from observations of social insect colonies, such as bees and ants. The process of stigmergy is described as:

"The production of a certain behaviour as a consequence of the effect produced in the local environment by previous behaviour" [19].

Due to the lack of individualism, hierarchies and social relations between agents in insect colonies are virtually non-existent. In mammalian societies, hierarchies and social relations play a fundamental role in the organisation of such a society. Cooperation models are often based on hierarchical and role-based models. Insect colonies, as a cooperation model, were and still are very attractive for applications in robotics [5][98]. The main advantages of swarm robotics are that the insect-like robots are fairly simple to construct; cooperation between such robots should be an emergent property of such a system. Swarm robot teams are usually fault-tolerant (to a degree, because if a sufficient number of agents fails then the whole team might fail).

Coordination in insect-like multi-robot systems was initially done through interaction with the environment, using the principle of stigmergy [19]. The stigmergy-based coordination mechanism is very limited and although emergent cooperative behaviour was observed [168], it has imposed limitations on cooperation methods. The need for a more capable communication mechanism, even in insect-like societies, has been recognised relatively early in research and it has lead to various communication-capable behaviour based multi-robot systems [113]. A summary of the differences between the agent models is given in table 4.

Agent Characteristic	Insect Colonies	Mammal Societies
Communication	Sparse, localised	Complex
Individualism	No	Yes
Learning ability	No	Yes

Table 4. Differences between two biology-inspired agent models

6.3 Organisational Sciences-Based Approach

Since the emergence of more complex work-related structures, it has become a necessity to better organise such structures. This necessity for better organisation gave birth to a range of disciplines under the research field of organisational sciences. Researchers in organisational sciences have concentrated mainly on the two most popular approaches, namely the market-based approach and the hierarchical approach.

Each of these two approaches has its own advantages and disadvantages and each of them has been tried as a coordination technique in the field of MASs. The remainder of this section overviews market-based and hierarchical approaches.

6.3.1 Market-Based Approach

Markets are based on the voluntary exchange of commodities between parties at an agreed price. Market-based coordination is based on the same premise. Markets have many properties but from the point of view of its applicability to robotics, the following properties are of primary interest:

- **Self-organisational property**: Markets are self-organising through a pricing mechanism. This property is highly desirable as it helps with the social approach to MAS design, where agents are viewed as a self-organising society.

- **Demand-supply relationship**: Supply and demand are inseparable and self-regulating. This relationship assumes the existence of two entities: a buyer and a seller.

- **Scalability:** Theoretically there is no limit to the number of participants in a market.

Ideally, when the market functions properly, there is an equilibrium price. The equilibrium price is the fairest cost of the transaction and coordination is nearly optimal. The idea of using market-based coordination, in MAS in particular and in AI in general, has lead to the development of various auction based algorithms. One of the most widely used coordination mechanisms used in AI is the Contract Net Protocol (CNP) [175]. CNP assumes the existence of a buyer, a seller and a price.

With regard to robotics, auction based coordination has mainly been applied to simulated multi-robot teams. An example of such an application is given in [51]. Recently, the first real embodied agent systems that use auction based coordination have appeared [77]. It may be too early to judge the success of such coordination based MASs in real embodied agent applications, but there are concerns with the future of a purely auction based approach.

Firstly, the method for awarding a bid must be determined. It is usually a metric or fitness function that is used to determine a winning bidder. In the case of a task that has not been done before, it is uncertain how to determine a winning bidder.

Secondly, the auctioning mechanism relies on accuracy of the task details that is used by bidding agents to calculate the bid. The information that is submitted to bidding agents is not necessarily accurate or complete.

Thirdly, it is unclear how a purely auction based approach can handle a scenario when a task exceeds the capabilities of each individual bidder. One of the main strengths of MASs is the ability to solve problems that exceed the capability of individual agents.

6.3.2 Hierarchical Approach

Probably the simplest way of coordinating agents is by establishing a relatively strict hierarchical architecture that prescribes the roles for each agent. Such an approach often assumes a globally coordinated and optimised multi-robot system. The coordination task is done either by a specialised agent [10] or by an agent that has been temporarily assigned the coordination role [103]. The hierarchical approach often uses a symbolic based planner that can provide the optimal solution based on its symbolic environment model.

The hierarchical approach has a number of problems including:

- The inability to create an accurate world model performance [197], and
- it is somewhat contrary to the idea of autonomous agents. Instead of being fully autonomous agents, the agents in a hierarchical approach system are in effect controlled from a central agent.

The hierarchical approach also leads to easier specialisation of agents as agents can have different physical and logical characteristics, adjusted to their specific tasks. Agent specialisation can prohibit optimal load balancing. One of the extreme specialisations is that of MACTA [11][10], where the deliberative (or central planning) coordination agent is a desktop PC, while the team members are physical robots. The side effect of specialisation is that redundancy is reduced. A team member cannot easily replace another team member that has failed because they have different physical and reasoning characteristics.

Specialisation has a positive side: agents can be designed according to the role they are assigned to perform. Specialisation in this context can lead to lowering an agent's complexity and cost as only the required functionality is implemented.

If a centralised coordination mechanism is employed, the role of a reliable communication channel is crucial. In the case of failure of either the global coordinating agent or the communication channel, the hierarchical multi-robot system will fail. The single point of failure characteristic is not desirable. In certain environments (e.g. deep level mining, underwater exploration, electronics emissions saturated battlefield etc.), communication channels can be limited and unreliable. In such environments, it is improbable that a hierarchical approach would be an effective approach to coordination, as it requires reliable communication channels.

6.4 Social Networks

Traditionally, societies are organised according to socio-economic structures such as markets, hierarchies and networks. Only relatively recently, social networks have been identified as social organisational structures, with one of the first formal definitions of social networks given by Mitchell:

"A (social) network is generally defined as a specific type of relation linking a defined set of persons, objects or events" [125].

Using social networks analysis, social networks can be used to explain why a society, as an entity, functions the way it does. From a social network analysis point of view, a society can be expressed as patterns of relationships between interacting units [195]. Social network analysis can also give insights into emergent patterns, relationships and their implications to a society.

Societies with well-developed patterns of social networks have many advantages. Before exploring these advantages, an introduction to the field of social networks and its terminology is necessary.

6.4.1 History of Social Networks Analysis as a Science

It is outside the scope of this thesis to give a full detailed history of the development of social networks analysis as a science. The reader is referred to [167] for more details. For the purpose of this thesis, only a brief overview is presented next.

Social network theory did not appear suddenly as a unified, complete theory. Instead, social network theory has evolved from the works of various scientists over a period spanning almost a century. Initially, mainly behavioural and organisational scientists were interested in it. At the beginning of the 20th century behavioural scientists have posed the question "how much do we preconceive objects and concepts and how much do we really perceive them" [167]. One of the answers to this question was the *gestalt* theory by Kohler [97]. Kohler proposed that our perception is defined by organised patterns through which humans interpret the real-world. Kohler's gestalt theory was positively accepted and a number of researchers have expanded on his work, albeit mainly in the field of social psychology.

One of the scientists that has embraced the *gestalt* theory was Moreno [167]. Moreno's research focus was to determine the influence of structures, what he has termed "social configurations", on the psychological well-being of an individual. Examples of such social configurations are concepts of friendship, attraction, repulsion, etc. Moreno's main contribution to the field was the introduction of the concept of "sociometry". Sociometry is a metric function for social relations. Furthermore, he has introduced a "sociogram", which is basically a directed graph representing "social configurations". The improved versions of sociogram are still frequently used to describe social relations. In fact, the sociogram has provided a foundation for graph theory applications to sociological sciences [195]. One way of representing social networks is through such graphs.

Almost at the same time, a group of scientists at Harvard started their work on defining cliques, clusters or blocks within a society in the late 1920s. The most prominent amongst the researchers were Mayo and Werner [167]. Their research focus was different from that of Kohler and Moreno. While, as sociologists, Kohler and Moreno were interested in the application of social networks (or their humble beginnings) to social psychology, Mayo and Werner were interested in applying social networks to anthropology. This diversity of origins of modern social networks analysis emphasises the fact that from the early beginnings, social networks analysis was seen as an interdisciplinary technique.

The Harvard group proposed communities within societies and often informal mechanisms that govern them [167]. By the late 1930s, basically all components of modern social networks analysis were in place, namely relationships, actors, cliques etc. However, a unified social networks theory was still decades away.

In 1969, a Manchester University researcher, J.C. Mitchell, published his work that is widely seen as the foundation of modern social network theory [125]. The majority of concepts introduced by Mitchell are still applicable and in use. Concepts and notations related to social networks, and relevant to this thesis, are defined next.

6.4.2 Social Networks Analysis Concepts

For the purpose of this thesis, only a selection of social networks analysis concepts is presented here. The selected concepts are relevant to the research presented in this thesis and sufficient to support the work done. The selected concepts are used in comments on results presented in chapters 7, 8 and 9.

The selected concepts and definitions are:

- Actors that act semi-independently. Actors are autonomous, yet they are defined and embedded within the society through the existence of social networks. Actors can be seen as nodes in a graph that represents social networks. In MASs in general, and in INDABA in particular, actors are agents. Therefore, the remainder of this thesis uses the term agent instead of the term actor.
- Relationships that link agents to each other. The relationships can either be positively or negatively weighted, and are directed. The social relationships can be seen as indices of a graph that represents a society.
- Social network. A social network is a set of agents and a distinct relationship among the agents.
- Agent society. A society is a representation of a complete set of social networks.
- Cliques or clusters, that are sets of agents defined by existence of strong relationships. A clique is a sub-set of a society, or in graph terms, a sub-graph.

To illustrate the described concepts, consider a society with only one type of relationship, namely the frequency of cooperation between individuals A, B, C, D and

E on a specific project. Table 5 summarises the frequency of cooperation between the agents in the society.

Cooperated	A	B	C	D	E
A	0	27	25	5	0
B	27	0	31	0	3
C	25	31	0	8	7
D	5	0	8	0	35
E	0	3	7	35	0

Table 5. Matrix representing a social network

Based on the relationships given in table 4, figure 13 illustrates a graph that describes the resulting society.

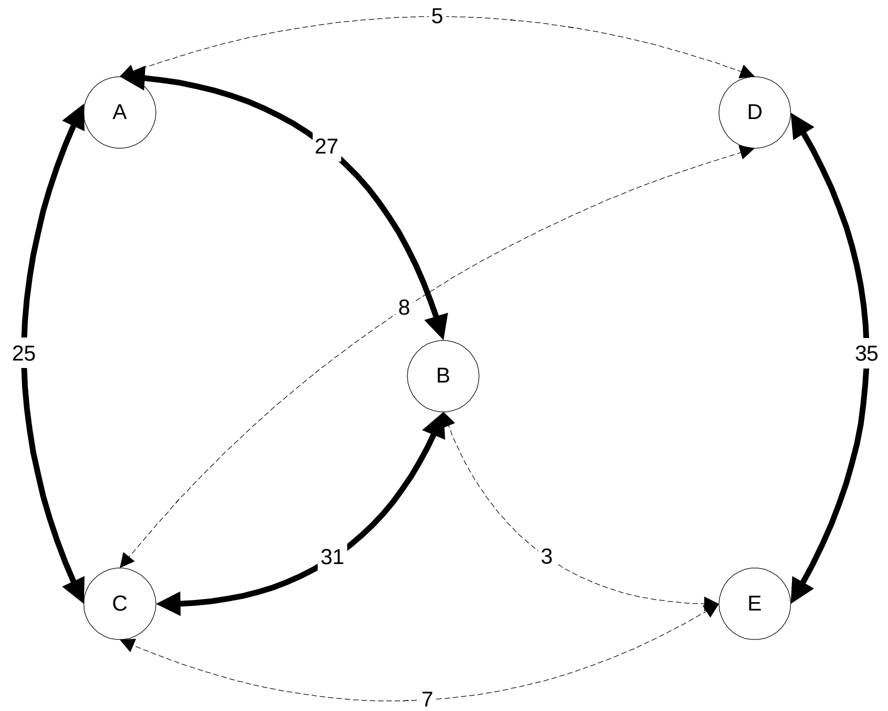

Figure 13. An illustration of a social network representation

Individuals A, B, C, D and E are agents, represented as nodes in the graph. The indices associated with links represent the frequency of cooperation between the agents on the project. The existence of strong relationships between certain members is an indication that those members form a cluster (or a clique). Figure 13 illustrates two clusters (as indicated by bold lines): one consisting of individuals A, B and C and the other consisting of individuals D and E.

Membership to a social cluster or a social relationship is not exclusive: a member can be linked to other members through multiple social relationships. The agents that are linked through a social relationship can be seen as the members of a social group. Examples in human society abound: a person can be a member of a sports club, a university study group and a family. In human societies, social networks are present in everyday interactions but they are not always simple to express and quantify. Key questions applicable to social networks are how the social relationships that define the social networks are formed and how they are maintained [180]. In more complex animal societies, concepts of kinship and trust form a fundamental, but not exclusive, role in the creation of social networks.

It is important to note that this is a simplified representation of social networks in comparison with social networks as observed in the real-world. Real-world social networks are often more complex and have attributes such as direction, durability and intensity [102].

6.4.3 The Importance of Uncertainty in Multi-Robot Teams

Uncertainty about task details is unfortunately one of the realities of implementing any MAS and specifically multi-robot teams in real-world environments. The problem of uncertainty is more evident in robotic applications operating in previously unexplored environments. Those environments are difficult to model, due to the uncertainty about the environment attributes. Furthermore, there is often no previous history of similar applications. Interplanetary robotic exploration is an example of such environments.

The majority of "robots" that were used for interplanetary exploration, starting with the early Soviet Lunokhod series [191] up to the recent NASA's Mars Rovers [187], are not agents or robots in the true sense of the definition of an agent (refer to section 2.2.3). These vehicles are not autonomous, but tele-operated from Earth, leaving just a basic interaction with its environment to their internal mechanisms.

While "robots" within the inner Solar system can be tele-operated due to the fact that the delay caused by the finite speed of radio signal propagation is within (barely) acceptable limits, the same method of exploration will not be possible for the further reaches of the Solar system and beyond. Interplanetary exploration is just one of the problem domains that will benefit from the evolution of more autonomous, self-organising robotic systems.

On other hand, biological systems, such as teams of animals or humans, generally cope with uncertainty. In other words, teams of animals or humans, when put in different environments and faced with unfamiliar tasks, generally achieve their goals. The view proposed in this thesis is that one of the contributing factors is the existence of social networks that define a team and the structures within it.

6.4.4 The Applicability of Social Networks to Multi-Robot Teams

A major advantage of societies with multiple, well-established social network is that they are flexible enough to allow the best team for the task to be selected by using the most appropriate social relationship that in turn defines the social group. For example, if a task involves participation in some sport, the member belonging to the sports club that practices that sport should be used to form a team. Affinities between social groups can also play a significant role. If there is uncertainty about the task and no social group satisfies the demands of that task, then the group with the highest affinity for the task should be selected. The existence of such affinity relationships between social groups is very important. Affinity between social groups is especially important when there is uncertainty (lack of detailed information) about the task or in the case where the best candidates for a task are already allocated to another task. In the case that the best candidates are not available, the "next best" candidates must be selected. The "next best" candidates are the members of the social group with affinity to the optimal social group.

It has been noted in section 6.4.2, that social networks are sets of agents and relationships between them. Agents are members of a society and robots in multi-robot teams can be viewed as members of a society. Zini, for example, defines a MAS as a society of agents [204].

With the previous definition in mind, a multi robot-system can be defined as:

"…a social and cognitive entity with a relatively identifiable boundary, that functions on a relatively continuous basis through the coordination of loosely interdependent, cognitive and autonomous agents" [22].

By considering a multi-robot team as a society, social networks between robots can be identified and analysed. The knowledge obtained from the social networks analysis can then be used to describe multi-robot teams, as well as to predict and coordinate the behaviour of the team as a single entity.

Social relationships within a society are often very complex and can be multidimensional. More often than not, there can be more than one social relationship between two agents. In human societies, it is easy to grasp the wealth of relationships with all members of societies being linked to others.

For the purpose of this thesis, the approach taken is to initially consider higher mammalian societies and to isolate only a few applicable relationships. This is by no means an exhaustive approach, but rather an exploratory approach. The social network-based approach, as presented in this thesis, is by no means limited to the number of relationships that each agent can have. However, the implementation of the social networks based approach developed for the purpose of the simulations presented in this thesis is limited to two social relationships only, namely kinship and trust.

6.5 Related Work

Higher mammalian societies, and specifically human societies, have inspired research related to the applicability of the coordination mechanisms in MASs. During the 1990s, a number of researchers were involved with various society-inspired MAS-related research.

Social networks are an integral part of such societies and although there is no directly related work done on utilising social networks for task allocation as presented in this thesis, it is important to overview existing research efforts.

The social network-related research efforts can be broadly divided into three main categories:

- Research interested in social hierarchies as coordinating mechanisms.
- Research in modelling higher mammalian societies in order to better understand the subtle relationships that exist in them, with a view to be possibly used as coordination mechanisms.
- Research into the use of social networks for trust propagation in MASs.

6.5.1 Social Hierarchies and MAS Applications

Social hierarchies have been of interest to researchers from the early days of DAI research. The early work related to decentralised AI with application to multi-robot teams can be traced back to the work of Luc Steels in 1990 [177], although Steels was mainly interested in societies which were less complex than mammalian societies (insect societies).

The higher mammalian societies, such as packs of wolves and troops of chimpanzees, have also been investigated [149]. The hierarchies within these societies have inspired research at MIT [69] to explore the benefits of hierarchies in MASs for the purpose of tasks such as streamlining inter-agent negotiations and forming alliances between agents.

6.5.2 Modelling Societies

Using agents to model societies is becoming increasingly popular. Agent systems have been used to model the societies of primates [149], using tools such as MACACA [100], and even early human societies [55].

Modelling of social relationships between agents in a MAS, specifically between robots in multi-robot teams has been the focus of research headed by Dautenhahn [47][48]. A survey of socially interactive robots can be found in [69].

6.5.3 Social Networks for Trust Propagation In MAS

The concept of trust is very important in interaction between agents in a MAS. If coordination and/or cooperation is required, an agent makes a decision based on its perception of other agents' capabilities. The ideal situation is that each agent in a MAS is fully aware of all other agents' capabilities and their current status. More often than not this is not possible and in these situations an agent must trust the other agents' estimates of their own capabilities. For a detailed survey of trust-related research in the field of MASs the reader is referred to [163].

One of the methods for establishing trust in an agent's capabilities is through trust propagation via the process of querying trusted agents about the capabilities of the agent whose credentials need to be established. Social networks provide a mechanism for trust propagation. More on utilisation of social networks in trust propagation can be found in the work of Yu *et al* [201] and Schillo *et al* [164]. It is important to note that trust as defined in work of Yu is related to information systems security issues, while in this thesis the trust is related to agent capabilities.

6.6 Social Networks Based Approach

The new social networks based approach to coordination presented in this thesis uses task allocation as a coordination mechanism. The basic concepts and origins of the social networks based approach for coordination are discussed next.

6.6.1 The Biology Origin

The social networks based approach has its foundations in the observed similarities between higher mammalian societies and multi-robot systems. In both biological (e.g.

a wolf pack) and artificial (multi-robot team) systems, there is often the need for cooperation. In this section, a conceptual comparison is given between a multi-robot team and a pack of wolves.

Wolves are social animals that are organised into packs, governed by strong male and female animals (alpha male and alpha female) [133]. A wolf pack is characterised by the existence of a strong social hierarchy [133]. A wolf pack is usually a family unit, reflecting the existence of a strong kinship relationship between the pack members. A wolf pack is a very effective hunting team and can bring down prey much bigger than an individual wolf could.

The comparison between a wolf pack and INDABA will be made in relation to the five steps of the cooperative problem-solving approach as proposed in INDABA (refer to chapter 5), namely potential recognition, team formation, plan formation, plan execution, task evaluation and recognition.

In terms of the adopted robotic MAS taxonomy (refer to section 4.1), a wolf pack and INDABA can be compared as in table 6 (note that the characteristics of INDABA are described as the upper limit, not as a particular implementation, as INDABA is only a framework):

Team Characteristic	Wolf Pack	INDABA
Size of Team	LIM	INF
Communication Range	NEAR	INF
Communication Topology	BROAD	GRAPH
Communication Bandwidth	MOTION	INF
Collective Reconfigurability	DYN	DYN
Processing Power of a Team Member	TME	TME
Collective Composition	HET	HET

Table 6. Comparison of a wolf-pack and INDABA

6.6.1.1 Potential Recognition

When successful task completion requirements exceed the capabilities of an agent, a need for cooperation is recognised by members of the society. In the case of a wolf pack, wolves often hunt prey that is too big to be hunted by a single member. The need for cooperation is recognised and a wolf pack hunts as a team.

In the case of multi-robot teams, tasks that exceed the capabilities of a single robot abound. Box pushing [199][107] is one of the well-known problems, as well as foraging under a time constraint, where a single robot cannot complete the task in a prescribed time period, while a multi-robot team can.

For both biological (wolf pack) and artificial systems (multi-robot team), the need for cooperation is recognised. The potential recognition leads to the next step, i.e. team formation.

6.6.1.2 Team Formation

In a society a team is formed according to the relationships between its members. Considering a wolf pack, the distribution of labour is according to the hierarchical structure of the pack. The hunt is lead by the alpha male and alpha female, as they are the most efficient hunters.

If neither the alpha male nor the alpha female is capable of leading the hunt, the next most capable members will lead the hunt (a beta male or beta female will assume their position).

Considering heterogeneous multi-robot teams, similarities with a wolf pack are many. The most capable members of a multi-robot team are selected for a task. When the most suitable robots are not an acceptable choice (they might be cost-prohibitive or unavailable or malfunctioning etc.), the next most suitable robot will be selected, according to a social network relationship (e.g. “next of kin” or another member of a social group).

6.6.1.3 Plan Formation and Plan Execution

The plan formation step is not always applicable to packs of animals and multi-robot teams, and is largely ignored for the purpose of this comparison. It is sufficient to note that members of a society can have a specialised role that they perform and that a plan formation should take such specialisation into account.

The same applies to the plan execution step of a cooperative problem-solving process. Plan execution is not relevant to this comparison of similarities between biological and artificial societies as the focus is on task allocation.

6.6.1.4 Task Evaluation and Recognition

It might not be obvious, but a form of reward and punishment mechanism can be found in both a wolf pack and multi-robot teams (if it is so designed and programmed).

Considering a wolf pack, if the hunt was successful, the social hierarchy will be updated by strengthening the existing relationships. Furthermore, because the feeding order is dictated by the social hierarchy, the alpha male and female will eat first and eat the best parts of the hunted animal, in turn maintaining their physical supremacy over the rest of the pack.

In a multi-robot team (if so programmed) the agents that succeed in task execution will be rewarded and their affinity to the task will be increased. The team leader will strengthen its affinity to the task and maintain its team leader status.

However, if the hunt fails, probably nothing will happen immediately for the wolf pack. However, if the alpha male fails to feed the pack for extended periods of time, its social position can deteriorate to the extent that it is challenged by a beta male.

The same principle is applied to multi-robot teams. If a team leader repeatedly fails, its affinity to a task is decreased. It can happen that at a certain point in time the team leader is no longer the top-scoring agent in the team and it stops being the team leader. In other words, it may happen that over time, the team leader's bid to secure a task may be insufficient, in which case a different robot may win the bid and select a new team.

6.6.2 Comparison to Other Task Allocation Coordination Mechanisms

Approaches other than INDABA have been developed to use task allocation as a coordination mechanism, for example MURDOCH [77] and BLE [196].

BLE is based on the subsumption architecture and uses a port-inhibiting strategy for task allocation, where a robot can decide that it is the best eligible for a task. If a robot is eligible for the task, the robot can inhibit a communication port, effectively seizing control. MURDOCH, a market-based approach, uses a more traditional approach, an auctioning mechanism that governs task allocation, again based on a robot's own estimate of its capabilities. An extensive review of multi-robot task allocation mechanism can be found in [76].

While each robot is an autonomous agent in MURDOCH and BLE, the agents are basically unaware of other members of the society. The social networks approach presented in this thesis is different, because it relies on agents to belong to social groups and that agents maintain social links among themselves.

Task allocation in INDABA consists of selecting a team leader. This can be done by either using an auctioning mechanism or agent's historical performance on same or similar tasks. Once a leader is selected, a team is formed based on the strength of agents' relationships to the team leader and the task. A team is formed using a scoring system which takes into consideration all applicable relationships to the task in question. Based on relationships and task affinity each agent is given a score. The society members with the highest scores form the team together with a team leader.

6.6.3 Definitions and Notification

To describe the social networks based approach to task allocation in more formal terms, the following definitions are necessary.

Let T_k be a task that needs to be allocated, where $k = 1,..,K$ with K the number of tasks. If there are n known attributes of task T_k , then task T_k can be represented by an n-tuple, $(T_{k1}, T_{k2}, \ldots, T_{kn})$, where $T_{k1}, T_{k2}, \ldots, T_{an}$ are the n attributes that define task T_k .

The value of attributes $T_{k1}, T_{k2}, \ldots, T_{kn}$ can either be binary, discrete or continuous valued.

Let A_x be an agent, whose suitability to task T_k needs to be evaluated, where $x = 1,..,X$ with X the number of agents in the society. If there are m known attributes of agent A_x, then agent A_x can be represented by an m-tuple $(A_{x1}, A_{x2}, \ldots, A_{xm})$ where $A_{x1}, A_{x2}, \ldots, A_{xm}$ are the m attributes that define agent A_X .

The value of attributes $A_{x1}, A_{x2}, \ldots, A_{xm}$ can either be binary, discrete or continuous valued.

Let A_{lk} be the agent whose applicability to task T_k is the highest. Then agent A_{lk} is a team leader. Each leader has m known attributes, $A_{lk1}, A_{lk2,\ldots,} A_{lkm}$. The leader is represented by a m-tuple $(A_{lk1}, A_{lk2}, \ldots, A_{lkm})$ where $A_{lk1}, A_{lk2,\ldots,} A_{lkm}$ are the m attributes that define agent A_{lk} .

Let there be I relationships between agents in the society. Then relationships between agents A_{lk} and A_x in relation to task T_k are denoted as $R_i\,(A_{lk}, A_x, T_k)$ for $i = 1,.., I$. R_i is a function normalised to the interval [0,1], that is, $R_i : (A_{lk}, A_x, T_k) \rightarrow [0,1]$.

It is important to note that not all functions that model relationships require T_k as an input. For example, kinship is independent of a task under consideration, and therefore the kinship relationship depends only on the two agents involved.

With these definitions in mind, it is possible to define a fitness[2] (or scoring) function for an agent A_X and the team leader A_{lk} in relation to given task T_k as

$$F_{xk}\,(A_{lk}\,,\,A_x,\,T_k) = \Sigma_{i\,=\,1...I}\,(1\text{-}k_i)\,R_i\,(A_{lk}\,,\,A_x,\,T_k). \qquad (6.1)$$

where $\Sigma_{i\,=\,1...I}\,k_i = 1$. It is interesting to note that if A_{lk} is omitted from equation (6.1), then the remaining function $F_{xk}{}'$, given as

$$F_{xk}{}'\,(A_x) = \Sigma_{i\,=\,1,...,I}\,(1\text{-}k_i)\,R_i\,(A_x,\,T_k) \qquad (6.2)$$

is used by agent A_x to estimate its own eligibility to task T_k.

6.6.4 The Social Network Task Allocation Algorithm

The social networks task allocation algorithm is outlined below in general terms. A specific implementation that uses two social relationships is presented in greater detail in section 7.1. The social networks task allocation algorithm can be seen as an enhanced or augmented auction based task allocation algorithm. The bid is a function of the strength of social networks. The algorithm itself is surprisingly simple, and the key to its efficiency is in keeping the relationships up to date. The relationships can be stored either in a central repository or they can be distributed, with each agent maintaining its own social networks.

The advantage of a central repository system is that it is simpler to implement, but on the other hand it does require reliable communication channels between all the agents in the society and the central repository.

The distributed model is more applicable to robot teams and, in general, closer to the true notion of an agent. It does, however, require more complex implementation than a central repository system.

[2] Please note that in the context of this thesis, the notion fitness function is different from the concept of fitness function as used in evolutionary computing, i.e. it does not influence the survival of the agents in the society.

In the new proposed architecture, INDABA (refer to chapter 5), each agent maintains its own relationships data. The relationships data is created, stored and maintained in the interaction layer. From the robotic application point of view, the advantage of this approach is that an agent can rely on its internal social network to find the best candidates in its local environment, even if there is no complete and reliable communications with the rest of the agents.

The algorithm consists of four main steps:

- task detail propagation component
- the selection of a team leader most suitable to the task
- selection of the remaining team members
- task evaluation and social network maintenance

Each of these steps is described next.

6.6.4.1 Task Details Propagation Component

Once the task details are known, they are propagated to all available agents in the society. An external party, such as a user of the system, can either give the details of the task, or the task details can be obtained through the agent's exploration of its environment.

An example of the latter approach is a heterogeneous robot team where a specific robot performs the role of a scout and collects the information about the environment. The scout collects all the details about the environment using its own sensor suite and sends the details to the rest of the team for task allocation.

The task details T_k, represented as the n-tuple, $(T_{k1}, T_{k2}, \ldots, T_{kn})$, needs to be propagated to all participating members of the team. The propagation of task details may utilise any available communication protocol or method. The implementation of task propagation is not prescribed by INDABA and can take any form, from a simple binary coded string of predetermined length (where values correspond to the sensor readings of a scout) to much more flexible approaches such as KQML [66][99] or XML[186].

For truly unknown environments, where even metadata about the environment is not available, KQML together with a semantic descriptive language, for example Knowledge Interchange Format (KIF) [75], would be an advised approach. To illustrate, consider a scout in an unknown environment. If the scout discovers new concepts, those newly discovered concepts (metadata) can be described using KIF and propagated to the rest of the agents in the society using KQML.

Propagation of task details can be done either by a centralised entity (such as an external supervisory program, an approach similar to a supervisor in MACTA) or an agent can initiate the propagation of task details, in which case the agent becomes a managing agent. While the agent-initiated approach is advisable, it is more complex to implement. The implementation of propagation of task details is made according to the environment requirements and agent capabilities.

6.6.4.2 Team Leader Selection

Team leader selection can again proceed in at least two ways: either an agent can submit its own task affinity evaluation or an agent can be evaluated by an external supervisory entity. Considering leader selection, once the task details T_k are received, scoring takes place based on agent attributes and scoring function F_{xk}' (refer to equation (6.2)). The agent with the highest F_{xk}' is selected as the team leader.

Team leader selection is not social network related in its true sense (links with the other members of the society are not examined or utilised in selection), but it can rely on either direct matching of attributes to the task details (if possible) or to historical data (if available). If direct matching is not possible nor historical data available (which is the case when the task is executed for the first time), then an alternative selection method must be used, using a different scoring function.

As an example of an alternative team leader selection method, the leader can be selected randomly. The main advantage of a random selection method is that all agents are given an equal chance for task selection. However, using a random selection method, there is no guarantee that the selected team leader is capable of executing task of team leader. Alternatively, the cheapest member, determined using a cost function (if implemented), can be selected.

Both the social networks team leader selection method and an alternative team leader selection method can be combined in a single algorithm. If there is no historical data, an alternative selection method that uses a different scoring function $F_{xk}''(A_x)$ can be utilised, otherwise social networks selection based on scoring function $F_{xk}'(A_x)$ is used, as illustrated in algorithm 7.

$A_{lk} = A_1$
If historical data available or direct matching possible
  $F_{xk}(A_X) = F_{xk}'(A_X)$
Else
  $F_{xk}(A_X) = F_{xk}''(A_X)$
End If
For all agents A_x *in society S*
    If $F_{xk}(A_x) > F_{xk}(A_{lk})$
      $A_{lk} = A_x$
EndFor

Algorithm 7. Team leader selection in social networks based approach

6.6.4.3 Team Selection

Once a team leader has been selected, the rest of the team is selected. At this stage social networks play a crucial role.

The team member candidates are not only evaluated in relation to the task, but also based on their relationships to the team leader. This may look counter-intuitive, but it is not enough that a team member has an affinity to the task, the candidate must also be capable of working together with the team leader.

The relationships to the team leader are the crucial part of the algorithm and form the premise of the whole social networks-inspired approach. Agents are not individual, independent entities, but are defined in relation to the other members of the society. For team selection, team member candidates are also evaluated according to their ability to cooperate with the team leader, based on previous history (trust) and similarity to the team leader (kinship). In analogy to human societies, an agent is evaluated on how good it is as a "team player".

Team selection can be done by a team leader either according to its existing relationships to the other team members or the relationships can be recalculated prior to team selection. For both cases the algorithm is basically the same as summarised in algorithm 8.

While team T less than TeamSize
  $A_N = A_1$
  For all agents A_x *in society S*
    If A_x *not allocated to team T and* $F_{xk}(A_{lk}, A_x, T_k) > F_{xk}(A_{lk}, A_N, T_k)$
      $A_N = A_x$
  EndFor
  Add A_N *to team T*
EndWhile

Algorithm 8. Team selection in social networks based approach

6.6.4.4 Social Networks Maintenance

Once task execution finishes (successfully or not), the social networks need to be updated. Each member of the team needs to be evaluated and its relationships updated.

The exact method of updating the relationships is not prescribed by INDABA and can take the form of simply increasing a counter of successful or unsuccessful executions related to a particular task, or towards a particular team member. More complex methods can also be implemented.

Assuming that methods of updating agent relationships are given as *strengthen* and *weaken*, the social networks maintenance algorithm is summarised in algorithm 9.

For all agents A_x *in team* T
  For all R_i (A_{lk}, A_x, T_k)
    If T_k *completed*
      strengthen R_i (A_{lk}, A_x, T_k)
    ElseIf
      weaken R_i (A_{lk}, A_x, T_k)
    EndIf
  EndFor
EndFor

Algorithm 9. Social network maintenance

6.7 Summary

This chapter started with an overview of biologically inspired approaches to coordination. This discussion was followed by an overview of the two main approaches to coordination in MASs, which are based on organisational sciences, namely the market-based and hierarchical approaches.

The remainder of the chapter introduced the concept of social networks and presented a new approach to coordination in MASs. The new approach was based on social networks, and a modification for application in multi-robot teams was presented.

The next chapter presents the implementation of the social networks based approach within the INDABA framework, applied to simulated robots in an abstract simulated environment.