Cooperative Robotics: A Survey

(HS-IDA-MD-00-002)

Nicklas Bergfeldt (nicklas@ida.his.se)

Department of Computer Science
University of Skövde, Box 408
S-54128 Skövde, SWEDEN

Supervisor: Tom Ziemke

Cooperative Robotics: A Survey

Submitted by Nicklas Bergfeldt to the University of Skövde as a dissertation towards the degree of M.Sc. by examination and dissertation in the Department of Computer Science.

2000-10-16

I certify that all material in this dissertation which is not my own work has been identified and that no material is included for which a degree has already been conferred upon me.

Cooperative Robotics: A Survey

Nicklas Bergfeldt (nicklas@ida.his.se)

Abstract

This dissertation aims to present a structured overview of the state-of-the-art in cooperative robotics research. As we illustrate in this dissertation, there are several interesting aspects that draws attention to the field, among which 'Life Sciences' and 'Applied AI' are emphasized. We analyse the key concepts and main research issues within the field, and discuss its relations to other disciplines, including cognitive science, biology, artificial life and engineering. In particular it can be noted that the study of collective robot behaviour has drawn much inspiration from studies of animal behaviour. In this dissertation we also analyse one of the most attractive research areas within cooperative robotics today, namely RoboCup. Finally, we present a hierarchy of levels and mechanisms of cooperation in robots and animals, which we illustrate with examples and discussions.

Keywords: Cooperation, Collective Behaviour, Autonomous Robotics, Artificial Intelligence, Mobile Robots, Communication

Acknowledgments

There are several people I would like to thank for their different kinds of support during my work on this dissertation. First and foremost, I thank my supervisor Dr. Tom Ziemke for his valuable support, comments and help because without him, this dissertation would not exist. I would also like to thank my family and friends for bearing with me during this time, and last but certainly not least, my deepest gratitude to my fiancée Ida Jonegård for her love, support and patience.

Thank You!

Table of Contents

Table of Figures

1 Introduction

During the 1990s the interest for multiple, cooperating robots has increased dramatically. As we will see in this dissertation, this primarily depends on the fact that cooperative robotics attracts many different disciplines into the same domain and that the research in cooperative robotics can be used in many ways. In nature, cooperation is not unusual among animals, which often cooperate to achieve goals that would have been unreachable otherwise. A very popular species to bring up when talking about cooperation among animals and insects is ants (Kube and Bonabeau, 2000; Grassé, 1959; Wagner and Bruckstein, 1995). They have been the targets for numerous researchers and several robot experiments as well as theoretical analyses have been based on ants and their behaviour. For instance, ants cooperate to bring home food and building blocks. Another example is related to wild dogs, which are herd animals that hunt in flocks since the preys they hunt are relatively larger than themselves. It is this cooperative behaviour that is the reason for researchers to make comparisons and to take inspirations from animals, and in particular humans, since the cooperation and intelligence found there seems like a good first step towards artificial counterparts. This dissertation will present a collection of works from the cooperative robotics field and which disciplines that can benefit from this area.

First, in Section 1.1, the aims of this dissertation will be stated and in Section 1.3 an introduction to the cooperative robotics field will be presented. This chapter will end with an outline of the dissertation, which can be found in Section 1.4.

1.1 Aims of this dissertation

As we will see in this dissertation, there are a lot of contributors and many different perspectives in the field of cooperative robotics. Hence, it is difficult to form an overall picture of the field. The aim with this dissertation is to present a structured overview of cooperative robotics research, with special interest in the fundamental concepts and the current research topics identified in the field. Additionally, some open research areas will be suggested and various problems that are still open in the field will also be presented and discussed. The overview will for instance include some commonly used concepts and the similarities, as well as the distinction between them, will be emphasized with respect to other concepts.

This survey of the cooperative robotics field will include both physical, mobile robots as well as simulated ones. For instance, the survey will include cooperative robot pushing, clustering, team games (in particular robotic soccer), exploration, etc. Furthermore, this survey will also explore software issues of robot groups, for instance 'awareness' of other robots (which is discussed in more detail later in this dissertation). However, this dissertation will not include issues like human-robot cooperative systems[1] or competitive robots that are run one-on-one, like predator-prey co-evolution[2]. The reason for this is simple, since in the case of human-robot cooperation one cannot talk about cooperative robotics as in cooperation between robots, simply because there is only one robot. The same argument applies to robot systems that are run one on one; there is no cooperation involved since there is only one robot in every team.

1.2 Related work

There have been a few earlier surveys of the cooperative robotics field (Cao *et al*., 1997; Asama, 1992; Dudek *et al*., 1993), which address different aspects of mobile cooperative robotics, but not any of them has covered any research after 1995. Since the area has so rapidly grown in the last couple of years this dissertation will follow up on the previous works and extend them with recent research in the field.

1.3 Introduction to Cooperative Robotics

Researchers within the artificial intelligence community had put up a goal that should be achieved before the year 2000. That was to beat the current human world champion in chess with a computer program. This goal was achieved in 1997 (cf. Schaeffer, 1997) and now a new goal has been set:

> "By the mid-21st century, a team of autonomous humanoid robots shall beat the human World Cup champion team under the official regulations of FIFA." - (Asada and Kitano, 1999).

[1] For this kind of work, see for instance Kuniyoshi and Nagakubo, 1997, where they train a humanoid robot with visual observation of a human who is performing the desired task. Another article on robot-human interaction is by Dautenhahn and Werry, 2000, where they discuss the dynamics involved and how mobile robots can play a therapeutic role in the rehabilitation of children with autism.

[2] There are a number of articles that address the issue of *competitive* robot co-evolution (see for instance Floreano, Nolfi and Mondada, 1999).

Nowadays robots are playing soccer in teams against other robots. This is still pretty far of from humanoid robots but the research is continuing and the first steps towards soccer playing humanoid robots have been taken (Mataric, 2000; Kuniyoshi and Nagakubo, 1997; Cheng and Kuniyoshi, 2000). Since 1997 world championship tournaments in robot soccer have been held annually (Asada and Kitano, 1999), so one might think that cooperative robotics research already has come very far and that the area has been investigated completely. However, a closer look at successful teams reveals that in many cases the robots actually do not really cooperate, but they merely 'coexist' in the same physical arena, and in the rare cases where cooperation exists it is at a trivial level. The cooperation has a long way to go if we intend to give human soccer players a challenge.

Earlier, most robotics research focused on developing single robots that should accomplish a certain task. Now, it looks like we have come to a point where the need for multiple, cooperating robots is increasing. Surprisingly, this can in some cases be done rather easily; take a couple of "dumb" robots, and put them together. The individual robots are not capable of performing many tasks but as a colony they can show a more intelligent behaviour (Parker, 1999; Werger, 1999; Kube and Zhang, 1997). This is much like the behaviour we see in ants; a single ant may seem to be very random in its behaviour and has apparently no goal in life, whereas an ant colony can accomplish rather intelligent things together (Kube and Bonabeau, 2000).

We can relate the above advantage to humans, who often cooperate in many and varied ways; they help each other to accomplish goals that could not have been achieved otherwise. The cooperation humans are doing with other people often involves exchanging resources for mutual benefit. For instance, one person cannot move a piano a very long distance but if there are four people helping each other, they can. In the long run this may be seen as a form of mutual benefit since the participating humans at a later time may request help back. This of course also applies to robots, as well as the aspect that more agents can together accomplish goals that otherwise would have not been feasible (Kube and Bonabeau, 1997). With all of this in mind it is easy to see that cooperation among robots can be very beneficial to accomplish. However, the number of articles that address this is growing rapidly and thus the necessity for an overview of the field of cooperative robotics becomes more and more increasing.

In this dissertation, we present an overview with detailed examples of what has been done and what techniques are being used in the field. Common concepts regarding cooperation among mobile robots will also be listed and analysed. This dissertation will identify areas in cooperative robotics where there has been little work done yet.

1.4 Dissertation Outline

In Chapter 1 an introduction to the field of cooperative robotics is presented along with the aims of this dissertation. Chapter 2 will give an overview of some of the different perspectives found in cooperative robotics research along with a motivational background for doing this kind of work. In Chapter 3 we present some of the most commonly appearing concepts in the cooperative robotics literature and various definitions of those will be stated. A taxonomy of the current research field is presented in Chapter 4, and in Chapter 5 we discuss what it is that has been done in the cooperative robotics field and what seems to be missing. We also emphasise the contributions of this dissertation.

2 Perspectives on Cooperative Robotics research

This chapter will present some examples where the common ground is cooperative (mobile) robotics. As we will see in this chapter, there are many different disciplines involved in cooperative robotics research, e.g. biology, engineering, cognitive science and AI. In this chapter we will present a couple of examples from each discipline, and the link to cooperative robotics will be emphasized.

We roughly see the field of Artificial Intelligence, and in particular the field of Cooperative Robotics, as two major interest areas; namely so-called ‘Life Sciences’ and ‘Applied AI’. According to Pfeifer (1995) there are two main motivations for AI research: firstly that of science, the understanding of natural systems, and secondly that of engineering, the concern of building ‘useful’ systems. Similarly, we can distinguish between cooperative robotics research motivated by the cognitive science, biology, ethology, etc. which is mostly concerned with the use of robots as tools in the investigation of the mechanisms underlying animals’ collective behaviour. Secondly, we can distinguish another group that is more concerned with building ‘useful’ systems, namely applied or engineering-oriented cooperative robotics research. However, the engineering-oriented systems may never the less certainly be inspired by ‘Life Sciences’ but the main concern is still that of engineering and not biological or cognitive plausibility.

‘Life Sciences’ are mostly concerned with biological and cognitive plausibility in their systems whereas ‘Applied AI’ mostly are concerned in building ‘useful’ systems. However, the usefulness should not be disregarded in ‘Life Sciences’ even if that is often just a side-affect of the primary goal. The examples that will be presented are only there to point out that cooperative robotics can attend so many different disciplines. The reader should thus note that the examples are not fully described (for more details we refer to the original articles). What we will see from these examples is that since there are so many different disciplines involved in this research area, there is some degree of conceptual confusion and that the approaches and goals can tend to be quite different.

First some low-level, basic behaviour is analysed and then we go on with more complex forms of cooperation activities like language evolution (we try to present the experiments in an order that is somewhat easy to manage, since the order in a way reflects the complexity in the different experiments). From the small overview presented in this chapter, we will see that the field is very dispersed and that many different disciplines have a lot to gain when using cooperative robotics (and probably, this is why they use it).

This chapter is divided as follows; first some examples of work from different disciplines in the field will be presented, followed by a discussion of the important issues in the different disciplines that are related to cooperative robotics. In Section 2.1 various works that are biologically or cognitively inspired will be presented. In Section 2.2 a more practical example will be presented that is more related to engineering-oriented cooperative robotics, and in Section 2.3 a short summary of this chapter will be presented.

2.1 Biological and Cognitive plausibility

As mentioned above, much cooperative robotics research is motivated by cognitive and behavioural sciences and are concerned with understanding the behaviour of living organisms and its underlying mechanisms. This section briefly illustrates some works investigating principles of natural collective behaviour in robotics experiments and this will illustrate some of the usefulness of cooperative robotics for disciplines like for instance cognitive-, biology- and behaviour sciences.

The following sections will include some examples that deal with cooperative robotics in some way and where the main goal with the systems is not purely engineering-oriented. First, in Section 2.1.1, the concept of stigmergy is introduced and briefly described (a more detailed description along with a discussion about the concept can be found later in the dissertation). This stigmergy example reflects some very low-level behaviour seen in, for instance wasps, ants and bees. Since stigmergy is a very common cooperative mechanism that is often found in nature, it is not surprising that much research has addressed this. Two important research areas in cognitive- and behaviour sciences is imitation and language (Dautenhahn, 1995) and in Section 2.1.2, some imitation examples are described and some language related examples are presented in Section 2.1.3.

2.1.1 Stigmergy

In research related to cooperation among robots, many parallels are made with different animals and how they use cooperation to achieve different tasks. Animals can be very intriguing and from simple behaviour, complex cooperation can emerge.

Grassé first defined the concept of stigmergy in 1959, when he was studying nest building in a colony of termites. Beckers *et al*. later (1994) defined stigmergy as follows:

> "The production of a certain behaviour in agents as a consequence of the effects produced in the local environment by previous behaviour."

Stigmergy is a notion for changing something in the environment and that the change then affects another agent.

Stigmergy is also called "cooperation without communication" by some researchers (e.g. Beckers *et al*., 1994) because the cooperation achieved is not based on direct communication but rather some indirect variant of communication that uses the environment to transfer the messages from one agent to another. For instance, one agent may move a box and another agent notices that the box has been moved and reacts to that.

Several robotics researchers have used the concept of stigmergy in their research (e.g. Beckers *et al*., 1994; Kube *et al*., 2000). In one experiment by Beckers *et al*. (1994) a group of mobile robots should gather 81 randomly distributed objects. The robots in this experiment used stigmergy in the sense that when one robot moved an object, that affected the others' behaviour since they sensed the object at a new position.

Another experiment that has been conducted along the lines of stigmergy is a work by Kube *et al*. (2000) where multiple mobile robots are used to accomplish cooperative transport. In this particular experiment the robots' task is to move a large box towards a lightened area. The box that the robots shall move is too heavy for one robot to push and it is thus required that several robots push the box simultaneously from the same direction, in order to make the box move. Here the robots sense when the box is being pushed in one direction and reacts on that, as well as on other sensing in the environment (like for instance light change). The environment is intentionally constructed this way because cooperation is of concern here and by doing this, the robots must cooperate in order to accomplish the task. This kind of cooperation can also be observed in for instance ants (Kube *et al*., 2000), where they cooperate when transporting a large and heavy prey.

With these examples in mind it is easy to understand that cooperative robotics can aid in biological research in that way as when trying to investigate the cooperation seen in for instance some insects. If we observe the cooperation in nature and then try to build a model based on those experiences, the model and the hypotheses can be validated by simply building our own colony of robots and then investigate the cooperation seen there.

2.1.2 Learning by Imitation

In cognitive research and behaviour sciences, one major area involves research concerned with social skills, such as imitation (Dautenhahn, 1995; Billard and Dautenhahn, 1999). Billard and Dautenhahn use imitation to train a group of agents in a simulator. In previous articles by Billard and others, they trained a single robot with imitation so that the robot would learn the vocabulary of the robot teacher. The vocabulary contained words that could be used to describe the surroundings and the task for the learner robots were to be able to use this vocabulary to describe the surroundings of the robot. Several experiments have been done regarding for instance perception of objects (Billard and Dautenhahn, 1998) and internal perceptions of movement and orientation (Billard and Hayes, 1998) and inclination (Billard and Dautenhahn, 1997). In those experiments there was only one learner and one teacher and the hypothesis was that this would scale up to several learner robots.

In the environment used by Billard and Dautenhahn (1999), there were one teacher robot and eight learner robots and the task for the teacher robot was to teach all learner robots its vocabulary (i.e. the correct association between phrases that describes the environment and corresponding sensory input). All robots wander around in the environment and the teacher describes the surroundings as they walk (see Figure 1 for an illustration of the simulated environment). In the environment, there are areas that have different colours and the learned vocabulary could be used to distinguish the different areas by location and colour. During a second experiment the robots used the learned vocabulary to investigate a large area. The robots would wander around in the environment and when coming across a coloured patch they would tell the other robots of its location and colour. In this way the robots covered a larger area by cooperation that would have been possible with one robot in the same time.

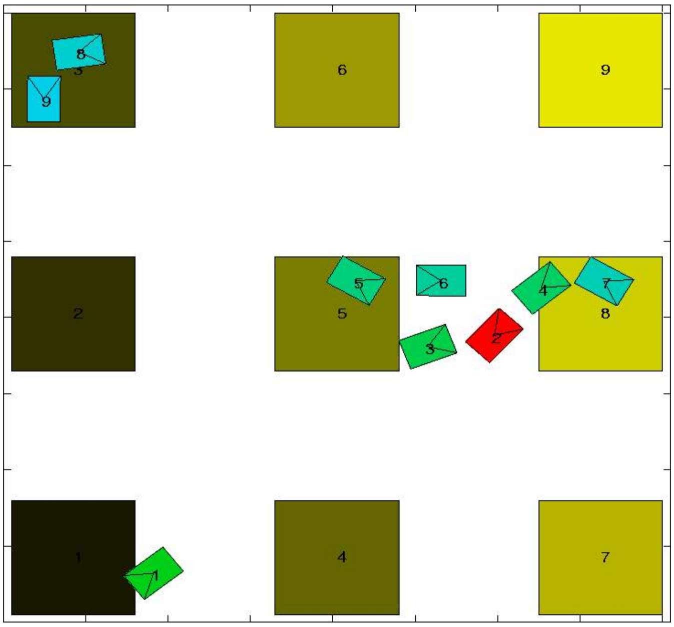

Figure 1 - Agents in Simulator used by Billard and Dautenhahn (1999). With courtesy of Billard.

Here we see that cooperative robotics can be very useful when investigating cognitive and behaviour aspects like imitation and learning.

2.1.3 Language

Another major area in cognitive research involves research on language, including its origin and evolution through time. In more recent years, some research has been focusing on physical, mobile robots that communicate with each other in an environment by using some sort of language (Billard and Dautenhahn, 1997; Steels and Vogt, 1997; Vogt, 2000). Some research that has been done in this area lately is focused on letting the robots evolve their own language (see for example Steels, 1996; Steels and Vogt, 1997; Vogt, 2000) from the perceptions that the robot gets from its surroundings and by communication between robots in the colony. This research is focused on the hypothesis that languages emerge from cultural evolution, and is not purely genetically determined. Other research is based on supplying one robot with a language and then the other robots in the colony should learn from that robot (see for example Billard and Dautenhahn, 1997; Billard and Hayes, 1998; Billard and Dautenhahn, 1999), and from this teacher-learner scenario, the colony of robots will evolve a common language and thus a common understanding of their environment.

Steels and Vogt have done some experiments with a couple of physical, mobile robots (see for example Steels, 1996; Steels and Vogt, 1997; Vogt, 2000). In the experiment from 1997 the main research interest was focused on how language can become common in a group of individuals using so-called adaptive language games. The game used in their experiment was a naming game. The robots walked around in an enclosed environment and when two robots found each other the game started. The robots should then focus on an object in the environment and 'talk' about this, i.e. agree on a common word for that object. When the robots have agreed on what word to use for this object they scatter and try to find other robots. In this experiment physical, mobile robots where successfully used to investigate two fundamental questions regarding the origins of cognition (Steels and Vogt, 1997):

> "(1) How can a set of perceptual categories (a grounded ontology) arise in an agent without the assistance of others and without having been programmed in (in other words not innately provided)."

> "(2) How can a group of distributed agents which each develop their own ontology through interaction with the environment nevertheless develop a shared vocabulary by which they can communicate about their environment."

In this example we easily see the usefulness of cooperative robotics when investigating social and cognitive aspects like language evolution, and how it may be possible to learn representations of objects in the world.

2.2 Engineering-oriented Cooperative Robotics

There are a lot of works that do not really care about biological or cognitive plausibility, but the main concern here is to make the robots complete the task in time (cf. 'Applied AI'). In this section we will give an example of such a work. Donald *et al*. (2000) use three cooperating robots to manipulate and wrap a rope around three boxes in various ways and the main concern here is not biology or cognitive plausibility, but merely achieving the goals.

Donald *et al*. performed three different experiments with the robots:

Binding: Tying a rope around the boxes (see Figure 2).

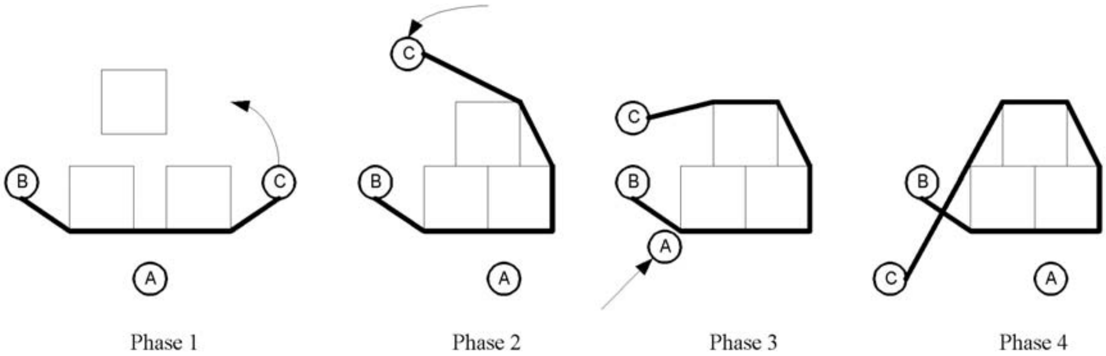

Figure 2 - Three robots tying a rope around three boxes, redrawn from Donald *et al*. (2000).

Flossing: Affecting rotations of the tied boxes by pulling the rope in various ways (see Figure 3).

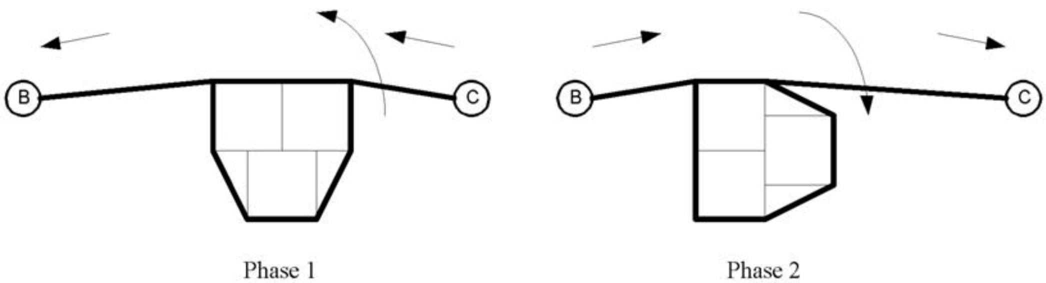

Figure 3 - Two robots rotating three boxes, redrawn from Donald *et al*. (2000).

Ratcheting: Affecting translations of the tied boxes (i.e. moving the tied boxes along a straight line, see Figure 4).

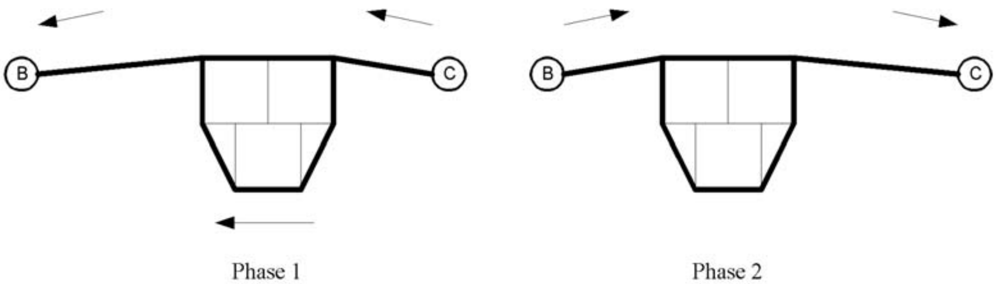

Figure 4 - Two robots moving three boxes, redrawn from Donald *et al*. (2000).

In these experiments the aim is not the cognitive or biological plausibility, but merely the aim of achieving the specified task reliably. The steering mechanism used in these experiments is purely mechanical and there is no learning in the system (to solve the task that is not necessary). However, the cooperation used in this task is necessary for the robots to be able to accomplish the goals.

2.3 Summary

As shown in the previous sections, the use of physical, mobile robots can be used in many different ways and with a lot of different intentions. There are many different disciplines that come together in the mobile, robotic domain. For AI researchers it is an intriguing and challenging domain to make the robots cooperate and to make them solve the tasks that are put before them, and as mentioned before, in some cases it may be necessary to have the robots cooperate. Cognitive scientists can use this domain to find out, for example, how language can evolve and how / why and when cooperation is established between different agents. Moreover, they can use robots to model underlying mechanisms for intelligent behaviour and thus test different hypothesis about various cognitive aspects. Obviously, when dealing with robotics, there usually are robots, and thus engineers have the opportunity to both build robust robots and to engineer new and better technology for building robots. Of course, since this is actually a real-time domain, real-time researchers also can be interested in this, as well as others can benefit from algorithms and solutions from the real-time domain. Finally, biologists can use this domain to test theories and hypotheses about animal behaviour in certain contexts. As we can clearly see, the cooperative robotics domain can very well attract a lot of different researchers from different disciplines.

3 Key concepts

This chapter presents some of the concepts that were found in the cooperative robotics literature. As mentioned before, the cooperative robotics field is very dispersed (since there are so many and different disciplines involved) and this will become very clear in this chapter. The various definitions will be discussed and related to other definitions of the same concepts and the similarities, as well as the distinction between the concepts, will be emphasized with respect to other concepts.

First, in Section 3.1 the distinction between a distal and a proximal view will be discussed. Then in Section 3.2 the concepts of cooperative and collective behaviour are discussed and the distinction between these two concepts are emphasized. In section 3.3 different levels of communication and their complexity is presented. Section 3.4 brings up the issue of awareness among robots and the chapter is ended with Section 3.5 where some social aspects in cooperative robotics are discussed.

3.1 Distal and Proximal view

The distinction between a distal and a proximal view (cf. Sharkey and Heemskerk, 1997; Harnad, 1990; Nolfi, 1997) is quite important to realise when dealing with cooperative robotics. Since the facts that are being considered when looking from a purely distal view is only what can be observed from the outside, the real underlying mechanisms may be neglected.

From a distal perspective it may seem as the robots are cooperating, but when analysing the underlying mechanisms it becomes clear that the robots maybe for instance have no notion of the other robots in the environment, and can thus not be aware of any team members, even though it *looks* like that is the case. The robots only act *as if* they were aware of each other; actually they are not (cf. Melhuish *et al*., 1999).

In most works the concepts of, for instance, social- and cooperative behaviour have been defined from a distal perspective. All that has been done is that the researchers have observed the robots from the outside and then drawn the conclusion that the robots are cooperating or that they are performing social interactions with each other. Then again, since we yet only can observe human social behaviour from the outside it becomes clear that defining what social behaviour is and then saying that these robots are interacting socially is very hard, since it is hard to say what it really is (at a proximal level) to socially interact among humans.

3.2 Cooperative and Collective behaviour

What is cooperative behaviour and is there any distinction from collective behaviour? When we say that a group of agents are cooperating, then what do we really mean – what is it to have collective behaviour in a group of agents?

The distinction between cooperative and collective behaviour seems to be a little vague or somewhat ambiguous in the literature. Some researchers speak of collective behaviour (e.g. Beckers *et al*., 1994) when others would have called it cooperative behaviour (e.g. Kube and Bonabeau, 2000) and this could lead to misunderstanding in some cases since we believe (and other researchers with us, e.g. Cao *et al*., 1997; Touzet, 2000) that there is a quite noticeable distinction between these two concepts. Cooperative behaviour can be seen as a subclass of collective behaviour, which has some specific attributes that do not exist in its superclass collective behaviour. What these specific attributes are that distinguishes cooperative behaviour from collective behaviour depends on which author you ask, for instance Cao *et al*. (1997) state the following regarding cooperative behaviour:

> "Given some task specified by a designer, a multiple-robot system displays cooperative behavior if, due to some underlying mechanism (i.e., the "mechanism of cooperation"), there is an increase in the total utility of the system." – (Cao *et al*., 1997)

It may be argued that this definition is a little vague since the authors speak of "some underlying mechanism", i.e. "the mechanism of cooperation", which is not discussed in the definition. However, they do mention earlier in the article "Cooperative behaviour is a subclass of collective behaviour that is characterised by cooperation" but this still does not give a satisfying definition of cooperative behaviour. They say that cooperative behaviour intuitively should give some performance gain over naive collective behaviour and this seems to be the key issue in most definitions.

Regarding the "mechanism of cooperation" Cao *et al*. (1997) state the following:

> "The mechanism of cooperation may lie in the imposition by the designer of a control or communication structure, in aspects of the task specification, in the interaction dynamics of agent behaviors, etc." – (Cao *et al*., 1997)

Thus, Cao *et al*. mean that it may be the case that the designer of a cooperative robotics system is enforcing the cooperation by introducing some sort of control structure or by defining the task in a certain way as to impose cooperation. These design choices will thus have an impact on the cooperation being performed by the system and to what extent this cooperation will emerge, and thus in the end it is the designer who controls what kind of cooperation that is being performed by the system and in what way this cooperation is being manifested.

Mataric (1994) gives the following definition of collective behaviour and by this states that it is only an observer-subjective definition:

> "Collective behaviour is an observer-subjective definition of some spatial and/or temporal pattern of interactions between multiple agents."

This can be related to Brook's "Intelligence is in the eye of the observer" (Brooks, 1991), and in this particular case the observer see some interaction between the agents and draws the conclusion that this is collective behaviour, all from a distal perspective (cf. Section 4.1).

In order to explain cooperation Mataric (1994) gives the following definition:

> "*Cooperation* is a form of interaction, usually based on communication."

It can be argued that Mataric's definition of cooperation is rather vague, since the definition is based on two other not-well-defined concepts (namely "interaction" and "communication"). Mataric has defined these concepts in her thesis but even with those additional definitions the meaning of cooperation is still vague. Mataric tries to explain the definition by further divide the concept of cooperation into two sub domains:

- *Explicit cooperation* – is defined as a set of interactions, which involve exchanging information or performing actions in order to benefit another agent.
- *Implicit cooperation* – is defined as a set of actions that are part of the agent's own goal-achieving behaviour repertoire, but have effects in the world that help other agents achieve their goals.

With these additional definitions of cooperation it is now easier to see the similarity with Cao *et al*. (1997). According to Mataric the agents cooperate when they benefit other agents (i.e. help them achieve their goals, intentionally or unintentionally) and if the agents goals are fulfilled faster then the performance of the system would naturally increase. Thus, from Mataric's definition of cooperation we see that it is based on performance gain as with Cao *et al*. (1997).

Surprisingly to some, in the cooperative robotics literature the definition of cooperation and collective behaviour is very sparse. If this is due to that these concepts are in fact well defined and commonly understood in the community, then that does not show in any of the articles that we have surveyed. Jung and Zelinsky (2000) make a rather interesting comment regarding cooperation, it is only a "label for a human concept" and it only refers to a category of animal behaviour and that humans have made the category up. Jung and Zelinsky does not state anything about what *cooperation* is, but they do however discuss the *relation* between communication and cooperation and state that those two concepts are closely tied:

> "Communication is an inherent part of the agent interactions underlying cooperative behaviour, whether implicit or explicit."

Jung and Zelinsky (2000) divide the notion of cooperation into terms of communication, which they in turn divide in several sub categories. According to Jung and Zelinsky (2000), cooperation and communication is very closely tied, and thus the explanation of cooperative behaviour resides in the definitions of communication. Jung and Zelinsky classify communication into four categories and then use that classification to examine some examples of cooperative behaviour seen in some biological systems (see Section 3.3 for a more detailed discussion about these categories).

In the cooperative robotics literature, the usual approach is to let cooperation include everything that is cooperative from a distal perspective (regardless of what the internal states are in the individual agents, and thus awareness of the cooperation is often not discussed). Furthermore, there are not many attempts in trying to define what it really is to cooperate. One thing that is very common though, is that many researchers (see for instance Grassé, 1959; Kube and Bonabeau, 2000; Jung and Zelinsky, 2000) compare different levels of cooperation with the different kinds of cooperation seen in nature among for instance animals and insects.

Although definitions of cooperation seem to be a little different, the main idea underlying the different definitions seem to be based on the simple measure of performance gain. If there is an increase in performance, compared with simple and naive collective robotics, then the system is said to cooperate.

Cooperation seems to be hard to define precisely, and in many articles where it is used there are no definitions at all. The cooperative behaviour that is mostly referred to in the literature is when the authors of the articles say that the agents cooperate, not according to any specific definition but just that *they* think that the agents cooperate, thus agreeing with Mataric's (1994) subjective view as mentioned earlier (of course there are several definitions which were discussed in this section, we only want to stress that definitions of cooperation and collective behaviour are surprisingly sparse in the literature).

Despite the differences in the definitions, the relation between the two concepts cooperative and collective seems however to be commonly understood, cooperative behaviour is seen as a subclass of collective behaviour. Thus, there is some additional agenda that have a positive impact on performance when dealing with cooperative- and not purely collective robotics.

3.3 Communication - Stigmergy and Language

Many articles use some sort of communication in their experiments, and since communication is the most common means of interaction among intelligent agents (according to Mataric, 1994), this is not surprising. Many researchers have proposed classifications for the types of communication found in biological and artificial systems (see for instance Arkin and Hobbs, 1992; Dudek *et al*., 1993; Balch and Arkin, 1994; Kube and Zhang, 1994; Cao *et al*., 1997; Mataric, 1997; Jung and Zelinsky, 2000). The complexity of the communication found in various works is rather different and varies quite a lot. Some articles only use the simplest forms of interaction between agents (i.e. stigmergy) to achieve their goals, whilst others try to have their robots learn to communicate about different objects, physical impacts and even representations of other robots actions and intents.

Cao *et al*. (1997) characterise the following three major types of interactions that can be supported by a multi-robot system:

- Interaction via Environment – is the simplest, most limited type of interaction that occurs when the environment itself is the communication medium, and there is no explicit communication or interaction between the agents. This is also called stigmergy by other researchers.
- Interaction via Sensing – is the local interactions that occur between agents as a result of agents sensing one another, but still without explicit communication. Of course this kind of communication requires the agents to be able to distinguish between objects and other agents in order to be able to react accordingly.
- Interactions via Communication – is the explicit communication with other agents. The agent knows that it is sending some message but it does not have to know the recipient(s) – they may be known in advance but that is not necessary for this kind of communication. Here the agent knows that it is sending some message but whether or not this implies the intent to communicate is still an open issue.

Mataric (1994) divided the notion of communication into two sub domains:

- *Direct communication* – the communication is directly aimed at a particular receiver(s), which is/are identified.
- *Indirect communication* – is based on the product of other agents' behaviour and this is often called stigmergy by other researchers.

Mataric's 'Indirect communication' is actually identical with 'Interaction via Environment' presented by Cao *et al*. (1997), the only thing that differs is the choice of words. However, Mataric's 'direct communication' is actually more specific than the highest level presented by Cao *et al*. since Mataric requires the receiver to be known in advance, and hence it can be argued that Mataric's definition of direct communication requires a higher level of complexity by the agents and that the intent of sending a message to a specific receiver is somewhat more present.

Jung and Zelinsky (2000) define communication as follows:

> "A communicative act is an interaction whereby a *signal* is generated by an *emitter* and 'interpreted' by a *receiver*."

The complexity of interpreting a message can of course vary much between different agents and Jung and Zelinsky also discuss this in their article. They split up communication into the following four characteristics and the complexity of the interpretation is included as the fourth category:

- Interaction distance – The distance between the involved agents.
- Interaction simultaneity – The period between the signal emission and reception.
- Signalling explicitness – This is an indication of how much the emitter is changing the signal (e.g. due to an evolved or learnt behaviour or just an implicit signal emission).
- Sophistication of interpretation – This is an indication of the complexity of the interpretation process that gives meaning to the signal (e.g. it is possible that a signal can have very different meaning to the emitter and receiver).

Billard and Dautenhahn (2000) state a somewhat more straightforward definition of communication:

> "… two agents are communicating once they have developed a similar interpretation of a set of arbitrary signals in terms of their own sensor perceptions."

This statement can be related to Jung and Zelinsky (2000) where they state that a signal is *interpreted* by a receiver and this interpretation should, according to Billard and Dautenhahn (2000), result in a similar categorisation both in the sending agent and the receiving. However, when also considering Jung and Zelinsky's (2000) 'Sophistication of interpretation' this becomes more ambiguous, since they state that the emitting agent and the receiving agent could very well interpret the signal differently. But this contradicts Billard and Dautenhahn's (2000) statement, since they think that the two agents should achieve a *similar* categorisation of the set of sensory perceptions.

As mentioned earlier in this section there have been many attempts to categorize the notion of communication into various levels and the result of these efforts has been a number of rather different categorisations, although still with a common notion of something that is being transmitted between two or more agents. However, it is easy to see the common low-level in almost every categorization: stigmergy. This category is often included as the lowest level in some form.

3.4 'Awareness' of cooperation

The issue about whether or not the robots are aware of the cooperation being performed, or the awareness of other team members, is very rarely discussed or even mentioned. What is it to be aware of cooperation and is it beneficial in some way to be conscious of other team members and what they can achieve in different situations, or is that just not necessary?

As we will see later in this section, there are some articles that do performance comparisons and the results are unambiguous: awareness should be addressed in some way, since even with very simple forms of awareness the performance gain could be noticeable (of course this depends on the task being performed by the robots but the possibility of awareness among the robots should always be in the mind of the designer of the system).

The awareness issue is often somewhat overlooked and often the robots are expected to see other robots as obstacles rather than team mates. However, there are some articles that address the issue of awareness of cooperation (or awareness of other robots in the environment, as most of the researchers implement it).

Mataric (1992) shows that the ability to distinguish between other robots and the rest of the objects in the world provides an advantage when training robots. Mataric divides the notion of recognising other robots as robots into three different levels:

- *Ignorant Coexistence* - The robot in question thinks it is the only robot in the environment; all other robots in the environment are treated as obstacles. This is the simplest case of coexistence and it is very common in the cooperative robotics literature (see for instance Beckers *et al*., 1994; Cao *et al*., 1997; Kube and Bonabeau, 2000). Everything that the robot encounters in the environment are treated as obstacles (e.g. walls) and this in particular applies to other robots.
- *Informed Coexistence* - The robots have the ability to sense the presence of each other, i.e. to discriminate between two types of objects in the world: obstacles and other robots. With this form of coexistence there can only be reactive behaviour upon sight of other robots, when the other robots leaves the field of vision the concept of those robots is no longer available.

- *Intelligent Coexistence* - This is an expanded version of the Informed Coexistence. In this case the robots not only sense other robots in front of them (i.e. in their field of view) but also within a specified radius all around the robot (Mataric (1992) used a radius of 36-inch). This case of coexistence thus includes the ability to have an internal representation of other robots in the environment (although only within a specified radius), since the robots do not have to be directly in front of each other in order to be recognised. This can be seen as a simplified variant of the ability humans have to internally represent a person in the room, even if we close our eyes we know that the person is still there. In this case the internal representation is managed and simplified by adding extra sensory information. It may be argued that the so-called 'intelligence' here is nothing more than additional sensory inputs and that one might have expected a little more advanced level of awareness in the level called 'Intelligent Coexistence', but on the other hand it may be argued that the resulting behaviour can be seen as more sophisticated than purely 'Informed Coexistence'.

Besides Mataric's article there are not many articles that address the issue of recognising other robots, nor the ability to be aware of the cooperation in progress. However, there is another article that addresses a similar issue and where the results are somewhat similar. Parker (1995) investigates how awareness of team members actions can affect performance in a group of mobile robots and the results point in the same direction as Mataric's; the more the robots are aware of each other's presence and actions the more beneficial for the whole task. The results that Parker presents show that it may be more beneficial however to increase the number of robots instead of just adding awareness of other robots in the group. Of course this depends on the task and thus the designer of such a robot system should think about what would be more beneficial for the particular task – awareness or just adding more robots.

Parker used a puck-moving mission to investigate the impact of robot awareness of team member actions. The task for the robots is to move all pucks to the desired final position as fast as possible.

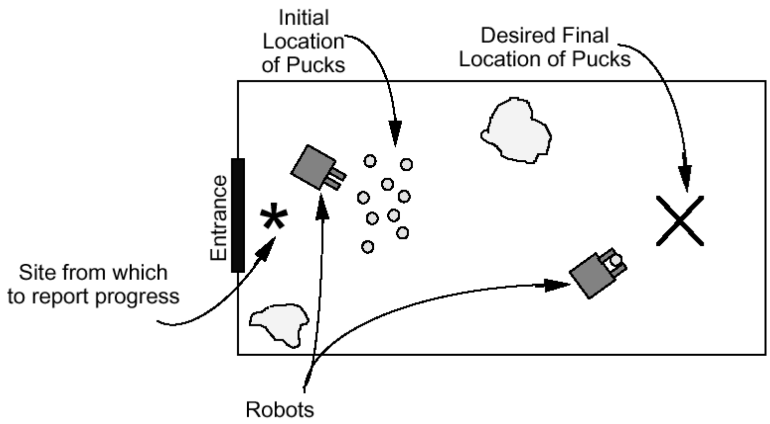

Figure 5 - Environment used by Parker (1995), with courtesy of Parker.

To address the issue of awareness in the robots Parker used communication as a means for the robots to sense each other and to be able to exchange information about certain positions. Thus, when considering and comparing with non-aware robots the ability to communicate was the only thing that was different. A number of comparisons were made between these and other groups and the results were unambiguous; awareness increased the performance in all cases (for this particular task).

Touzet (2000) takes another approach to implement awareness and uses the ability to sense the other robots. With this Touzet wanted to emphasise the fact that a robot may become aware of a team member's actions without the use of any explicit communication, for instance through passive action recognition. Touzet shows that awareness of other robots can be used instead of communication to achieve better performance when trying to cover a large search space.

But if (according to Parker, 1995) communication can be used to achieve awareness, and (according to Touzet, 2000) awareness gives better performance than only using communication, then the communication that Touzet speaks about must be very limited since it obviously cannot achieve awareness. From this it is easy to deduce that Touzet does not think that the ability to communicate is a sufficient prerequisite for awareness.

Furthermore, Touzet states that awareness is necessary when dealing with cooperative robot learning (there are other researchers that state that awareness is not a necessary condition for cooperation, see for instance Jung and Zelinsky, 2000). Touzet gives the following definition of awareness:

> "Awareness encompasses the perception of other robot's locations and actions"

If we compare this definition with Mataric's "Informed Coexistence" or 'Intelligent Coexistence', we see that they are practically identical (since Touzet does not mention anything about a field of vision, just '…the perception of other robot's locations and actions') and thus this awareness can be related to purely reactive behaviour based on the environment as well as other robots as some sort of triggers for different behaviours.

Touzet confirmed the mathematical statements using simulation and he stated that this does neither impact on the legitimacy for awareness in cooperative robotics, nor the generality of the method.

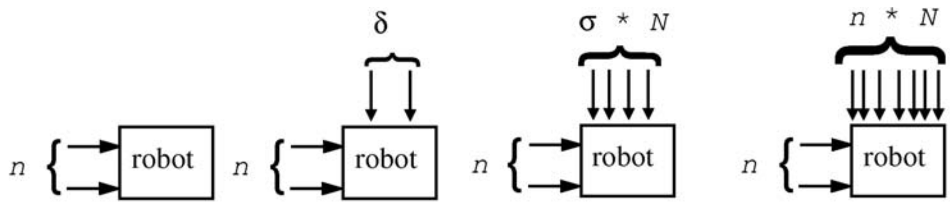

Figure 6 - Four awareness levels, with courtesy of Touzet.

Furthermore, Touzet presents the following four different levels of awareness. These are illustrated in Figure 6, which shows the difference in how much information each robot gets in each time step. The number of sensors used to perceive the world situation is represented by '*n*' and '*N*' is the number of robots in the environment. A fixed set of additional inputs ('δ') was used to represent the knowledge about all the other members of the group. Since the amount of additional information to consider can grow very large when considering all the other robots sensory information, a more simplified variant was represented by 'σ'. This is a set of additional inputs where the number of extra inputs is lower than using all sensory information from the other robots, this could for instance be the speed and direction of the other robots. The four levels of awareness are presented from left to right (as seen in Figure 6).

- *No awareness* - This is the simplest case and all interactions are made through the environment (cf. stigmergy). Although no robot is aware of the other robots in the environment, they still act *as if* they were aware of each other. Here the only input the robot gets is its own sensing of the environment.

- *Restricted awareness* - Here the robots have a fixed set of additional inputs to represent knowledge about the other members of the group (how to obtain such knowledge was not discussed in the article).
- *Awareness of all* - On this level all robots are taken into account, but the amount of input from other robots that is considered is reduced. In the example by Touzet, only mobile robots were used and thus orientation and distance to other robots were used as additional inputs, in addition to the ordinary sensory information used to perceive the environment.
- *Complete communication* - This is the most advanced and complex level of awareness and here the objective is to get as much information as possible from the other robots in the environment. This is obtained by sharing a number of inputs of the other robots and thus the robot get as additional input all the other robots sensing of the environment.

It may be argued that the awareness defined by Touzet (2000) is very different from the kind of awareness presented in the other examples. The kind of awareness stated in the definitions is somewhat biologically acceptable (that a robot should be able to become aware of a team member's actions without any use of explicit communication). This can be related to humans, they can certainly be aware of another person's actions without that person telling them anything about it. However, the awareness presented in the examples is rather different. There the implementation of awareness is achieved by just adding extra inputs to the robots and if we look among animals that are aware of each other and each other's actions then it is safe to say that they do not share any sensory information with each other in the same way that is presented in these examples. If we again compare with humans, then it is easy to understand that we do not get extra information about what another person *actually* sees; we can only internally simulate what we *believe* that person sees.

It has been shown in a couple of articles that the awareness of other robots in the environment very well could be beneficial for the performance (see for instance Parker, 1995 or Touzet, 2000). There are however a couple of different approaches of how to make the robots aware of each other, as we have seen in this section.

3.5 Social Robotics - Social Behaviour

In the cooperative robotics community there nowadays seem to be an increasing interest for ‘social activities’ and ‘social behaviour’ within robot societies. What do we then mean with ‘social’, the few definitions in the cooperative robotics literature are not very precise, where they do exist at all. One thing that seems to be commonly understood (although it may seem a little trivial) is that social behaviour requires more than one entity (i.e. in this case when considering social robotics, there must be at least two robots). If we, for instance, look in ‘The Concise Oxford Dictionary’ the definition of ‘social’ is as follows:

> “Living in companies or organized communities; not fitted for or not practising solitary life; interdependent, co-operative, practising division of labour; existing only as member of compound organism, (of insects) having shared nests etc., (of birds) building near each other in communities”

This seems to be a very broad definition but since there are few definitions in the cooperative robotics literature, it must be the case that the researchers in the community are accepting this commonly understood definition. Hemelrijk (1997) states the following about social behaviour:

> “… any study trying to explain the complexity of social behaviour should at the same time ask what part of it must be coded explicitly as capacities of the individuals and what part is determined by interactions between individuals.”

This illustrates the key issues with social behaviour, where is the social interaction taking place and what is it? The social behaviour could be inside the agents or between the agents, or maybe more possibly both inside and between.

Billard (1997) states:

> “Communication is a social skill and as such would be desirable for artificial autonomous agents that may be expected to interact with other agents or humans”

Here communication is a social skill and in the experiments letting a teacher learn another robot, rather than studying how communication can be used to establish a relationship between the two robots constructs a ”social relationship”. Thus, the term 'social' here includes teaching of another robot and apparently this constitutes a 'social relationship'.

Billard and Dautenhahn (2000) state the following regarding social interactions:

> "… communication is an interactive process between the two communicative agents and as such is a social interaction."

In this statement we can see that an interactive process between two agents is a social interaction (according to Billard and Dautenhahn, 2000) and in this particular case the social interaction includes communication.

However, the statement by Billard (1997) can be related to Dautenhahn (1995) where Dautenhahn state that the investigation of social intelligence may be a necessary prerequisite for those scenarios in which autonomous robots are integrated into human societies. Dautenhahn states that 'social skills' cannot be defined as the 'rational manipulation' of others but are strongly related to individual feelings, emotional involvement and empathy. Dautenhahn further states that social intelligence might be interesting even if the robots are used in applications where they are nearly all the time dealing with non-social tasks since this would benefit the whole learning process. Based on results from the study of natural societies and especially influenced by the social intelligence hypothesis (which derives from primatology research and states that primate intelligence originally evolved to solve social problems and was only later extended to problems outside the social domain; see Chance and Mead, 1953; Jolly, 1966; Humphrey, 1976 and Kummer and Goodall, 1985 for details), Dautenhahn formulates the following issues that should be addressed when designing social multi robot systems:

- Robots should, comparable to the normal development of a human child, 'grow up' in a social context. Dautenhahn means that research aiming at the development of intelligent artefacts should not only focus on solving problems with the dynamics of the inanimate environment, but it should also take into account the social dynamics.
- If artificial agents should resemble living beings, they should not have a reset-button. Dautenhahn argues that artificial agents will never show an elaborated 'mental life' if they do not have the chance to have an individual ontogeny like

all natural agents. Although it is well known that different species develop at different speeds in nature, little effort has been made to take this into account in robotics research.

- The embodiment criterion (see for instance Brooks, 1991) is hard to fulfil unless we do not include the ontogeny of body concepts. Dautenhahn argues that no robot will ever evolve a concept of a body or personal self, unless the robots themselves actively use their bodies and develop a kind of ‘body conception’.

Dautenhahn states that there are two mechanisms that are crucial for the development of individual interactions and social relationships: imitation and the collection of body ‘images’. For imitation, she stresses that it is crucial that the imitation is used as a ‘social skill’ (i.e. the imitator should use imitation to ‘get to know’ the robot it is imitating, instead of just implementing imitation behaviour in order to let the model learn a specific task). The collection of body ‘images’ is related to the ‘body conception’ discussed earlier and she stresses that it is not the mere presence of a physical body that is necessary.

There seem to be some agreement in the cooperative robotics literature, namely that social behaviour involves some kind of interaction that should lead to something more than purely goal-specific accomplishment. Then again, what this actually should accomplish is yet undefined.

4 Research areas and issues

This chapter will elaborate on several aspects of cooperative robotics research. A very large and still growing area of the cooperative robotics field will be described in more detail in Section 4.1; namely RoboCup (cf. Kitano *et al*. 1995). Various groups that can be used to divide systems exhibiting cooperative behaviour have been identified and will be presented along with clarifying examples in Section 4.2. Additionally, Section 4.3 presents some issues that are specific when considering learning in cooperative robot groups. In Section 4.4 research areas related to cooperative robotics will be presented and the chapter ends with a summary in Section 4.5.

4.1 RoboCup

In the cooperative robotics field there is one field that is really growing fast nowadays, and that is the RoboCup domain. The Robot World Cup Initiative (RoboCup) is an attempt to advance AI and robotics research by providing a standard problem where a wide range of technologies can be integrated and examined under similar conditions. For this purpose, the RoboCup committee has chosen to use robotic soccer, and they organize an annual world championship (Kitano, 1998).

The interesting thing with RoboCup is that it involves many different disciplines (including for instance engineering-, real-time-, cognitive-, biology- and computer-science and of course artificial intelligence research in general) and can thus be used to help researchers to see other fields and be inspired by other approaches.

Despite the fact that RoboCup may seem very practically oriented, the theoretical use of RoboCup is very significant (Asada and Kitano, 1999), as we will see in Section 4.1.7.

This section is divided as follows, first some aspects of RoboCup will be presented that will show how RoboCup could be beneficial for all those disciplines previously mentioned. Section 4.1.1 contains an introduction to RoboCup, where the ideas and motivations underlying RoboCup will be presented. After that, in Sections 4.1.2 - 4.1.6, the different leagues will be described more thoroughly and a discussion about the current status regarding the cooperation in these leagues will be done. The amount of cooperation among team mates in the various teams is varying quite much from league to league (and of course from team to team), as we will see. In Section 4.1.7 we will see that the RoboCup domain is not only about robots competing in soccer against each other but there are also several other areas that are of great concern in RoboCup.

4.1.1 Introduction

Robot World Cup (RoboCup) was proposed by Kitano *et al*. 1995 where soccer would be used to have robots compete against each other in teams. The proposal of RoboCup was intended as a complement to the annually competitions held by the American Association for Artificial Intelligence (AAAI), where a *single* robot should solve a specified problem (see for instance Nourbakhsh, 1993).

The soccer domain was chosen because of the diverse problem areas found there, for instance various technologies must be incorporated including: design principles of autonomous agents, multi-agent collaboration, strategy acquisition, real-time issues and sensor-fusion. Within RoboCup, teams of multiple moving robots can compete against each other and since it is in a standardized environment, algorithms and approaches can be compared and evaluated. Of course there are certain rules that imply restrictions but according to the RoboCup committee these rules should not impact on the research areas of the domain, but only to insure fair play.

RoboCup offers a wide problem domain, which makes it an interesting one for many different disciplines (including engineering-, real-time-, cognitive-, biology-, computer science and of course artificial intelligence research in general); also it forces the designers to build robots, which reliably perform the task and cope with uncertainty and noisy environments.

Currently, there are four different leagues in RoboCup (Asada and Kitano, 1999; Christensen, 1999) and those will be described in the following subsections. There is one simulation league and a couple of real-robot leagues and the leagues will be analysed in how they work and what the current status is in relation to cooperative robotics. Furthermore, a new league will be introduced into the domain till 2002. This league is called “Humanoid League” and is described in Section 4.1.6.

4.1.2 Simulation League

In the simulation league the robots are simulated on a computer and separate threads control the team mates in the game, i.e. there is not a single point of control so the robots are autonomously controlled, although the same computer controls all robots. Communication to other robots is allowed and achieved by sending messages through the server. What the players see (the field of view) is of course also simulated and is achieved by a message being sent every 300 millisecond to every player. This message contains information about how far away the robot is from different flags, the goal and the ball. Along with this information are degrees so the actual positioning of the robot in the field is left to own calculation, i.e. there is no information in the message about the robot's position in the playing field (see Figure 7 for a view of the playing field). As in real human soccer the team consists of eleven members, one goalkeeper and ten field players. Beside the ordinary simulator, the RoboCup federation is developing a more advanced one with "RoboCup Advanced Simulator" as the project name (Asada and Kitano, 1999). The intent with this simulator is to allow simulation of physical properties of robots as well as to give a detailed simulation of physical aspects of their environment like gravity and sensory inputs. The ultimate goal is to allow the development of software in the simulator and then downloading it onto a real physical robot, where it should work well with little modification.

Figure 7 - Simulation League.

Major Research Issues

Since this league is simulated, the main research issue here is purely oriented around the software of the agents, and not to any physical aspects. However, given that the software actually adds noise to the signals, some consideration has to be made to this in the incoming data from the "sensors", but other physical aspects like hardware failure and slippery floors are not a relevant problem in this league. Because of this, the researchers can focus more on actual teamwork in their teams (see for instance Stone *et al*. 2000), thus addressing issues like passing the ball to a team mate or real-time adaptation to the opponent's strategies.

Discussion

The preconditions for this league to perform well, with respect to cooperation among team mates, is very good since the agents always have information about where the opponents goal is and the location of team mates can be easily maintained by communication. Even if there is noise in the data distributed to the agents they still *know* that they actually get all the data, maybe noisy but they know that the object in focus will be there somewhere (in contrast to physical robots, where not all data is available at all times, as we will see later).

It might be argued that the agents do not cooperate at all. Since they are all run on the same computer, there has to be a single seat of control that manages all the agents (and then it cannot be difficult to make the agents "cooperate" since everyone is controlled from the same location). But in a sense one can argue that the agents are autonomous, since each agent is controlled by its own thread and perceives its own view of the environment and reacts upon this in conjunction with own internal states. Thus, one can argue that we have autonomous agents that cooperate in soccer playing and there is no single seat of control that manages all the agents.

There are several works that discuss the advantages and disadvantages when using a simulator to emulate physical robots (e.g. Steels, 1994). For instance, one motivation for using real, physical robots that is often discussed is the fact that when dealing with physical robots one must not calculate as much as in a simulator. In a simulator, in order to make a realistic world model, you must calculate a lot of things; for instance gravity, sensory responses, line of sight and much more. This makes reality an appealing domain since the calculations of these things is already taken care of by the laws of physics. Another point of view in favour of use of a simulator is of course that one does gain a lot of time, since running several tests in a simulator usually does not require too much work (compared with running the same amount with real, physical robots, which often becomes impossible due to time constraints).

4.1.3 Small Robot League

In this league the robots are physical and they play on a table tennis-sized field. The game is played with 5 robots in each team and the main restriction is the so called "18cm rule", which states that no robot may have a diameter larger than 18 cm, i.e. the robot must fit in a cylinder with 18 cm in diameter. The robots are allowed to use wireless communication between each other and/or a single host residing beside the playing field. The vision systems allowed include both a centrally placed, global camera and a distributed variant where each robot has its own camera. The type of vision system used must be announced in advance.

Major Research Issues

Naturally, this league includes physical aspects (for instance faulty sensors, slippery floor and motor malfunction) since it is a physical league. Beside this, the league is still quite different from the simulation league. First of all, in this league all robots (often) use the same camera over the playing field (whereas in the simulation league all robots got their own view of the field from their own current location in the field) and secondly, the rules say nothing about autonomous robots (only that no human intervention can occur) (whereas in the simulation league each robot should be controlled by a separate thread and that all communications between team mates should be made through the server). This means that a single program may control all robots in the Small Robot League, thus eliminating some of the problems found in the simulation league (e.g. the need to make independent agents to cooperate disappears and is replaced by the problem of strategic planning for a whole team at a time), while introducing some new problems when adding the physical aspects.

Because of the issues mentioned above the major research issues in this league are more concerned with strategic planning of consecutive actions (see for instance Veloso *et al*., 1998a) and of course more individual, mechanical skills like for instance being able to aim at targets and shooting the ball in a specified direction (see for instance Asada and Kitano, 1999 or Veloso *et al*., 1998a). For instance, Veloso *et al*. (1998a) presents a vision-processing algorithm that is able to track multiple moving objects and even predict their trajectories.

Discussion

In the simulation league it could be argued that the robots did not cooperate at all since it was all run on the same computer, this argumentation can instead apply here. In this league all robots have one single seat of control (i.e. the same program controls all the robots) and all robots get the same view from above the playing field. Even if one can use the distributed vision system according to the rules, there are very few teams that actually use this type of vision system (for instance, in RoboCup-97 three out of four teams used the global vision system and in RoboCup-98 it was ten out of eleven teams, see Asada and Kitano, 1999 and Bräunl and Graf, 1999).

This league is probably designed to be a kind of “in between” and thus are not all issues addressed here, which we think is good since research not only has to be broad, it also has to be focused and narrow.

When the global vision system is used the main research issues with this league is more concentrated on strategic planning within a single unit that has several actuators available (see for instance Veloso *et al*. 1998a). But when the distributed vision system is used the research focuses more on cooperation among several, independent agents where the individual robots really are autonomous and are capable of making their own decisions (see for instance Bräunl and Graf, 1999).

4.1.4 Middle Robot League

In the middle league the robots are also physical and the game is played on a field that is approximately 10 meters long and 7 meters wide. The use of a global vision system is not allowed in the middle league, which means that every robot in the team must be able to perceive its own surroundings. Every team can at most have four players, and one of these must be the goalkeeper. The size restrictions of the robots are usually met if the robot fits within a circle of 50 cm in diameter. The communication methods are not restricted in any way, except that no central computer may be used, and of course there must not be any human intervention.

Major Research Issues

In this league, the complexity has been increased to include all aspects of robotic soccer, for instance robots must be fully autonomous (i.e. no central steering, central computer or global vision may be used), physical aspects that where discussed earlier are all incorporated, and of course there are still the design of the software of every robot (which for instance include tracking and strategic planning). This league is one of the most interesting ones when considering cooperation among autonomous agents, since this is one of the few leagues where the agents actually are fully autonomous.

Discussion

The cooperative game playing in this league is not very good (in fact, it is often not there at all). From a distal view, it is like this league is filled with individual players that are said to play together but still play individually (i.e. treat other robots as obstacles) (see for instance Shen *et al*., 1998).

It is interesting to note that the cooperation among agents is so very different from the simulation league; actually it is only two things that differ. First of all, there is of course the physical aspects that are present in the middle league, but apart from that the only thing that is changed, is that in the simulation league all robots got information about where they are in relation to a couple of markers on the playing field, including information about the ball location and the opponents' goal.

It is either the physical aspects that have made such a strong impact on the cooperation achieved or the extra information provided in the simulation league is so crucial, or maybe a combination of them both (which we find most likely).

4.1.5 Sony Legged Robot League

As the name implies, this league is all about using Sony's AIBO™ dog-like robot in the soccer games (see for instance Veloso *et al*., 1998b). The teams are composed of three robots each and the robots are not to be changed physically in any way, this means that the only thing that is allowed to change is the robot's program. The three robots in the team include the goalkeeper, and the game field size is three meters long and two meters wide.

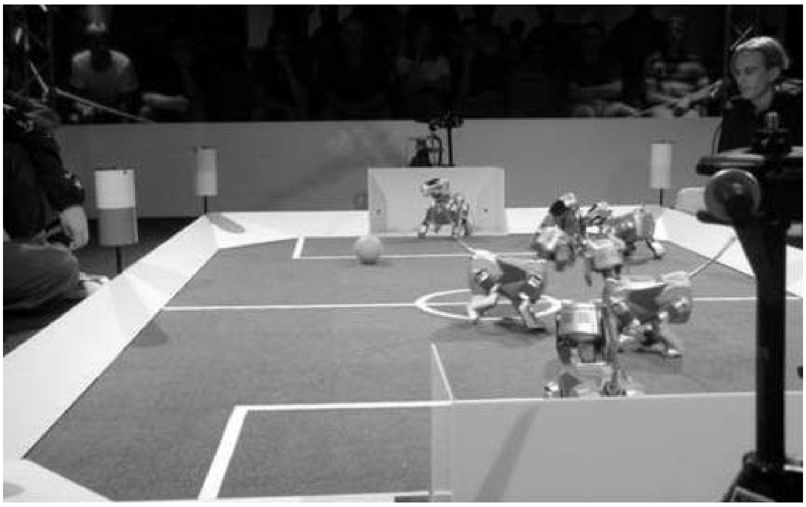

Figure 8 - Sony Legged Robot League.

Major Research Issues

In this league the major research areas are purely software oriented since the robots are pre-built and are not to be changed in any way. Of course, as a consequence of a physical implementation, there are still aspects like sensor malfunction and noisy data, but the researchers do not have to focus any attention on how to build the robots, since this is already done. As with the middle robot league, the robots here are fully autonomous and since it is not allowed to change the hardware the research here tends to focus on the pure software aspects of cooperation between autonomous agents.

Among the most common research areas in this league we can for instance see machine-learning algorithms for automated colour calibration and detection (see for instance Veloso and Uther, 1999).

Discussion

When looking at this league when the robots are playing, it may very well be argued that the robots do not cooperate at all (as with the middle league). In a sense those two leagues are quite similar. The only thing that really differs from the middle league is that in the Sony league the robots are pre-built and it is not allowed to change them physically in any way. However, one has to have in mind when discussing this league that it has not existed for so long (it started just a couple of years ago) so research in this field has just begun to emerge.

4.1.6 Humanoid League

The RoboCup Federation will start a humanoid league in RoboCup-2002. Prior to the official kick-off some extensive exhibitions are going to be held at RoboCup-2000 and RoboCup-2001. The league will be divided into three different categories and it is assumed that the participating humanoids are fully autonomous in the sense that after the kick-off the robots shall play soccer by them selves, without human intervention. Note that it is allowed to use off-board computers, but there must not be any remote control by humans. The categories that will be present are “Small Size Humanoid” with a maximum height of 60cm, “Middle Size Humanoid” with a maximum height of 120cm and “Large Size Humanoid” with a maximum height of 180cm. Initially there will only be three robots in each team but the number of players is supposed to increase as the technology matures.

In addition to the humanoid league in robotic soccer there will also be competitions is pure walking/running of the robots where the speed of the humanoids are important, this will be called RoboCup Humanoid Athletic.

4.1.7 Other domains in RoboCup

RoboCup is not all about competitions and robotic soccer playing but there are also many other interesting research areas. Apart from the practical competitions in RoboCup various conferences about the approaches and problems are arranged in order to make the field develop (Asada and Kitano, 1999).

Within RoboCup there is also the *RoboCup Challenge*, which stands as a common test bed for robots, this can be used as a benchmark to be able to see the advances made in both design of software and hardware (Asada and Kitano, 1999). Currently there are three different challenges in RoboCup:

- RoboCup Synthetic Agent Challenge (Kitano *et al*., 1997).
- RoboCup Physical Agent Challenge (Asada *et al*., 1998).
- RoboCup Natural Language Challenge (Asada and Kitano, 1999).

The Synthetic and Physical Agent Challenge differs from the ordinary RoboCup competitions in that robots participating in the challenges do not have to be able to play a complete game of soccer. The challenges are set up in a way as to only encourage development in a specific area, these challenges can thus be seen as sub-goals to the complete, autonomous robot soccer-playing problem. While the Synthetic challenges are focused on simulation the Physical challenges are realized with real, physical robots (and thus dealing with real-world noise, sensor perceptions etc.). In addition to these, there is another challenge, RoboCup Natural Language Challenge. This challenge addresses the non-robotics community since the task here is to have a computer commentary of the game in progress, i.e. the RoboCup Automatic Commentator. This commentary should be made in natural language in real-time during the soccer game.

4.1.8 Summary

We have presented the different leagues in RoboCup and discussed what their main research areas include.

For the cooperative robotics community, we think that it is very good that RoboCup has emerged, since this is apparently helping the research forward (see for example Kitano *et al*., 1999; Asada and Kitano, 1999). We think that it is mainly due to two things that RoboCup has grown so much in the last few years: the widely dispersed domain, and not to forget, RoboCup offers a funny and entertaining way of testing hypotheses and theories. Many people are familiar with the human version of soccer and can without difficulty identify with and understand the problem. Much because of this, the RoboCup domain becomes a common ground for researchers to exchange ideas and to help each other (since all these problem domains intersect with each other, researchers from one area may already have done something that others are trying to do, but information flows slowly across borders and that is why it is good to have a common research platform where everybody have the same goal).

We believe that simulations of multi-agent societies will have more influence in the near future since more and more processing power is being available. In addition to that, the knowledge is also growing and now it is easier to build a more realistic simulator with friction, gravity, wind and different kinds of noise. There are a lot of advantages with using a simulator instead of making the experiments in reality, e.g. the ability to simulate a very large society containing several hundreds of individuals are becoming rather manageable nowadays. This type of experiment would probably not be feasible to do in reality and since the area of focus in the field of cooperative robotics is going towards larger and larger environments with more and more robots that interact with each other, the need for simulation is clear. See for instance the RoboCup Advanced Simulator (Asada and Kitano, 1999), the ultimate goal there is to be able to develop a robotic team in simulation and when the ground training is done it should be easy to transfer the control software onto real robots, without too much modifications.

When we compare the robots that are playing soccer today with human players we realise that we have quite a long way to go. In the physical leagues one barely see any robot running across the field to put itself in a better position, where other robots can pass it so that it can score. This kind of long-term tactical plans we only, but hardly, see in the simulation league and the small league. In the small league much more elaborated plans are being built and passing of the ball to another player that is in a better position is not uncommon (see for instance Veloso *et al.*, 1998a). There are even algorithms for robot control that are used in the small league that also take into account the possible intents of other players, in order to be able to respond faster (i.e. in advance) (see for instance Stone *et al.*, 2000 or Veloso *et al.*, 1998a).

In the other physical leagues this kind of advanced cooperation is still in the future, hopefully.

4.2 Homogeneous and Heterogeneous Robots

There are two major groups that the robotics field can be divided into, namely homogeneous robots and heterogeneous robots. The two groups reflect in a way what kind of experiments that have been conducted. This chapter will give a more thorough description of what the two groups include and what research that has been done in the areas.

In section 4.2.1 homogeneous robots will be described, what is distinct with those and what kind of experiments that has been done in that field. Section 4.2.2 will discuss heterogeneous robots and some of the experiments they have been used in.

It is noticeable that there are researchers who claim that the either/or division into homogeneous and heterogeneous robot teams is too distinct. Balch (1997) proposes a metric to measure the diversity in multi-robot systems he calls *social entropy* – which also recognizes physically identical robots that differ only in their behavioural repertoire, this will be discussed in more detail in Section 4.2.3.

4.2.1 Homogeneous robots

Since the commonly understood definition of homogeneous (see for instance Balch, 1997 or Jung and Zelinsky, 2000) includes that the software must also be identical, then this area becomes rather small since there often are differences in the learned control system.

Melhuish *et al*. (1999) are conducting an ongoing project that has been going on since 1996, and this is investigating how ”lots of dumb robots can collectively do something smart". In those experiments there are homogeneous robots and the only means of communication that is allowed is by using stigmergy (i.e. to communicate using only the environment as the communications medium). Figure 9 shows a picture of a clustering task, where the robots should gather the coloured frisbees into clusters (in this experiment there are yellow and red/black/white frisbees).

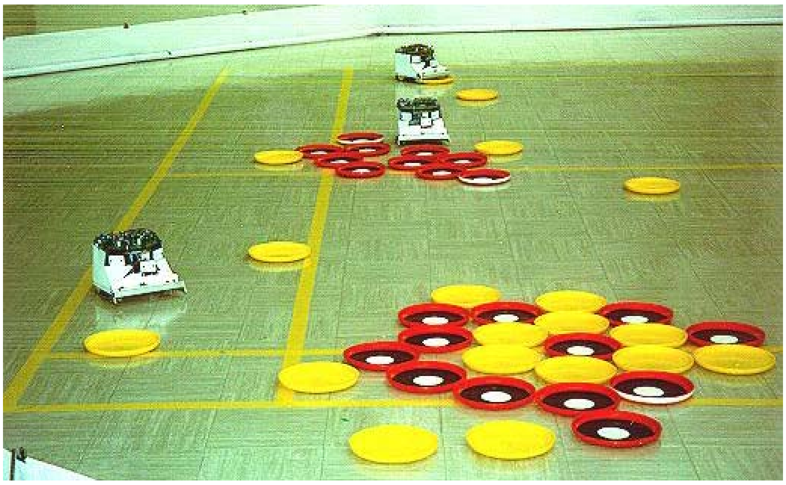

Figure 9 - Robots used by Melhuish *et al*. (1999). With courtesy of Melhuish.

If we look in the RoboCup domain there are a lot of soccer teams that are composed by homogeneous robots. In most cases there is some difference between the field players and the goalkeeper, but if we only look at the field players then there is a lot of homogeneous teams in this domain, and in particular in the physical leagues. It may be argued that by only using homogeneous robots one cannot build an effective team, but if we take a look at today's RoboCup teams we see that such simple control structures can be quite effective. For instance there are the 'Spirit of Bolivia', which is a robotic soccer team that competed in RoboCup-97 (Werger, 1999). The robots in 'Spirit of Bolivia' are entirely controlled by a 'Behaviour Based System' (cf. Brooks, 1986) and there is no learning of the control system in the robots. Because of this, all the robots are identical, both regarding physical construction and regarding control software and thus the team is homogeneous (when only considering the field players, the goalkeeper in 'Spirit of Bolivia' had another control software as well as different physical architecture).

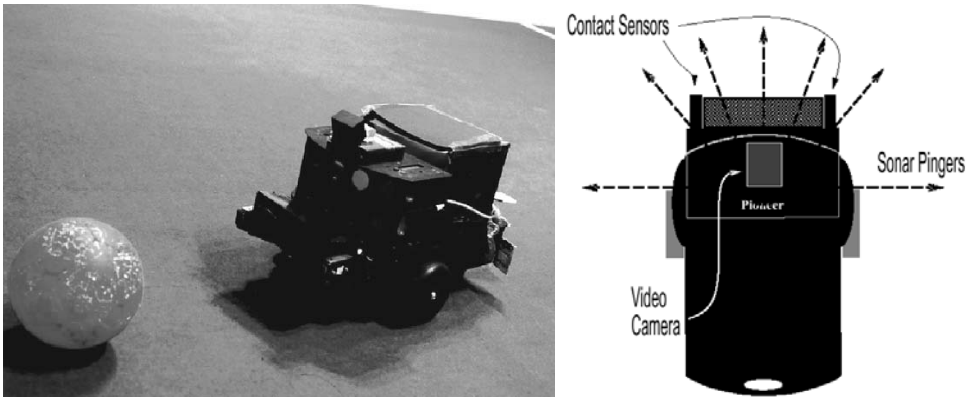

Figure 10 - Field player in 'Spirit of Bolivia' (Werger, 1999). With courtesy of Werger.

4.2.2 Heterogeneous robots

A robot system is classified as heterogeneous if one or more agents are different from the others (cf. Balch 1997). When the software is also included in this definition then one can easily understand that heterogeneous is the area where most work has been done. Since by the time a robot learns something about the environment it is situated in, it has probably changed its behaviour and representation about its surroundings and thus the robot is different from the others (regarding software).

Jung and Zelinsky (2000) use two autonomous mobile robots that perform cooperative cleaning to investigate the impact of using symbolic communication when cooperating. Since they were interested in heterogeneous cooperation, they equipped each robot with a different set of sensors and actuators (see Figure 11 for pictures of the robots). One robot has a vacuum cleaner mounted between the drive wheels, and as a result of that it cannot vacuum close to walls or furniture. The other robot has a brush tool that is dragged over the floor to sweep distributed litter into larger piles that the vacuum cleaning robot can pick-up, since the vacuum cleaning robot cannot walk near walls or furniture the piles have to be made in open floor space.

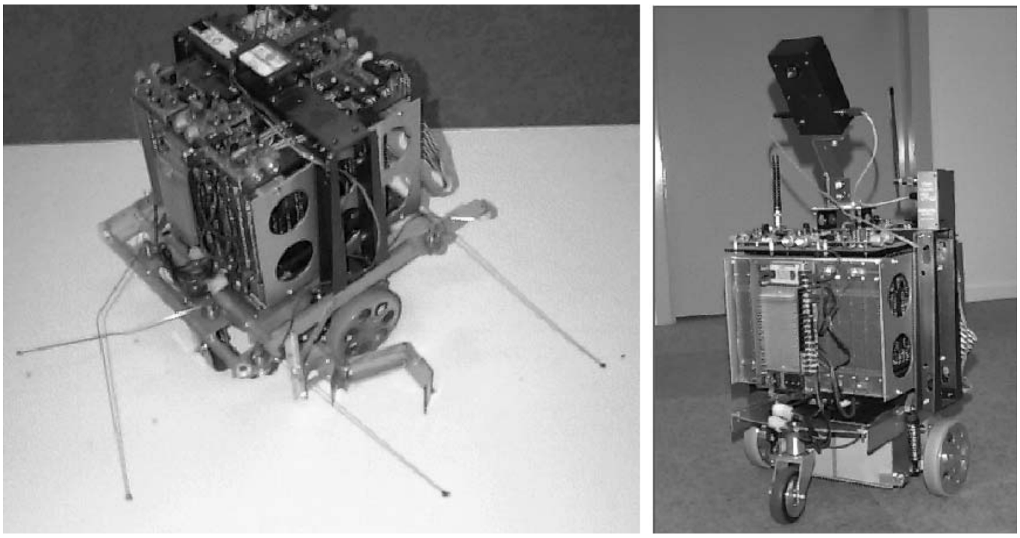

Figure 11 - Cleaning robots used by Jung and Zelinsky (2000), on the left we see the brush robot and on the right the vacuum robot. With courtesy of Jung.

Another work that has been done with cooperative heterogeneous robots is by Sukhatme *et al*. (1999). Here they use three heterogeneous robots for surveillance and patrol of a predefined area. There were two autonomous mobile robots and one helicopter (see Figure 12), which were flown by a human pilot during these experiments for purposes of safety and speed. This is also a very clear example of a heterogeneous robot group since the differences is not only in the control software. Picture courtesy of the USC Robotics Research Laboratory Autonomous Helicopter Project (http://www-robotics.usc.edu/~avatar).

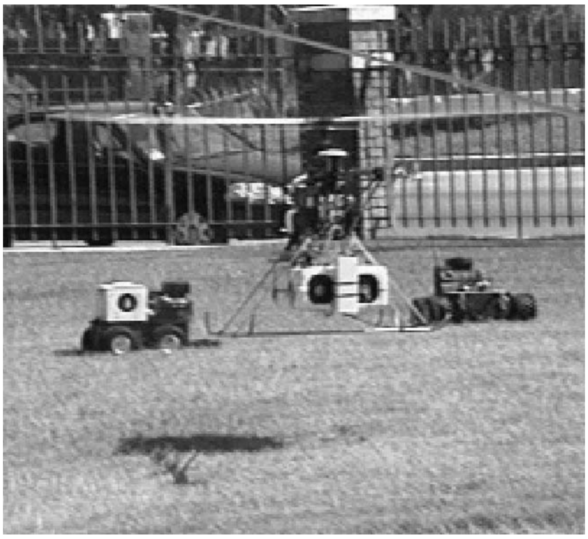

Figure 12 - Two mobile robots and the helicopter used by Sukhatme *et al*. (1999).

4.2.3 Social Entropy

In the previous subsections we presented the most commonly used way of categorizing various robot groups, i.e. into either homogeneous or heterogeneous. Here, we will now describe a metric proposed by Balch (1997) that should replace these groups. This metric should be able to more accurately divide different teams of robots (and thus cope with different scales of diversity so that one team of robots can be said to have more diversity than others) instead of the more commonly used homogeneous/heterogeneous division, which is binary. The metric Balch proposes, *Social Entropy*, is inspired by Shannon's Information Entropy (see Shannon, 1949).

The main idea with Social Entropy is to give a value of how much diversity there are in a society. Balch states the following goals that should be fulfilled in order to have a metric, which captures the behavioural diversity in a society:

- The least diverse society is one in which all agents are equivalent.
- The greatest diversity in a society is achieved when no agent is equivalent to any other agent.
- A society in which one agent is different and all the rest are equivalent is slightly more diverse than the society in which all agents are the same.

- If two societies have uniformly-sized groups, the one with more groups is more diverse.

To measure the diversity in a society, Balch partitions the robots into several groups, where *all* members of a given group have *similar* behaviours. When the partitioning is done the Social Entropy is calculated with the following formulae:

$$Het(R) = -\sum_{i=1}^{C} p_i \log_2(p_i) \quad \text{where} \quad p_i = \frac{|C_i|}{\sum_{j=1}^{C} |C_j|}$$

Het(R) is the function that calculates the heterogeneity of the robot society *R*, and p_i calculates the proportion of robots in the corresponding subgroup C_i to *R*.

With this metric, it is possible to calculate a value of *how* diverse a robot society is (on the contrary from the binary division we saw in homogeneous / heterogeneous), and it is also possible to compare this value with other society's diversity and thus come up with the conclusion that society *A* is more diverse than society *B*. This metric is based on the number and size of the groups that make the society up. Balch (1997) however states that even though the metric may seem to be focused on the level where only the external observable behaviours are taken into account, it could very well be applied to other levels. For instance by measuring the behavioural difference between the robots by comparing the real-valued sensor inputs and the actuator outputs, then we would have a metric of how diverse a robot society is with respect to not only the externally observable behaviours but also with respect to the internal differences.

4.2.4 Discussion

We have presented a couple of example experiments that have been conducted, and which can be divided into the groups homogeneous and heterogeneous robot systems. We do however believe that the distinction is too strong and that the use of for instance Balch's (1997) Social Entropy is preferable when discussing different kinds of robots. As we can see from the examples, the differences can vary quite a lot and the either/or distinction in homogeneous and heterogeneous robots becomes unfeasible when discussing the differences between the robots.

It may be argued that using homogeneous multi-robot systems is preferable since they will more easily cope for robot malfunction since every robot can do everything and it will be much easier to just introduce more robots into the environment if the demand is increased (since if they are homogeneous they will not notice if one more comes along). For instance, Werger (1999) prefer simplicity and uses a homogeneous robotic soccer team for just those reasons.

On the other hand there are some researchers that state that homogeneous multi-robot systems will not play a large role in the cooperative robotics domain in the long run (see for instance Jung and Zelinsky, 2000). The main argumentation for this is that their definition of heterogeneous multi-robot systems includes difference not only in physical aspects but also in the software, and that the robots shall learn about their environment in order to be effective.

One should notice here however, that the choice of whether or not to use homogeneous robots of course depends on the task that the robots should do and what is to become from the experiments.

4.3 Learning in cooperative robot groups

There are some researchers who believe that when dealing with group robotics it is very favourable to use behaviour based control systems (see for instance Sukhatme *et al*. 1999 or Billard and Dautenhahn, 2000), since practice in the last several years has favoured behaviour-based systems for group robotics.

As stated earlier in this dissertation, there are roughly two major areas of interest in cooperative robotics, ‘Life Sciences’ and ‘Applied AI’, and this is important to have in mind when discussing issues like learning in cooperative robot groups. In ‘Life Sciences’ the use of evolutionary algorithms and learning is very common but when looking at works more oriented at ‘Applied AI’ the learning takes another approach. ‘Applied AI’ mostly deals with designing and construction of the underlying behaviours that may lead to cooperative behaviour. Since the goal here is to construct a functional system and not care so much about biological or cognitive plausibility. Correspondingly, when considering ‘Life Sciences’ the learning approaches taken there are often inspired by nature and the ways various animals acquire knowledge.

Robot learning in general is known to be a very difficult problem (see for instance Mahadevan, 1996), and since multi-robot learning imposes even more areas of complications, this has been found to be a difficult problem. Mahadevan compare four different techniques for robot learning: Inductive Concept Learning, Explanation-based Learning, Reinforcement Learning and Evolutionary Learning.

Since multi-robot learning is not that different from single robot learning we will not discuss any details regarding training and learning in multi-robot groups. However, there are some aspects that must be considered when addressing multi-robot groups: the credit assignment problem, the trade-off between different goals and the modelling of other agents.

Credit Assignment Problem – This problem arises when the use of delayed reward/punishment is used for training robot teams, and since this is the most commonly used technique the problem is very often encountered. If we compare with for instance social insects like ants and birds, they are born with the cooperative behaviour whereas robot teams has to learn to cooperate and it is here the problem arises. This results in for instance the following scenario; when a team of robots has completed a task and should receive credits for that, the knowledge of which robot to assign credits to is not known, since it could have been Robot A or it could have been Robot B. To further complicate things it could also be the interaction *between* the robots that made the final call and so, it can be very difficult to know which robot to give credits for a particular task. This is known as the *Credit Assignment Problem* and there has been quite much research in this area (see for instance Minsky, 1963; Weiss, 1996; Versino and Gambardella, 1997; Salustowicz *et al.*, 1998).

Mataric (1996) uses communication to cope with the credit assignment problem in one of her experiments. The robots used by Mataric communicate their internal states to each other in order to synchronise their behaviour. In doing so the robots know what kind of actions the other robot has executed and hence they are able to share the reward or punishment with each other.

Another way of dealing with the credit assignment problem, which by the way is often used in evolutionary algorithms, is by using Implicit Credit Assignment (see for instance Moriarty *et al.*, 1999). In this technique the evaluation and credit assignment is based on the total performance of the system or subtask and thus there will be implicit credit assignments on those subtasks that often occur in succeeding solutions.

There are many articles that address the issue of the credit assignment problem but there are no articles yet that has claimed to solve the problem, and since the problem is task-specific that is not surprising. The credit assignment problem is something that every designer must bear in mind when doing experiments that could be affected by issues like those presented here.

Trade-off between individual and group goals – When considering teams of robots and when there are some global goal that the group should accomplish, the individual goals of the individual robots are often not mentioned. If we for instance relate robotic soccer with human soccer we easily see the issue. In human soccer there of course are the global goal of winning the match but the individual players also have their own goals, these could for instance be the will of being "Player of the match" or just some emotional feeling such as anger or fear. These individual drives will of course impact both positive and negative on how team play is performed and thus to what extent the team will cooperate and how they will cooperate. In robotic soccer there are no individual goals in the robots, the only goal that is present is the global goal spanning the whole team. As a result of this there will of course be differences in the soccer playing performed by robots and the kind of play being performed by humans. So when addressing this issue there will of course be some trade-off between the individual goals and the group goals since the group goal may not be in sync with every individual's goal.

Now one may argue that since we do not have any individual goals in the robots there is no need to cope with individual goals, since they are not there. The only thing we need to take into account is the global goal of the whole system. But when there are several robots their goals may very well be conflicting and the designer of such a multi-robot system must bear this in mind since it may have a great impact on how the system performs.

We have not found any article that addresses this issue so the problem remains open, what to do when there are conflicting goals among the robots.

Modelling and prediction of other agents' behaviour – There has been some research in this area (see for instance Veloso *et al*., 1998a), and now with RoboCup there is a standardised way of testing the algorithms and to see what the strengths and weaknesses are. In order to be able to extend the notion of cooperation among several independent individuals it is useful to be able to predict what other agents in the environment will do in the near and far future. This will help every individual agent when trying to build a more complex plan of what to do next. Further more, the communications requirements can also be lowered if the ability to model other agents is present (see for instance Touzet, 2000).

Cao *et al*. (1997) defines modelling of other agents as follows:

> "Modelling requires that the modeller has some representation of another agent, and that this representation can be used to make inferences about the actions of the other agent."

Cao *et al*. mean that when modelling of other agents should be accomplished, there has to be something more than purely implicit communication via the environment. The perception of another agent and the internal representation of that agent are very important.

4.4 Research areas related to cooperative robotics

According to Cao *et al*. (1997), the most critical disciplines for the growth of cooperative robotics are distributed artificial intelligence, biology and distributed systems.

Distributed Artificial Intelligence – The main concern in DAI is the study of distributed systems of intelligent agents. When seen from this view, the relevance with cooperative robotics is easy to understand. Bond and Gasser (1988) define DAI as "the subfield of artificial intelligence (AI) concerned with concurrency in AI computations, at many levels."

Since DAI comes from both traditional symbolic AI and the social sciences, there are two major sub-disciplines: Distributed Problem Solving (DPS) and Multi Agent Systems (MAS). DPS is mostly concerned with solving a single problem using many agents, whereas MAS research is not bound to these specifications. In MAS there could for instance be several goals among the agents and these can be both common and conflicting.

Since DAI mostly deals with distribution among independent entities and the research have concentrated on domains where uncertainty is not as much of an issue as it is in the physical world, the influence on cooperative robotics research has been limited.

Biology – The biological influences in cooperative robotics is huge. The majority of existing work has cited biological systems as inspiration or justification for various choices. In nature, there are many examples of insects that perform complicated tasks with very simple agents (for instance ants, wasps and bees). From these simple agents the complex behaviour *emerge* out of the interactions between the agents, which individually only are following simple rules. This can be found in the cooperative robotics community as a 'bottom-up' approach (see for instance Brooks, 1991; Agah and Bekey, 1997; Steels and Vogt, 1997; Jung and Zelinsky, 2000) in which individual agents are more like ants – they follow simple rules and cooperative behaviour emerge from the interaction among the robots.

Distributed Systems – In fact, a multiple robot system is a special case of a distributed computing system. Thus, the field of distributed systems is a natural source of ideas and solutions. There are many aspects that are applicable to both domains, for instance distributed computing, message passing and resource allocation to name a few. Much research that has been made in distributed systems can thus be applicable to cooperative robotics.

4.5 Summary

In this chapter we have presented a brief taxonomy of the cooperative robotics research and we see that there are many disciplines that may take advantage from cooperative robotics research. The fact that so many theories and disciplines have been brought together in cooperative robotics shows the fascination and the attraction of the field. We have presented a domain within cooperative robotics that brings many different disciplines together, namely RoboCup. Within RoboCup there are several leagues, which were presented and the major research areas for these were discussed. We also elaborated on the issue of how to divide different societies of robots into groups, which can be done in a couple of ways. Either one can make an either/or distinction using homogeneous/heterogeneous robot groups, or a more fine-grained scale can be used like for instance Balch's (1997) Social Entropy. We also discussed some issues particularly regarding learning in cooperative robot groups and presented some examples that illustrated why the various problems arise. Finally some related research areas to cooperative robotics were presented along with a discussion about what it is that brings the different fields together.

5 Summary and Conclusions

This chapter presents the achievements and the contributions of this dissertation. In Section 5.1 we will first give a summary of what research has been done in the field of cooperative robotics and within that summary we give an outline of the various levels of cooperation found in works nowadays. Section 5.2 states what is left to do with respect to the former outline, and in the end some final remarks are made in Section 5.3.

5.1 What has been done

There are many redundant articles that address the issue regarding cooperative mobile robotics in the last few years. One major way to address problems in this area is by building things bottom-up and that is precisely what has happened in the research that has been conducted. A lot of work has been done on low-level cooperation and simple interacting tasks where the robots mostly see each other as obstacles (cf. discussion of stigmergy in Sections 2.1.1 and 3.3). Also, quite a lot of work has been done in a little more advanced cooperation areas like flocking and herding where the robots do sense each other to some degree but the behaviour is still purely reactive and still no complex plans are being constructed by the robots.

The concept of sensing other robots as other things than obstacles is a very hot topic nowadays and there are some articles that have already addressed this issue (cf. for instance the discussion of awareness in Section 3.4). These experiments however are still very simple in nature and often only comparisons are made with systems where no awareness of other team mates is present and where the so-called awareness is very positive with respect to performance. Despite the big interest around the awareness among robots in certain articles, the number of articles addressing this issue is still relatively small but we expect a further increase in the upcoming years.

Communication among the robots is another thing that is being used more and more nowadays, and the number of articles addressing communication among autonomous robots is growing every day. We believe that this much depends on the fact that awareness is beginning to grow in the community, and one way of dealing with awareness is to let the robots communicate with each other.

Although no expanded team planning at the individual level is being performed in any experiments or articles, there has been however a lot of work regarding planning of group behaviours when there are access to some sort of central control or common vision system. This can be related to for instance the simulation league and of course the small league in RoboCup. Other areas that has been doing this kind of research is for instance formation following (Brock *et al*. 1992).

In the cooperative robotics literature we have roughly found a couple of levels of cooperation that the various experiments can be divided into. In this section we will present an outline of those levels as well as elaborate on the kinds of works that belong to them by giving some examples. It may be argued that the outline presented here is not complete or that it is not entirely correct. The outline is in no way entirely ready since research in the community has not reached the top-level yet, and hence it cannot be complete. This section will only illustrate the various levels of cooperation seen in works nowadays. Roughly, the following six levels were found in the literature:

- No cooperation
- Stigmergy
- Distinguishing each other
- Explicit communication
- N/A
- Planned behaviour

No cooperation – It may seem peculiar to include this level but firstly, every scale has to have a beginning. Secondly, if we look back in Section 3.1, we discussed the difference between pure collective- and cooperative behaviour and this level includes those experiments where no cooperation (i.e. there are only collective behaviour) is being performed by the participating agents. Some may for instance argue that since there are robots in RoboCup that does not cooperate at all, those should belong to this level. However, when dealing with robotics there is (almost) always some kind of sensing of the environment that the robots are operating in, and if there are sensing of the environment then the robots can be affected through the environment. Thus, according to the definitions of stigmergy found in Section 3.3 those robots are indeed using stigmergy. If we for instance look at the RoboCup leagues, there is actually one common link between all robots, and that is the ball. If we compare this with the box-pushing example discussed in Section 2.1.1 there are actually not so much that differ. In the box-pushing experiment the robots pushed the box and the other robots sensed that the box moved in a certain direction and reacted upon that. The same applies to robotic soccer, where a robot kicks the ball ahead and the others sense that the ball is moving / is on a new position.

Stigmergy – The distinction between the lowest level and this one is very clear. Here the robots do have some notion of other things in their environment and they can react upon that. This can be used to accomplish a kind of communication through the environment and in this level the agents can *seem to be* very intelligent and cooperative. However, when looking at the underlying mechanisms in their behavioural repertoire it become obvious that the robots cannot distinguish other team mates from obstacles in the environment or what it is to cooperate.

As we saw in Section 4.2.1, the soccer team "Spirit of Bolivia" is a purely reactive system and does not have any notion of the other team mates. Yet it *looks like* the robots are aware of each other and that they are cooperating, but this is only from a distal view (cf. Section 4.1). When looking at the underlying mechanisms (i.e. the proximal view) of the behaviours one can easily see that there are no notion of other team mates as team mates, they are only treated as obstacles and that the cooperation that emerge is only due to cleverly designed underlying behaviours.

Now, it may be argued that many of the robotic soccer teams do not use *any* cooperation at all, and this particularly applies to the middle-sized league where, as we recall from Section 4.1.4, practically there is no cooperation. However, the robots are indeed using stigmergy even when just shooting the ball ahead – the robots are affecting each other's environment since the robots do sense the ball.

Another example that was discussed earlier in Section 4.2.1, clearly illustrates what kind of robots are included in this level, is the ongoing experiment by Melhuish *et al*. (1999). The robots in Melhuish *et al*.'s experiment are purely reactive upon their environment and for those experiments performed there that is all that is needed. In Figure 13 we see a close-up of one of the robots used by Melhuish *et al*. and the different kinds of frisbees used. The frisbees differ only in colour but since this is enough to distinguish them into groups, the robots can now for instance cluster the frisbees into two different piles.

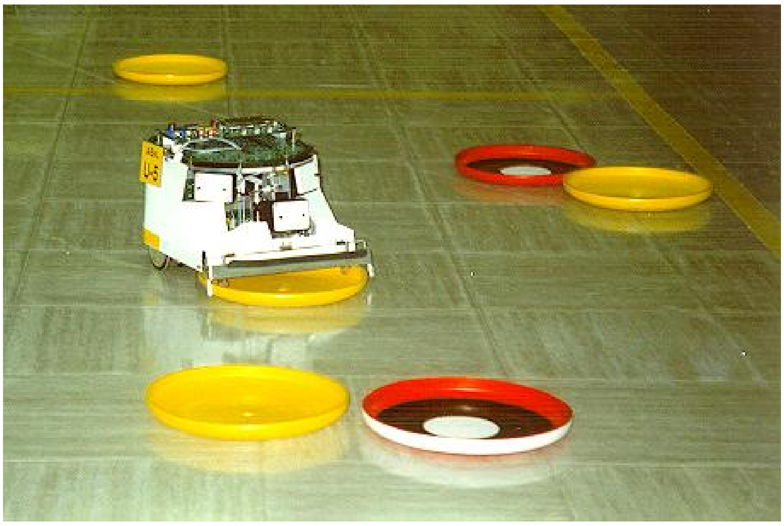

Figure 13 - Close-up of a single robot used by Melhuish *et al*. (1999). With courtesy of Melhuish.

Distinguishing each other – This constitutes the third level that we have found and the distinction from the second level is rather obvious. Here the robots not only sense their environment but they are also able sense other robots as robots and not only as obstacles, and hence react on other robots as well as the environment. The sensing of other robots as robots have been achieved in many and varied ways in the cooperative robotics literature (see for instance Parker, 1995; Brogan and Hodgins, 1997; Veloso *et al*. 1998a; Touzet, 2000) and one way to achieve awareness of other team mates is to mark the robots in a special way and thus having the robots distinguish themselves with the aid of this special marker. Another way is to have the awareness over some kind of communication channel, like for instance letting the robots communicate with each other via radio and thus being able to react upon the other robots behaviours.

Apparently, this level includes flocking, herding and swarming behaviours that can be found in many works in the cooperative robotics literature (see for instance Parker, 1995 or Touzet, 2000).

For instance, Brogan and Hodgins (1997) use simulated robots to investigate how a flock of robots moves around an obstacle that is put in their way. In this experiment the robots have the ability to sense the other robots and to distinguish them from the obstacles, and this leads to reactive behaviour not only to the environment but also to reactive behaviour on the other robots in the vicinity.

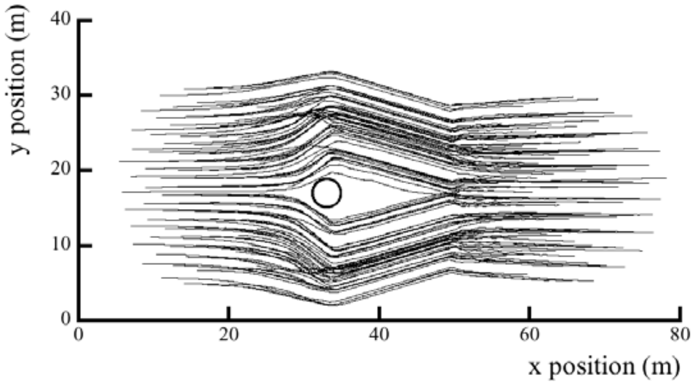

Figure 14 - Simulated robots used by Brogan and Hodgins (1997). With courtesy of Brogan.

Explicit communication – The distinction in this level from the lower ones is the use of explicit communication (cf. 'Explicit communication' and 'Direct communication' in Section 3.3). In the lower levels the use of communication is frequently used (i.e. implicit communication) but here there are something more to it. The emitting of a message is somewhat 'intentional', unlike in the lower levels where the use of communication was a purely reactive behaviour and there where no 'intention' in sending the message.

Steels and Vogt (1997) conducted an experiment that clearly illustrates this kind of cooperation (cf. Section 2.1.3). Here, the robots wondered around in the environment and occasionally ran into each other – then the communication started. The robots *agreed* on an object in their surroundings to talk about (i.e. the two robots explicitly told the other one what they wanted to talk about, and from this an agreement was established on an object). Here, the distinction from the lower levels is pretty obvious; the robots use explicit communication to communicate only with the robot in front of them. However, it may be argued that the robots are still only using a purely reactive behaviour and that the so-called explicit communication is only because of the reactive behaviour from sensing another robot in the vicinity. This is of course true, since the robots are reacting on their environment, but what makes this so different is because of the *intentional* sending and the there after following *discussion* with the other robot (cf. Sections 3.3 and 3.4).

In the cooperative robotics literature there are a very obvious gap between the former and the top level, which is rather easy to stipulate:

Planned behaviour – This level can be compared to the kind of planned behaviour seen among humans and it may be argued that this is the highest level of cooperation seen yet. Here the *intentional planning of the future* is one of the things that distinct this level from the lower ones. To some extent this can be seen in RoboCup when the aid of a central team control or a global camera are available. Especially one can see this kind of future planned behaviour in the small- and simulation league where the robots may be able to predict the future and take that into account when planning what to do next. For instance, Veloso *et al*. (1998a) not only use the environment perceived locally and by communication with other team mates, they also use the *predicted* behaviours observed in the opponents in order to faster cope with rapid changes and to build more complex plans. For instance, Veloso *et al*. used an Extended Kalman Filter (for details, see Veloso *et al*. 1998a) that was used to predict the forthcoming position of the ball. Another technique used in this team is the evaluation of weather to shoot at the opponent goal or weather it may be better to pass the ball to another team mate. Here the robots take into account the distance the opponents have to the calculated trajectory of the ball if it will be shoot in the specified direction.

The outline presented in this section, which shows the various levels of cooperation found in the cooperative robotics literature nowadays, not only illustrates what kind of work that has been done in the field, but it also points out which areas that need more attention.

5.2 What needs to be done

As we can see in the outline presented in Section 5.1, the simpler forms of cooperation have been covered rather well in the literature and thus it seems rather natural to move on forward onto the next step, which includes awareness of other robots in the environment. This issue is beginning to materialise it self in the community but there are still much to do. If we for instance compare current world champions in robotic soccer it is sad to say that the robots there do not use any kind of cooperation (in practice) at all and that a very common approach is to let the robot treat each other as obstacles. Sure, this can lead to quite complex behaviour with respect to what the abilities of the robots are, but we do not se any complex planning like we for instance see in human soccer playing. If we of course ignore those systems with a global vision or some form of global control (where the main issue is not to individually make up a team, but instead construct a team controller that is able to guide a group of robots).

In the outline presented in Section 5.1 we see that there are much left to do in the transition between the 'Explicit communication' and the 'Planned behaviour' subdivisions. There are some works that can be put pretty high in this scale, but since they use some restrictions like a global control or some sort of global sensing, they do not reach to the top (see for instance the small league in RoboCup, there are some high level planners available but since all is done with the aid of a central camera and a central control system, this cannot be seen as the kind of high level planning seen among groups of humans).

From the outline presented in Section 5.1 it is easy to see where we need to make more experiments and elaborate more on various techniques in order to achieve higher level planning and cooperation.

5.3 Final remarks

We have presented a survey of the cooperative robotics literature and we have investigated some aspects of the research that is going on. Many different disciplines have been identified as relevant to the cooperative robotics research and examples have been presented. Roughly two major areas of interest for cooperative robotics have been discovered, 'Life Sciences' and 'Applied AI', whose research interests differ quite a lot. Some key concepts were identified and analysed, where the various definitions were presented and any possible anomalies and of course similarities where emphasized and discussed. Thereafter some distinct groups in cooperative robotics were presented, among which one of them were more emphasized, namely RoboCup. Some open issues in multi-robot learning where also discussed and some works that have tried to tackle these problems were presented.

Finally we presented an outline that clearly illustrated what parts of the cooperative robotics field that has a lot remaining work to be done. This outline also showed that there has been a lot of work done in the simpler forms of cooperation, but that much work remains to be done in more complex forms before we can see for instance human-like cooperation being performed by robots. Within this outline we classified existing systems and works in the cooperative robotics field and the parts that demands more research was identified in the outline.

In the area of cooperative robotics, this dissertation has clearly shown which branch of the research field that still has a lot of work to be done.

References

Agah, A., Bekey, G. A. (1997). Phylogenetic and Ontogenetic Learning in a Colony of Interacting Robots. *Autonomous Robots*, volume 4:1, pp 85-100.

Arkin, R. C., and Hobbs, J. D. (1992). Dimensions of Communication and Social Organization in Multi-Agent Robotic Systems. *From animals to animats 2*: Proc. 2nd International Conference on Simulation of Adaptive Behavior, Honolulu, HI, Dec. 1992, MIT Press, pp 486-493

Asada, M., and Kitano, H. (1999). The RoboCup Challenge. *Robotics and Autonomous Systems*, volume 29, pp 3-12.

Asada, M., Kuniyoshi, Y., Drogoul, A., Asama, H., Mataric, M., Duhaut, D., Stone, P., and Kitano, H. (1998). The RoboCup Physical Agent Challenge: Phase-I. *in Applied Artificial Intelligence (AAI) Journal*.

Asama, H. (1992). Distributed autonomous robotic system configurated with multiple agents and its cooperative behaviors. *Journal of Robotics and Mechatronics*, volume 4:3.

Balch, T. (1997). Social Entropy: a New Metric for Learning Multi-robot Teams. In Proceedings of 10th International Florida Artificial Intelligence Reseach Symposium, pp 272-277. Florida AI Research Society.

Balch, T., and Arkin, R. C. (1994). Communication in Reactive Multiagent Robotic Systems. *Autonomous Robots*, volume 1, pp 27-52, Kluwer Academic Publishers, Boston.

Beckers, R., Holland, O.E., and Deneubourg, J.L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In *Proceedings A-Life IV*, MIT Press.

Billard, A., and Dautenhahn, K. (1997). Grounding Communication in Situated, Social Robots. Proceedings of TIMR 97, Towards Intelligent Mobile Robots Conference, (Technical Report Series UMCS-97-9-1), Dept of Computer Science, Manchester University.

Billard, A., and Dautenhahn, K. (1998). Grounding communication in autonomous robots: an experimental study. In *Robotics and Autonomous Systems*, special issue on Quantitative Assessment of Mobile Robot Behaviour, volume 24:1-2, pp 71-79.

Billard, A., and Dautenhahn, K. (1999). Experiments in learning by imitation – Grounding and Use of Communication in Robotic Agents. To appear in *Adaptive Behavior*, volume 7:3-4.

Billard, A., and Dautenhahn, K. (2000). Experiments in social robotics: grounding and use of communication in autonomous agents. To appear in *Adaptive Behavior*, volume 7:3-4.

Billard, A., and Hayes, G. (1998). Transmitting Communication Skills through Imitation in Autonomous robots. in *Learning Robots: A Multi-Perspective Exploration*, Birk, A., and Demiris, J. (editors), LNAI Series, Springer-Verlag, 1998.

Brock, D. L., Montana, D. J., and Ceranowicz, A. Z. (1992). Coordination and Control of Multiple Autonomous Vehicles. Proc. of the IEEE Conference on Robotics and Automation, Nice, France.

Brogan D. C., and Hodgins J. K. (1997). Group Behaviours for Systems with Significant Dynamics. In *Autonomous Robots*, volume 4:1, pp 137-153.

Brooks, R. A. (1986). A Robust Layered Control System for a Mobile Robot. *IEEE Journal of Robotics and Automation 2*.

Brooks, R. A. (1991). Intelligence without Reason. in Proceedings of the 12th International Conference on Artificial Intelligence, Morgan Kaufmann, San Mateo, CA, pp 569-595.

Bräunl, T., and Graf, B. (1999). Autonomous Mobile Robots with OnBoard Vision and Local Intelligence. In Proceedings of Second IEEE Workshop on Perception for Mobile Agents, Colorado.

Cao, Y. U., Fukunaga, A., Kahng, A., and Meng, F. (1997). Cooperative mobile robotics: Antecedents and directions. In *Autonomous Robots*, volume 4:1, 1997, pp 7-27.

Chance, M. R. A., and Mead, A. P. (1953). Social behaviour and primate evolution. *Symp. Soc. Exp. Biol VII (Evolution)*, pp 395-439.

Cheng, G., and Kuniyoshi, Y. (2000). Complex Continuous Meaningful Humanoid Interaction: A Multi Sensory-Cue Based Approach. *Proceedings of IEEE International Conference on Robotics and Automation* (ICRA 2000), pp 2235-2242, San Francisco, USA, April 24-28, 2000

Christensen, H. (1999). RoboCup-99. in *Robotics and Autonomous Systems*, volume 29, pp 103-105.

Dautenhahn, K. (1995). Getting to know each other – Artificial Social Intelligence for Autonomous Robots. *Robotics and Autonomous Systems*, volume 16, pp 333-356.

Dautenhahn, K. (1999). Embodiment and Interaction in Socially Intelligent Life-Like Agents. In: Nehaniv, C. L. (editor), Computation for Metaphors, Analogy and Agent, Springer Lecture Notes in Artificial Intelligence, volume 1562, Springer, pp 102-142.

Dautenhahn, K., Werry, I. (2000). Issues of Robot-Human Interaction Dynamics in the Rehabilitation of Children with Autism. To be published in Proc. *FROM ANIMALS TO ANIMATS*, The Sixth International Conference on the Simulation of Adaptive Behavior (SAB2000), 11 - 15 September 2000, Paris, France

Donald, B., Gariepy, L., and Rus, D. (2000). Distributed manipulation of multiple objects using ropes. In Proceedings of IEEE International Conference on Robotics and Automation, pp 450-457.

Dudek, G., Jenkin, M., Milios, E., and Wilkes, D. (1993). A taxonomy for swarm robots. In *IEEE/RSJ IROS*, pp 441-447.

Floreano, D., Nolfi, S., and Mondada, F. (1999). Co-evolution and ontogenetic change in competing robots. In *Robotics and Autonomous Systems*.

Grassé, P. P. (1959). La reconstruction du nid et les coordinations inter-individuelles chez *Bellicositermes natalensis* et *Cubitermes sp*. La théorie de la stigmergie: Essai d'interpretation des termites constructeurs. *Insects Sociaux,* volume 6, pp 41-83.

Harnad, S. (1990). The Symbol Grounding Problem. *Physica D 42,* pp 335-346.

Hemelrijk, C. K. (1997). Cooperation Without Genes, Games Or Cognition. Husbands, P., and Harvey, I. (editors), Proceedings of the 4th European Conference on Artificial Life (ECAL97). Cambridge: MIT Press, pp 511-520.

Humphrey, N. (1976). The social function of intellect. In Bateson, P. P. G., and Hinde, R. A. (editors), *Growing points in ethology*, pp 303-317. Cambridge University Press.

Jolly, A. (1966). Lemur social behaviour and primate intelligence. *Science*, volume 153, pp 501-506.

Jung, D., and Zelinsky, A. (2000). Grounded Symbolic Communication between Heterogeneous Cooperating Robots. In *Autonomous Robots*, volume 8:3, July 2000.

Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., and Osawa E. (1995). RoboCup: The Robot World Cup Initiative. Proceedings of IJCAI-95 Workshop on Entertainment and AI/Alife, Montreal.

Kitano, H. (1998). *RoboCup-97: Robot Soccer World Cup I*. Springer-Verlag, Berlin, Heidelberg, New York.

Kitano, H., Tadokoro, S., Noda, I., Matsubara, H., Takahashi, T., Shinjou, A., Shimada, S. (1999). RoboCup-Rescue: Search and Rescue for Large Scale Disasters as a Domain for Multi-Agent Research. *Proceedings of IEEE Conference on Man, Systems, and Cybernetics (SMC-99).*

Kitano, H., Tambe, M., Stone, P., Veloso, M., Coradeschi, S., Osawa, E., Matsubara, H., Noda, I., and Asada, M. (1997). The RoboCup Synthetic Agent Challenge 97. *In Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*. San Francisco, CA. Morgan Kaufman.

Kube, C. R., and Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, volume 30, pp 85-101.

Kube, C. R, and Zhang, H. (1994). Collective Robotics: From Social Insects to Robots, *Adaptive Behavior*, volume 2:2, pp 189-218.

Kube, C. R., and Zhang, H. (1997). Task Modeling in Collective Robotics. *Autonomous Robots*, volume 4:1, pp 53-72.

Kummer, H., and Goodall, J. (1985). Conditions of innovative behaviour in primates. *Phil. Trans. R. Soc. Lond. B 308*, pp 203-214.

Kuniyoshi, Y., and Nagakubo, A. (1997). Humanoid Interaction Approach: Exploring Meaningful Order in Complex Interactions. Proceedings of the *International Conference on Complex Systems*.

Mahadevan, S. (1996). Machine learning for robots: A comparison of different paradigms. In *Workshop on Towards Real Autonomy*, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-96), Osaka, Japan.

Mataric, M. J. (1992). Minimizing Complexity in Controlling a Mobile Robot Population. in Proceedings, IEEE International Conference on Robotics and Automation, Nice, France, pp 830-835.

Mataric, M. J. (1994). Interaction and intelligent behavior. MIT EECS PhD Thesis AITR-1495, MIT AI Lab, August 1994.

Mataric, M. J. (1996). Using Communication to Reduce Locality in Distributed Multi-agent learning. Brandeis University Computer Science Technical Report CS-96-190, Nov 1996.

Mataric, M. J. (1997). Using Communication to Reduce Locality in Distributed Multi-Agent Learning, *Proceedings, AAAI-97*, Providence, Rhode Island, July 27-31, pp 643-648.

Mataric, M. J. (2000). Making Humanoids Move and Imitate, *IEEE Intelligent Systems*, Jul 2000.

Melhuish C., Welsby J. & Edwards C. (1999). Using Templates for Defensive Wall Building with Autonomous Mobile Ant-Like Robots. Technical Report Series. Manchester University. U.K.

Minsky, M. (1963). *Steps toward artificial intelligence*. In Feigenbaum, E. A., and Feldman, J. A. (editors), Computers and Thought, pp 406-450. New York: McGraw-Hill.

Moriarty, D. E., Schultz, A. C., and Grefenstette, J. J. (1999). Evolutionary Algorithms for Reinforcement Learning. Journal of Artificial Intelligence Research, volume 11, pp 199-229.

Nolfi, S. (1997). Using Emergent Modularity to Develop Control Systems for Mobile Robots. *Adaptive Behavior*, volume 5:3-4, pp 343-363.

Nourbakhsh, I., Morse, S., Becker, C., Balabanovic, M., Gat, E., Simmons, R., Goodridge, S., Potlapalli, H., Hinkle, D., Jung, K., and Vactor, D.V. (1993). The Winning Robots from the 1993 Robot Competition. *AI Magazine*, volume 14:4, 1993.

Parker, L. E. (1995). The Effect of Action Recognition and Robot Awareness in Cooperative Robotic Teams. In IEEE/RSJ IROS, pp 212-219.

Parker, L. E. (1999). A case study for life-long learning and adaptation in cooperative robot teams *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II*, 1999, volume 3839, pp 92-101.

Parker, L. E. (1999). Adaptive Heterogeneous Multi-Robot Teams. *Neurocomputing*, special issue of *NEURAP '98: Neural Networks and Their Applications*, volume 28, pp 75-92.

Pfeifer, R. (1995). Cognition – Perspectives from autonomous agents. *Robotics and Autonomous Systems*, volume 15, pp 47-70.

Salustowicz, R. P., Wiering, M. A., and Schmidhuber, J. (1998). Learning Team Strategies: Soccer Case Studies. *Machine Learning*, Kluwer.

Schaeffer, J. (1997). Kasparov versus Deep Blue: The Re-match. *ICCA Journal*, volume 20:2, pp 95-102.

Shannon, C. E. (1949). *The Mathematical Theory of Communication*. University of Illinois Press.

Sharkey, N. E., and Heemskerk, J. (1997). The neural mind and the robot. In Browne, A. J., editor, *Neural Network Perspectives on Cognition and Adaptive Robotics*, pp 169-194. IOP Press, Bristol, UK.

Shen, W-M., Adibi, J., Adobbati, R., Cho, B., Erdem, A., Moradi, H., Salemi, B., and Tejada, S. (1998). Towards Integrated Soccer Robots. AI Magazine, Spring 1998.

Steels, L. (1996). The Spontaneous Self-organization of an Adaptive Language. In: Muggleton, S. (editor), *Machine Intelligence 15*. Oxford University Press, Oxford.

Steels, L., and Vogt, P. (1997). Grounding adaptive language games in robotic agents. *Proceedings of the fourth European Conference on Artificial Life*. Cambridge MA and London. The MIT Press.

Steels, L. (1994). Building agents out of autonomous behavior systems. in *The "artificial life" route to "artificial intelligence". Building situated embodied agents*, Steels, L., and Brooks, R. (editors). Lawrence Erlbaum, New Haven.

Stone, P., Riley, P., and Veloso, M. (2000). Defining and Using Ideal Team mate and Opponent Agent Models. To appear in the Twelfth Innovative Applications of AI Conference (IAAI-2000).

Sukhatme, G. S., Montgomery, J. F., and Mataric, M. J. (1999). Design and implementation of a mechanically heterogeneous robot group. In *Sensor Fusion and Decentralized Control in Autonomous Robotic Systems*, Proceedings of SPIE, pp 111-122.

Touzet, C. F. (2000). Robot awareness in Cooperative Robot Learning. *In Autonomous Robots*, volume 2, pp 1-13.

Veloso, M., Stone, P., and Han, K. (1998a). Prediction, Behaviors, and Collaboration in a Team of Robotic Soccer Agents. In Proceedings of International Conference on Multi Agent Systems (ICMAS'98), Cite des Sciences - La Villette, Paris, France.

Veloso, M., Uther, W., Fujita, M., Asada, M., and Kitano, H. (1998b). Playing soccer with legged robots. In Proceedings of IROS-98, Intelligent Robots and Systems Conference, Victoria, Canada.

Veloso, M., and Uther, W. (1999). The CMTrio-98 Sony Legged Robot Team. In Asada, M., and Kitano, H. (editors), *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin.

Versino, C., and Gambardella, L. M. (1997). Learning real team solutions. In Weiss, G. (editor), *DAI Meets Machine Learning,* Vogt, P. (2000). Grounding Language About Actions: Mobile Robots Playing Follow Me Games SAB2000 Proceedings Supplement Book, Meyer and Bertholz and Floreano and Roitblat and Wilson (editors), International Society for Adaptive Behavior. Forthcoming.

Wagner, I., A., and Bruckstein, A., M. (1995). Cooperative Cleaners: A Study In Ant-Robotics.

Weiss, G. (1996). Adaptation and learning in multi-agent systems: Some remarks and a bibliography. In Weiss, G., and Sen, S. (editors), *Adaptation and Learning in Multi-Agent Systems*. Volume 1042 of *Lecture Notes in Artificial Intelligence*, pp 1-21. Springer-Verlag, Berlin Heidelberg.

Werger, B. B. (1999). Cooperation without Deliberation: A Minimal Behavior-based Approach to Multi-robot Teams. *Artificial Intelligence*, 110, pp 293-320.