Experiments in Realising Cooperation between Autonomous Mobile Robots

David Jung, Gordon Cheng and Alexander Zelinsky
Robotic Systems Laboratory, Department of Systems Engineering
Research School of Information Science and Engineering
Australian National University
Canberra, ACT 0200, Australia
http://syseng.anu.edu.au/rsl

Abstract: This paper describes experiments in cooperation using autonomous mobile robots to perform a cleaning task. The robots have heterogeneous capabilities and the task is designed so that cooperation is required. Each experiment increases the sophistication of the cooperation scheme to assess the effect on task performance. The experiments range from using emergent cooperation with no communication to explicit cooperation and communication. We also propose an action selection mechanism that can also be used for distributed planning of joint actions inspired by the way primates co-construct joint plans. *Keywords: Cooperation, Vision, Touch, Mobile*

1. INTRODUCTION

The research effort into multi-robot systems is driven by the assumption that multiple agents have the possibility to solve problems more efficiently than a single agent. Agents must therefore cooperate in some way. Before we can begin to consider cooperation between robots we must define exactly what is meant by cooperation. Next we will investigate how cooperation can be achieved within a behaviour-based framework, by implementing cooperative behaviours using mobile robots to perform a concrete task.

The basic control design approaches can be broadly divided into four types. These are defined by [Mataric 92a] and we briefly reiterate them and their shortcomings here. The *purely reactive* approaches use a mapping from sensor sets to associated actions; a set of rules [Brooks 87]. The *planner-based* strategies originated with the symbolic AI community and employ a sense-plan-act cycle. The plan stage uses cognitive techniques to reason about a symbolic world model. There also exist hybrid systems which employ reactive components beneath planner-based systems to provide the benefits of both. Another approach is *behaviour-based*, which uses a set of interacting distributed concurrent behaviours, each of which may incorporate memory and learning of environment representations.

The purely reactive approach achieves robust real-time performance, but the tasks that can be achieved are limited because of the lack of any cognition. The planning approach suffers from a number of problems including slow interaction with the environment due to slow processing, the frame problem and the symbol grounding problem. The hybrid approaches attempt to marry two incompatible philosophies and still suffer from many of the problems of the planning approaches. For more depth on these problems refer to [Pfeifer 95]. The behaviour-based approach has the potential of real-time response and cognitive processing in a uniform manner. We have chosen to implement our cooperative behaviour within the behaviour-based philosophy.

2. OUR AIMS

As research into multi-agent cooperation is still a relatively new field there are no standard formalism's for describing cooperation nor benchmarks for measuring the performance of techniques. The aim of our research is to determine the defining characteristics of cooperative behaviour; examine the effect of these characteristics on performance and to propose a scheme for implementing cooperation within the framework of behaviour-based robotic systems.

In order to assess the performance of different levels of sophistication of cooperative behaviour we have focused on a concrete application. The task we have chosen is for two autonomous mobile robots to clean the floor of our laboratory. The 'Yamabico' robots [Yuta 91] shown in Figure 1 each have different tools and sensors such that neither can accomplish the task alone.

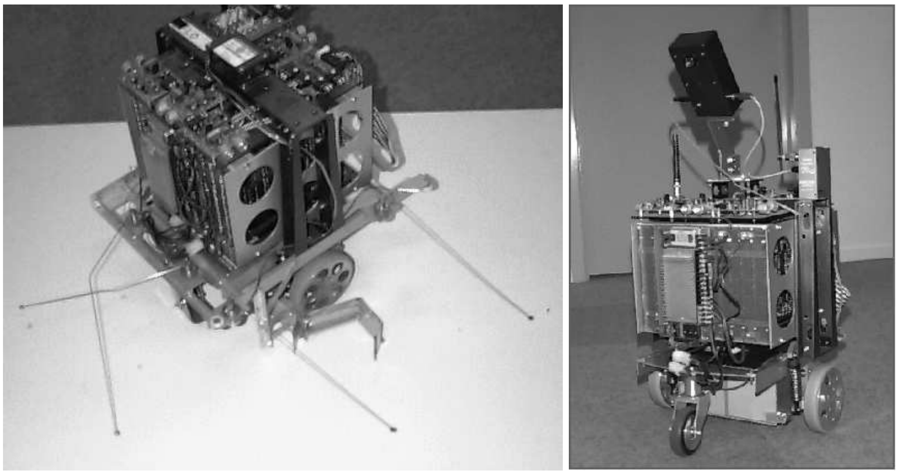

Figure 1 - (a) Yamabico Flo (b) Yamabico Joh

The robot *'Joh'* navigates by vision, and has a vacuum cleaner that can be turned on and off via software. Joh's task is to vacuum piles of litter from the laboratory floor. It cannot vacuum close to walls or furniture. It has the capability to 'see' piles of litter using its vision system, but not fine particles scattered over the floor.

The other robot *'Flo'*, has a brush tool that is dragged over the floor to sweep distributed litter into larger piles for Joh to pick up. Flo uses whisker sensors to navigate (see Figure 1a).

The task is to be performed in the unmodified indoor environment of a laboratory. Our laboratory is cluttered and the robots have to contend with furniture, other robots, people, opening doors, changing lighting conditions, equipment with dangling cables and other hazards.

3. WHAT IS COOPERATION?

The word 'cooperation' has been applied to behaviour between robots and also between humans and robots, without much thought having been given to its exact meaning. Although its literal reading - simultaneous operation - is quite general, the word has historically been used primarily to refer to the joint behaviour of humans, and sometimes animals. As such its traditional meaning may be loaded with anthropomorphic assumptions which do not hold in the modern robotics context. Before we consider an appropriate definition for robot cooperation we need to examine cooperative behaviour in humans and animals. As cooperation is a social phenomena we may look to sociology and social psychology for insight into its mechanisms [Tinbergen 53][Wilson 75][Wright 95].

3.1 The origins of Altruism

Given that Charles Darwin's theory of natural selection implies that individuals behave in a completely selfish manner to increase their own fitness, an obvious question is 'why do biological organisms cooperate at all?'. This very question perplexed Darwin who introduced a mechanism he called group selection to account for the cooperation he observed - without any idea as to why group selection should exist. The answer to the question came once the gene was discovered and was identified as the basic selectionist unit, not the individual. Modern sociobiology and evolutionary psychology is based on the application of Darwin's theory at the level of the gene instead of the individual. This selfish-gene approach has been popularised across disciplines in a number of books by Richard Dawkins [Dawkins 96]. It is interesting to note that according to the *social intelligence hypothesis* intelligence in primates "*originally evolved to solve social problems and was only later extended to problems outside the social domain*" [Byrne 95][Dautenhahn 96a].

The evolutionary path from early primates to humans was a long and convoluted one. It was considerably accelerated once performance in the social context became the dominant driving force behind selection rather than just survival in the natural environment. This was largely due to the beginnings of language. A detailed account of this is given in [Wright 95] and [Diamond 91]. What we should keep in mind from this is that the dynamics of cooperation between humans implicitly includes a history all the peculiarities of our particular path through the evolutionary space of possibilities.

3.2 Characterising Cooperation

The robotics community has investigated various characteristics of cooperation. Using the terms defined by [Cao 95], which we can relate to the social organisation of biological systems, some of the characteristics that define the *group architecture* of a cooperative system are:

- Centralisation/Decentralisation - of planning for cooperative actions
- Differentiation - homogeneous or heterogeneous capabilities
- Communication Structures - none, implicit, explicit, dialogue for planning
- Modeling of Other Agents - no awareness, awareness, modeling of others

A more detailed description of these and a classification of current research into these categories is given in [Jung 96].

Our research focuses on systems which are distributed, allow heterogeneous individuals, explicit communication, and modelling of other agents. Individuals are autobiographical - learn from experience, engage in explicit communication, and dynamically plan cooperative actions. Although cooperation is dynamically planned, the mechanisms are determined at design time - not learned via social or cultural learning. The learning will be limited to learning the environment. The robots will also be given an implicit model of some actions of other agents.

4. REALISING COOPERATION

In the sections to follow we will describe four experiments we have designed around the cleaning task. The first will be a simple scenario not involving any explicit communication or learning, and will serve as a benchmark against which to assess the performance of the other three. They will increase the level of cooperative sophistication. Before we describe the experiments we will briefly describe the robot hardware we are using.

4.1 Hardware

We have the two Yamabico robots mentioned above, each equipped with basic locomotion and four ultrasonic range sensors. The robot Joh also has a CCD camera and video transmitter that sends video to our *Fujitsu MEP tracking vision system*. The vision system does template correlation, and can match about 100 templates at frame rate. The vision system can communicate with the robot, via UNIX host, over a radio modem to close the loop. We have implemented a vision-based navigation system that is capable of landmark based navigation and can operate safely in dynamic environments at speeds up to 600 mm/sec [Cheng 96].

The robot Flo has a brush and some novel proportional whisker sensors we have developed for accurately sweeping close to walls and furniture. Flo has two whisker sensors mounted on its left side for wall following and two whiskers in front for collision detection (see Figure 1a). The basic behaviours we have implemented on Joh and Flo are discussed below. First we will outline the three experiments we have designed to carry out the cooperative cleaning task and assess performance.

4.2 Experiments

4.2.1 Emergent Cooperation

The simplest experimental case, which will serve as a benchmark against which to access the following three, involves no awareness of the other robot and hence no explicit cooperation or communication. Any communication is implicit via interaction through the environment. The cooperation can be described as emergent. We have found that very little needs to be added to the base competency of each robot to achieve this kind of cooperation.

Flo uses wall following and obstacle avoidance behaviours to brush around the perimeter of the laboratory close to the walls. Periodically, Flo decides to deposit the litter collected so far into a pile in clear space farther away from the wall where Joh can vacuum it.

Joh navigates our laboratory using visual landmarks. A foraging behaviour is executed in which it wanders about the laboratory, avoiding obstacles, people and Flo while searching for piles of litter. The piles of litter are identified using the vision system. To conserve power Joh only turns its vacuum on when its sees litter. Occasionally Joh may attempt to vacuum artefacts on the floor that aren't litter, such as a patch of sunlight or a piece of equipment. In the case of sunlight no harm is done, while in the case of a piece of equipment the bump sensor will stop the robot.

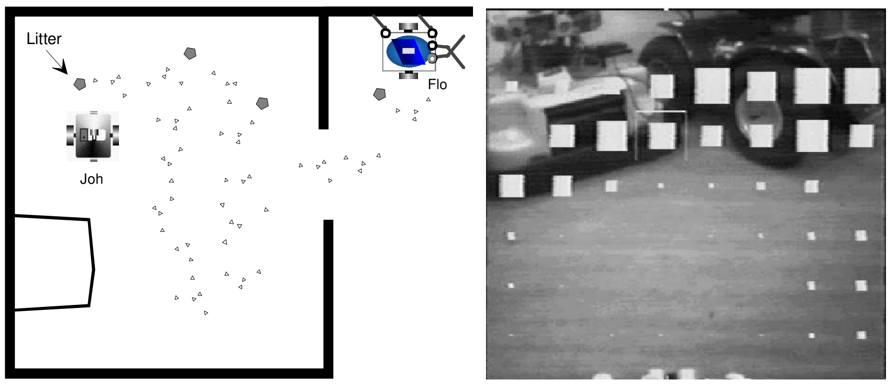

Figure 2 - (a) Cleaning by emergent cooperation (b) Visual obstacle avoidance

Using these simple behaviours Flo and Joh clean the floor in an inefficient manner. Neither robot requires the capacity to purposively navigate around the laboratory using a learned map in this case.

Since we required a close high speed wall following capability in Flo for sweeping close to the walls, after investigating existing technologies, we decided to develop a unique proportional whisker sensor [Jung 96a]. We utilised a Kalman filter to estimate the robot position relative to the wall by fusing the information from the side whiskers with the wheel odometry. Flo is also fitted with a simple scoop so that wall following at an appropriate distance causes litter to be scooped from the wall. Flo then periodically drives into free space away from the wall and reverses leaving a pile of litter from scoop behind.

The wall following behaviour was extended to follow rough contours around the laboratory by using the front whiskers to detect corners and obstacles and to make an appropriate turn. Flo also has other miscellaneous simple behaviours such as reflex stopping, door traversal, tracking along straight trajectories and others [Jung 96a].

Joh has the ability to visually distinguish between the carpet on our laboratory floor and other obstacles at video frame rate (30Hz). This is accomplished using our template matching vision system and has been discussed in [Cheng 96]. This provides a good free space wandering and obstacle avoidance behaviour as shown in Figure 2b. Joh also needs the ability to detect piles of litter left by Flo. For this we use an 'interest' operator that segments areas of non-carpet surrounded by carpet (see Figure 3a). The 2D array of template correlation values must be normalised to compensate for camera lens distortion first. Once a possible pile of litter has been detected a visual servoing behaviour moves the robot over it. The vacuum can be turned on and off via software control. If the object was an obstacle not litter, then since Joh is fitted with a row of bump sensors on the front, it will be stopped from trying to vacuum it.

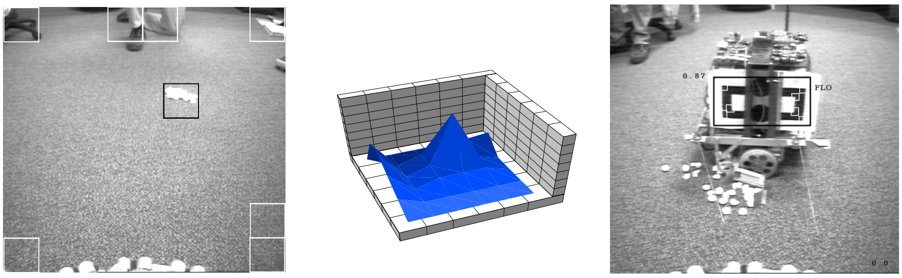

Figure 3 - (a) Pile of litter and Interest operator output (b) Flo from Joh's camera

The beauty of this approach to litter detection is that it is very simple. It requires little computation, no model, is very robust, and delivers enough information for the task. A scheme that used sophisticated classical computer vision techniques, perhaps using a model of litter, would be computationally expensive, difficult to implement and would still fail some times. Hence almost nothing would be gained since the bump sensors would still be required as backup.

4.2.2 Cooperation by Observation

The second experiment, a more interesting case, utilises explicit cooperation using implicit communication by passive observation (cf. [Kuniyoshi 94]). In this case Joh uses the vision system to identify Flo by matching a unique geometric pattern. By observing Flo's actions Joh can determine the approximate location of the litter deposited by Flo. This new behaviour augmented the existing foraging behaviour and improved the efficiency of the cleaning task. This case requires awareness of Flo's existence by Joh.

In particular Joh needs to be able to identify and track the motion of Flo in the live video. This is accomplished by placing a unique pattern on the sides of Flo and tracking templates from it with geometric constraints (Figure 3b). The match correlation value for each template is fused using a network of Kalman filters. Other work in our laboratory used this technique with the vision system for robust human head tracking and is discussed in [Zelinsky 96]. Because we know the size and shape of the pattern we can also easily estimate the distance and heading of Flo relative to Joh.

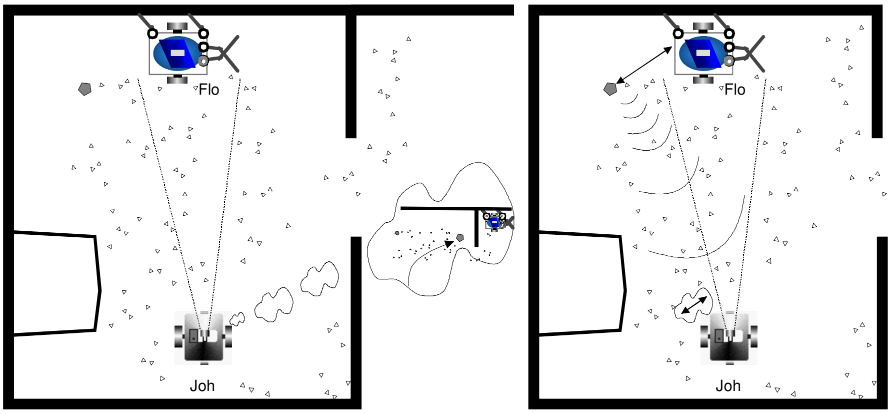

Figure 4 - (a) Cooperation by observation (b) Cooperation with communication

4.2.3 Cooperation with Communication

The third experiment, by including explicit communication between the robots, can further improve the efficiency of the task. When Joh can see Flo it communicates this fact. In response Flo communicates the location where it next plans to dump a pile of litter, specified relative to its own location. Since Joh knows Flo's location relative to itself, the dump location can also be calculated relative to itself. Joh then navigates to the approximate dump location and initiates a visual search for the litter pile. If it is acquired a visual servo behaviour vacuums over it. This case only requires some simple communicative behaviours over the last experiment. Since it is not always possible to predict Flo's actions from observation alone, the communication removes ambiguity and improves performance. For example, visually determining if, when and where Flo plans to dump litter by observation alone is somewhat uncertain.

The communication between the robots is achieved using two pairs of radio modems between each robot and a UNIX host in combination with some networking software we developed for the custom robot operating system.

4.2.4 Co-construction of Joint Plans

All the capabilities required to implement the first three experiments have been described. The fourth experiment requires a significant increase in the sophistication of the basic behavioural capabilities of the robots. In particular they must be able to learn a map of their environment and purposively navigate by it. A mechanism for this is currently being implemented, and is briefly described below.

The are two possible dialogues in this experiment. When Joh sees Flo it chooses a name for the location where it sees Flo and labels the location in its internal map. It then tells Flo to label the location in its map. Joh and Flo have different representations for their maps. The map is represented in terms of the behavioural and sensory space of the robots, which means Joh will have visual landmarks in its map while Flo will have whisker based landmarks in its. Hence there is no possibility for communicating an absolute location except where the location has been labelled as just described.

The second dialogue occurs when Flo dumps litter, in which case it communicates the litter location in terms of a relative position from the nearest named location. In this case Joh simply navigates to the location using its map and visually searches for the litter to vacuum. Alternatively Joh could just note the location and vacuum the litter pile when it next comes near the location.

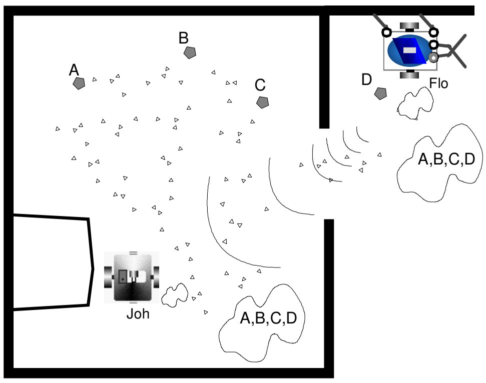

Figure 5 - Joint planning (eg. communication of rendezvous points)

The communication is completely grounded in the behaviour-sensor space of the robots even though each has a different such space. There is no communication of symbolic concepts that have been anthropomorphically designated by the designer. For example, we may consider that Flo can detect 'doors' using its whiskers and Joh can detect 'doors' visually. Hence we might imagine some communication using the concept of a 'door'. However, since a door to Flo is really just a pattern of whisker movements and to Joh a geometric arrangement of matched templates, there will be situations where the implicit identity assumption in our anthropomorphic designation will break down. As Connah and Wavish state "*communication between them (robots) will have to be flexible and natural; growing out of the perceptions that are part of their general behaviour patterns*" [Connah 90].

The mechanism we have developed for purposive navigation uses a spatial and topological map representation that is integrated into our action selection mechanism, which incorporates learning. This mechanism is briefly described below because it is also the basis of our proposed scheme for the dynamic co-construction of joint cooperative action plans, which is currently being implemented.

One major problem to be solved in robotics is, given a robot with a repertoire of basic behaviours - which behaviour should be selected next? Various action selection mechanisms have been proposed, such as Rodney Brooks' subsumption architecture [Brooks 87]. We have developed a mechanism based loosely on Pattie Maes' spreading activation scheme [Maes 90a], extended to add integrated learning and adapted to a behaviour based framework. This is beyond the scope of this paper, please refer to [Jung 97b].

Briefly, in simplified form, behaviours are connected in a network, where each behaviour has a set of preconditions. The preconditions are the outputs of feature detectors that are basically virtual sensors. Only one behaviour is active at once, and it must have all its preconditions satisfied. Behaviours are also connected to those feature detectors that become satisfied as a result of its execution - its postconditions. Behaviours become active according to an activation level - which is a result of a spreading activation algorithm that depends on the connections. For an example see Figure 6. The precondition-behaviour connections are designed, but the behaviour-postcondition connection are learnt according to correlations between behaviours and feature detector outputs. The spreading activation achieves distributed planning of action sequences. It can handle multiple hypotheses and contingencies.

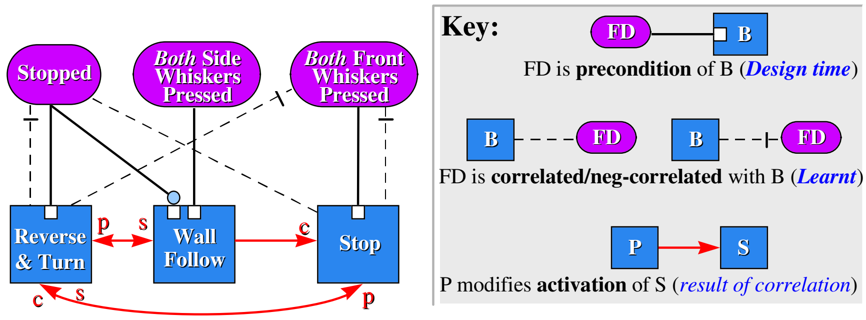

Figure 6 - Behaviour network implementation of a simple perimeter follower

The important point to note is that action sequences are always linked via a condition that is directly sensed from the environment. If a behaviour doesn't have the desired effect on the state of the environment then its successor behaviour will not be blindly activated.

We have designed a scheme for using this same mechanism for representing topological environment maps. These are landmark feature detectors linked by behaviours to take the robot from one landmark to the next. We have also incorporated spatial mapping into the scheme by using feature detectors for locations which self-organise like nodes of a Kohonen self-organising-map.

Primates are very social animals. As Bond writes in reference to vervet monkeys "*They are acutely and sensitively aware of the status and identity of other monkeys, as well as their temperaments and current dispositional states*" [Bond 96]. Humans, as other primates, have the ability to co-construct plans with more that one interacting person, and flexibly adapt and repair them all in real time.

Bond goes on to describe the construction and execution of joint plans in monkeys. He defines a *joint plan* as a conditional sequence of actions and goals involving the subject and others. In order to achieve interlocking coordination each agent needs to adjust its action selection based on the evolution of the ongoing interaction. The cooperative interaction will consist of a series of actions - including communication acts - where each agent attempts different plans, assesses the other agents' goals and plans, and alters the selection of its own actions and goals to achieve a more coordinated interaction where joint goals are satisfied. This model of interaction is similar to that proposed by human conversation theorists [Goodwin 81].

The action selection mechanism outlined above can be effortlessly applied to planning this sort of joint plan. This is because it is irrelevant which robot causes a change in the environment that triggers the precondition of the next action of a sequence. For example a cooperative interaction may consist of a sequence of actions by each robot interleaved (see Figure 7).

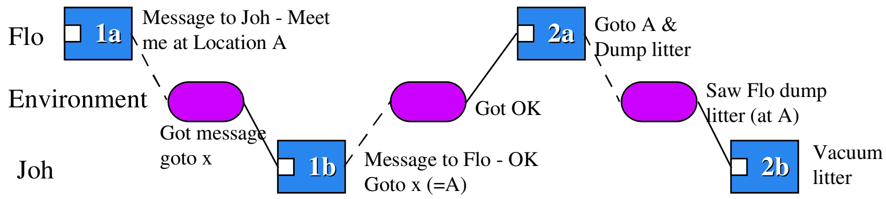

Figure 7 - Interleaved cooperative action sequence

This type of cooperative interaction, where the repertoire is essentially fixed at design time, is adequate for realising the fourth experiment. If we consider communicative utterances as behaviours and heard utterances as sensed features then this mechanism is also capable of planning dynamic communication dialogues. We may investigate this in future research.

5. Conclusion

We have presented the design and implementation of some experiments in cooperative cleaning behaviour between autonomous mobile robots. In particular we:

- Implemented a novel proportional whisker sensor and data fusion technique for close wall following.
- Implemented basic robot behaviours for emergent robot cleaning, including whisker and vision based navigation capabilities.
- Improved the performance of cleaning by adding, first implicit - by observation, and then explicit, communication between the robots.
- Proposed an action selection mechanism for joint planning of cooperative actions, including the possibility of planning communication dialogue.

Although this was a slightly contrived task one can easily imagine applications where there exists a trade-off between using one or a few very complex robots or a larger number of much simpler and less expensive robots to perform a task. For example, you may design a large complex and expensive automated cleaning robot, but in the event of failure the task cannot be performed. If instead there were numerous smaller, cheaper and simpler robots performing the same task cooperatively, then the reliability of any single robot will be higher. In addition, because they are cheaper to manufacture it may be possible to have redundancy so that the failure of a single unit does not render the whole system useless. In addition the time required to develop a few robots with simple behaviours may be less than that required for a single robot with complex behavioural requirements.

We also believe that the approach to action selection we are proposing will prove effective for implementing high-level behaviour in simple behaviour-based robots with meagre perceptual abilities, including the possibility of distributed planning of cooperative behaviour and dialogue for communication.

Videos demonstrating our experiments are available from our web site.

REFERENCES

[Bond 96] Bond, Alan H., "**An Architectural Model of the Primate Brain**", *Dept. of Computer Science, University of California*, Los Angeles, CA 90024-1596, Jan 14, 1996.

[Brooks 87] Brooks, Rodney A. and Connell, Jonathan H., "**Asynchronous Distributed Control System for a Mobile Robot**", *SPIE's Cambridge Symposium on Optical and Opto-Electronic Engineering Proceedings*, Vol. 727, October, 1986, 77-84.

[Byrne 95] Byrne, R., "**The thinking ape, evolutionary origins of intelligence**", *Oxford University Press*, 1995.

[Cao 95] Cao, Y. Uny, Fukunaga, Alex S., Kahng, Andrew B. and Meng, Frank, "**Cooperative Mobile Robotics: Antecedents and Directions**", *IEEE 0-8186-7108-4/95*, 1995.

[Cheng 96] Cheng, Gordon and Zelinsky, Alexander, "**Real-Time Visual Behaviours for Navigating a mobile Robot**", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol 2. pp973. November 1996.

[Connah 90] Connah, David and Wavish, Peter, "**An experiment in Cooperation**", *Decentralized A.I.*, Yves Demazeau & Jean-Pierre Müller (Eds.), Elsevier Science Publishers B.V. (North Holland), 1990.

[Dautenhahn 96a] Dautenhahn, K., "**Social modelling - on selfish genes, social robots and the role of the individual**", **submitted to**: *Int Journal in Computer Simulation*, Special issue on the Simulation of Social Behavior, August 1996.

[Dawkins 96] Dawkins, Richard, "**Climbing Mount Improbable**", *Penguin Books*, 1996. ISBN 0-14-026302-0.

[Diamond 91] Diamond, Jared, "**The rise and fall of the third chimpanzee**", *Vintage Books*, 1991.

[Goodwin 81] Goodwin, Charles, "**Conversational Organization : interaction between speakers and hearers**", *Academic Press*, New York and London, 1981.

[Jung 96] Jung, David and Zelinsky, Alexander, "**Cleaning: An Experiment in Autonomous Robot Cooperation**", *Proceedings of the World Automation Conference (WAC) - International Symposium on Robotics and Manufacturing* (ISRAM), Montipellier, France, March, 1996.

[Jung 96a] Jung, David and Zelinsky, Alexander, "**Whisker-Based Mobile Robot Navigation**", *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol 2. pp497. November 1996.

[Jung 97b] Jung, David, **"An Architecture for Cooperation between Autonomous Mobile Robots**", *PhD Thesis, Intelligent Robotics Laboratory, University of Wollongong*, submitted December 1997. For further information on this research see the web site `http://pobox.com/~david.jung`.

[Kuniyoshi 94] Kuniyoshi, Yasuo, Rougeaux, Sebastien, Ishii, Makoto, Kita, Nobuyuki, Sakane, Shigeyuki and Kakikura, Masayoshi, "**Cooperation by Observation [The framework and basic task patterns]**", *Proceedings of the 1994 IEEE International Conference on Robotics and Automation*, San Diego, California, May 8-13, 1994.

[Maes 90a] Maes, P., "**Situated Agents Can Have Goals.**", *Designing Autonomous Agents*. Ed: P. Maes. MIT-Bradford Press, 1991. ISBN 0-262-63135-0. Also published as a special issue of the Journal for Robotics and Autonomous Systems,Vol. 6, No 1, North-Holland, June 1990.

[Mataric 92a] Mataric, Maja J., "**Behavior-Based Systems: Key Properties and Implications**" in *Proceedings, IEEE International Conference on Robotics and Automation*, Workshop on Architectures for Intelligent Control Systems, Nice, France - May 1992, 46-54.

[Parker 95] Parker, Lynne E., "**The Effect of Action Recognition and Robot Awareness in Cooperative Robotic Teams**", *IEEE 0-8186-7108-4/95*, 1995.

[Pfeifer 95] Pfeifer, Rolf, "**Cognition - Perspectives from autonomous agents**", *Robotics and Autonomous Systems 15*, pp47-70, 1995.

[Steels 90] Steels, Luc, "**Cooperation between Distributed Agents through Self-Organisation**", *VUB AI Lab*, Pleinlaan 2, 1050 Brussels Belgium, IEEE *IROS 90*.

[Steels 96a] Steels, Luc, "**Synthesising the origins of language and meaning using co-evolution, self-organisation and level formation**", Artificial Intelligence Laboratory, Vrije Universiteit Brussel. In: Hurford, J. (ed.) *Evolution of Human Language*, Edinburgh University Press. Edinburgh, July 1996.

[Steels 96b] Steels, Luc, "**Self-organizing vocabularies**", In: Langton, C. (ed.) *Proceedings of Alife V*, Nara Japan, 1996.

[Tinbergen 53] Tinbergen, N., "**Social Behaviour in Animals**", *Methuen & Co. Ltd.*, London, 1953.

[Wilson 75] Wilson, E. O., "**Sociobiology: The New Synthesis**", *Harvard*, 1975.

[Wright 95] Wright, Robert, "**The Moral Animal : Why We Are the Way We Are : The New Science of Evolutionary Psychology**", *Vintage Books*, Luann Walther (Ed.), ISBN:0679763996, 1995.

[Yuta 91] S. Yuta, S. Suzuki and S. Iida, "**Implementation of a small size experimental self-contained autonomous robot - sensors, vehicle control, and description of sensor based behavior**", *Proc. Experimental Robotics*, Tolouse, France, LAAS/CNRS (1991).

[Zelinsky 96] Zelinsky, Alexander and Heinzmann Jochen, "**Real-Time Visual Recognition of Facial Gestures for human-Computer Interaction**", *Proceedings of the 2nd International Conference on Automatic Face and Gesture Recognition*, IEEE Computer Society Press, Killington, Vermont, October 1996.