Improving performance in a multi-robot task through minimal communication

Jose Carmena* **and John Hallam****

Institute of Perception, Action and Behaviour
Division of Informatics, University of Edinburgh
5 Forrest Hill, EH1 2QL Edinburgh
Scotland, UK

Abstract. It is a well known result from economics and mathematical theory that communication can increase group performance. In the real world there are many collective tasks which need a terminating signal to inform the agents about the end of the task. In this work we show an example of how using minimal communication improves task performance in a set of mobile robots while performing a retrieval task. The task is performed by a set of LEGO robots inside an arena, of which a small area is designated the nest, with some polystyrene targets dispersed about the rest of it. The task is to collect the targets and bring them to the nest. The robots are provided with a radio-communication unit which allows them to communicate with each other. Two different experiments are done: A collective task and a cooperative task. For the cooperative task minimal information is shared by the robots using the radio communication unit. Performance of the robots was compared and it was found that increasing the number of robots improved the results obtained for both tasks. Also, the cooperative task was significantly better than the collective task in all the cases tested.

1 Introduction

Within the field of distributed systems, multi-agent systems research has increased during recent years because of the necessity of achieving tasks which can not be achieved by a single agent [5]. Multi-robot systems are particularly different from the rest because of the implication of the interaction with the real world during the execution of any task.

But, why using a group of robots? As everybody knows, there are lots of either human or animal tasks which need multiple agents to be achieved, either because of the size of the task or because of the necessity of cooperation between individuals, *e.g.* building a house, playing a group game (such as football), building a termite nest, etc.

We find the same requirement for robots. There are some tasks which could be very complex or impossible for a single robot to accomplish, or in which performance benefits such as time spent in the task can be gained from using multiple robots [5].

Thinking of different types of these tasks we find the following characteristics: need to be finished at some point, having an extra agent will improve group performance, and the necessity of a terminating signal to inform the agents when the task is finished. Examples of these tasks are:

- **exploration tasks** of hazardous environments, unknown places, etc.
- **foraging tasks** such as ant-like behaviour.
- **cleaning tasks** of big places such as airports, public buildings, etc.

For such tasks, it is a well known result from economics and mathematical theory that communication can increase group performance. This has been tested in robots by many researchers with successful results as seen in [9, 15, 16, 17]. But, what kind of information can be shared by a communicating group of robots? We can classify this information from lower to higher levels as follows:

* Email: jose@dai.ed.ac.uk
** Email: john@dai.ed.ac.uk

- **kernel level:** how long should the underlying process run
- **sensory level:** readings from any kind of sensors
- **behaviour level:** which behaviours to activate/inhibit
- **territorial level:** suggesting specific places to go, e.g. "I'm going here"

Also, increasing the number of agents performing the task improves performance [15] but this improvement will start decreasing when a critical number N of agents is reached, since too many robots will be sharing the same space on the arena as well as the radio bandwidth, generating a "traffic jam" problem. This is addressed by Cao, Fukunaga and Kahng in [5] as a "resource conflict" stating that "when a single indivisible resource is requested by multiple robots, *resource conflict* arises... With multiple robots, resource conflict occurs when there is a need to share space, manipulable objects or communication media."

2 The role of communication in common tasks

Common tasks can be performed in two different ways according to the agents' behaviour: collectively and cooperatively. First of all, it is worth differentiating between them. Because of the several possible contexts in which these terms can be used (as seen in [5]), in this paper I will use the definitions proposed by myself in [6]. In that work, I defined collective and cooperative behaviour for a group of agents as follows:

- **Collective behaviour:** When multiple agents co-exist and treat each other as obstacles while going about a common task.
- **Cooperative behaviour:** When multiple agents co-operate during the performance of a collective task by exchanging some type of information.

A clear example of a collective behaviour task is the work done by Beckers, Holland and Deneubourg [2]. In such a task, the robots are not provided with global sensing or communication. The environment consists of an arena crowded of pucks, and the task is to collect them into piles of increasing sizes. As the robots perform, small piles become bigger until all the pucks have been grouped. There is no explicit knowledge inside the robots about what they are doing but there is a kind of intelligent collective behaviour in the performance of the group.

Another example is Altenburg's work [1]. He experimented with ant-like robots (ALRs) investigating their performance in collective search and retrieval type tasks.

Many examples of cooperative behaviour using robots can be found in the literature, being the most representatives (because of the size of the colony used) those done by Mataric as seen in [13, 14, 15]. Most of Mataric's cooperative experiments are focussed on control and learning group behaviour, using variations of Reinforcement Learning algorithms. She explodes the cooperativeness by territorial divisions along the arena in which the robots perform the task. Mataric states that there are two main problems in learning in multi-agent systems: hidden state and credit assignment. She addresses both problems by simple communication based on local broadcast and exploits message-passing as a broader form of sensing and reinforcement.

Another representative example of cooperative behaviour is the work done by Sugawara *et al.* [18] in which a quantitative analysis of the effectiveness of multirobot group behaviour is done.

As in nature, where the velocity of a herd is determined by the velocity of the slowest individual, collective behaviour is influenced by the behaviour of the worst individual. In terms of a group of robots performing a collective task such as those mentioned above, the worst individual is the slowest and/or less accurate robot. Cooperative behaviour, however, is not so influenced. Instead of being limited individually, task performance is limited collectively.

The important point of all this is that in the collective and cooperative examples seen above, a *terminating signal* is needed in order to make the agents to stop, otherwise how could they know when is the task finished? In the collective case, each agent has its own terminating signal where in the cooperative case this signal is inferred by all the agents through communication.

In the example described in next section, we will see a group of robots performing a common task under collective and cooperative behaviour. This cooperativeness arises by sharing only the minimal communication needed by the agents in order to infer this terminating signal. Hence, the point of the story could be summarised as the fact that even this tiny amount of communication buys an improvement in performance as we will see in the next section.

3 Example case: Experiments with LEGO robots

Real robots

The necessity of building physical agents, such as mobile robots, instead of using software agents in a simulated world in order to know how they would perform when interacting with the real world is well known from the literature [3, 4, 10, 11]. Hence, we must build robots if we really want to know about them, their performances, capabilities, constraints, that is, the real problems existing in a physical world rather than in a virtual world.

The most important consequence of using these physical agents (LEGO robots in our experiments) is that every agent, even though they have been identically built (*i.e.* identical shape and sensor distribution), performs different than the others. The reason for this is because of the un-repeatability of the mechanical and electronic domains [12]. The former one includes gears, wheels, joints, and any LEGO bit. In the electronic domain, sensors and actuators (motors) are the most relevant factors. These are the consequences of being physical agents (*situatedness*) interacting with the real world (*embodiment*) [4].

Therefore, the robots will perform the task according to the capabilities given to them by hardware (shape of their bodies, and sensor and actuator locations) and software (a set of processes and their activation and inhibition). A complete description of these capabilities can be found in [6].

Two different views of one of the robots used in the experiments can be seen in figures 1 and 2.

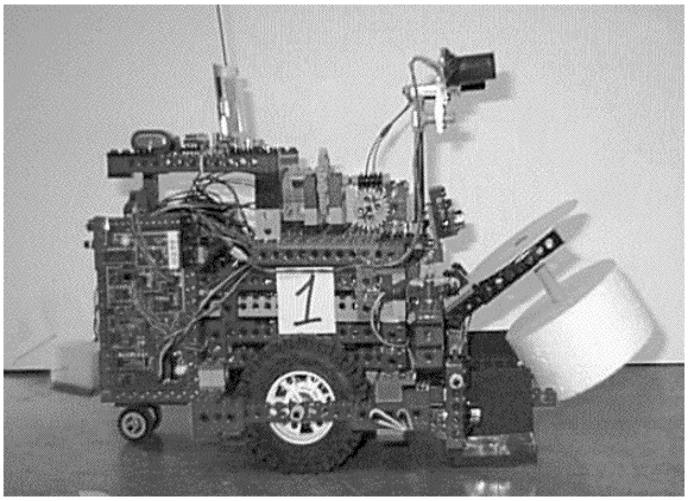

Fig. 1.: Side view.

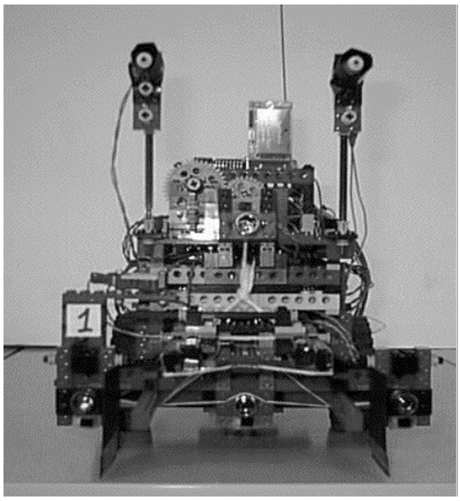

Fig. 2.: Front view.

The task

The task described next is inspired by how insects work together in a common task. In our case, the common task will consist of searching for 'food' and once found bringing it to the nest.

This task will be performed by a group of up to four LEGO Robots inside an arena, of which a small region is designated the nest, and some polystyrene targets (as food pieces) will be dispersed about the rest of the arena.

The main objective for each robot is to find and place on the nest the number of targets necessary for their survival. This number of targets will depend on how many robots are performing the task.

The robots will be performing the task during a period of time whose maximum will be constrained by the batteries' life-time. After placing the required number of targets, a robot will stop and it will be removed from the arena[3]. The task will be finished when the group of robots has collected the required number of targets.

The task will be done by the robots under two different conditions of performance: Collective behaviour (without cooperation among the robots) and cooperative behaviour (sharing information about the task with the other robots through a radio-communication unit placed on the top of each robot). The only information shared during the cooperative task is a short message broadcast[4] by an agent as it delivers a target, from which all agents keep track of the total targets on the nest. In this case, the goal of the task is to place on the nest the total number of targets needed for the group survival. This is an important difference compared with the collective task in which a robot only takes care of itself. During the cooperative task, a single robot could bring to the nest more targets than other robots, as a consequence of being more accurate, fast or agile than the rest. This is the important difference between the two ways of performing the task. In figure 3 a schematic diagram of the task can be seen.

As soon as the needed number of targets are placed on the nest, every robot will stop and the task will be finished. One of the important results obtained from the pilot tests is the speed of each robot compared with the rest. We will use this measure to classify the robots in a decreasing speed order, starting with robot 2 as the fastest, followed by robots 4 and 3 and ending with robot 1 as the slowest.

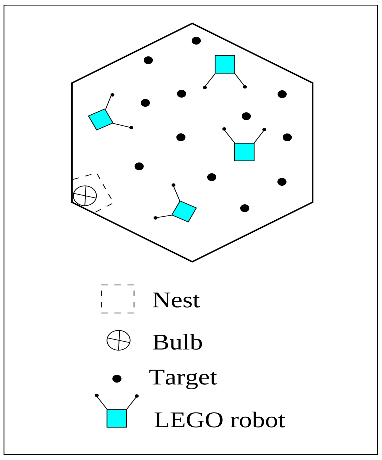

Fig. 3.: A schematic diagram of the task.

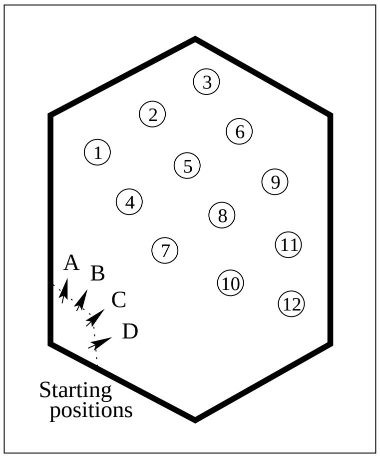

Fig. 4.: Target distribution and starting positions.

Description of the experiments

A single experiment will consist in four runs in each of which twelve targets must be collected before a time-out. This time-out has been fixed at thirty minutes to accord with the batteries' life time. The main factor to measure is the time spent by the robot/s to accomplish the task. Also, for each of the robots we will measure time spent in finding a target, the time spent in carrying that target to the nest and the number of the fetched target.

Experiments are done with one, two, three and four robots in both modes: collective and cooperative[5]. In the collective experiments, the number of targets to be collected by each robot depends on the number of robots performing the experiment: twelve targets in the case of one robot, six targets in the case of two robots, four in the case of three robots and three targets in the case of four robots. In the cooperative experiments there is no fixed number, only 12 in total. The distribution of the targets about the arena can be seen in figure 4. In figures 5 and 6 we can see a set of four robots performing the cooperative task[6].

[3] This will be done only in the collective task since in the cooperative task all the robots finish at the same time.
[4] This half-duplex communication is produced by a radio communication unit placed on the top of each robot. Detailed information about these units can be found in [6] and [7].
[5] For the cooperative case experiments will be done with two, three and four robots since cooperation can't be done

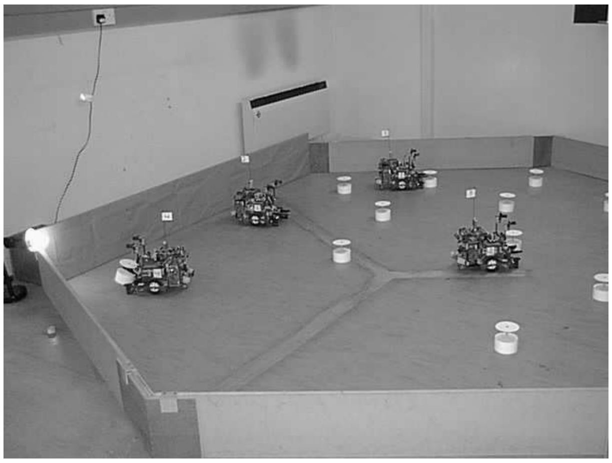

Fig. 5.: Four robots performing the cooperative task.

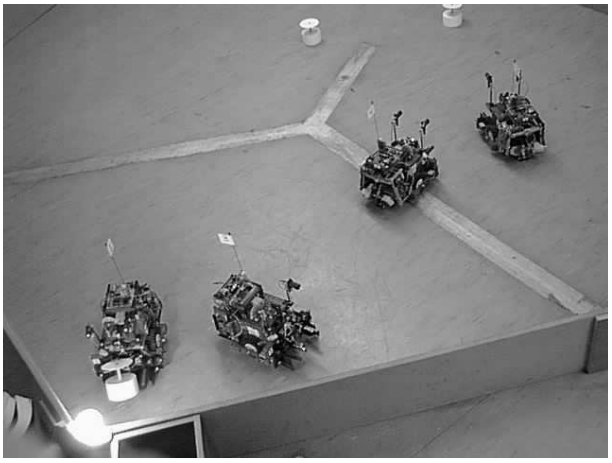

Fig. 6.: A robot bringing a target to the nest while the rest are leaving.

4 Results

First we will analyse the individual performance of each robot while performing the whole task by their own. Robots start in position C as shown in figure 4. In figure 7 we can see an histogram for each robot representing the mean and standard deviation of the time spent in searching for and gripping each target. The full data collected from the experiments can be seen in [6].

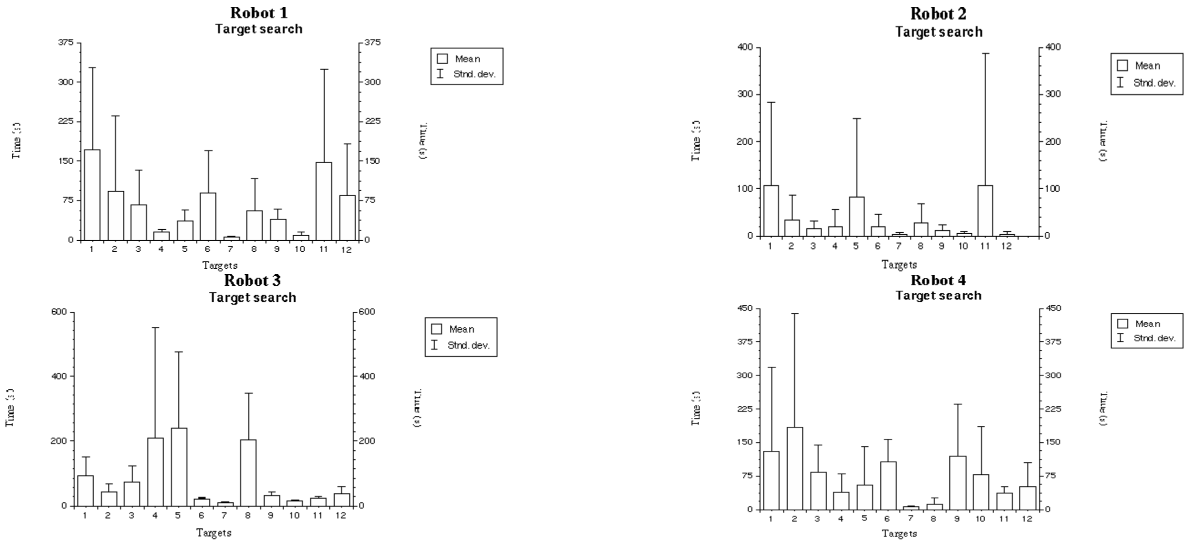

Fig. 7.: Mean & standard deviation of the time spent by each robot to find the targets.

In these histograms we see that there are targets which are more easily found[7] than others. Such is the case of targets 7 & 10 compared with target 1. However, it is important to clarify that the reason for targets 7 & 10 being so easily found is because of their location in front of the starting position C. The nearer a target is from the nest, the faster it is found.

with one single robot.

[6] Actually, no differences between collective and cooperative behaviour are visible in the pictures since, as we have seen, the difference between them is in the use of communication.

[7] The time for a target to be found starts counting with the robot positioned at the nest, either because of bringing a target or because of starting the task from its initial position.

As well, the standard deviation of time to find most targets is large, as a result of the unpredictable trajectory of the robots after placing a target on the nest. Another important result is the time spent by some robots in finding particular targets. The fewer targets left in the arena, the greater the amount of time needed for a robot to find one of these targets.

The mean and standard deviation of the time spent in carrying each target to the nest is shown in figure 8. As we can see in this figure, in most of the cases the time needed to bring a target from the place where it was found to the nest is significantly lower than the time needed to find such target. This happens because the robot goes directly towards the nest (after raising a target and detecting the light bulb) interrupting its performance only if an obstacle is detected before arriving to the nest. Therefore, the time needed is in most cases less than in the case of searching for targets.

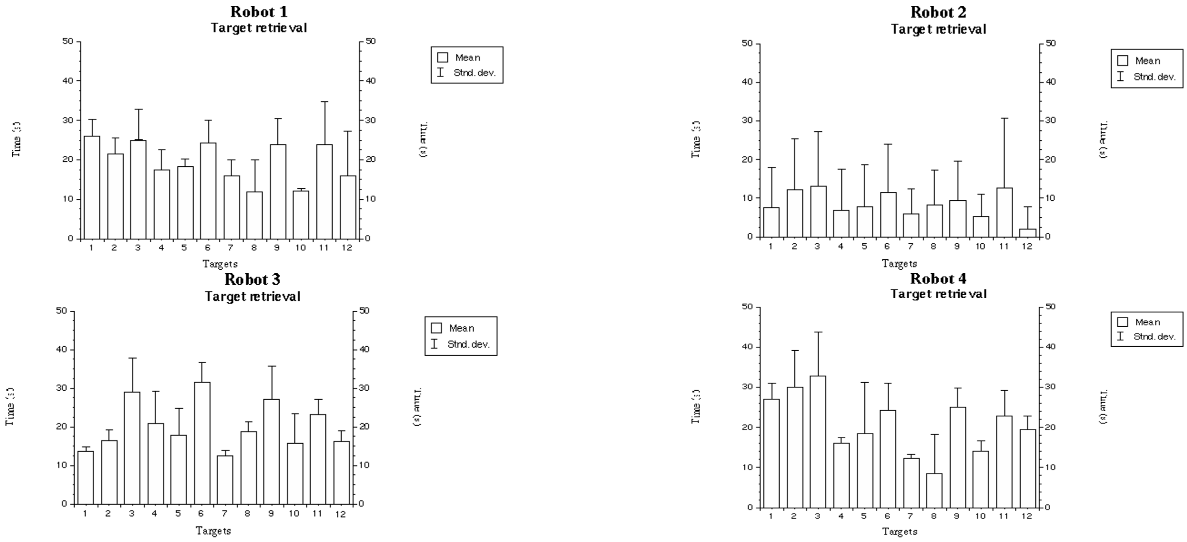

Fig. 8.: Mean & standard deviation of the time spent by each robot to place the targets on the nest. Note the $\frac{1}{10}$ (approx.) vertical scale relative to figures 7 & 9.

Finally, figure 9 shows the mean and standard deviation of the total time spent by a robot handling each target, that is, the sum of the target search and retrieval.

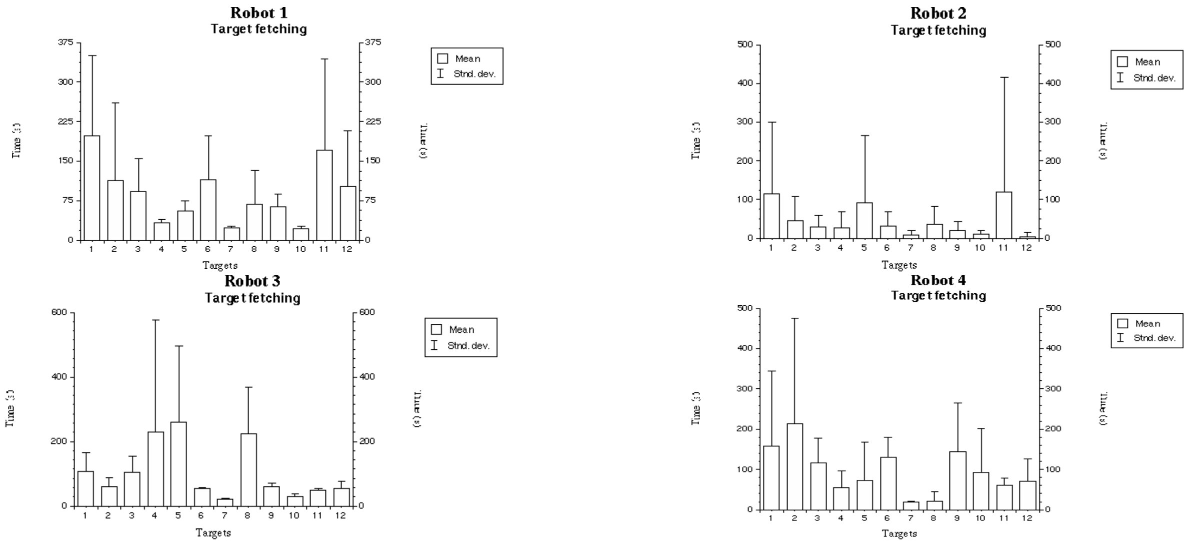

Fig. 9.: Mean & standard deviation of the total time spent by each robot to find and place the targets on the nest.

Statistical analysis of the results

We are interested in comparing the time spent by the robots accomplishing the collective task and the cooperative task with minimal communication. As we have a relatively small sample size we will analyse the results applying the t Test in two different versions: the Two-Sample and the Paired Sample [8]. For each of these tests, we will consider a p value between 10% and 5% as evidence of significant difference between the pair of samples. For p values $\leq 5\%$ we will consider that there is a strongly significant difference between the sampled pair. For $p > 10\%$ we will consider the pair of samples as similar in performance.

We use the Two-Sample test for comparing performance between different collective experiments, and the Paired Sample test to compare each collective experiment with its matching cooperative one.

Collective vs. Cooperative behaviour

The results of doing the Two-Sample t Test with the collective set of experiments is shown in table 1. The key for the results is the following: "**-**" No significant difference ($p > 0.1$), "**s**" Significant difference ($p < 0.1$) and "**S**" Very significant difference ($p < 0.05$)

Comparison	t value	N	p value	Result
1 Robot vs. 2 Robots	0.9265	24	> 10%	-
2 Robots vs. 3 Robots	0.9386	16	> 10%	-
3 Robots vs. 4 Robots	1.5417	12	10% - 5%	s
1 Robot vs. 3 Robots	3.5179	24	0.1% - 0.05%	S
2 Robots vs. 4 Robots	2.2643	12	2.5% - 1%	S
1 Robot vs. 4 Robots	3.4941	20	0.5% - 0.1%	S

Table 1.: Collective task Two-Sample t Test results

	1 Robot	2 Robots	3 Robots	4 Robots
1 Robot	X	-	S	S
2 Robots	-	X	-	S
3 Robots	S	-	X	s
4 Robots	S	S	s	X

Table 2.: Collective task comparison diagram.

As we can see on the table, in the case of 1 Robot vs. 2 Robots and 2 Robots vs. 3 Robots, there is no significant difference in performance. However, as we continue increasing the number of robots (up to four), performance continues improving monotonically as an evidence of significant difference. Also, there are comparisons between 1 Robot vs. 3 Robots, 2 Robots vs. 4 Robots and 1 Robot vs. 4 Robots. In all these cases, there are very significant differences in performance. A performance comparison diagram of the collective task can be seen in table 2.

For cooperative behaviour we can state the same as in collective behaviour that is, increasing the number of robots improves performance. The key is the same as in the collective case. As we can see in table 3, there is more significant difference in performance in the case of 2 Robots vs. 3 Robots than in the case of 3 Robots vs. 4 Robots. This differs from the collective case as we saw before.

Comparison	t value	N	p value	Result
2 Robots vs. 3 Robots	3.6383	16	0.5% - 10%	S
3 Robots vs. 4 Robots	1.7896	12	10% - 5%	s
2 Robots vs. 4 Robots	5.1730	12	< 0.05%	S

Table 3.: Cooperative task Two-Sample t Test results

	2 Robots	3 Robots	4 Robots
2 Robots	X	S	S
3 Robots	S	X	s
4 Robots	S	s	X

Table 4.: Cooperative task comparison diagram.

A performance comparison diagram of the cooperative task commented above can be seen in table 4, which gives a visual representation of the statistical significance for each of the experimental settings.

Looking at table 5 we can see that the most significant difference in performance between collective and cooperative behaviour corresponds to the case of 3 Robots (1% - 0.5%). Nevertheless, there is significant difference in all cases which can be understood as a better performance of the cooperative behaviour with respect to collective behaviour for the task performed.

Case	t value	N	p value
2 Robots	2.9195	8	2.5% - 1%
3 Robots	3.4062	8	1% - 0.5%
4 Robots	4.1791	4	2.5% - 1%

Table 5.: Comparison of matching collective vs. cooperative tasks.

5 Conclusions

As we saw in section 2, communication between agents is used when performing cooperative behaviour. This can be extended to this work since, due to the task designed, the robots need to know when a task is finished, that is, the terminating signal which tells the robots when the 12 targets are placed on the nest. This couldn't be done without communication. As we stated before, the information shared by the robots in our task is minimal. From the results obtained we can draw the following conclusions:

- Having more robots improves performance in our experiments. However, performance will eventually start decreasing with a number of robots greater than four.
- Cooperative behaviour is more efficient than collective for the task studied in this work.
- Minimal communication is needed to perform the cooperative task (*i.e.* no versatile inter-robot communication system is needed).

As a general conclusion we can state that a small amount of communication, as the target retrieval information shared by the robots in our experiments, buys an improvement in performance in a common task.

6 Further work

Increasing the number of robots to perform the task is one of the things left to do. This will allow us to determine the maximum number of them before their performance starts decreasing.

Also, it might be possible to use a similar strategy to improve the performance of the individual agents working on the task, for instance by focusing activity in areas of need such as target localisation and density. This can be done by incrementing the information shared among the robots by, for example, determining their positions while navigating about the arena, and sending these positions to the rest of the robots when relevant. It is important to notice, however, that this increment in communication will still needing a terminating signal in order to make the common task effective and finish at some point.

Another possible improvement is to design a learning task, that is, experiment with learning techniques such as reinforcement learning and compare the results obtained with ours. The learning task could consist in a robot being rewarded or punished according to its actions inside the arena, and reporting the information obtained to the rest of the group which will perform the learnt task.

7 Acknowledgments

Thanks to Esther Dura for comments on previous versions of this paper. This work was supported by the European Union TMR Network SMART2 under contract number FMRX-CT96-0052.

References

1. K. Altenburg and M. Pavicic. Initial results in the use of inter-robot communication for a multiple, mobile robotic system. In *Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots*, pages 96–100, Chambery, France, 1993.
2. R. Beckers, O.E. Holland, and J.L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, 1994.
3. R.A. Brooks. A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, 2(1):14–23, 1986.
4. R.A. Brooks. Intelligence without reason. In *Proceedings of the Twelfth International Conference on Artificial Intelligence*, volume 1, pages 569–595, 1991.
5. Y. Uny Cao, Alex S. Fukunaga, and Andrew B. Kahng. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4(1):7–27, March 1997.
6. Jose M. Carmena. Cooperative tasks need versatile inter-robot communication systems? Master's thesis, Department of Artificial Intelligence, University of Edinburgh, 1998.
7. Jose M. Carmena, John Hallam, and Juan Domingo. A versatile inter-robot radio communication system. DAI Technical Paper (in preparation), Department of Artificial Intelligence, University of Edinburgh, 1999.
8. Paul R. Cohen. *Empirical Methods for Artificial Intelligence*. MIT Press, 1995.
9. Jim Dallas. Co-operative search behaviour in a group of lego robots. MSc dissertation, Department of Artificial Intelligence, University of Edinburgh, 1990.
10. John Hallam. Autonomous robots: from dream to reality. DAI Technical Paper 526, Department of Artificial Intelligence, University of Edinburgh, 1991.
11. John Hallam. Intelligent sensing & control course notes: Architectures section. Department of Artificial Intelligence, University of Edinburgh, 1997.
12. Chris Malcolm. Intelligent assembly systems lecture notes. Department of Artificial Intelligence, University of Edinburgh, 1998.
13. Maja J. Mataric. Interaction and intelligent behaviour. PhD thesis, MIT, EECS, 1994.
14. Maja J. Mataric. Reinforcement learning in the multi-robot domain. *Autonomous Robots*, 4(1):73–83, March 1997.
15. Maja J. Mataric. Using communication to reduce locality in distributed multi-agent learning. *Journal of Experimental and Theoretical Artificial Intelligence*, 10(3):357–369, Jul-Sep 1998.
16. L.E. Parker. Heterogeneous multi-robot cooperation. PhD thesis, MIT, 1994.
17. Peter Stone and Manuela Veloso. Communication in domains with unreliable, single-channel, low-bandwith communication. Submitted to ICMAS'98, 1997.
18. Ken Sugawara, Masaki Sano, Ikuo Yoshihara, and Kenichi Abe. Cooperative behaviour of interacting robots. *Artificial Life and Robotics*, 2:62–67, 1998.

This article was processed using the TEX macro package with SIRS99 style