Probabilistic Modelling of a Bio-Inspired Collective Experiment with Real Robots

A. Martinoli, and F. Mondada[*]

[*]*Microcomputing Laboratory, Swiss Federal Institute of Technology*
IN-F Ecublens, CH-1015 Lausanne

Abstract. This paper describes the implementation and modelling of a biologically inspired collective behaviour. The experiments are concerned with the gathering and clustering of randomly distributed small cylinders. Each experiment has been repeated ten times in a simulated environment (parametric simulation) and carried out five times with a group of ten Khepera miniature mobile robots. The simulated and experimental results are compared, quantified and discussed showing the advantages and the weaknesses of both approaches.

Key Words. Collective Bio-Inspired Robotics, Modelling, Autonomous Robotics

1   Introduction

Bio-inspired collective robotics favours decentralised solutions, i.e. solutions where coordination is *not* taken over by a special unit using private information sources, or concentrating and redistributing most of the information gathered by the individual robots. Inspired by the so-called collective intelligence demonstrated by social insects (Bonabeau and Theraulaz, 1994), bio-inspired collective robotics studies robot-robot and robot-environment interactions leading to robust, goal-oriented, and perhaps emergent group behaviours. Such a bio-inspired approach in collective robotics seems to be a promising way to solve problems which are hard to tackle using classical control methods.

Let us now address the state of the research in a particular well-suited experiment in bio-inspired collective robotics: the gathering and clustering of randomly distributed objects. This experiment is inspired from the studies of cluster emergence with social insects. Deneubourg (Deneubourg *et al.*, 1991) showed that a simple mechanism involving the modulation of the probability of dropping corpses as a function of the local density, was sufficient to generate the observed sequence of the clustering of corpses.

In (Beckers *et al.*, 1994; Maris and te Boekhorst, 1996) similar experiments were carried out with real robots architectures based basically on reactive behaviour. In these two papers, a precise statistical analysis was carried out but no modelling of the experiment was presented.

In (Martinoli and Mondada, 1995) we also presented a similar experiment. However, due to the difficulty for the recognition algorithm to distinguish between a small cylindrical object (which will be referred to as "seed" from now on) and another robot, it was possible to analyse quantitatively the data, but the high rate of destructive interferences and of experimenter interventions prevented the creation of an adequate probabilistic model in simulation which could have generated similar results.

The experiments presented in this paper are carried out with a more reliable distinguishing algorithm. We aim to show that the improvements in the control architecture and a systematic measuring procedure of the team performances will allow us to compare the experimental results with those obtained in simulation with a Markov chain model. It is worthwhile to mention that the main purpose of the parametric simulation (as opposed to the "classical" simulation, where as exhaustive as possible feature sets of the robot and of the environment are used) is not to deliver data which fit as close as possible to the experimental data. The simulations aim to outline the parameters of the real set-up which play a crucial role in the evolution of the collective performances. A better comprehension of collective mechanisms, such as interference among the robots or stigmergic communication, will help to evaluate the expected collective performances of a given number of robots, a given amount of work, a given work area and a pre-established control architecture.

2 Materials and Methods

2.1 Experiments with Real Robots

a)   b)

Figure 1   a) Seed scattering at beginning of the experiment and b) after about 2 hours, at the end of the longer experiment.

2.1.1 Experimental Set-Up   Khepera is a miniature mobile robot developed to perform "desktop" experiments (Mondada *et al.*, 1993). Its distinguishing characteristic is a diameter of 55 mm. Each robot can be extended with a gripper module, which can grasp and carry objects with a maximum diameter of 50 mm. The experiments are carried out with a group of 1 to 10 Kheperas and 10 to 40 seeds (see fig. 1 as an example). The seeds have a cylindrical form, with a diameter of 15 mm and a height of 25 mm. We use two square arenas with different sizes, the largest

Figure 2   a) Geometrical representation of the cluster construction probability. The ratio between the identification perimeter and the total detection perimeter of the cluster represents the probability to increment the cluster size of one seed. b) Geometrical representation of the cluster destruction probability. The robot, in order to decrement the size of the cluster by 1 seed first has to detect the cluster as in figure 2a and then grasp a seed. c) Plots of both probabilities. They are calculated from the perimeter ratios (external perimeter of the grey zones divided by the total cluster perimeter) illustrated in Fig. 2a and b.

having double the surface of the smallest ($80 \times 80\,$cm and $113 \times 113\,$cm). The initial scattering of the seeds and the starting position of the robots are arbitrarily predefined and differ from replication to replication. Several experiments which differ in the number of scattered seeds, the number of robots, and the working surface are performed and the team performances are measured and compared. As performance measurement we chose the same as in (Martinoli and Mondada, 1995), that is, the mean cluster size expressed in number of seeds at a given time. The experiments terminate when a pre-established time lapse is over, in our case approximatively 20 minutes, and are repeated 5 times. A further experiment, which lasted about 2 hours and was replicated 3 times, is also presented as reference.

It is worth emphasising that in all the experiments the robots operate completely autonomously and independently; all sensors, motors and controls are on-board, and there is no explicit communication (IR or radio link) with other robots or with the experimenters. The only possible interactions among robots are the reciprocal avoidance of collisions and an indirect form of messages, which arise from the modifications of the environment (i.e., for instance the cluster geometry).

2.1.2   Control Algorithm   We can summarise the resulting robot behaviour with the following simple rules: the robot moves on the arena looking for seeds. When its sensors are activated by an object, the robot starts the discriminating procedure. Two cases can occur: if the robot is in front of a large obstacle (a wall, another robot

Figure 3 The dynamics of the clustering experiment represented as a Markov chain.

or an array of seeds), the object is considered as an obstacle and the robot must avoid it. In the second case, the object is considered as a seed. If the robot is not already carrying a seed, it grasps the seed with the gripper; if the robot is carrying a seed, it drops the seed it is carrying close to the one it has found; then, in both cases, it begins searching again.

The discriminating behaviour is basically similar to that reported in (Martinoli and Mondada, 1995): the improvement is based on an increased number of spatial and temporal samples. A test of reliability has shown that this algorithm correctly discriminates objects with a probability close to one.

Figure 2 illustrate the geometrical and numerical results considered in the calculations of the construction and destruction probabilities of a cluster of a given size when the robot has found it.

2.2 Parametric Simulation with a Markov Chain

In this section we describe a simple probabilistic model of the clustering experiment. We take into account the building and destruction probabilities of a cluster, its probability to be found by a robot (proportional to its detection surface) and the interference probability due to the encounter with other teammates. Interference can occur during different activities: by the search or by the cluster construction-destruction. If another teammate lies in the sensor range, the robot has to recognise

Figure 4   a) The first random process: the clusters are represented with their
detection area scaled with the total surface of the environment; the
probability that a robot encounters a cluster is proportional to the
detection area of the cluster. b) The second random process: the robot
can modify the size of the cluster, incrementing or decrementing it by 1
seed, only if the robot runs into the cluster from the positions described
in Fig. 2. The two bars represent the whole set of clusters scattered at a
given moment on the arena: if the random cursor (set on the left bar at
0.47) fails on a grey zone, the selected cluster will be decremented (left
bar) or incremented (right bar) by 1 seed. The third stochastic process,
the encounter of other teammates, is always overlapped to both the
former processes: interference can occur during the search as well as
during seed pick up or drop activity.

and avoid it. The parametric simulation takes in account the time needed for this
action as well as the one required to modify the size of a cluster. The coordinates
of the robots are randomly assigned at the beginning of each program iteration (see
the first random process depicted in Fig. 4a). There are two fundamental approxi-
mations in our simulation:

- the robot is not moving in the environment: the simulation calculates the
  global probability of finding a cluster or another teammate based on their
  detection area and the arena surface;

- in order to convert the number of iterations into time, we assume that the
  clusters and the robots are scattered homogeneously on the arena and that
  we can calculate the minimal mean distance between two of them (robots or
  clusters, see Fig. 4a); using the mean distance and the experimental mean
  velocity of the robots (in these experiments 120 mm/s), we can compute the
  minimal time to move from one cluster to another, which is set equal to two
  simulation iterations (a simulated robot needs at least two iterations to pick up

a seed from a cluster and to drop it in the same or in another one); if another teammate is found, the robot searching behaviour is frozen for a couple of further iterations, depending of the calculated time out based on the real robot interaction.

The clustering process can be modelled as a Markov chain (see Fig. 3). The chain has as many states as possible clusters with different sizes, which actually corresponds to the number of seeds scattered on the arena. The transition probabilities from state to state are calculated as a function of the total construction and destruction area represented by all the clusters with the same size and as a function of the number of other teammates on the environment. As a consequence, the rules to calculate the transition probabilities are pre-established by the geometrical constraints of the set-up but their values are updated every time that the number of clusters of a given size changes. Notice that, for a cluster of size n, if there is no cluster of size n-1, its building probability is zero.

Every robot can increment or decrement the size of a cluster by one seed at a time. The building and destruction probabilities (Bi and Di of Fig. 3, with i between 1 and the number of scattered seeds) are conditioned by three stochastic processes which are explained in Fig. 4. First, a random position in the environment is assigned to the robot. If this position is inside the detection area of a cluster, the second random process is started. According to the state of the robot (carrying or not carrying a seed) the size of the found cluster is incremented or decremented by one seed if the number delivered by the random process is within the construction or destruction region (calculated with the values of Fig. 2c). The third random process, the interference with the other teammates, is overlapped to the former ones. Each random process is repeated for each robot independently before the next iteration of the program is started. Notice that both the first two random processes always consider the whole actual set of clusters and are very lightly coupled: the former delivers only a boolean value "true" if a cluster is found but neither its size nor its position in the arena. As a consequence, the second process, if enabled by the first one, also takes into account the whole actual set of clusters, which is a vector containing the building and destruction probabilities of each cluster.

3   Results and Discussion

As mentioned in the previous section, we use the mean size of the clusters as collective performance measurement of the group of robots. On the following pairs of plots, the left figure displays the real team fitness and the right displays the fitness obtained by simulation. In all of figures, only the average fitness is plotted: 10 runs have been performed for each simulated experiment and 5 for the ones with real robots (only 3 replications for the longer experiment, see Fig. 8).

In order to estimate the repeatability of the experiments and the reliability of the model we have calculated the relative error on the basis of the data variance of each experiment for all the time frames: the experimental results were affected by a relative error always smaller than 25% (35% with only 3 replications in the longer
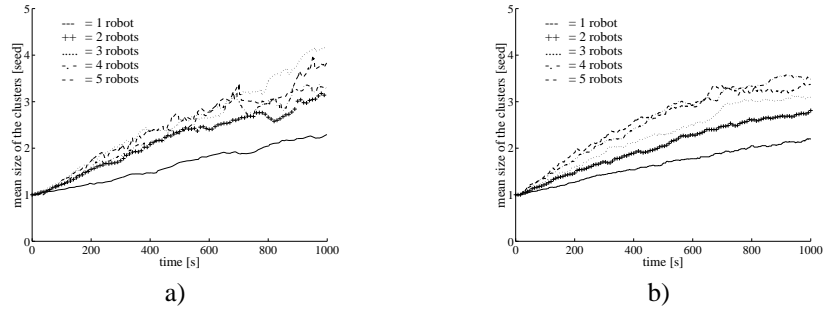
Figure 5 Fitness of the group with increasing number of robots (1 to 5) on the arena of 80x80 cm and 20 seeds to be gathered. (a) Results of the experiments with real robots. (b) Results of the simulations.


Figure 6 Fitness of the group with increasing number of robots (1,3 and 5) on the arena of 80x80 cm and 20 seeds to gather in a) and fitness of the group with the double of teammates (2,6 and 10), the double of seeds (40) and an arena two times bigger (113x113 cm) in c). The corresponding simulated fitnesses are depicted in Figure b) and d).

experiment) and the simulated results by one smaller than 15%. The experimental relative error could be decreased by a more extended set of replications for a given experiment. However, this could drastically increase the time spent for experimenting, if the same team fitness is measured.

Figure 5 shows the team fitness for a group of 1 to 5 robots. Although the team fitness of the simulated group of robots is slightly smoother than that of the real ones (it is namely the average of twice the number of the experimental replications), the two plots shown are quite similar. The main difference is that the performances of the group of 3 real robots is less rapidly saturated than the one in the simulations. Although the territorial division is slightly taken into account in the parametric simulation by the conversion of the number of iterations into time, the approximation introduced by considering only static robots could explain this difference.

Fig. 6 compares the team performances where the number of robots, the surface,

a)   b)

Figure 7   Mean fitness of the single teammate within a group composed by an increasing number of robots (1 to 5) on the arena of 80x80 cm and 20 seeds to gather. (a) Results of the experiments with the robots. (b) Results of the simulations.

a)   b)

Figure 8   The evolution of cluster mean size during a) 2 hours with real robots and b) about 7 days in simulation.

and the number of seeds to gather is doubled. The purpose here is to demonstrate that robot- and seed-density (meant as amount of work to do) are two key parameters of the experiments and that we can obtain the same results in the team fitness with the same density of robots and seeds (this is very important because of the rare availability of a greater number of robots). If we compare Fig. 6a with Fig. 6c we can conclude that, considering the repeatability of the experiments, the differences between the two pictures are negligible. This tendency is confirmed by the simulations (compare Fig. 6b with Fig. 6d).

Fig. 7 shows that in both simulation and experiments with real robots there is no superlinearity in the team performances. On the contrary, in the experimental and simulated results with 4 and 5 robots there is a substantial sublinearity because of the destructive interferences.

Finally, Fig. 8a shows the experimental results obtained using a special tool de-

veloped at our laboratory to extend the autonomy of the Khepera robots (Martinoli *et al.*, 1997). Although a single experiment replicated only three times can not be representative, the figure shows that after 3000 seconds the team fitness becomes saturated due to interferences. In comparison, we have performed the same experiment but longer in simulation. The simulation results are shown in Fig. 8b. Since the probability to build a cluster is consistently greater than that to destroy it, the seeds can be gathered in a single cluster if enough time is available. At the moment, we do not have enough experience with long experiments to guess whether it is possible to obtain this result with the real robots and the set-up described in this paper. Remember also that the resulting growth probability (building probability minus destroying probability) is ever smaller. As a consequence, before all the seeds can be gathered together, the average of the cluster size could reach a saturation zone where interference and building gradient contributions are in equilibrium.

4 Conclusion

The results of this paper have shown that it is helpful, in collective robotics, to compare simulation results with real robot experiments in order to better understand mechanisms such as interference or stigmergic communications. The previous section has shown that the parameters chosen for the simulation (number of robots, number of seeds, arena size and seed-robot interaction geometry) play a crucial role in the evolution of the team fitness. The influence of these key parameters has also been verified in the experiments with real robots, changing one parameter at a time.

Both approaches, simulation and experimentation, have also shown their limitations. It is very difficult to simulate robots considering all the mechanisms and the noise of the real world: the optimal key parameters found in simulation could be quite different from those of the real environment. Therefore, we advocate a parametric simulation of the robots, which will not suppress the need to experiment with real robots but will help to understand the mechanisms found in the real world. On the other hand, the limited energy autonomy of the Khepera robot as well as the whole procedure needed for preparing the set-up represent a great handicap to perform many replications of the same experiment. A first improvement in this direction has been done by introducing the extended autonomy tool mentioned in Fig. 8b.

The results show that in this kind of experiments with no adaptivity in the robot control, a greater number of robots does not necessarily help to increase the fitness. If we compare the performances of each teammate with that of the single robot, we see that the robot working alone always achieves the best performance (linearity or sublinearity of the fitness).

The introduction of adaptivity could, for instance, allow the single robots to switch from an active phase to an inactive one when the ratio between the amount of work (in our case the seed finding rate) and the interference (in our case the encounter rate with other teammates) decreases under a given threshold. Similar mechanisms are supposed to play a crucial role in ant colonies (Pacala *et al.*, 1996).

5 Acknowledgements

6 REFERENCES

Beckers, R., O.E. Holland and J.L. Deneubourg (1994). From local actions to global tasks: Stigmergy and collective robotics. In: *Proceedings of the Fourth Workshop on Artificial Life* (R. Brooks and P. Maes, Eds.). MIT Press. Boston, MA. pp. 181–189.

Bonabeau, E. G. and G. Theraulaz (1994). *Intelligence Collective*. Hermès. Paris, France.

Deneubourg, J. C., S. Goss, N. Franks, A. Sendova, A. Franks, C. Detrin and L. Chatier (1991). The dynamics of collective sorting: Robot-like ant and ant-like robot. In: *Simulation of Adaptive Behavior: From Animals to Animats* (J. A. Mayer and S. W. Wilson, Eds.). pp. 356–365. MIT Press.

Maris, M. and R. te Boekhorst (1996). Exploiting physical constraints: Heap formation through behavioral error in a group of robots. In: *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems IROS-96*. Vol. 3. Osaka, Japan. pp. 1655–1660.

Martinoli, A. and F. Mondada (1995). Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In: *Proceedings of the Fourth International Symposium on Experimental Robotics* (O. Khatib and J. K. Salisbury, Eds.). Springer Verlag. Stanford, U.S.A.. pp. 3–10.

Martinoli, A., E. Franzi and O. Matthey (1997). Towards a reliable set-up for bio-inspired collective experiments with real robots. In: *Proc. of the Fifth International Symposium on Experimental Robotics ISER-97* (A. Casals, Ed.). Springer Verlag. Barcelona, Spain. In press.

Mondada, F., E. Franzi and P. Ienne (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In: *Proceedings of the Third International Symposium on Experimental Robotics* (T. Yoshikawa and F. Miyazaki, Eds.). Kyoto, Japan. pp. 501–513.

Pacala, S. W., D. M. Gordon and H. C. J. Godfray (1996). Effects of social group size on information transfer and task allocation. *Evolutionary Ecology* **10**, 127–165.