

Division of Labor in a Group of Robots Inspired by Ants' Foraging Behavior

THOMAS H. LABELLA and MARCO DORIGO

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

and

JEAN-LOUIS DENEUBOURG

CENOLI, Université Libre de Bruxelles, Brussels, Belgium

In this article, we analyze the behavior of a group of robots involved in an object retrieval task. The robots' control system is inspired by a model of ants' foraging. This model emphasizes the role of learning in the individual. Individuals adapt to the environment using only locally available information. We show that a simple parameter adaptation is an effective way to improve the efficiency of the group and that it brings forth division of labor between the members of the group. Moreover, robots that are best at retrieving have a higher probability of becoming active retrievers. This selection of the best members does not use any explicit representation of individual capabilities. We analyze this system and point out its strengths and its weaknesses.

Categories and Subject Descriptors: I.2.9 [Artificial Intelligence]: Robotics

General Terms: Algorithms, Performance

Additional Key Words and Phrases: Adaptive systems, ant algorithms, bio-inspired systems, adaptation

1. INTRODUCTION

Nature is full of examples of animals that can cooperate efficiently in big groups. Ants are probably the most cited: they can collectively retrieve big prey, adapt

This work was supported by the "SWARM-BOTS project", funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant IST-2000-31010. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication. Marco Dorigo acknowledges support from the Belgian FNRS, of which he is a Research Director, through the grant "Virtual Swarm-bots", contract no. 9.4515.03, and from the "ANTS" project, an "Action de Recherche Concertée" funded by the Scientific Research Directorate of the French Community of Belgium. Jean-Louis Deneubourg acknowledges support from the Belgian FNRS, of which he is a Research Associate.

Authors' address: T. H. Labella, IRIDIA—CP 194/6, Université Libre de Bruxelles, Av. F. Roosevelt 50, B-1050, Bruxelles, Belgium; email: h1abella@ulb.ac.be.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 1515 Broadway, New York, NY 10036 USA, fax: +1 (212) 869-0481, or permissions@acm.org.

© 2006 ACM 1556-4665/06/0900-0004 \$5.00

to the richness of the environment, or find the shortest path to food [Hölldobler and Wilson 1990]. Other examples are termites that are able to build big and complex nests, and bees that collectively choose a new location for their nest. Amazingly, direct communication is seldom observed: these insects use only locally available information, exploit features present in the environment, and use indirect communication, for instance *stigmergy* [Grassé 1959].

Such collective insect behavior tends to be robust in the face of uncertainty and environmental change. Also noteworthy is the fact that in most cases the individual group members seem to lack an internal map of the environment. Nor do the individuals in any sense seem to have a global plan to which they adhere. Indeed the individuals have a limited knowledge and perception of their environment. The results of collective insect behavior are usually well beyond the capabilities of a single insect. The key mechanism at work tends to be that of self organization [Camazine et al. 2001]. Researchers in the field of Multi-Robot Systems look with interest at such examples because they might suggest new ideas for developing robust and adaptive robotic control algorithms. At the same time, biologists see robotics as a useful experimental tool for checking and validating their theories [Webb 2000].

In this article, we study the collective behavior of a group of robots, both real and simulated, performing an object search and retrieval task. The control algorithm of each individual was inspired by a model of ants' foraging, developed while observing the behavior of *Pachycondyla apicalis* (previously named *Neoponera apicalis*). The robots we used for our experiments are quite simple. They have very limited computational power, they do not communicate with each other and they are equipped with simple sensors. The sensors are too simple to allow them to build a map or any other model of the environment. Nevertheless, we show in this article that they are able to cooperate in order to increase the efficiency of the group. The control algorithm we use induces division of labor at the group level. The controller also encourages a selection mechanism at the group level—robots more suited to the task are more likely to carry out the task than less capable robots. These results are achieved in a distributed fashion, without the use of explicit models either of the environment or of the capabilities of the robots.

This work is part of a larger research agenda. We started from implementing in real and simulated robots a model of ants' foraging. This is studied, to the best of our knowledge, only theoretically in the literature. The first steps, discussed in this article, are the validation of the theoretical model and the analysis of its features. Some of the results might be easily foreseen from the model, but they need a formal validation to be useful. The results of the analysis will then be used to perform comparisons with similar algorithms that are found in the literature and eventually to improve and test them in other contexts. These last steps are ongoing work and are not presented here.

The analytical methodology used in this work makes extensive use of statistical tools and concepts (e.g., nonparametric tests and the blocking design of experiments). In fact, our experiments were designed in such a way as to obtain statistically significant results even when using a limited number of runs. This is useful especially when experimenting with real robots, a highly

time-consuming activity. While such concepts are extensively used in other disciplines, we are aware of very few works that make use of this methodology in robotics or that at least explicitly mention it. We hope that this article can be an inspiration for other researchers to improve the efficiency of their experimental work.

The next section describes the task of our robots. Section 3 fixes some issues with the terminology, Section 4 explains the methods that we used for our experiments that are reported and discussed in Section 5. Section 6 summarizes related works in the literature, and Section 7 draws some conclusions and indicates the directions of our future research.

2. THE APPLICATION DOMAIN: PREY RETRIEVAL

The test application we chose for our experiments is object search and retrieval: a group of robots has to look for objects, termed *prey*, spread in the environment and retrieve them to a special area, termed *nest*. In this article, we refer to this task as *prey retrieval* to emphasise the similarity with the corresponding behavior observed in ants.

Prey retrieval is often used as a model for other real-world applications (such as toxic-waste cleanup, search and rescue, demining or collection of terrain samples in unknown environments) and is among the canonical tasks for collective robotics [Cao et al. 1997]. The main scientific question is whether there is an actual performance gain in using more than one robot since the task can be accomplished by a single one [Cao et al. 1997]. Stated in another way, the question is whether more robots are more efficient.

There are two components that must be taken into account in a prey retrieval task in order to measure the efficiency of the group. On the one hand, retrieved prey may be seen as an *income* for the group. On the other hand, searching and retrieving also has drawbacks that can come from dangers in the environment, from the interferences among robots (such as blocking the way to others, or collisions that slow down their speed), or from the fact that robots spend energy while moving. All these are *costs* for the group.

Income and costs depend on the number, X , of robots that are exploring the environment. Both income and costs increase with X but not in the same way. The income saturates when X becomes too high (robots can not retrieve more prey than their actual number in the environment), but costs potentially increase without limit. A straightforward definition of *efficiency* is

$$\eta = \frac{\text{income}}{\text{costs}}. \quad (1)$$

We indicate as *performance* of the group the number of retrieved prey items, that is, the income.¹

¹Note that the words “performance” and “efficiency” have been used with different meanings in the robotics literature. For instance, in some works, performance refers to the time it takes to retrieve all the prey in the environment, while in others it refers to the number of retrieved prey. The choice of which definition to use depends, in our opinion, on the particular application the researcher has in mind: for instance, time is an important factor when considering search and rescue applications, but the number of collected items is a more interesting measure in the case of spatial exploration.

We see three possible strategies for increasing the efficiency of a group of robots performing prey retrieval: increase the income by providing better sensors or better search strategies; decrease the costs, for example, by introducing communication in order to reduce interference; use an optimal number of robots.

The third strategy, the one that we follow in this article, comes from the observation that if $X = 0$, the income is null and so is the efficiency (costs are never null because they also include the energy spent by the robots in the nest). When X becomes too high, costs increase faster than income, and the efficiency tends towards 0. Therefore, there is a value \hat{X} that maximises η . We refer to the mechanism that tunes the number of the robots involved in the retrieval task as *division of labor*.²

The optimal number of robots can be estimated a priori if the characteristics of the environment are well known and fixed as done, for instance, by Hayes [2002]. Otherwise, the group of robots should use some form of adaptation in order to cope with uncertain and dynamic environments. The robotic literature already offers algorithms for learning and adaptation (some of them are described in Section 6), but we are at the moment interested in studying the properties of the biological model described in the following section.

2.1 A Model of Ants' Foraging

Although many aspects of ants' foraging behavior are still under study, we can summarize its main features as follows [Detrain and Deneubourg 1997; Hölldobler and Wilson 1990]:

- (1) Ants randomly explore the environment until one of them finds a prey. If the prey is not too heavy, the ant tries to pull it to the nest; otherwise, it tries to cut it or to use short or long range recruitment (by spreading some chemical substance in the environment or going back to the nest while leaving a pheromone trail).
- (2) The prey is pulled straight to the nest (pushing is never observed), both in the case of individual or collective retrieval.³
- (3) After the retrieval, the ant returns directly to the location where it found the prey.

It has been noted that the foraging behavior of a single ant may be influenced by several factors like age or genetic differences. Among the few authors that pointed out that learning might play a key role, we refer to the work by Deneubourg et al. [1987]. They model each ant of the species *Pachycondyla apicalis* as an agent that has a probability to leave the nest P_1 that changes

²In the robotics literature, the terms “division of labor” and “task allocation” are often used synonymously. We see, however, some important differences in their use. Division of labor is more typical for bio-inspired studies and the problem is usually “how many robot should perform a single task?”. Task allocation is often used in more traditional robotic studies. In this case, there are many tasks which can be executed by one or more robots and the problem is to find the best assignment robots/tasks once the robots' qualities are known (see Gerkey and Mataric [2004]).

³In this article, we consider only the simple case in which one robot can retrieve one prey without either recruitment of other robots or dividing up (cutting) the prey.

according to previous successes or failures. If an ant retrieves a prey, its P_1 increases by a constant Δ . If it spends too much time without retrieving a prey, its P_1 decreases by the same constant. P_1 is always kept between a minimum and a maximum value. They show, through numerical simulations, that this model can explain the division of labor observed in ants and how the colony is able to adapt to the environment.

The model that Deneubourg et al. [1987] propose emphasizes the role played by learning during the lifetime of an individual. No direct communication, not even signaling, is necessary for the colony to adapt to the environment.

The simplicity of this division of labor model indicates that it might be well suited for use in the context of group robotics. This is the key idea we study in this article. We have two goals. To date, the model has only been tested using purely numerical simulations. Our first goal is, therefore, to subject the model to further testing in a more real environment. Our second goal is to provide a deeper level of analysis with a view to identifying the salient features of the model that can then be reused in other contexts.

3. TERMINOLOGY

There are a few terms that we need to clarify before proceeding with our discussion. In fact, these terms are used with different meanings by roboticists and biologists. Moreover, some of the definitions are still a subject of debate in the research community. It is out of the scope of this work to argue for one definition or another: here we limit ourselves to clearly stating what we mean when we use these terms in this article.

We have so far used the word “learning” only when discussing ants. For a biologist, learning is in fact a behavior observed in animals. It produces a “durable modification of [another] behavior in response to information acquired from specific experiences [within a given time scale]” [Alcock 1995]. Researchers in robotics tend to associate learning with Reinforcement Learning (RL), a field which studies “how to map situations to actions so as to maximise a numerical reward signal” [Sutton and Barto 1998].⁴ As the setup for our experiments and the algorithm we study are far simpler than those typically studied in RL, we use the term “adaptation”, which seems more appropriate and avoids confusion.

In this article, we use the concept “indirect” communication, but unfortunately there is still a lot of debate on how to define and classify the different forms of communication. By indirect communication, we mean that there is nothing in the robots’ controller that an observer can interpret as an act of transmitting information to other robots. In our experiments, robots use neither wireless communication nor signaling. Nevertheless, they do communicate by modifying the environment, for instance, by retrieving a prey.

It might also be argued that what we show is more a form of “specialization” than of division of labor. The former has indeed a meaning in biology that is different from what we show in Section 5.2, but we use it in a robotics context to refer to the process of tuning the number of robots performing a task. Moreover,

⁴RL is in fact a particular mechanism of implementing learning in the biological meaning.

some of the definitions given for specialization⁵ imply some form of individual adaptation (behavioural or physical) in order to better perform a task. This is not what we show either: some of our robots will perform a task more and more often. However, how well the individual robot performs the task does not change.

4. METHODS

This section describes the tools and the methodology that we used in our experiments. Before doing so, we briefly explain the rationale behind our choices. We used two different set-ups: real robots, built with Lego MindstormsTM, and simulated *s-bots*.⁶ This choice was driven by the concepts explained later and by the tools available at the time of experimentation.

Different set-ups allow us to draw conclusions that are less dependent on particular experimental conditions and that are potentially more generally valid. The conclusions have a more qualitative than quantitative nature, given that different platforms with different characteristics are compared.

As we wrote previously, we need to validate a theoretical model. Real robots are preferable for this purpose in order to avoid erroneous validation due to ill-chosen simulation models or parameters. Once we know that the model is valid through experiments with real robots, we can use simulation after verifying it against the empirical evidence brought by the real robots. The advantage of simulation is that it speeds up analysis because it can produce more data in shorter time.

We designed the experiments in order to exploit the best characteristics of each set-up. With real robots, we validate the theoretical model and define the path for further analysis. With simulation, we first crosscheck the results with those obtained with real robots before continuing with more in depth analysis.

4.1 Real Robots

The robots, which we call *MindS-bots* (Figure 1(a)), were built using Lego MindstormsTM. The CPU is a Hitachi H8300-HMS 1MHz, with 32Kb for both the operating system and the user program. Figure 1(a) shows a top view of a *MindS-bot* and its sensors. Two light sensors are used to sense prey, which are black cylinders, and the nest, which is spotted in the experimental arena by a light above it. Two bumpers are used for obstacle avoidance. The CPU controls two tracks and a gripper. Random numbers were generated using the random number generator of BrickOS (the operating system of the robots), which is “inspired by Press et al. [1992, 279–281].⁷ Detailed instructions to replicate the *MindS-bot* can be found in Labella [2003].

⁵For example, the following is taken by the Merriam-Webster English dictionary: “structural adaptation of a body part to a particular function or of an organism for life in a particular environment”.

⁶*S-bots* are small autonomous robots produced within the SWARM-BOTS project [Dorigo et al. 2004; Mondada et al. 2004]. See also the Web site of the project at <http://www.swarm-bots.org>.

⁷From BrickOS source code. See <http://brickos.sourceforge.net/>.

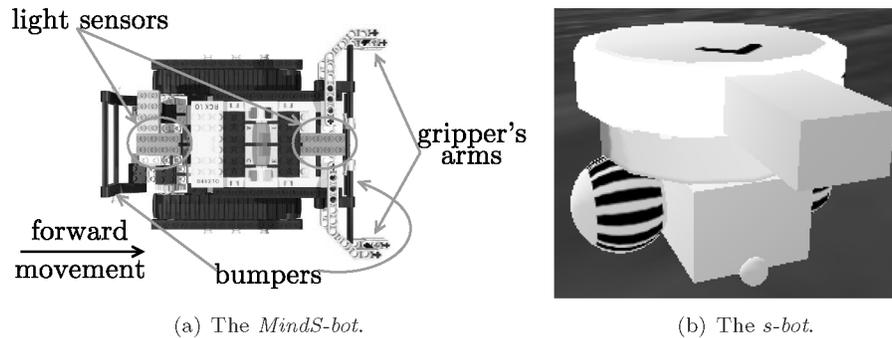


Fig. 1. The robots, real and simulated, used for the experiments. Left: top view of a *MindS-bot*. The circled bricks are light sensors and are used to locate prey and nest. There are two bumpers, one on the front and one on the back, to avoid obstacles. The gripper is made of two arms that close by rotating around two vertical axes relative to the *MindS-bot* (coming out of the picture). Right: an *s-bot*. It has simulated light sensors placed around the main circular body and an omnidirectional camera (not shown in the picture). The frontal parallelepiped is the gripper, which can stick to the prey for retrieval.

4.2 Simulation

The simulator used for the experiments, named *swarmbots3d*, was developed by the members of the SWARM-BOTS project and is described by Mondada et al. [2004] in more detail. It is a dynamics simulator, based on Vortex, a commercial dynamics engine.⁸ It simulates the *s-bots*, robots that can explore the environment autonomously and connect to each other whenever the environmental contingencies require it. Among other features, *swarmbots3d* can simulate the *s-bots* with different levels of detail, ranging from a simple cylinder with wheels to a nearly exact replica of the real *s-bot*. Obviously, the more detailed the simulation is, the slower it runs. Figure 1(b) shows the model used in our simulations.

Simulated light sensors placed around the main body and an omnidirectional camera, which has limited perceptual range, are used to locate prey and nest. Infrared sensors, also around the main body, are used to detect and avoid obstacles. The simulation of the gripper is somewhat simplified but still functional: the parallelepiped in front of the *s-bots* can stick and unstick to the prey by dynamically creating and deleting a joint between the *s-bots* and the prey. Random numbers are generated using the Mersenne Twister algorithm in the implementation of the GNU Scientific Library.⁹

4.3 Control

Figure 2 depicts the control program both of the *MindS-bots* and of the *s-bots*. Different states are the different phases of prey retrieval, that is, the subtasks in which the overall prey retrieval task is decomposed. These subtasks are as

⁸<http://www.cm-labs.com>.

⁹<http://www.gnu.org/software/gsl/>.

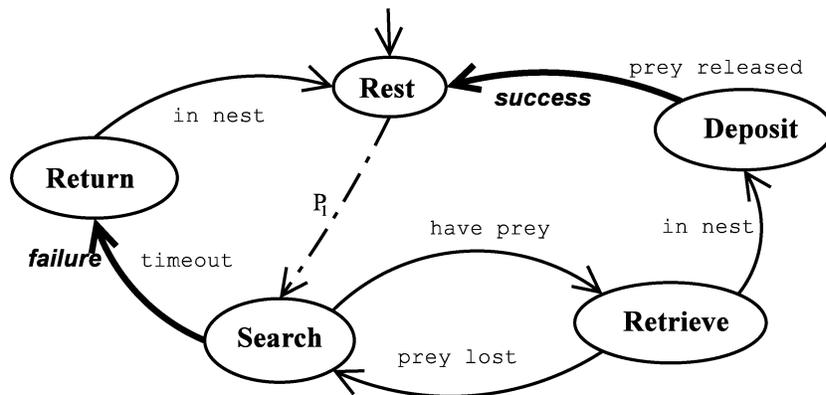


Fig. 2. Sketch of the control system of the robots. The states represent different phases of retrieval (see text). The labels on each edge represent the conditions under which the robots change state. The robots start in the **Rest** state. The transition from **Rest** to **Search** (dash-dotted) is based on the probability P_1 . The transition from **Deposit** to **Rest** (bold arrow) represents a successful retrieval (P_1 is increased), the one from **Search** to **Return** (also a bold arrow) is a failure (P_1 is decreased).

follows:

- Search** The robot looks for a prey, randomly exploring the environment and changes direction when it encounters an obstacle. If a prey is found, the robot grasps it. If a timeout occurs before the robot has grasped a prey, the robot gives up searching.
- Retrieve** The robot looks for the nest and pulls the prey toward it.
- Deposit** The robot leaves the prey in the nest and turns to the point from which it came (to mimic ants' behavior).
- Return** The robot looks for the nest and returns to it.
- Rest** The robot rests in the nest.

Transitions between states occur when the labels on the edges in Figure 2 are true except the one from **Rest** to **Search** which occurs with probability P_1 every second. More details about the implementation of the behaviors are given in Labella [2003].

4.4 Adaptation

The robots change from **Rest** to **Search** with probability P_1 each second whose value is updated during the transitions from **Search** to **Return** (henceforth called *failure*) and from **Deposit** to **Rest** (henceforth called *success*). The probability to leave the nest P_1 is adapted in a fashion similar to the one described by Deneubourg et al. [1987] and is shown in Algorithm 1. While Deneubourg et al. [1987] increment and decrement P_1 by a fixed constant Δ , in our algorithm, Δ is multiplied by the number of consecutive successes or failures. Also in our case, P_1 is bounded in the range $[P_{\min}, P_{\max}]$.

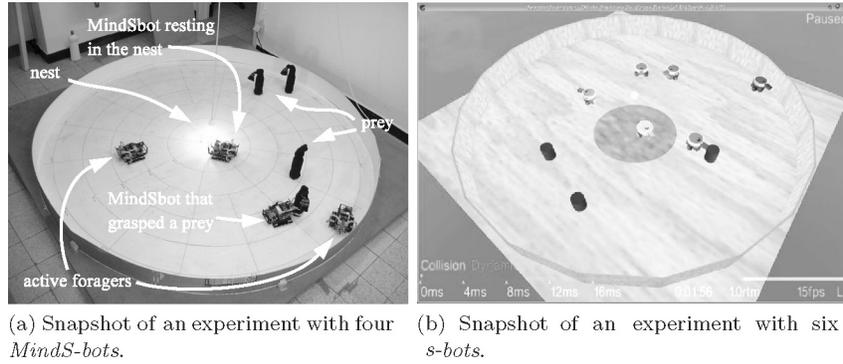


Fig. 3. Set-up of the experiments. The nest is indicated by a light in the center. The robots in the nest are resting and not active. The other robots are searching the environment, and one in each picture has found and is retrieving a prey.

Algorithm 1 Adaptation rule of P_1 , the probability to leave the nest. The variables *succ* and *fail* are the number of consecutive successes and failures.

```

succ ← 0; fail ← 0; P1 ← Pinit;
if success then
    succ ← succ + 1
    fail ← 0
    P1 ← min{Pmax, P1 + succ · Δ}
fi
if failure then
    succ ← 0
    fail ← fail + 1
    P1 ← max{Pmin, P1 + fail · Δ}
fi

```

It is out of the scope of this article to discuss different methods for adapting P_1 , but the origin of the difference between the method studied by Deneubourg et al. [1987] and ours deserves some explanation. In fact, some preliminary tests showed that the dynamics of adaptation would have been too slow if the original algorithm had been used [Labella 2003]. With our modification, individuals adapt faster, although the experiments with real robots still took 40 minutes each.¹⁰

4.5 Experimental Set-Up

For the experiments, we used a circular arena (Figure 3) with a diameter of 2.40 meters. A light bulb is used to signal the position of the nest in the center of the arena. Walls and floors are white in the experiments with the real robots, prey are black. The search timeout is fixed to 228 seconds for the *MindS-bots* and 71.2 seconds for the *s-bots*.¹¹ P_{\min} is set to 0.0015, P_{\max} to 0.05, P_{init} to

¹⁰The length of the experiments is generally affected not only by the speed of adaptation, but also by other factors such as the speed of the robots and how effective their search for prey is. The latter are in fact the most limiting factors. For instance, the mean time for a *MindS-bot* to find and retrieve a prey is nearly 300 seconds. As an approximation, under the hypothesis that there is no timeout, that the *MindS-bot* is always successful, and that it never rests in the nest, the average number of prey retrieved in 40 minutes is 8. Therefore, in each experiment, P_1 can be updated a maximum 8 times.

¹¹These values are the estimated median times needed by a single robot to find one prey when it is alone in the arena. Notice that these values do not depend only on the speed of the robots, but

0.033, which correspond to a mean time spent in the nest of approximately 11 minutes, 20 seconds and 30 seconds respectively. Δ is set to 0.005. Prey appear randomly in the environment during the experiments. The probability with which this happens each second is referred to as *prey density* and changed between experiments. A new prey is placed randomly in the arena so that its distance from the center is in [0.5 meters, 1.1 meters]. We chose the values of these parameters based on some a priori reasoning about the dynamics of the system. The purpose was to find the set of parameter values that best shows the effects of adaptation. More detail about the choice of the parameters are given in Labella [2003].

4.6 Efficiency Index

It is not possible to directly measure the value of η (as given by Equation 1) because the costs as we defined them are hard to be (or cannot be) quantified. In fact, they comprise too many factors, some of which are generally unknown. Therefore, we decided to use the *group duty time* as an estimate of the costs. The group duty time is the sum of the time that each robot spent in searching or retrieving, that is, the time it was on duty. The group duty time is directly related to the costs: the higher it is, the higher the probability that some robot gets lost or breaks down, the higher the energy consumption, and so forth. Thus, as an efficiency index we used

$$\nu = \frac{\text{performance}}{\sum_{\text{robots}} \text{duty time}} , \quad (2)$$

where *performance* is the number of retrieved prey. It is easy to see that if ν increases, η increases too.

4.7 Experimental Methodology

The results of our experiments are strongly stochastic. The randomness induced by our particular adaptation algorithm adds to the natural noise coming from the sensors. The final result of an experiment is therefore a random variable with an unknown distribution. In order to draw sound conclusions, we need to give statistical relevance to the results.

Some of the experiments described in Section 5 are designed in order to identify the impact of varying a single experimental parameter (e.g., adapting vs. nonadapting robots). The experiments were designed according to the concept of block design [Montgomery 2000]. We generated a set of random *instances* before running the experiments. An instance, parametrized by prey density, describes where and when prey appear in the environment. Then, we tested different set-ups against the same instances. Finally, the results were compared using preferentially nonparametric statistical tests [Siegel and Castellan Jr. 1988]. This procedure often allows statistical significance to be demonstrated

also on their sensors. The *s-bots* can detect a prey in their surroundings with the omnidirectional camera more easily than the *MindS-bots* which need to have the prey in front of their light sensor to perceive it.

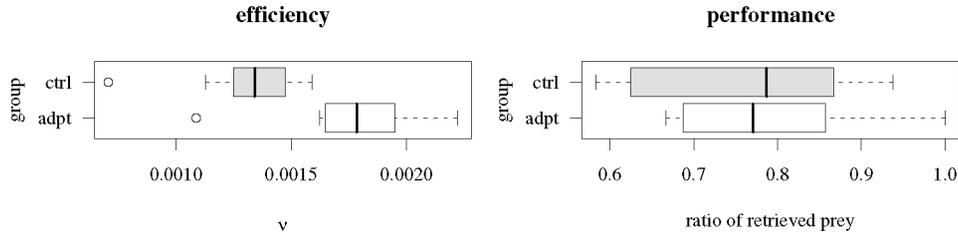


Fig. 4. Left: Value of ν (Equation (2)) when P_1 is adapted (bottom) and when it is not (top) in the *MindS-bots*. Right: relative performance of the adapting and control group (1 on the x axis means 100% of prey retrieved in an experiment). The right and left limits of a box extend from the first to the third quartile of the distribution of the results. The line in the box shows the median value. The whiskers extend to the most extreme data point which is no more than 1.5 times the interquartile range from the box. Circles are considered outliers.

with fewer experiments. This, in turn, greatly speeds up research, especially when working with real robots.

5. EXPERIMENTS AND RESULTS

We now analyze the effects of individual adaptation on the group of robots. We focus mainly on three features: efficiency improvement (Section 5.1), division of labor (Section 5.2) and selection of best individuals (Section 5.3). We also describe some of the limitations that we observed and these will be used as a basis for further research.

5.1 Efficiency

5.1.1 Real Robots. We used groups of four *MindS-bots* chosen out of a pool of six. The four robots were changed after each experiment.¹² Each trial lasted 2400 seconds (40 minutes). We created ten instances with prey density set to 0.006 s^{-1} . We repeated each experiment with a control group made of the same robots with P_1 fixed to 1 and using the same instances.¹³ Figure 4 summarizes the results: on the left side are the values of ν for both the adapting and the control group; on the right side, we show the ratio between the number of retrieved prey and the total number of prey that appeared during the experiment. When the robots used adaptation, there were 2.57 active robots and 2.44 prey on average in the arena in the period between 1000 seconds and 2400 seconds. In the control experiments, there were 3.63 active robots and 3.49 prey.

5.1.2 Simulation. The simulation experiments used groups which varied from 2 to 8 *s-bots* with increments of two units. The groups were tested with prey density equal to 0.005 s^{-1} , 0.01 s^{-1} , 0.02 s^{-1} , and 0.04 s^{-1} . We generated fifty instances to be used for each combination of prey density/group size. As a base result, we also tested a single robot on the same instances. The experiments lasted 2400 seconds (simulated time). As we did with the *MindS-bots*, the

¹²The choice was not completely random but was biased by the status of the *MindS-bots* after each experiment. For instance, those with low battery or those which needed some maintenance were taken out and new ones were added to the group.

¹³In both original and control experiments, a computer next to the arena signaled to the experimenter the time and the position of the new prey.

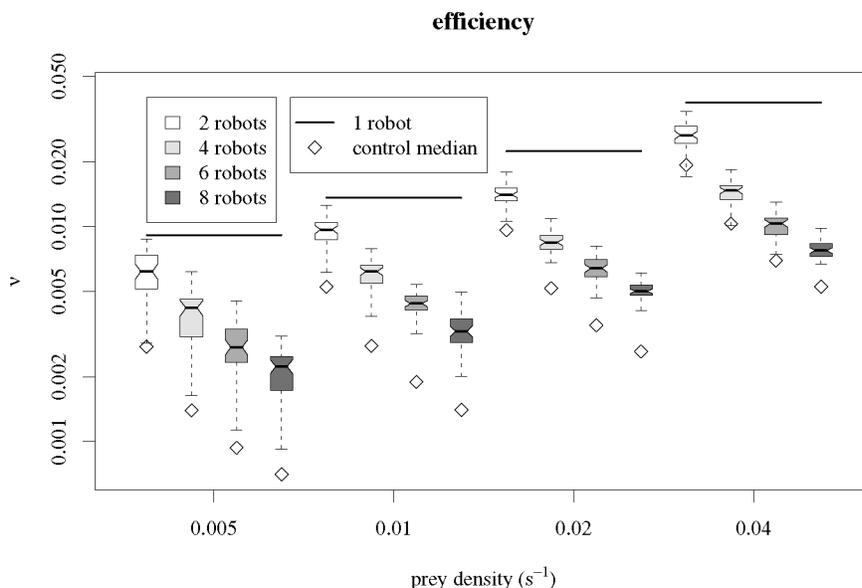


Fig. 5. Effects of prey density and group size on the efficiency of retrieval in the simulated *s-bots*. The plots report the results of fifty experiments for each prey density (on the x axis), for each group dimension (different filling of the boxes). Both the x and the y axes use a logarithmic scale. The horizontal black line shows the median value of the efficiency of one adapting robot tested in the same conditions as the other groups. The diamonds show the median value obtained in the control experiments. We show only the median values and not the whole distribution for the sake of readability of the plot.

experiments were replicated using the same instances but with a control group that did not use adaptation. The results are summarized in Figure 5 which reports the final distribution of the values of ν for different combinations of prey density and group size. Figure 6 reports the final distribution of the performances of the different groups in each environment.

5.1.3 Discussion. The group that uses adaptation is significantly more efficient both in the case of the real robots (after 1400 seconds)¹⁴ and of simulation¹⁵ (confidence level 95%).

There is no statistical difference in the performances between the two groups of *MindS-bots*,¹⁶ while in simulation, the control group performs better.¹⁷ In the latter case, the average difference of retrieved prey is 3.4 units, that is, a negligible amount with respect to the total.¹⁸ The group size seems not to have any effect on the performance of the group, but this is mainly due to the fact

¹⁴Sign test for paired data [Siegel and Castellan Jr. 1988, 80-87]: Null hypothesis: ν is the same in the two colonies. Alternative hypothesis: ν , and therefore the efficiency, improves with adaptation.

¹⁵Wilcoxon rank sum test. The null and alternative hypotheses are the same as in footnote 14.

¹⁶Permutation test, null hypothesis: the performances are the same. Alternative hypothesis: the performances are different. The p -value is 0.53.

¹⁷Wilcoxon rank sum test. The null and alternative hypotheses are the same as in footnote 16.

¹⁸The only exception is for prey density 0.04 s^{-1} and 2 *s-bots*, where the control group retrieved on average 8.9 more prey.

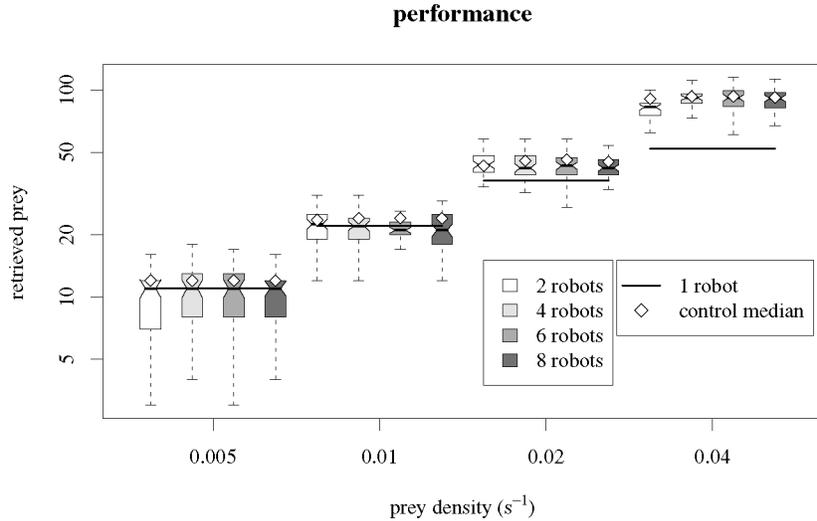


Fig. 6. Performance of different set-ups in simulation. The x and y axes are in logarithmic scale. The value plotted is the number of retrieved prey. For the meaning of the boxes and the other symbols, see Figure 4 and Figure 5.

that the relative number of retrieved prey, that is, the number of retrieved prey divided by the total number of prey, is on average already close to 1.

Both in simulation and with the real robots, the differences in performance, if any, are not enough to explain the differences in efficiency. Therefore, the improvements are due to the decrease of the group duty time. We show in Section 5.2 that this is achieved by division of labor.

We see in Figure 5 that the gap between the adapting group and the control group tends to decrease when the prey density increases. This is not surprising because it is better to use all available robots, as the control group does, in rich environments. It is important to observe that ν also decreases with the group also when adaptation is used. One possible explanation is that our adaptation algorithm is not good at reducing the number of explorers. For instance, if the optimal number of explorers is 2 for a given prey density, our robots might end up with 2.5 for a group of 4 *s-bots* and 3.5 for 8 *s-bots*. An alternative explanation is related to the way we measure ν , which depends on the group duty time. The latter depends on the group size. In fact, all the robots in a group spend some time in searching because none of them can have $P_1 = 0$ (remember that $P_1 \in [P_{\min}, P_{\max}]$ and $P_{\min} > 0$), so each robot contributes to the final group duty time. These two explanations are not mutually exclusive, but it is important to test if the first case does occur since this would demonstrate a limitation of our adaptation algorithm. Section 5.2 provides a partial answer to this issue.

5.2 Division of Labor

We showed that adaptation improves the efficiency of the group. The differences of the performances in both simulation and hardware when present are

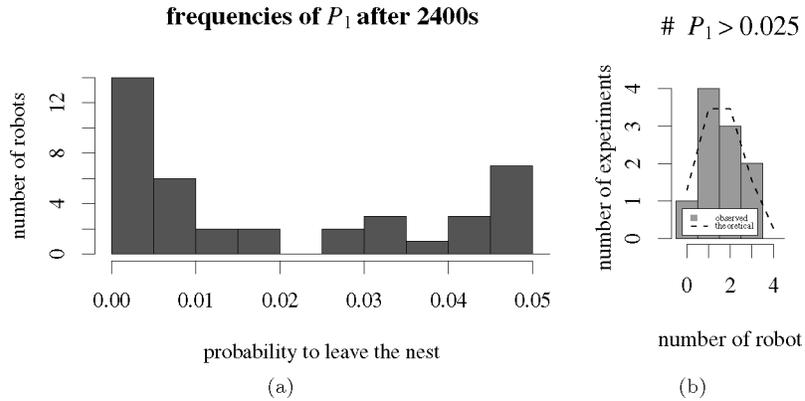


Fig. 7. (a) Frequency of P_1 observed in the *MindS-bots* 2400 seconds after the beginning of experiments. The two peaks demonstrate the occurrence of the division of labor. 40% of the observations are above 0.025. (b) Distribution of the number of *MindS-bots* with $P_1 > 0.025$ observed in each experiment compared with the theoretical binomial distribution with $p = 0.4$.

not enough to explain the improvement. Therefore, we deduce that adaptation reduces the group duty time. There are two ways in which the group can achieve this. In the first case, all robots end up having the same, albeit low, P_1 so that the mean number of robots in the foraging area is reduced. In the second case, only a few robots are active foragers with high P_1 and the others have low P_1 . Obviously, the robots with high P_1 would spend more time in searching than the others, therefore we could observe division of labor.

At any given instant t after the beginning of the experiment, the value of P_1 in a robot is a random variable which has different values for every robot and every experiment. Whether the group uses division of labor or not can be observed in the distribution of P_1 : in case division of labour occurs, then at the end of the experiments, the distribution of P_1 will have two peaks; otherwise it will have only one peak.

5.2.1 Real Robots. During the experiments of Section 5.1.1, we recorded the value of P_1 for each *MindS-bot* in order to estimate the distribution. The result after 2400 seconds plotted in Figure 7(a), clearly shows two peaks. Figure 8 reports the distribution of P_1 over time.

5.2.2 Simulation. We used the data from the experiments of Section 5.1.2 to analyze the effects of group size and prey density on the division of labor. We observed that the evolution of the distribution of P_1 is similar to Figure 8, for each combination of prey density and group size (data not shown). We also noticed a wider gap between the two peaks of the distributions.

We classify the robots into three classes: *foragers*, *loafers* and *undecided*. Foragers are those *s-bots* whose P_1 is bigger than 0.042, while loafers have P_1 lower than 0.007, and the rest are undecided (notice that this last group spans a range of values for P_1 that is five times bigger than the others). Figure 9 plots

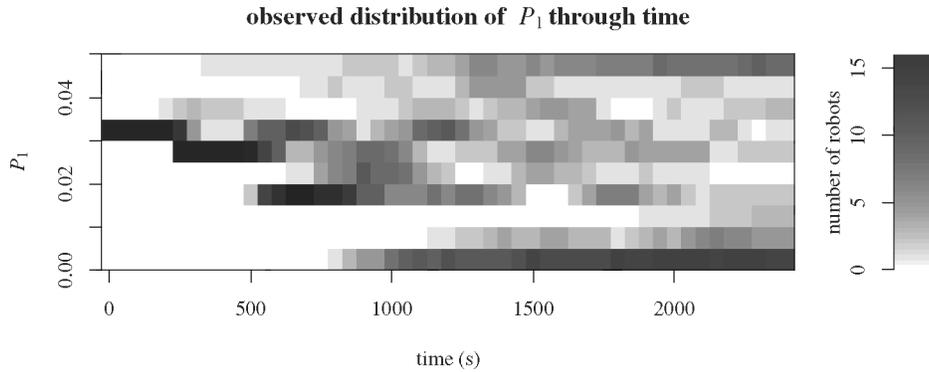


Fig. 8. Dynamics of the observed frequency of P_1 in the *MindS-bots*. The darkness of a cell in position (t, p) is proportional to the number of *MindS-bots* with $p = P_1$ after t seconds from the beginning of the experiment. The relationship is given by the bar on the right. At $t = 0$, all the *MindS-bots* have $P_1 = 0.033$ (see the black stripe on the left). After 1000 seconds the number of *MindS-bots* with low P_1 drastically increases (see the dark stripe on the bottom). Similarly, after 1500 seconds, the number of robots with high P_1 increases, although slowly and reaching a lower value than that of the loafers (top-right part of the plot).

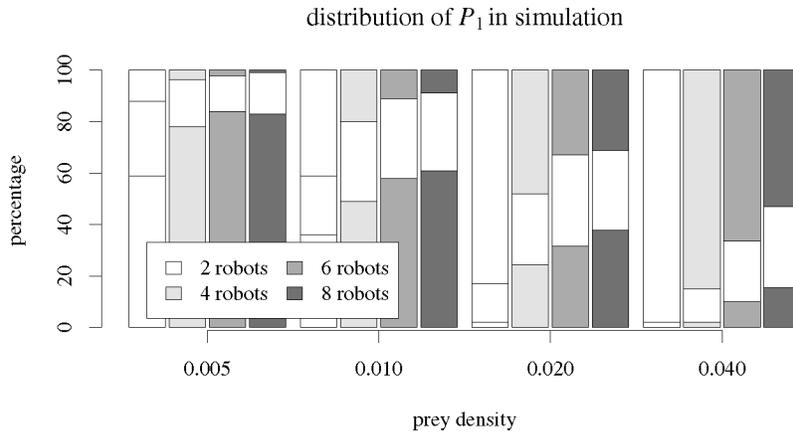


Fig. 9. Division of labor in the *s-bots*. Each group of four columns refers to different environments with increasing prey density. Each bar refers to a group size (see the legend). Each bar is divided into three parts whose height is proportional to the ratio of robots belonging to the following groups: foragers ($P_1 > 0.043$) on the top, loafers ($P_1 < 0.007$) at the bottom, and undecided ($0.007 \leq P_1 \leq 0.043$) in between. For example, if the top part is 25% of the total height of the bar for a group of 8 robots, it means that on average 2 robots were foragers.

the proportions of *s-bots* belonging to each class at 2400 seconds. The graphs clearly show a strong division of labor in the colonies. Individuals tend to have either high or low P_1 but seldom values in between.¹⁹

¹⁹The undecided category spans a broader range of P_1 than either the forager or loafer categories. Note that in Figure 9, therefore, even in cases where the percentages of undecided *s-bots* is similar to the percentage of foragers or loafers, the undecided *s-bots* are still more sparsely distributed over P_1 .

5.2.3 *Discussion.* One might object that the right peak of Figure 7(a) could be the result of a few experiments in which all the *MindS-bots* happened to have high P_1 . To see that this is not the case, it is enough to look at the number of *MindS-bots* with $P_1 > 0.025$ at the end of the experiments and how this number is distributed. From the data in Figure 7(a), we know that 40% of the population has $P_1 > 0.025$. Therefore, we expect that the number of *MindS-bots* with $P_1 > 0.025$ in each experiment follows a binomial distribution with $p = 0.4$. Figure 7(b) shows that the profiles of the theoretical and the observed distributions are indeed very similar.

The evolution of the distribution of P_1 over time (Figure 8) shows that the *MindS-bots* with high P_1 appear later than those with low P_1 (the former at 1500 seconds, the latter at 1000 seconds). It may be hypothesized that the presence of robots with low P_1 are necessary for the group with high P_1 to appear. All *MindS-bots* start with the same P_1 , as can be seen from the black stripe for $t = 0$ and $0.03 \leq P_1 \leq 0.035$. After some time, some *MindS-bots* reduce their P_1 because they are not successful (see the black stripes that start at 250 seconds and 500 seconds for $P_1 = 0.025$ and $P_1 = 0.015$), while the others alternate successes with failures (indicated by the region in the middle of y range that remains dark till 1500 seconds. The number of *MindS-bots* in the arena decreases, that is, there are less competitors for those that managed to keep their P_1 high enough. Less competitors implies more and easier retrievals, which increase the P_1 of the remaining foragers.

As we expected, the ratio of foragers in the *s-bot* group increases with higher prey density for a fixed group size. More interestingly, for most prey densities, the proportions of foragers for groups of six and eight *s-bots* are nearly the same, and thus there are on average more foragers in the latter group.²⁰ This phenomenon could explain the loss in efficiency when increasing the group size even when *s-bots* adapt. Because of the particular set-up we used, we could not test colonies with more than eight individuals: if more robots had been used, there would not have been room enough in the nest for all of them. We speculate that the particular rule we implemented, or the set of parameters we used, can be effective only to a certain extent and that other rules or other parameter settings could work better in such conditions.

5.3 Best-Individual Selection

The adaptation mechanism we are studying is based only on individual successes or failures. If one robot, for any reason, is better than the others for the task of retrieving, then it will be more successful, and therefore it is more likely to become a forager. We might, in fact, have come to this conclusion intuitively. It is still important, however, to validate it experimentally.

Generally speaking, the differences can be artificially created or intrinsic in the robots. In the first case, for instance, some robots can be intentionally designed for the task of retrieving, while the others are designed to explore the environment to find and mark dangerous spots. In the second case, the

²⁰This is true also for prey density 0.04 s^{-1} , where there are on average 4.23 foragers in a group of 8 robots vs. 3.98 in a colony of 6.

Table I. For each *MindS-bot*, identified by a unique name, the total number of experiments in which it was used and the number of times it was a forager ($P_1 > 0.025$) are reported. Data refers to ten experiment, four *MindS-bots* per experiment.

ID	Tot. Exp.	#foragers	ID	Tot. Exp.	#foragers
<i>MindS-bot1</i>	6	5	<i>MindS-bot4</i>	9	4
<i>MindS-bot2</i>	3	2	<i>MindS-bot5</i>	3	0
<i>MindS-bot3</i>	9	1	<i>MindS-bot6</i>	10	4

differences come from the imperfections of the robots' components which can never be identical (e.g., one motor that is less powerful than another). Any mechanism for division of labor should take into account this type of heterogeneity of the group.

This section shows that individual adaptation can be effective to select the best suited individuals for the retrieving task. We want to stress the fact that the algorithm we are studying does not take into account the presence of other robots. In fact, a robot neither knows how many nestmates are present nor whether it is working in a group or alone. There is no explicit model of the environment or of the robot's own capabilities.

The *MindS-bots* are built identically. The only differences come from the components. With the simulated *s-bots*, we artificially introduced some heterogeneity. In the following, we are interested in those robots whose P_1 is greater than 0.025. Slightly modifying our earlier definition, we refer to them as foragers.²¹ Given the stochastic nature of the experiments, we can model the fact that a robot i is a forager at the end of an experiment as a random event.

5.3.1 Real Robots. We took data from the experiments of Section 5.1 where we use groups of 4 robots selected out of a pool of $N = 6$ robots. Table I reports the number of times each *MindS-bot* was observed to be a forager at the end of the experiments.

5.3.2 Simulation. We created six different *s-bots* that differ in their maximum speed. More precisely, we set the maximum speed of the first *s-bot* to half the speed of the *s-bots* used in Section 5.1.2 and 5.2.2. The speed of each of the other five robots was the speed used in Section 5.1.2 and 5.2.2, scaled respectively at 0.7, 0.9, 1.1, 1.3, and 1.5. The six robots were combined into all possible colonies of four robots, forming fifteen different groups. Each group was tested in the same fifty instances randomly created with prey density 0.01 s^{-1} . The groups were simulated for 2400 seconds, and we counted how many times each *s-bot* in each group ended up being a forager. The results are not shown for reasons of space.

5.3.3 Discussion. If adaptation takes into account differences among individuals, the probability that each robot becomes a forager at the end of an experiment either depends on the other robots in the same group or is different

²¹Note that both with *MindS-bots* and *s-bots* the robots are not aware of such concepts as being a forager or being undecided. These are categories defined a posteriori by us, that is, they are arbitrary definitions, whose purpose is to help us in the discussion. Therefore, we can modify them, if this helps to better explain the results of our experiments.

for at least one robot. Adaptation does not consider differences among robots if the probability of becoming a forager is constant for all the robots.

For the *MindS-bots*, the data in Table I allows us to reject the hypothesis that the probability of becoming a forager is constant with confidence 95%.²² Thus one of the other two conditions can be assumed to be true. Both of them prove that adaptation considers individual differences for the division of labor.

This way of proving our thesis might seem complex but it has a great advantage. It allows us to reach significant conclusions with only 10 experiments. Consider the simpler approach of testing all possible groups with different instances. It would require 15 experiments (all the combinations of 4 robots out of 6) for each instance.

A two-way ANOVA test on the data from simulation shows statistically relevant differences both among robots and groups (confidence level 95%); therefore we can conclude that both individual characteristics and average abilities of the group are crucial to the selection of the best individuals. Not surprisingly, the probability of becoming a forager is always higher for faster robots.²³

6. RELATED WORK

When working with several robots, interference between them reduces the efficiency of the group [Balch 1999; Goldberg and Matarić 1997]. To reduce interference, smarter behaviors can be implemented as in Goldberg and Matarić [1997]. The authors estimate where the interferences occur the most by counting the number of collisions or manoeuvres to avoid other robots. With this data, they design a control algorithm that avoids the most problematic zones. Schneider-Fontán and Matarić [1996] reduce interference by assigning a predefined part of the arena to each robot and each robot to one particular area.

It is not surprising that cooperation and task allocation are extensively studied in the robotic literature as a means of reducing the interference problem. Jin et al. [2003] and Flint et al. [2004] give some examples of distributed control and task allocation for Unmanned Air Vehicles (UAV). Gerkey and Matarić [2004] recently proposed a taxonomy of task allocation problems and analyzed some of the known solutions which use intentional cooperation [Parker 1998]. In these examples, however, robots have enough information in order to create explicit models of the environments and of their own capabilities unlike our case.

Our working conditions are more akin to those usually found in swarm robotics [Dorigo and Sahin 2004] or minimalistic robots as in Jones and Matarić [2003]. They suggest a method for division of labor when robots have to collect items of two different types. Their robots collect one type or the other with a probability that changes according to estimates of the relative number of items in the environment and of the relative number of robots already focusing on one type. Krieger and Billeter [2000] show an example of division of labor in a

²² χ^2 test. Null hypothesis: all the probabilities are equal. Alternative hypothesis: there is at least one robot with different probability.

²³Paired Wilcoxon tests among the data for each robot, applying Bonferroni correction for multiple tests.

group of robots using the activation-threshold model [Bonabeau et al. 1996]. The thresholds of the robots are assigned at the beginning and do not change during the experiments. Their work is interesting because it is one of the first showing that complex group behavior can be produced using simple control algorithms.

Among the works that use adaptation or learning, we would like to highlight a few which show some similarity with our work. In the L-ALLIANCE framework [Parker 1998], each robot has a set of behaviors associated with each task. Two values, *impatience* and *acquiescence*, are associated with each set. Impatience increases every time a robot is not performing a task, until it reaches a particular threshold. Then the robot executes the task, and its acquiescence starts growing. The robot gives up performing the task when the acquiescence is too high. Individual impatience and acquiescence rates of growth are adjusted by each robot according to their experience and to the information coming from other robots. The system can, however, work without communication.

Ijspeert et al. [2001] showed how robots can collaborate in order to pull a stick out of a hole without using communication. The behavior of the robots is characterized by a *gripping time parameter* (GTP) which sets the time to wait for help from other robots. Li et al. [2004] proposed an adaptation algorithm to adjust the GTP in order to improve the collaboration rate. They tested it only in simulation and in an extended version of the problem where k robots are needed. Their algorithm works by incrementing (or decrementing) the GTP by a variable value. Both local information (i.e., the rate of personal successful collaborations) and global information (i.e., the swarm collaboration rate evaluated globally and broadcast to all the robots) are used by each robot to decide how to modify the GTP.

Agassounon et al. [2004] propose an architecture, based on the threshold model to allocate workers in a puck clustering application. Their algorithm is also characterized by a time constant, that is, the time spent without finding a puck to transport. If this time expires, a robot returns to the nest and does not exit any more. They use a form of adaptation which consists of estimating the density of pucks at the beginning of the experiment and setting the time constant accordingly.

The three works described show some fundamental differences with respect to ours. Unlike L-ALLIANCE, our approach neither needs intentional communication nor a model of other robots. Unlike the two other works, our actions are not deterministic. It is still an open question whether the other adaptation algorithms differ from ours in terms of performance or efficiency. Unfortunately, the applications used by the authors and the set-ups are too different to draw any scientifically sound conclusion.

7. CONCLUSIONS

We implemented and analyzed an algorithm for division of labor which is inspired by a model of ants' foraging. The model stresses the role that learning plays in the collective behavior of ants' foraging. A simple form of adaptation which occurs in each individual and which uses only information locally available can improve the efficiency of a group of robots. Efficiency improves by

means of division of labor which implicitly takes into account heterogeneity among robots. The robots do not need to communicate with each other, and, in fact, their control algorithms do not need to know how many other robots, if any, are present.

We are convinced that our work can be useful both for biologists and roboticists. Biologists tend to downplay the importance of learning in collective behaviors. Many biologists still think that complex phenomena (like prey retrieval) must imply some form of communication or some form of social hierarchy. Our work is a strong argument against this point of view. Roboticists can take advantage of the methodology that we used. For instance, to the best of our knowledge, very few works exploit (at least explicitly) control groups for the statistical assessment of hypotheses. In this way, we obtained statistical significance with a few experiments.

We also pointed out some limitations of our adaptation algorithm. We showed the negative effects of group size on the global efficiency and noticed that the ratio of foragers seems to reach a fixed value for some prey density. We also note that adaptation occurs quite slowly (the experiments with *MindS-bots* took forty minutes). This still might not be a problem if the group is designed to work for a long time as in the case of planetary explorations.

Future work will try to identify the causes of the problems in order to improve the efficiency of the group and the division of labor. We will also compare our approach with similar ones like those described in Section 6. The comparison will be done by porting the different algorithms to a common application and by analyzing the result in the same way as we did in this article. Hopefully, we will be able to determine if different algorithms yield different results and, if so, their strong and weak points are.

Finally, another future research line will investigate the effect of combining our adaptation algorithm with signaling and more complex communication. We expect that the the group will be more efficient, although it is an open question whether communication and adaptation will interact synergically or destructively. The results will integrate the conclusion already given by Balch and Arkin [1994] who considered only the effect of communication on the group performance.

ACKNOWLEDGMENTS

The authors want to thank the anonymous reviewers for their precious comments on this article.

REFERENCES

- AGASSOUNON, W., MARTINOLI, A., AND EASTON, K. 2004. Macroscopic modeling of aggregation experiments using agents in teams of constant and time-varying sizes. *Autonomous Robots* 17, 2–3, 163–192.
- ALCOCK, J. 1995. *Animal Behavior* 5th Ed. Sinauer, Sunderland, MA.
- BALCH, T. 1999. The impact of diversity on performance in multi-robot foraging. In *Proceedings of the 3rd International Conference on Autonomous Agents (Agents'99)*, O. Etzioni, J. Müller, and J. Bradshaw, Eds. ACM Press, New York, NY. 92–99.
- BALCH, T. AND ARKIN, R. 1994. Communication in reactive multiagent robotic systems. *Autonomous Robots* 1, 1, 27–52.

- BONABEAU, E., THERAULAZ, G., AND DENEUBOURG, J.-L. 1996. Quantitative study of the fixed threshold model for the regulation of division of labor in insect societies. *Proceedings of the Royal Society of London, Series B-Biological Sciences* 263, 1565–1569.
- CAMAZINE, S., DENEUBOURG, J.-L., FRANKS, N., SNEYD, J., THERAULAZ, G., AND BONABEAU, E. 2001. *Self-Organisation in Biological Systems*. Princeton University Press, Princeton, NJ.
- CAO, Y., FUKUNAGA, A., AND KAHNG, A. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots* 4, 1, 7–27.
- DENEUBOURG, J.-L., GOSS, S., PASTEELS, J., FRESNEAU, D., AND LACHAUD, J.-P. 1987. Self-organization mechanisms in ant societies (II): Learning in foraging and division of labor. In *From Individual to Collective Behavior in Social Insects*, J. Pasteels and J.-L. Deneubourg, Eds. Experientia Supplementum, vol. 54. Birkhäuser Verlag, Basel, Switzerland, 177–196.
- DETRAIN, C. AND DENEUBOURG, J.-L. 1997. Scavenging by *Pheidole Pallidula*: A key for understanding decision-making systems in ants. *Animal Behaviour* 53, 537–547.
- DORIGO, M. AND SAHIN, E. 2004. Guest editorial. *Autonomous Robots* 17, 2–3, 111–113.
- DORIGO, M., TRIANNI, V., SAHIN, E., GROß, R., LABELLA, T., BALDASSARRE, G., NOLFI, S., DENEUBOURG, J.-L., MONDADA, F., FLOREANO, D., AND GAMBARDELLA, L. 2004. Evolving self-organizing behaviors for a *Swarm-Bot*. *Autonomous Robots* 17, 2–3, 223–245.
- FLINT, M., E. FERNÁNDEZ-GAUCHERAND, E., AND POLYCARPOU, M. 2004. A probabilistic framework for passive cooperation among UAV's performing a search. In *Proceedings of the 16th International Symposium on Mathematical Theory of Networks and Systems (MTNS'04)*, B. De Moor, P. Van Dooren, V. Blondel, and J. Willems, Eds. Leuven, Belgium.
- GERKEY, B. AND MATARIĆ, M. 2004. A formal analysis and taxonomy of task allocation in multi-robot systems. *Int. J. Robotics Resear.* 23, 9, 939–954.
- GOLDBERG, D. AND MATARIĆ, M. 1997. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings of the 14th National Conference on Artificial Intelligence (AAAI'97)*. MIT Press, Cambridge, MA, 637–642.
- GRASSÉ, P. 1959. La reconstruction du nid et les coordinations inter-individuelles chez *Belligositermes natalensis* et *Cubitermes*. La théorie de la stigmergie: essai d'interprétation des termites constructeurs. *Insectes Sociaux* 6, 41–83.
- HAYES, A. 2002. How many robots? Group size and efficiency in collective search tasks. In *Proceedings of the 6th International Symposium on Distributed Autonomous Robotic Systems (DARS'02)*, H. Asama, T. Arai, T. Fukuda, and T. Hasegawa, Eds. Springer Verlag, Heidelberg, Germany, 289–298.
- HÖLDOBLER, B. AND WILSON, E. 1990. *The Ants*. Springer Verlag, Heidelberg, Germany.
- ISPEERT, A., MARTINOLI, A., BILLARD, A., AND GAMBARDELLA, L. 2001. Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots* 11, 2, 149–171.
- JIN, Y., MINAI, A., AND POLYCARPOU, M. 2003. Cooperative real-time search and task allocation in UAV teams. In *Proceedings of the 42nd IEEE Conference on Decision and Control*. Vol. 1. IEEE Press, New York, NY, 7–12.
- JONES, C. AND MATARIĆ, M. 2003. Adaptive division of labor in large-scale minimalist multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*. Vol. 2. IEEE Press, New York, NY, 1969–1974.
- KRIEGER, M. AND BILLETER, J.-B. 2000. The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics Autonom. Syst.* 30, 1–2, 65–84.
- LABELLA, T. 2003. *Prey retrieval by a swarm of robots*. Thesis for the Diplôme d'Études Approfondies (DEA). Tech. rep. TR/IRIDIA/2003-16, IRIDIA, Université Libre de Bruxelles, Brussels, Belgium.
- LI, L., MARTINOLI, A., AND ABU-MOSTAFA, Y. 2004. Learning and measuring specialization in collaborative swarm systems. *Adaptive Behavior* 12, 3–4, 199–212.
- MONDADA, F., PETTINARO, G., GUIGNARD, A., KWEE, I., FLOREANO, D., DENEUBOURG, J.-L., NOLFI, S., GAMBARDELLA, L., AND DORIGO, M. 2004. *Swarm-Bot*: A new distributed robotic concept. *Autonomous Robots* 17, 2–3, 193–221.
- MONTGOMERY, D. 2000. *Design and Analysis of Experiments* 5th Ed. John Wiley & Sons, New York, NY.

- PARKER, L. 1998. ALLIANCE: An architecture for fault tolerant multi-robot cooperation. *IEEE Trans. Robotics Automat.* 14, 2, 220–240.
- PRESS, W., FLANNERY, B., TEUKOLSKY, S., AND VETTERLING, W. 1992. *Numerical Recipes: The Art of Scientific Computing* 2nd Ed. Cambridge University Press, Cambridge, UK.
- SCHNEIDER-FONTÁN, M. AND MATARIĆ, M. 1996. A study of territoriality: The role of critical mass in adaptive task division. In *From Animals to Animats 4, Proceedings of the 4th International Conference on Simulation of Adaptive Behavior (SAB'96)*, P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. Wilson, Eds. MIT Press/Bradford Books, Cambridge, MA, 553–561.
- SIEGEL, S. AND CASTELLAN JR. N. 1988. *Nonparametric Statistics for the Behavioral Science* 2nd Ed. Statistics Series. McGraw-Hill, New York, NY.
- SUTTON, R. AND BARTO, A. 1998. *Reinforcement Learning: An Introduction*. MIT Press, Cambridge, MA.
- WEBB, B. 2000. What does robotics offer animal behaviour? *Animal Behaviour* 60, 5, 545–558.

Received January 2005; revised July and November 2005; accepted April 2006