

Collaboration through the exploitation of local interactions in autonomous collective robotics: the stick pulling experiment

Auke Jan Ijspeert,^{1,2} Alcherio Martinoli,² Aude Billard,^{1,2}
and Luca Maria Gambardella¹

¹Istituto Dalle Molle di Studi sull'Intelligenza Artificiale (IDSIA), Lugano

²Laboratoire de Microinformatique (LAMI), EPFL, Lausanne*

*A.J.I. and A.B. are now at the Brain Simulation Laboratory, University of Southern California, Los Angeles, CA 90089, emails: ijspeert@usc.edu, billard@pollux.usc.edu. A.M. is at the Microsystems Research Laboratory, California Institute of Technology, Pasadena, CA 91125, email: alcherio@micro.caltech.edu. L.M.G.'s email is luca@idsia.ch.

Abstract

This article presents an experiment which investigates how collaboration in a group of simple reactive robots can be obtained through the exploitation of local interactions. A test-bed experiment is proposed in which the task of the robots is to pull sticks out of the ground—an action which requires the collaboration of two robots to be successful. The experiment is implemented in a physical setup composed of groups of 2 to 6 Khepera robots, and in Webots, a 3D simulator of Khepera robots.

The results using these two implementations are compared with the predictions of a probabilistic modelling methodology (Martinoli, Ijspeert, & Mondada, 1999a; Martinoli, Ijspeert, & Gambardella, 1999b) which is here extended for the characterization and the prediction of a collaborative manipulation experiment. Instead of computing trajectories and sensory information, the probabilistic model represents the collaboration dynamics as a set of stochastic events based on simple geometrical considerations. It is shown that the probabilistic model qualitatively and quantitatively predicts the collaboration dynamics. It is significantly faster than a traditional sensor-based simulator such as Webots, and its minimal set of parameters allows the experimenter to better identify the effect of characteristics of individual robots on the team performance.

Using these three implementations (the real robots, Webots and the probabilistic model), we make a quantitative investigation of the influence of the number of workers (i.e robots) and of the primary parameter of the robots' controller—the gripping time parameter—on the collaboration rate, i.e. the number of sticks successfully taken out of the ground over time. It is found that the experiment presents two significantly different dynamics depending on the ratio between the amount of work (the number of sticks) and the number of robots, and that there is a super-linear increase of the collaboration rate with the number of robots. Furthermore, we investigate the usefulness of heterogeneity in the controllers' parameters and of a simple signalling scheme among the robots. Results show that, compared to homogeneous groups of robots without communication, heterogeneity and signalling can significantly increase the collaboration rate when there are fewer robots than sticks, while presenting a less noticeable or even negative effect otherwise.

Keywords: collective autonomous robotics, swarm intelligence, collaboration, sensor-based simulation, probabilistic modeling.

1 Introduction

Swarm Intelligence (SI) (Beni & Wang, 1989) is an innovative computational and behavioral metaphor for solving distributed problems, that takes its inspiration from the biological examples provided by social insects - ants, termites, bees, and wasps - and by swarming, flocking, herding, and shoaling phenomena in vertebrates (Bonabeau, Dorigo, & Theraulaz, 1999). The abilities of such systems appear to transcend the abilities of the constituent individual agents. In most biological cases studied so far, the robust and capable high level group behavior has been found to be mediated by nothing more than a small set of simple low level interactions between individuals, and between individuals and the environment. The SI approach emphasizes distributedness and exploitation of direct (robot-to-robot) or indirect (via the environment) local interactions among relatively simple agents.

The main advantages of the application of the SI approach to the control of a group of robots are three-fold: first, scalability: the control architecture is kept exactly the same from a few units to thousands of units; second, flexibility: units can be dynamically added or removed, they can be given the ability to reallocate and redistribute themselves in a self-organized way; third, robustness: the resulting collective system is robust not only through unit redundancy but also through the unit minimalistic design (Boehringer, Brown, Donald, Jennings, & Rus, 1995). Although a formal and quantitative definition of minimalism has yet to be defined for collective systems, minimalistic design in SI implies an effort to keep the resources for computation, sensors, actuators, and communication as low as possible for each unit, while aiming at having an as smart as possible group behavior.

In the last few years, the SI control principles have been successfully applied to a series of case studies in collective robotics: aggregation (Beckers, Holland, & Deneubourg, 1994; Martinoli et al., 1999a; Martinoli, 1999) and segregation (Holland & Melhuish, 1999), exploration (Hayes, Martinoli, & Goodman, 2000), collaborative transportation (Kube & Bonabeau, 2000), work division and task allocation (Krieger & Billeter, 2000), and self-assembling (Hosokawa et al., 1998; Yoshida, Murata, Tomita, Kurokawa, & Kokaji, 1999). All these works have been performed using groups of 1 up to 12 simple, autonomous real robots, exploiting local communication forms among teammates (implicit, through the environment, or limited explicit, wireless communication), and fully distributed control. However, the lack of rigorous, scalable methodologies for designing and analyzing such fully distributed robotics systems has, for the moment, prevented a more extensive application of the SI approach to collective robotics. There are indeed many applications such as traffic regulation (Wang & Premvuti, 1995), waste cleaning (Gage, 1995; Parker, 1998), surveillance (Everett, Gilbreath, Heath-Pastore, & Laird, 1993), collective navigation (Mataric, 1994; Balch & Arkin, 1998), collaborative mapping (Yamauchi, 1999), or foraging (Mataric, 1994), which, so far, have been carried out using approaches which can not be classified as SI (either because they made extensive use of global communication, they had centralized control, or used sophisticated sensors based on global references), and for which, we believe, the SI approach would be very well suited because of the advantages cited above.

This article aims at contributing to research in swarm intelligence 1) by making a quantitative study of how collaboration in a group of simple reactive, autonomous robots can be obtained and controlled through the exploitation of local interactions, and 2)

by consolidating and extending a novel modeling methodology (Martinoli et al., 1999a, 1999b) for characterizing and predicting the collective behavior of (large) distributed robotic systems.

Collaborative robotics One way to increase the solving performance of a robot team without significantly modifying the robots capacities is *collaboration*. In particular, when collaboration is obtained with stigmergic mechanisms (i.e. implicit communication via the environment) or with simple explicit communication schemes such as binary signalling, the task accomplished by the team can be more complex and its performance enhanced without loosing autonomy or increasing in a relevant way the complexity at the individual level. In some cases, e.g. (Johnson & Bay, 1995; Boehringer et al., 1995; Mataric, Nilsson, & Simsarian, 1995; Ghanea-Hercock & Barnes, 1996; Khatib, 1999; Ota & Arai, 1999; Humberstone & Smith, 2000; Kube & Bonabeau, 2000; Wang, Kimura, Takahashi, & Nakano, 2000) and the experiment presented in this article, the task may even require collaboration to be successfully performed, with single robots not being able to carry out the task alone. Such types of experiments can be defined to be "strictly collaborative" (Martinoli, 1999).

The experiment presented in this article is the follow-up of initial tests presented in (Martinoli & Mondada, 1995). The task is to locate sticks in a circular arena and to pull them out of the ground.¹ The task is carried out by groups of two to six² Khepera robots (Mondada, Franzi, & Jenne, 1993) equipped with grippers and capable of distinguishing the sticks with their frontal sensors. Because of the length of a stick, a single robot is not capable of pulling it out of the ground alone, and collaboration between two robots is necessary for pulling a stick completely out. As the robots have only local sensing capabilities and do not use explicit communication (except in one experiment), there is no explicit coordination between robots. Coordination is purely probabilistic and happens based on local interactions (see the experiment description in Section 2). The experiment is not intended to represent a real-life application but to serve as a case study of the dynamics of collaboration and as an abstraction of future collaborative robotics applications.

A prediction tool for collective robotics One of the main difficulties in designing efficient robotic teams, is the problem of characterizing and predicting how the group behavior is affected by the hardware and software characteristics of the individuals forming the group. This is particularly true for large groups of robots controlled in a fully distributed way. As hardware costs are high and experimenting with real robots is time consuming, it is most useful to have prediction tools which allow one to determine the optimal number of robots or the optimal control parameters for an optimal team performance, for instance, before setting all the details of the experimental set-up. One popular approach is to use detailed sensor-based simulations which simulate as realistically as possible the sensor and actuator characteristics of the robots, including kinematic noise, as

¹Although the experiment is not intended to reproduce a biological system, the experiment presents several similarities with the matches extraction and transportation performed by some ant colonies (Chauvin & Janin, 1975).

²Six was the maximum number of real robots available for the experiments. Consistently with the SI approach, there is absolutely no limitation in the size of the team from a control architecture point of view.

well as the characteristics of the environment. Although useful, these types of simulation have the inconveniences that they are time-consuming both to develop, and to run when the group of robots is large. In addition, the number of parameters considered by these types of simulation is so large that it is then difficult to extract those which really play a crucial role on the team performance.

In this article, we present a simulation based on a probabilistic model which is significantly easier to implement and faster than a sensor-based simulation, while being able to predict group behavior, in particular the collaboration rate (the number of successful collaborations over time), with the same accuracy. The idea of the probabilistic model is that, with simple reactive autonomous robots and a distributed control scheme, a collective manipulation experiment is essentially a stochastic process based on simple geometrical considerations. Using the same methodology described in (Martinoli et al., 1999a, 1999b), in which it was used for the characterization of different clustering experiments with different robotic platforms, the probabilistic model represents the group of robots as a set of parallel processes which, instead of computing trajectories and sensory information, represent actions such as gripping a stick or encountering another robot as stochastic events.

Compared to previous applications of the methodology (Martinoli et al., 1999a, 1999b), the contribution of this article is four-fold. First, we demonstrate that the methodology can successfully be applied to a collective manipulation experiment in which the task needs strict collaboration between robots, where spatial and temporal coordination are needed for succeeding in the task. This requires, among other things, the introduction of probabilities which do not only depend on the current state of the environment (as for the clustering tasks), but also on the current state of the other robots. Second, while controllers were fixed in (Martinoli et al., 1999a, 1999b), we illustrate here how the probabilistic model can be used to make a detailed quantitative investigation of the effect of varying parameters of the robots' controller on the group behavior. Third, we demonstrate that the methodology can address heterogeneous as well as homogeneous groups of robots. Fourth, we investigate a case (collaboration using explicit signalling) in which the current methodology shows its limitations, and for which the model's predictions have to be corrected with the help of two additional free parameters such as to fit the results obtained with real robots and sensor-based simulations.

Three sets of experiments In the next sections, we present results of the stick pulling experiment in three different implementations: the real Kheperas, a sensor-based simulation (Webots (Michel, 1998)), and the probabilistic model. Three sets of experiments are carried out to quantitatively investigate the effects of variations of robot controller implementations and number of robots on group behavior, in particular, on the collaboration rate among robots, i.e. the number of sticks successfully taken out of the ground over time. In the first set, we investigate the case with homogeneous groups of robots, that is, groups in which all robots are identical from a hardware and software point of view. In the second set, we test whether introducing heterogeneity at the level of the controllers' parameters can improve the team performance compared to homogeneous groups. In the last set, which is only carried out with Webots and the probabilistic model, we analyze the effect of a simple signalling scheme among the robots.

2 The stick pulling experiment

2.1 The physical setup

The experiment is carried out in a circular arena (80 cm of diameter) delimited by a white wall. Four holes situated at the corners of a square with 30 cm edges, contain white sticks (15 cm long, diameter of 1.6 cm) which, in their lowest position, stick 5 cm out of the ground (Figure 1).




Figure 1: Physical set-up for the stick pulling experiment.

Groups of 2 to 6 Khepera robots, equipped with gripper turrets, are used to pull the sticks out of the ground. Because of their thinness, the sticks can be distinguished from the wall and from other robots³ using the six frontal IR proximity sensors of the Khepera. Two Kheperas are necessary for pulling a stick completely out of the ground. Collaboration is thus required, with a first robot taking the stick half out of the ground, until a second robot approaches the stick from the opposite direction and lifts the stick completely (see the right hand-side of Figure 1 and Figure 3). As described in the next section, the robots are able to determine whether another robot is holding the same stick using information about the gripper's arm position. After a successful collaboration, the stick taken out of the ground is released by the robot, and replaced in its hole by the experimenter.

2.2 The robots' controllers

The behavior of a robot is determined by a simple hand-coded program consisting of a loop through several functional blocks (Figure 2). The default behavior is to look for sticks, that is, to wander in the arena in a straight line until something is detected by the frontal proximity sensors, in which case the robot turns towards the detected object and starts a detection procedure. The detection procedure consists of taking multiple sensor samples of the same object with the robot turning on itself once to the left and once to the right (similarly to (Martinoli et al., 1999a)). A stick is recognized from obstacles (the wall or other robots) if, within these measurements, the number of proximity sensors

³To increase their reflectivity, robots have a belt of white paper, as well as a thin band of IR-reflective stickers in their back (not shown on the picture).

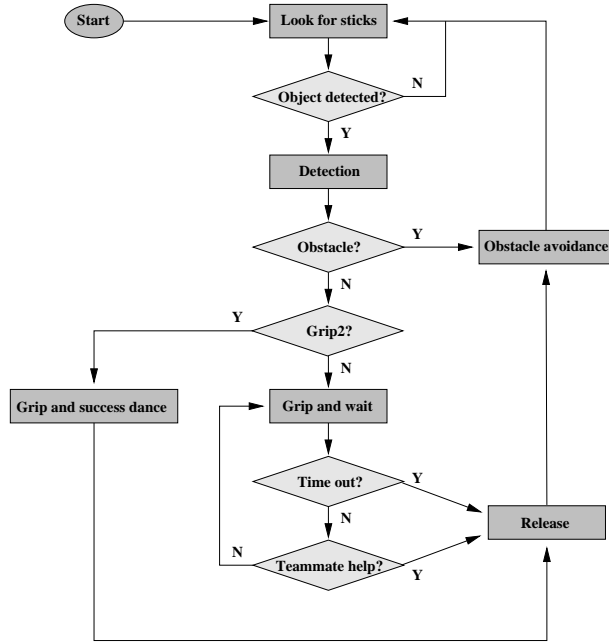


Figure 2: Flowchart of the robots' controller.

returning a significant signal does not exceed two. The multiple measurements enables an increased spatial resolution and a filtering of noise, leading to a perfect distinction of sticks and preventing, in particular, moving robots to be mistaken for sticks.

If the detected object is an obstacle, the robot turns away from it, performs obstacle avoidance for a few seconds, and returns to the looking-for-sticks procedure. If the object is a stick, the robot backtracks for a predefined distance (a few centimeters), grips the stick and pulls it up. During pulling, the robot can determine whether another robot is already gripping the same stick by measuring the speed of elevation of the gripper arm.⁴ If the elevation is fast, the robot assumes that the stick is free (no other robot holding it) and we call such a grip a *grip1*. If the elevation is slow, the robot assumes that another robot is already holding that stick and therefore “braking” the elevation. Such a grip is called *grip2*.

Robot making a grip1 When a robot makes a grip1, it holds the stick half out of the ground and releases it when either the duration of the grip exceeds a *gripping time parameter* (which is then considered as a failed collaboration) or another robot comes to make a grip2 (successful collaboration, see Figure 3). The robot can detect when another robot is making a grip2 because the force exerted by that robot on the stick leads to a slight elevation of its arm’s position compared to the arm’s programmed position. Once the stick is released, the robot turns away, performs obstacle avoidance for a few seconds, and returns to the looking-for-sticks procedure. The gripping time parameter therefore corresponds to the maximum duration a robot will wait with the stick lifted, from the moment the gripper elevation sensor indicates that the gripper is completely lifted. This parameter plays a primary role in determining the number of successful collaborations, and this role will be thoroughly investigated in the experiments of Section 4.

⁴The gripper turret is equipped with a sensor giving the arm’s elevation angle.




Figure 3: Collaborative sequence.

Robot making a grip2 When a robot makes a grip2, the pulling of the stick is temporarily braked until the robot making the grip1 releases its grip. The arm eventually reaches the programmed position, which corresponds to a complete lift of the stick out of the ground. To mark the successful collaboration, a small “success dance” (moving the arm up and down) is performed. Similarly to the end of grip1, the robot then releases the stick (which has to be replaced in the hole by the experimenter) and resumes looking for sticks.

Note that, because of the way sticks are recognized (i.e. only by their thinness), a stick which is held by one robot can only be recognized when approached from the opposite side, within a certain angle (approx. 125 degrees, see also Figure 6). In that case, the robot holding the stick is far enough for it not to be detected by the frontal proximity sensors. For the other angles of approach, both the stick and the robot are detected and the whole is therefore taken for an obstacle. This limits the probabilities for collaborations, but ensures that the second robot approaches the stick within an angle which allows it to grasp the stick without its gripper getting entangled with the first robot.

2.3 Simulations with Webots

In order to more systematically investigate the collaboration dynamics, we also implemented the experiment in Webots, a 3D simulator of Khepera robots. The simulator computes trajectories and sensory input of the robots in an arena corresponding to the physical set-up (Figure 4). The simulation is sufficiently faithful for the controllers to be transferred to real robots without changes,⁵ and for the robot behaviors to be very similar to those of the real robots (see Results). The mean acceleration ratio for this experiment with 5 robots between Webots and real time is about 15 on a Ultra Sun1 workstation.

⁵In Webots 2.0, Webots API commands can be directly crosscompiled into Khepera BIOS commands and downloaded into the real robots.




Figure 4: Implementation of the experiment in Webots, a 3D simulator of Khepera robots (Michel, 1998).

3 The probabilistic model

The central idea of the probabilistic model is to describe the experiment as a series of stochastic events with probabilities based on simple geometrical considerations. The states of the robots are defined by a program with exactly the same structure as that of the controllers of the real robots (Figure 2), but, instead of computing the detailed sensory information and trajectories of the robots, the change of states is determined by the throwing of dice (Figure 5). In other words, the probabilistic nature of interactions during an experiment is captured by the transformation of deterministic branch operators of the real robot/webots controllers into probabilistic branch operators in the probabilistic model.

Look-for-sticks mode Once a robot is in the look-for-sticks mode, it will, at each iteration, have probabilities P_N of encountering nothing, P_W of encountering a wall, P_R of encountering a robot, and P_S of finding a stick. Sticks can be distinguished between those that are available for a grip1 P_{G1} (if they are free) and those available for a grip2 P_{G2} (if another robot holds them), with $P_S = P_{G1} + P_{G2}$.

These different probabilities depend on the respective detection areas of the different elements in the arena, A_W for the surrounding wall, A_R one robot, and A_S one stick (Figure 6), relative to the area of the whole arena A_A . These detection areas in turn depend on the physical dimensions of each element, its surface features (e.g. IR reflectivity), the sensor range of robots, and the controller's parameters used in the detection procedure. These areas are measured from the corresponding detection distances using the real robots (Table 1).

The probabilities P_N , P_R , P_{G1} and P_{G2} also depend on the current state of the environment i.e. the number of robots N_R , and the number of sticks available for a grip1 N_{G1}

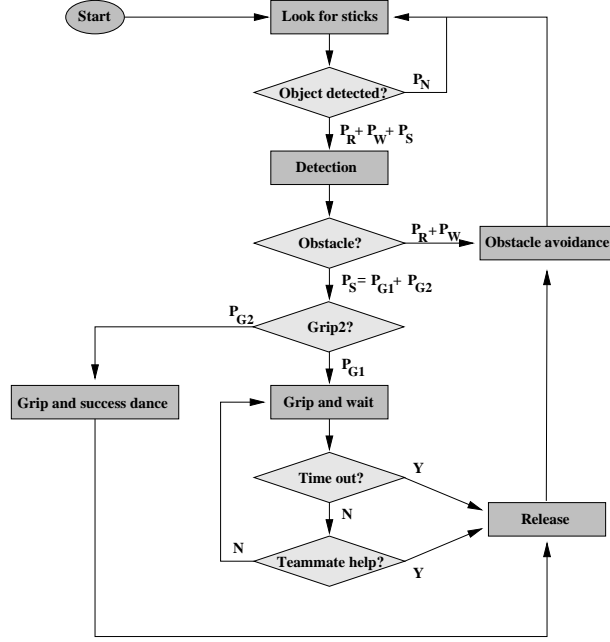


Figure 5: Flowchart of the robots' controller in the probabilistic simulation.

and for a grip2 N_{G2} (where $N_{G1} + N_{G2} = N_S$, the total number of sticks):

$$P_W = A_W/A_A \quad (1)$$

$$P_R = N_R \cdot A_R/A_A \quad (2)$$

$$P_{G1}(t) = N_{G1}(t) \cdot A_S/A_A \quad (3)$$

$$P_{G2}(t) = N_{G2}(t) \cdot R_{G2} \cdot A_S/A_A \quad (4)$$

$$P_N(t) = 1 - (P_W + P_R + P_{G1}(t) + P_{G2}(t)) \quad (5)$$

Variables which vary with time during an experiment are indicated by their dependency on current iteration t . Note that P_{G2} is computed by taking into account that a stick available for a grip2 (i.e. which is held by another robot) is only recognized when approached from an angle within an R_{G2} ratio (approximately 35%, see figure 6).⁶

The probabilities defined by equations 1-5 are based on the following assumptions: 1) robots move around the arena with a uniform distribution (i.e. they do not tend to stay longer or shorter in one part of the arena), and 2) the sticks are sufficiently spaced to be accessible from all sides. These assumptions will be discussed in Section 5.

Other modes and other robots It is worth noticing that the description level of both the robot controller's and the probabilistic process' flowcharts is arbitrarily defined by the experimenter. The level is chosen so that behavioral states that do not exert direct influence on the considered metrics (i.e., in this experiment, the collaboration rate) are simply summarized in blocks whose real time duration is taken into account in the probabilistic

⁶Note also that, while the number of robots is fixed for each experiment reported in this article, this is not a prerequisite. For instance, the number of robots could be varied during an experiment to investigate the effect on the collaboration dynamics of a sudden increase or decrease of the number of robots.

model. For instance, when the robot is in either the detection, obstacle-avoidance, grip-and-success-dance, or release mode, its behavior is frozen for a fixed number of iterations (corresponding to the time measured with a real robot, see below) in one of these states. The whole simulation consists of running several processes described in Figure 5 in parallel, with one process per robot, while keeping track of the state of the environment (i.e. $N_{G1}(t)$ and $N_{G2}(t)$, the numbers of sticks available for grip1 and for grip2, respectively). The different processes for the different robots influence each other indirectly by modifying $N_{G1}(t)$ and $N_{G2}(t)$, but also directly when a collaboration occurs: when a robot makes a grip2, one of the robots making a grip1⁷ is randomly chosen to release the stick it is holding.




Figure 6: Geometrical aspects considered for the probabilities' calculation.

Time-iterations transformation Similarly to the methodology proposed in (Martinioli et al., 1999b), the correspondence between iterations in the probabilistic simulation and time in the real experiment is obtained by linking the number of iterations and the time needed to systematically cover the whole arena in the probabilistic simulation and in the real experiment, respectively. In the probabilistic simulation $N = A_A/A_S$ iterations would be needed for the systematic search for sticks (i.e. without passing twice one the same position), while this would take a duration of $T = A_A/(V_R \cdot W_R)$ in the real experiment, where V_R and W_R are the robot's mean forward speed and detection width.⁸ The duration of one iteration therefore corresponds to $A_S/(V_R \cdot W_R)$ (in this case 1.15 seconds, see Table 1). Using this correspondence factor, it is possible to translate the different durations appearing in the real experiment, such as the gripping time, the duration of obstacle avoidance, and the duration of the detection procedure into numbers of iteration.

As the program requires little computation, it is very fast (at least 300 times faster than Webots). Table 1 gives all the parameters used in the simulation. These parameters

⁷Equation 4 ensures that a grip2 can only occur when at least one other robot is making a grip1. If no robot is making a grip1, $N_{G2}(t) = 0$.

⁸We define a systematic search as the search which takes the minimum time (or number of iterations) to discover *all* the elements (robots or sticks) in the arena. It therefore corresponds to traveling around the arena without passing twice on the same position (in this mental experiment, the robot is supposed to "pass through" the elements in the arena without having to detour). For the probabilistic model, the discretization of the arena area must therefore be realized with the detection area of the smallest of the elements in the arena, in this case the stick, for $N = A_A/A_S$ to be the minimum number of iterations for detecting *all* elements in the arena.

Table 1: Parameters used in the probabilistic model. The robot and stick detection distances are given from center (of the robot) to center (of the object). The interference duration is the time lost by a robot for recognizing another robot as an obstacle and turning away from it.

Variable	Value
Arena radius	40.0 cm
Robot detection width (W_R)	14.0 cm
Robot speed (V_R)	8.0 cm/sec
Wall detection distance	6.0 cm
Robot detection distance	10.0 cm
Stick detection distance	6.4 cm
Angle ratio for grip2 (R_{G2})	35%
Duration of one iteration	1.15 sec
Time for distinguishing and gripping a stick	10.0 sec (9 iterations)
Success dance duration	6.0 sec (5 iterations)
Obstacle avoidance duration	1.0 sec (1 iteration)
Interference duration	2.0 sec (2 iterations)

have all been measured from systematic tests with two real robots in the environment. This probabilistic model has therefore no free parameters.

4 Results

We present the results of several experiments implemented at the three different levels: the physical set-up, Webots, and the probabilistic simulation. In the first set of experiments, the influence of the gripping time parameter on the collaboration rate is tested with homogeneous groups of robots, i.e. groups of identical robots (same hardware and same controller). These experiments also evaluate how well the probabilistic model describes the collaboration dynamics. In a second set of experiments, we investigate whether the collaboration rate can be increased by using groups of heterogeneous, rather than homogeneous, robots.⁹ The heterogeneity is introduced at a software level, with robots differing from each other by their gripping time parameter. Finally, in a third set of experiments which are only carried out with Webots and with the probabilistic model, the benefits of a simple communication scheme —directional signalling— is investigated.

4.1 Homogeneous population

We carried out several experiments in order to quantify the influence of the number of robots and the gripping time parameter on the collaboration dynamics, in particular on the collaboration rate. All robots have exactly the same controller —the population is homogeneous— and the experiments are carried out in the same environment (fixed size and fixed number of sticks, see Section 2.1).

With the real Kheperas, a total of 20 runs are carried out with groups of 2 to 6 robots and time parameters equal to 5, 30, 100, and 500 seconds. Each run lasted approximately 20 minutes (the time for the batteries to discharge). With Webots, the influence of

⁹From a hardware point of view, the group is homogeneous (except, of course, for some minimal component differences in the real robotic platform). Heterogeneity in this article therefore only refers to differences in software control parameters.

the time parameter is tested more systematically, with time parameters varying between 5 and 1000 seconds.¹⁰ Each run lasts 30 minutes (simulated time) and is repeated 5 times. Finally, with the probabilistic model, the time parameter is varied between 5 and 1000 seconds by 5-second steps, and each 30-minute run (simulated time) is repeated 100 times. The choice of repeating a run a different number of times depending on the implementation (respectively, once, 5, or 100 times, with the Kheperas, Webots and the probabilistic model), is due to the significant differences in the (real) time necessary for carrying out each run in each implementation.

Figures 7, left and right, present the results of these different runs and illustrates the influence of the gripping time parameter on the collaboration rate and the relative collaboration rate per robot (i.e. the number of collaborations over time to which one robot participates by either making a grip1 or a grip2). Several observations can be made. First, the results with the three different implementations present a good correspondence qualitatively and quantitatively. In particular, almost all the collaboration rates with real robots are within one standard deviation of the mean values obtained either with Webots or the probabilistic model (two standard deviations at maximum). For Webots, this means that the sensor-based simulator faithfully reproduces the sensory information and trajectories of real robots. For the probabilistic model, it shows that, although it is very simple, it incorporates the essential characteristics determining the collaboration dynamics. The probabilistic model and Webots also present a very good quantitative agreement. While the averaged data are smoother with the probabilistic model than with Webots—they are namely the average of 100 runs instead of 5— they have very similar standard deviations, and their mean values \pm one standard deviation always overlap for all group sizes and all gripping time parameters. Figures 11 left and right in appendix, show that the probabilistic model also correctly predicts the rate of failed collaboration and the average time for a robot to find and grasp a stick. For a given group size, the rate of failed collaboration is found to decrease almost exponentially with the gripping time parameter, while the average time to find and grasp a stick remains approximately constant.

Second, the results demonstrate the importance of the influence of the gripping time parameter on the collaboration rate. As could intuitively be predicted, there is a different relation between the collaboration rate and the time parameter depending on the ratio between number of robots and number of sticks. When there are more robots than sticks, the collaboration rate increases monotonically with the gripping time parameter, until a plateau corresponding to the optimal collaboration rate. In other words, it is in this case a good strategy for a robot gripping a stick to wait a very long time for another robot to give a hand, because there will always be at least one “free” robot available. By contrast, when there are fewer robots than sticks, waiting a very long time becomes a bad strategy, as the few robots lose time holding different sticks. For instance (an extreme case), an infinite gripping time parameter would lead to a null collaboration rate with all robots eventually holding a different stick permanently. It is therefore important, when there are fewer robots than sticks, to adjust the gripping time parameter such as to optimize the collaboration rate.

The results also show that, at least with groups of up to 6 robots, the collaboration rate significantly increases with the number of robots, independently of the gripping time

¹⁰In Webots, the gripping time parameter is varied with 5-second steps between 5 and 200 seconds and 25-second steps between 225 and 1000 seconds.




Figure 7: *Left*: Collaboration rate as a function of the gripping time in homogeneous groups of robots. The large single markers correspond to the results with the real robots, the linked small markers to those with the Webots simulator, and the underlying continuous lines to those with the probabilistic simulation. *Right*: Relative collaboration rate per robot (i.e. the average number of collaborations over time to which one robot participates by either making a grip1 or a grip2). Errorbars correspond to \pm standard deviations of the results with the probabilistic model (thicker bars) and with Webot (thin lines). For reasons of clarity, only some errorbars are shown and results are only shown for gripping time parameters up to 600s.

parameter. Interestingly, this increase is super-linear; that is, increasing the number of robots not only increases the global performance of the group (the total collaboration rate, Figure 7 Left) but also the performance of each individual (the collaboration rate per robot, Figure 7 Right). However, as can logically be expected, increasing the number of robots will eventually lead to a diminution of the collaboration rate due to overcrowding and excessive interference. Only 6 Kheperas were available for experiments with real robots. The probabilistic model predicts that the increase of the collaboration rate remains super-linear for groups up to 7 robots, then becomes almost linear for groups up to 11 robots (Appendix, Figure 12, right). The model also predicts that the maximal collaboration rate is obtained with groups of 11 robots and that the collaboration rate then quickly decreases for larger groups due to excessive interference between robots (Appendix, Figure 12, left). For the group of 13 robots, the probability of making a grip is zero in the probabilistic simulation ($P_{G1} = P_{G2} = 0$) because the probability of encountering another robot, P_R , becomes equal to one (the total area of detection of a group of 13 robots indeed covers the whole area of the arena). Note that, for this last situation, the probabilistic model probably over-estimates this effect of interference compared to experiments with real robots, as, unless the arena is really too crowded, small movements within the group of robots can still allow, potentially, some robots to be sufficiently isolated to successfully grip a stick once in a while (i.e. P_{G1} and P_{G2} are not strictly zero). These boundary effects will be further discussed in Section 5.

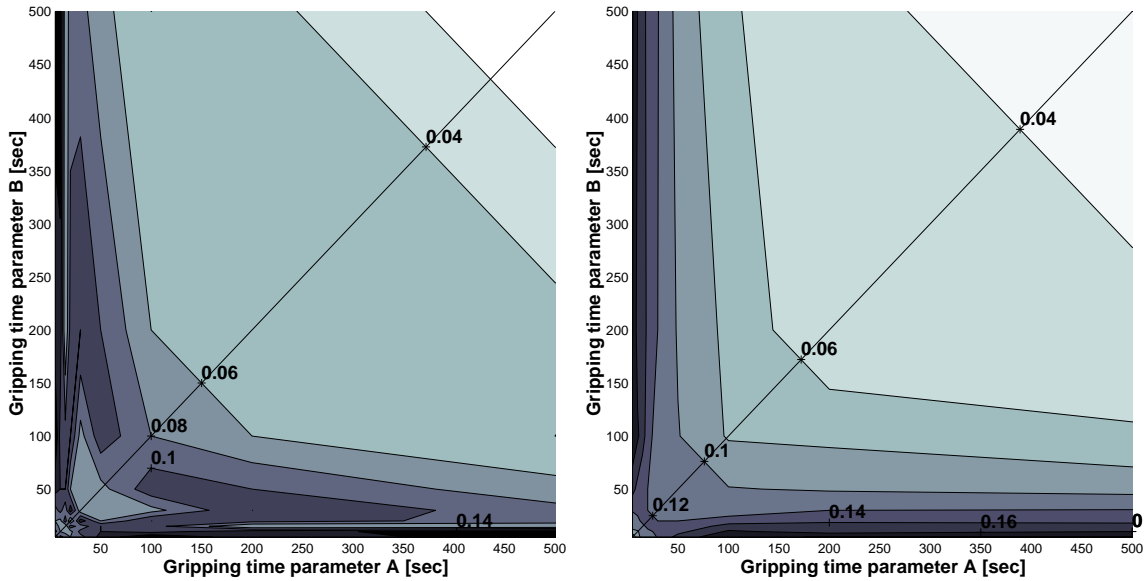


Figure 8: Collaboration rate as a function of the gripping time in a heterogeneous population of two robots. The collaboration rate is proportional to the darkness in the graph. *Left*: Results with the Webots simulator, *right*: results with the probabilistic simulation. In both graphs, the diagonal line corresponds to the collaboration rate in homogeneous groups of robots (time parameter A = time parameter B).

4.2 Heterogeneous population

The experiments carried out so far used groups of robots with identical controllers. In order to investigate whether heterogeneity could increase the collaboration rate of a group of robots of a given size, we carried out a series of experiments in which robots in a group have different gripping time parameters. The heterogeneity is therefore at a software level and concerns only a single parameter.

The experiments were mainly carried out with Webots and the probabilistic model (see below for the implementation with real Kheperas), with groups of 2 to 6 robots. The time parameters are varied between 5 and 1000 seconds in an approximately geometrical progression (5, 10, 15, 20, 30, 50, 100, 200, 500, and 1000 seconds). In order to reduce the number of possible combinations of different time parameters among a group of robots, groups are split into 2 subgroups,¹¹ with a given time parameter for each subgroup. All possible combinations of time parameters between the two subgroups are tested (55 runs). Each run lasts 30 minutes (simulated time) and is repeated 5 times in Webots and 100 times in the probabilistic model.

Figure 8 illustrates the collaboration rate as a function of the gripping time in a group of two robots, with both the Webots and the probabilistic model implementations. Although the collaboration rates are slightly higher with the probabilistic model, and the function representing the dependency of the collaboration rate on the different gripping time parameters smoother (it is namely the average of 100 trials instead of 5 with Webots), the probabilistic model gives a good prediction of the results with Webots.

¹¹The subgroups are of equal size if the number of robots is even, and closest to half the number of robots, if odd (i.e. a group of 5 robots is split in subgroups of 2 and 3).

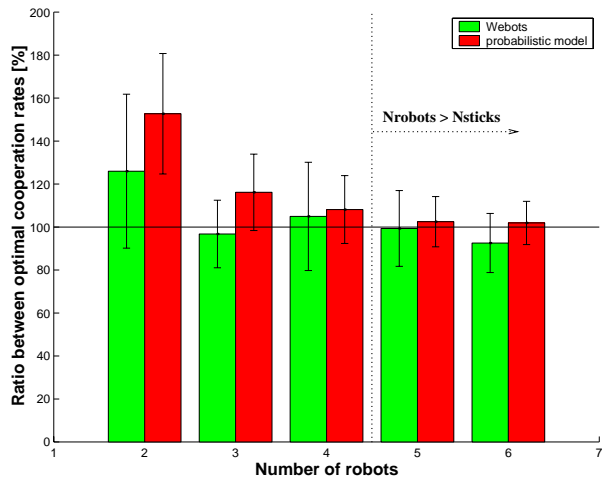


Figure 9: Ratio of optimal collaboration rates between heterogeneous and homogeneous groups of robots. The size of the error bars corresponds to the propagation of the standard deviations of the different runs with the optimal gripping time parameter(s): $\Delta(A/B) = \Delta A/\bar{B} + \Delta B/\bar{A}$, where $\Delta(A/B)$ is the size of the error bar, ΔA and \bar{A} are the standard deviation and average value of the heterogeneous runs with the optimal gripping time parameters, and ΔB and \bar{B} are the standard deviation and average value of the homogeneous runs with the optimal gripping time parameter.

The main outcome of these experiments is that heterogeneity can improve the collaboration rate when there are fewer robots than sticks. For instance, in the case of groups of 2 robots (Figure 8), the optimal collaboration rate in the heterogeneous group (gripping time parameter $A = 5s$, and gripping time parameter $B = 500s$) is approximately 50% better than with the homogeneous population (gripping time parameter = 30s). With groups of 2 or 3 robots, i.e. when the number of robots is less than the number of sticks, optimal collaboration is indeed obtained when one group of robots has a very small gripping time parameter, and the other, a large one (i.e. the dark areas close to the horizontal and vertical axes in Figure 8). This could be seen as the latter group specializing in performing the grip1s (and in waiting for a hand), and the former in performing grip2s, with short grips without waiting. In three initial experiments, we tested a similar specialization with groups of 2 to 4 real Kheperas. The groups were divided into 2 subgroups with gripping time parameters of 5 and 500 seconds. Although these single runs are by no means statistically significant, a 45% increase of the optimal collaboration rate compared to the homogeneous group was also observed with the group of 2 robots, while there was no significant increase for groups of 3 and 4 robots.

As illustrated in Figure 9 which summarizes the differences of optimal collaboration rates between homogeneous and heterogeneous groups, the benefits of heterogeneity disappear when the number of robots exceeds the number of sticks. In that case, the optimal strategy is having all robots waiting a long time once they found a stick, similarly to the homogeneous groups.

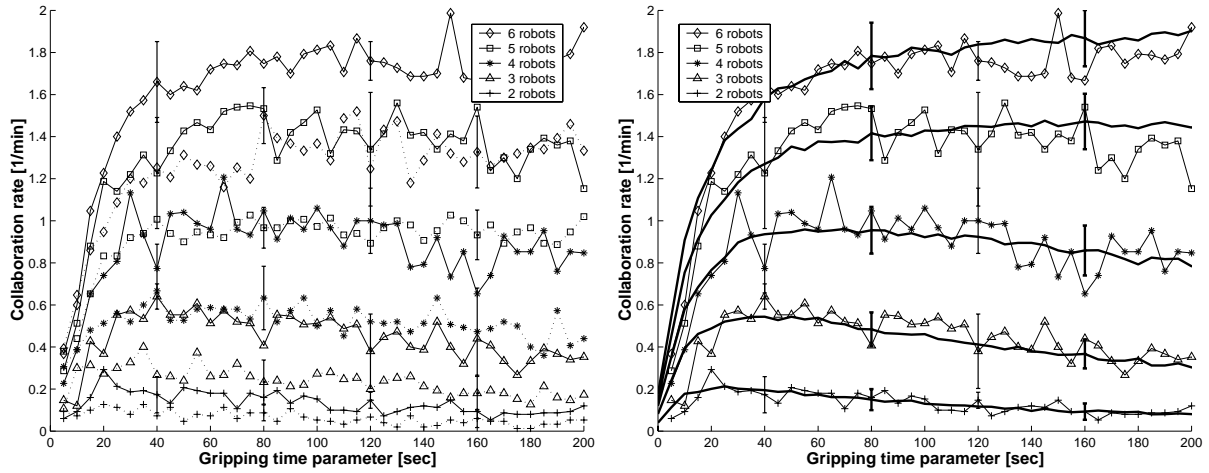


Figure 10: Collaboration rate as a function of the gripping time in a homogeneous population of robots with a simple signalling scheme. *Left:* Results with the Webots simulator. The results with (continuous lines) and without (dotted lines) signalling are superposed; the collaboration rate is systematically higher with signalling. *Right:* superposition of the results with Webots (linked markers) and the probabilistic simulation (continuous lines). Errorbars correspond to \pm standard deviations of the results with the probabilistic model (thicker bars) and with Webot (thin lines).

4.3 Communication

In this last section, we investigate whether introducing a simple communication scheme among the robots can increase the collaboration rate by introducing a less local and more explicit interaction between robots. The experiments are only carried out with Webots and with the probabilistic model, but are based on the IrDA communication turrets developed for the Khepera robots (Martinoli, Franzi, & Matthey, 1997). These turrets allow local communication through 4 directional IR emitters and receivers, separated by angles of 90 degrees.¹²

A simple signalling scheme is implemented as follows. Once a robot grips a stick, it emits a continuous signal in a 60° cone through its frontal emitter (“call for help”). The signal can be perceived by other robots within the whole arena (signalling range is fixed and larger than the arena’s size), as long as they are located within the emission cone. Robots which are in the looking-for-sticks mode and which sense that signal perform a phototaxis towards it until they detect an object, in which case they start the detection procedure as described in Section 2.2. As the emission is directional, robots moving towards the emitter tend to arrive to the caller robot with the right angle for making a grip2, unless they encounter another object on their way (a robot or another stick).

Figure 10 left shows the effect of the signalling scheme on the collaboration rate compared to the experiment without signalling (in the Webots implementation). The signalling systematically improves the collaboration rate, independently of the gripping time parameter. The effect is most visible with groups of few robots, for which the signalling nearly doubles the collaboration rate.

In its current version, the probabilistic model is not perfectly suited for including

¹²The emitters and receivers are on the front, the back, and the sides of the Kheperas.

such a signalling scheme, because the effect of signalling on the probabilities of finding sticks and robots cannot be determined from simple geometrical considerations as before. Difficulties arise partly because the area covered by the call depends significantly on the orientation of the caller¹³ and should be averaged, but mainly because signalling alters the distribution of robots over the arena, and therefore changes the probabilities of meeting other robots in a way which is hard to measure. However, by simply adding two free parameters which modulate the probabilities for making a grip2 and for finding robots, the model can be adapted to present a good fit of the results obtained with Webots. Figure 10 right presents a superposition of the results with Webots and those with the probabilistic model, when the probability of making a grip2 (P_{G2}) is increased by 80% and that of finding a robot (P_R) by 30%. These increases corresponds to what could intuitively be predicted: signalling significantly increases the probabilities of collaborations (i.e. of making a grip2), but also of encountering other robots, as several robots may move towards the same call, therefore increasing the chances of interferences between them. With these new probabilities for making a grip2 and for findings robots, the probabilistic model predicts that, for groups larger than 9 robots, this simple signalling scheme has a negative effect on the collaboration rate, and that groups without signalling perform better than with signalling (Figure 13 in Appendix).

In summary, a simple signaling scheme as presented here can significantly improve the collaboration rate as long as the increase of probability of collaboration outweighs the increase of interference between robots.

5 Discussion

This article presented an experiment in collaborative robotics with the motivations 1) to investigate the collaboration dynamics of a group of simple autonomous reactive robots, and 2) to develop a tool—the probabilistic model—for the characterization and prediction of such an experiment.

Probabilistic model As mentioned, the central idea of the probabilistic model is to represent the dynamics of a group of robots as a series of stochastic events without considering individual trajectories and sensory measurements. Similarly to (Martinoli et al., 1999a, 1999b) where the same approach was taken for the characterization of two different clustering problems, the benefits of the model are two-fold. First, it allows one to pin down the essential characteristics determining the collaboration dynamics. The approach is minimalistic—it includes very few parameters—and aims at capturing only the system parameters which play a relevant role for the metrics in which the experimenter is interested, in our case the collaboration rate. For instance, we took into account parameters such as the robots’ speeds, their detection width, and their gripping time parameter, and did not consider lower level parameters such as the PID parameters of the wheel controllers, the exact positions of the sensors, and the height of the robots, to name a few, as these parameters have either a negligible effect on the metrics we were interested in, or their effect was taken in account by the higher-level parameters that we

¹³The area covered by the call is much smaller when the caller is facing a near wall, than when it has its back to it.

had chosen.¹⁴ In our experiments, the surprisingly good fits between the results with the probabilistic model and those with the real robots and Webots, lead to think that the model represents correctly the dynamics of the experiment, and that it includes all the essential parameters defining it.

The second benefit is the development of a tool for prediction. The model requires only two robots for the setting of all its parameters (i.e. there are no free parameters in the simulation), and these parameters (timings and geometrical considerations) can easily be measured. Notice that the modelling methodology we used in this paper produces “micro-models”: at the local level, the robot-to-robot and robot-to-environment interactions are very precisely defined, with the granularity of details chosen by the experimenter; this micro-model can then be used to predict the collective behavior (the “macro-behavior”) with significantly bigger groups of robots via probabilistic simulations. Because of its speed, the probabilistic simulation is particularly well suited for experimenting with very large groups of robots, and therefore to test the behavior of real swarms. Of special interest, is the possibility of optimizing the number of robots, their hardware and software characteristics for a given objective, before developing dedicated hardware and setting up all the details of the experiment. In that respect, the advantages of the probabilistic model, compared to more complex, sensor-based simulations such as Webots, are the following. First, the probabilistic model is very easy to implement: it consists of a very limited number of lines of code, and is closely based on the flowcharts of the robots controllers. As described in Section 3, it only requires the transformation of local perception situations into probabilities, i.e., in our case, the transformation of deterministic branch operators of the robotic controller into probabilistic ones. A sensor-based simulation requires on the other hand the complete specification of the sensor and actuator properties of the robot, as well as accurate models of the environment. Second, it has few parameters which can readily be measured from a few robots (from two robots in this experiment): as mentioned above, the probabilistic model relies on high level parameters (such as the detection range of a stick, the robot detection width, i.e. parameters that summarize both physical characteristics of sensors, and control thresholds). Unlike sensor-based simulations, it does not require the detailed characterization of each individual sensor (e.g. their position, orientation, detection range, opening angle, or intrinsic noise). Finally, it is very fast: it requires few computations, and comprehensive characterizations of the experiment can be obtained in minutes. This is especially important for computation-intensive tasks such as optimization. Even if sensor-based simulation are faster than running experiments with real robots, making a systematic search for optimal multiple control parameters is often prohibitive with such a type of simulation. In our case, computations which might take almost a year with Webots, can be carried out in one day with the probabilistic model, assuming that the same computer is used.

Note that the probabilistic framework presented here can also serve as basis for a more explicit characterization of the experiment dynamics in terms of time-dependent probabilistic equations (macro-models). An example of such an extension of the approach to a collective exploration task is given in (Billard, Ijspeert, & Martinoli, 1999).

¹⁴The intuition of the engineer plays, of course, an important role in selecting the parameters she/he thinks play a major role on the metrics chosen for an experiment. An iterative process between defining the parameters to be included into the model and comparing the predictions with real data might be necessary to gradually identify the parameters. Note also that some experiments might not be suited to be modelled by this approach (see below).

Types of experiments to which the probabilistic model can be applied, and limitations The modeling methodology used in this paper was specifically designed to predict the dynamics of collective manipulation experiments using distributed control and autonomous, reactive robots. It therefore addresses problems which are pseudo-stochastic in nature, i.e. problems which involve groups of robots with limited navigation capabilities whose trajectories are pseudo-random because of their multiple interactions with other robots and the environment, and because of the noise in their sensors. In that case, trajectories, that is, correlated sequences of robot positions, can be approximated at the level of environmental changes by uncorrelated sequences of random positions, as illustrated by the good agreement of collaboration rates between the probabilistic model and the two other implementations. Unlike sensor-based simulations, the methodology can therefore only be applied to describe the average dynamics of a group of robots and average effect of the group of robots on its environment, a shared resource, rather than giving also examples of instances of an experiment (e.g. typical trajectories followed by the group of robots or which stick was pulled out by which robot).

The methodology currently relies on the assumption that the coverage of the arena by the groups of robots is uniform. In our case, this means that the probabilities of basic events (detecting a robot, a stick or a wall) only depend on geometrical considerations, and do not depend on time (except for our state variables, i.e. the number of grip1s and grip2s made at any time) nor on the positions/orientations of the robots. Uniform coverage might not always be the case, and depends on the environment —e.g. it could have bottle necks leading to higher density of robots in some parts of the environment —, on the robots’ configuration —e.g. their size, the positioning of their sensors —, and/or on the robots’ controllers—e.g. robots might be programmed to follow beacons or to move in flocks. See (Hayes et al., 2000) for probabilistic models that take into account these effects.

As illustrated by the experiments with signalling, experiments in which the probabilities of events are not constant in space and time, but depend on the behavior, position, and orientation of other robots are less straightforwardly implemented. They require some kind of averaging of the probabilities as was realized with the modification of P_R and P_{G2} for fitting the Webots results. This problem can however be solved by either adding free parameters, or extending the model to represent these specific effects more appropriately. For some situations, it might be worth introducing a non-uniform probabilistic density function for the robots’ positions. Experiments with robots that are programmed to generally remain close to walls, for instance, could include a density function in which robots have higher probabilities to be close to a wall.

In our experiments, the probabilistic model also relies on the assumption that the sticks are accessible from all sides for a grip1, which was the case in the real experiments with sticks being sufficiently spaced from each other and from the wall. In some situations, the exact position of a stick might become relevant, and the probabilistic model would need to be modified to take that into account. For instance, a stick which is close to the wall and only partially accessible would lead 1) to a reduction of its probability to be detected for a grip1 (this could easily be measured, or deduced from geometrical considerations), but also 2) to a more subtle effect on the probability of a grip2 as this probability would depend on the angle with which the grip1 is made (low probability if the robot making the grip1 is headed towards the wall, higher probability if the robot making the grip1 is parallel to the wall). Similar situations might occur when sticks are close to each

other in which case having a robot gripping one of the sticks could potentially reduce the probabilities of grips on the other sticks. In its current form, the probabilistic model is not well suited to describe these situations which depend on the trajectories, positions, and orientations of each robot.

As briefly discussed in Section 4, the probabilistic model also needs to be improved to better describe boundary effects such as, for instance, situations with many robots whose total area of detection comes close to, or exceeds, the total area of the arena. In such an “overcrowded” situation, the probabilistic model needs to be modified to take into account the fact that detection areas overlap (despite their obstacle-avoidance behavior, robots are forced to remain within each others detection range). There is no fundamental obstacle to consider this overlap in the modelling methodology, but the exact overlap might be difficult to assess as it intimately depends on details of the robots’ behavior during obstacle avoidance e.g. how much and how fast they turn away from an obstacle in comparison to how much time and space they need to perform a gripping operation.

Finally, some types of robot’s controllers might be difficult to be implemented into the probabilistic model. Creating the probabilistic model requires that the controllers (e.g. rule-based, behavior-based, neural networks-based controllers) present a relation from sensory space to action space which can be classified into a discrete set of actions triggered by well defined sensor-states. These different actions are then represented by branches in the probabilistic model, and probabilities are attached to each of these actions to represent how often they occur with the real robots.¹⁵ Controllers for which such a clear and discrete mapping does not exist, e.g. whose actions depend in a complex way on low level information from the individual sensors, are therefore not suitable to be modelled with this methodology. Note however that only actions which are meaningful for the chosen metrics need to be distinguished. In our case, for instance, obstacle-avoidance is implemented in a trivial neural network and is considered as a single action (i.e. turning away from an obstacle by 30 degrees or 60 degrees are not considered as different actions) which is then captured in the model only by its effective duration measured on a real robot. A last limitation is that the current methodology does not deal with adaptive controllers, and it remains to be seen how much adaptation it can integrate (see future works below).

A Swarm Intelligence approach towards collaboration The experiment presented in this article is one of the few experiments in *strictly* collaborative robotics implemented with real robots rather than only in simulation, and, more generally, one of the few experiments in collaborative robotics with groups of more than three robots. The experiment presented here is characterized 1) by the fact that it uses robots which are simple, reactive, and with local and noisy sensory information, 2) by the fact that the robots are autonomous and do not rely on an external supervisor or a robot leader, and 3) by the fact that collaboration is obtained by exploiting only local interactions (stigmergic or explicit communication mechanisms such as directional signalling). Collaboration occurs when there is a spatial and temporal coordination between two robots. In the experiments without signalling, this coordination happens randomly rather than being actively sought,

¹⁵Note that the controllers need not to be deterministic, the probabilistic model can easily integrate probabilistic action-selection, as long as the internal probabilities for each action given a sensori-state are known.

as robots do not coordinate their actions except through a *stigmergic* communication, i.e. by influencing the state of the environment (the state of the sticks).

Similar collaborative experiments with a SI approach were, for instance, presented in (Beckers et al., 1994; Holland & Melhuish, 1999; Kube & Bonabeau, 2000; Krieger & Billeter, 2000; Hosokawa et al., 1998; Yoshida et al., 1999). This is in contrast with many other experiments which, for instance, either relied on a central supervisor (Humberstone & Smith, 2000) or robot leaders (Wang et al., 2000) for generating guide lines and coordinating the group behavior, used extensive wireless communication for action coordination and compensating sensory weaknesses (Parker, 1994; Mataric et al., 1995), or used more sophisticated sensors (e.g. force and torque sensors) and models of the dynamics of the robot-to-environment interaction (Boehringer et al., 1995; Khatib, 1999; Ota & Arai, 1999).

Consistently with the SI approach, we minimized the complexity of the individual units by using the simplest possible sensing and communication capabilities, without centralized control. We therefore obtained good scalability of the system by avoiding global communication schemes and corresponding bottlenecks due to a limited bandwidth. The price to be paid for renouncing to any form of global information and global networking is system efficiency since action coordination and collaboration is based on probabilistic rather than deterministic rules.

Collaboration dynamics The collaborative essence of the experiment leads to the following observations.

First, the dynamics of the experiment —the relation between collaboration rate and gripping time parameter, for instance— differ significantly depending on the ratio between number of workers (i.e robots) and amount of work (i.e. number of sticks). This is due to the necessity for spatial and temporal coordination between at least two robots for successful collaborations. When there are few robots and a large amount of work (many sticks), special care must be given to prevent robots to disperse spatially and “temporally” (temporal dispersion corresponding to the situation where robots arrive at a same stick but with important time differences). Optimizing the time overlap for collaboration can be obtained by optimizing the gripping time parameters, as investigated with homogeneous and heterogeneous groups. Having homogeneous groups with a too high or especially a too low time parameter indeed leads to a strongly suboptimal performance. As observed in Section 4.2, the best performance is then obtained with heterogeneous groups and specialization. A natural continuation of this work would therefore be to implement adaptive rules for the robots to determine themselves their gripping time parameter, as, ideally, the robots should be able to adapt to the current work load and number of robots in the arena (as in (Parker, 2000; Touzet, 2000), for instance). It would also be worth investigating simple ways to prevent spatial dispersion (other than signalling) with, for instance, the implementation of a following-routine such that robots travel in pairs.

Although quantitative analysis is often missing in related research, we believe that the importance of the ratio between number of workers and amount of work can be generalized to a whole range of similar collaborative tasks, including object-pushing (Parker, 1994; Boehringer et al., 1995; Mataric et al., 1995; Wang et al., 2000), transportation (Johnson & Bay, 1995; Ghanea-Hercock & Barnes, 1996; Ota & Arai, 1999; Humberstone & Smith, 2000) and manipulation (Fujita & Kimura, 1998; Khatib, 1999), and that similarly strong differences of the “qualitative dynamics” of the experiment can be made between under-

staffed and adequately staffed groups of robots for a given amount of work. The exact dynamics depends, of course, on the details of task. In the box pushing experiment presented in (Kube & Bonabeau, 2000), for instance, the size and type of box is found to significantly influence the mean execution time of moving a box towards a goal for a given group size. When the influence of the group size is tested, it is found that pushing a box is done fastest with the smallest group of robots tested (with three robots). Had the box been heavier, the results would probably have been quite different (and closer to the experiment presented in this paper) with small groups being less efficient or even unable to move the box. As mentioned by the authors “... performance is dependent on some yet to be determined *task density* function”. In this paper, we experimentally defined such a function for a specific task in a specific environment. Our investigation is however not exhaustive (we did not, for instance, vary the number of sticks, see below), and further experiments to characterize such a task density in a broader sense will be mentioned below.

Second, the need for collaboration leads to a collaboration rate which, in groups up to a certain size, increases super-linearly with the number of robots. As mentioned, increasing the number of robots not only increases the global performance of the group, but also the performance of each individual. This is rarely the case in autonomous collective robotics which is not strictly collaborative.¹⁶ The increase of performance becomes sub-linear when the effects of interference outweigh the benefits of having more robots for collaboration.

Finally, the preliminary experiment with the directional signalling scheme showed how introducing a less local interaction between robots could significantly improve the collaboration rate, especially with small groups of robots. The interesting aspects of signalling are clear: they increase the probabilities of having the right spatial and temporal coordination for collaborations and they reduce the overall stochasticity of collaborations. In this simple case, where there is no agreement between robots receiving a call about which one should handle it, signalling also increases the interference between robots which means that for big groups of robots signalling leads to worse collaboration rates than no signalling. While a detailed analysis of the respective benefits of different communication protocols is out of the scope of the current article, the analysis of which communication protocol and range is most appropriate for a given task and a given number of robots is a central question in collective robotics (Balch & Arkin, 1994; Yoshida, Arai, Yamamoto, & Ota, 1998).

An interesting observation from these experiments is that for a given environment, a given task and a given initial group of robots, there are often multiple ways to improve the collaboration rate among the group. The collaboration rate can indeed be improved by, for instance, adding or removing robots, giving them better sensors, improving their controllers and adding specialization, or introducing signalling. Figure 10 left, for instance, shows that, given a group of 3 robots without communication, an almost identical improvement of collaboration rate can be obtained by either adding a fourth robot or giving the capacity to the 3 robots to communicate and use the signalling scheme described above. Each of these means has its advantages and drawbacks in terms of costs, time to implement, complexity (consistently with the minimalism idea of the SI approach), and system robustness. As mentioned above, having a prediction tool such as the probabilistic model is therefore most useful for investigating these issues and taking costs into consid-

¹⁶In (Mataric et al., 1995) another strictly collaborative experiment, superlinearity in the pushing time was also reported for the group of 2 robots.

erations before having to effectively buy and/or implement the different possibilities.

Future work The experiments presented in this article are by no means exhaustive, and need to be extended in several directions. First, the predictions of collaboration rates for groups larger than 6 robots need to be tested with Webots and the real robots. Of particular interest, is the predicted collapse of collaboration rate for groups larger than 11 robots. Second, more experiments need to be made with Webots and the real robots in order to assess how well the probabilistic model predicts the collaboration rate of the real robots compared to Webots from a statistical point of view. Current numbers of runs are too limited and too different from one type of implementation to the other to allow statistical tests to be used in a meaningful way. Further experiments are also required to characterize the task density function, not only as a function of the number of robots, but also of the number of sticks and the size of the arena. Experiments with very large number of sticks would be especially interesting in order to investigate how much heterogeneity in the gripping time parameters can improve the collaboration rate in a large group of heterogeneous robots compared to homogeneous robots. Finally, it would be interesting to test whether the methodology could be applied to adaptive controllers. The current probabilistic model could certainly help to develop an adaptive rule for each robot to determine its gripping time parameter given the number of successful and unsuccessful collaborations it makes as well as the number of obstacles it encounters, for instance. It remains to be investigated however how much the methodology is applicable when the whole controller organization is adaptive rather than just a few parameters.

6 Conclusion

This article investigated the collaboration dynamics among groups of simple, reactive, autonomous robots involved in a collaborative stick pulling task. In particular, a probabilistic model was developed for the characterization and the prediction of such dynamics. It was found that, by representing the experiment as a set of stochastic events with probabilities based on simple geometrical considerations, the probabilistic model was able to provide a very good prediction, both qualitatively and quantitatively, of the collaboration dynamics as a function of two main parameters studied in the experiment, the number of robots and their gripping time parameter.

Four observations can be made from the experiments. First, the experiments showed that there are two different dynamics depending on the ratio between the number of robots and the number of sticks. Second, with group sizes up to certain size (six robots), the collaboration rate increases *super-linearly* with the number of robots. Third, heterogeneity in the robots' controller parameters increases the collaboration rate when there are fewer robots than sticks while having no significant effect otherwise. Finally, the experiments showed that a simple signalling scheme can significantly improve the collaboration rate among robots unless the group of robots is large (larger than nine robots) in which case the increased interference between robots due to the chosen signaling scheme outweighs the benefits of increased coordination between robots.

Acknowledgments We gratefully acknowledge fruitful discussions with Jean-Louis Deneubourg, Francesco Mondada, and Guy Theraulaz, as well as insightful comments

by Marco Dorigo and Rodney Goodman on an earlier version of the article. We also acknowledge constructive comments by George Bekey and three anonymous reviewers. This project was funded by the Swiss National Science Foundation under a project entitled “A Methodology for Collective Robotics Design”.

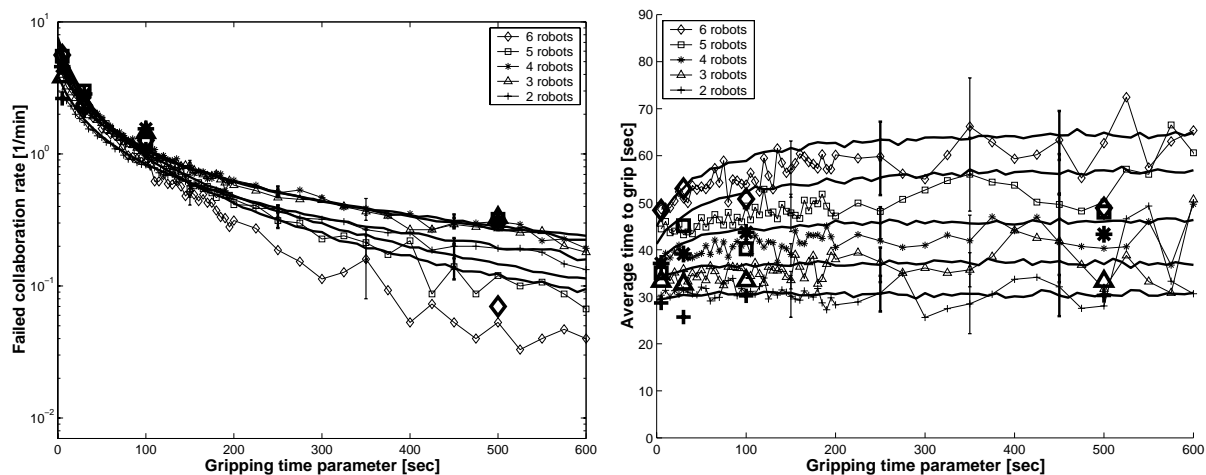


Figure 11: *Left*: failed collaboration rate with homogeneous groups of robots. *Right*: average time to find and grip a stick with homogeneous groups of robots. The large single markers correspond to the results with the real robots, the linked small markers to those with the Webots simulator, and the underlying continuous lines to those with the probabilistic simulation. Errorbars correspond to \pm standard deviations of the results with the probabilistic model (thicker bars) and with Webot (thin lines). Note that the correspondance between Webots and the probabilistic model on one side, and the real robots on the other side, is less good for the groups of 5 and 6 robots with gripping time parameter equal to 500sec. We believe this is due a statistical artifact which will disappear when more runs with the real robots will be carried out (see the future work Section).




Figure 12: Collaboration rates (left) and relative collaboration rates (right) predicted by the probabilistic model for homogeneous groups of robots with size larger than six.




Figure 13: Comparison of collaboration rates predicted by the probabilistic model for homogeneous groups of robots with size larger than six, with and without signalling.

References

- Balch, T., & Arkin, R. (1994). Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1), 27-52.
- Balch, T., & Arkin, R. (1998). Behavior-based formation control for multirobots teams. *IEEE Trans. on Robotics and Automation*, 14(6), 926-939.
- Beckers, R., Holland, O., & Deneubourg, J.-L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks & P. Maes (Eds.), *Proc. of the Fourth Workshop on Artificial Life* (p. 181-189). MIT Press.
- Beni, G., & Wang, J. (1989). Swarm intelligence. In *Proc. of the Seventh Annual Meeting of the Robotics Society of Japan, Tokyo, Japan, 1989* (p. 425-428).
- Billard, A., Ijspeert, A., & Martinoli, A. (1999). A multi-robot system for adaptive exploration of a fast changing environment: Probabilistic modelling and experimental study. *Connection Science*, 11(3/4), 359-379.
- Boehringer, K., Brown, R., Donald, B., Jennings, J., & Rus, D. (1995). Distributed robotic manipulation: Experiments in minimalism. In O. Khatib & J. K. Salisbury (Eds.), *Proc. of the Fourth Int. Symp. on Experimental Robotics, Stanford* (p. 11-25). Lecture Notes in Control and Information Sciences, Springer Verlag.
- Bonabeau, E., Dorigo, M., & Theraulaz, G. (1999). *Swarm intelligence: From natural to artificial systems*. SFI Studies in the Science of Complexity, Oxford University Press.
- Chauvin, R., & Janin, P. (1975). Facteurs de direction et d'excitation au cours de l'accomplissement d'une tache chez formica polyctena. *Insectes Sociaux*, 22, 199-206.
- Everett, H. R., Gilbreath, G. A., Heath-Pastore, T. A., & Laird, R. T. (1993). Coordinated control of multiple security robots. In *Mobile Robots VII* (Vol. 2058, p. 292-305). SPIE, Arlington, VA.
- Fujita, T., & Kimura, H. (1998). Tight cooperative working system by multiple robots. In *Proceedings of the 1998 IEEE/RSJ International Conference on Intelligent Robots and Systems* (p. 1405-1410). IEEE Computer Society.
- Gage, D. (1995). Many-robots MCM search systems. In A. Bottoms, J. Eagle, & H. Bayless (Eds.), *Proc. of the Autonomous Vehicles in Mine Countermeasure Symposium* (p. 9.55-9.63). Monterey, CA.
- Ghanea-Hercock, R., & Barnes, D. P. (1996). An evolved fuzzy reactive control system for co-operating autonomous robots. In P. Maes, M. J. Mataric, J.-A. Meyer, J. Pollack, & S. W. Wilson (Eds.), *Proc. of the Fourth Int. Conf. on Simulation of Adaptive Behavior: From Animals to Animats, Cape Cod, MA, September, 1996* (p. 599-607). MIT Press.

- Hayes, A. T., Martinoli, A., & Goodman, R. M. (2000). Comparing distributed exploration strategies with simulated and real autonomous robots. In L. Parker, G. Bekey, & J. Bahren (Eds.), *Proc. of the Fifth Int. Symp. on Distributed Autonomous Robotic Systems DARS-00, October, 2000, Knoxville, TN* (p. 261-270). Springer Verlag.
- Holland, O., & Melhuish, C. (1999). Stigmergy, self-organization, and sorting in collective robotics. *Artificial Life*, 5, 173-202.
- Hosokawa, K., Tsujimori, T., Fujii, T., Kaetsu, H., Asama, H., Kuroda, Y., & Endo, I. (1998). Mechanisms for self-organizing robots which reconfigure in a vertical plane. In T. Lueth, R. Dillman, P. Dario, & H. Wuern (Eds.), *Proc. of the Fourth Int. Symp. on Distributed Autonomous Robotic Systems, Karlsruhe, Germany, May, 1998* (p. 111-118). Springer Verlag.
- Humberstone, C. K., & Smith, K. B. (2000). Object transport using multiple mobile robots with pin joint endeffectors. In L. E. Parker, G. Bekey, & J. Barhen (Eds.), *Proc. of the Fifth Int. Symp. on Distributed Autonomous Robotic Systems, DARS-00, October, 2000, Knoxville, TN* (p. 417-426). Springer Verlag.
- Johnson, P., & Bay, J. (1995). Distributed control of simulated autonomous mobile robot collectives in payload transportation. *Autonomous Robots*, 2, 43-63.
- Khatib, O. (1999). Mobile manipulation: the robotic assistant. *Robotics and Autonomous Systems*, 26, 175-183.
- Krieger, M. B., & Billeter, J.-B. (2000). The call of duty: Self-organised task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, 30(1-2), 65-84.
- Kube, C. R., & Bonabeau, E. (2000). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1-2), 85-101.
- Martinoli, A. (1999). *Swarm intelligence in autonomous collective robotics: From tools to the analysis and synthesis of distributed control strategies*. Unpublished doctoral dissertation, PhD Thesis No 2069, EPFL.
- Martinoli, A., Franzi, E., & Matthey, O. (1997). Towards a reliable set-up for bio-inspired collective experiments with real robots. In A. Casals & A. de Almeida (Eds.), *Proc. of the Fifth Int. Symp. on Experimental Robotics, Barcelona Spain, June, 1997* (p. 597-608). Lecture Notes in Control and Information Sciences, Springer Verlag.
- Martinoli, A., Ijspeert, A., & Gambardella, L. (1999b). A probabilistic model for understanding and comparing collective aggregation mechanisms. In D. Floreano, F. Mondada, & J.-D. Nicoud (Eds.), *Proceedings of the Fifth European Conference on Artificial Life, ECAL99* (p. 575-584). Springer Verlag, Lecture Notes in Computer Science.
- Martinoli, A., Ijspeert, A., & Mondada, F. (1999a). Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, 29(51-63), 51-63.

- Martinoli, A., & Mondada, F. (1995). Collective and cooperative group behaviours: Biologically inspired experiments in robotics. In O. Khatib & J. K. Salisbury (Eds.), *Proc. of the Fourth Int. Symp. on Experimental Robotics* (p. 3-10). Lecture Notes in Control and Information Sciences, Springer Verlag.
- Mataric, M. (1994). *Interaction and intelligent behavior*. Unpublished doctoral dissertation, Dept. of Electrical Engineering and Computer Science, MIT, Cambridge.
- Mataric, M., Nilsson, M., & Simsarian, K. (1995). Cooperative multi-robot box-pushing. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems* (p. 556-561). IEEE Computer Society.
- Michel, O. (1998). Webots: Symbiosis between virtual and real mobile robots. In J.-C. Heudin (Ed.), *Proc. of the First Int. Conf. on Virtual Worlds, Paris, France*, (p. 254-263). Springer Verlag. (See also <http://www.cyberbotics.com/webots/>)
- Mondada, F., Franzi, E., & Ienne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In T. Yoshikawa & F. Miyazaki (Eds.), *Proc. of the Third Int. Symp. on Experimental Robotics* (p. 501-503). Lecture Notes in Control and Information Sciences, Springer Verlag.
- Ota, J., & Arai, T. (1999). Transfer control of a large object by a group of mobile robots. *Robotics and Autonomous Systems*, 28, 271-280.
- Parker, L. (1994). ALLIANCE: an architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *Proc. 1994 IEEE/RSJ/GI Int. Conf. Intell. Robot. Systems (IROS'94), Munich, Germany, September, 1994* (p. 776-783).
- Parker, L. (1998). ALLIANCE: An architecture for fault tolerant multirobot cooperation. *IEEE Trans. on Robotics and Automation*, 14(2), 220-240.
- Parker, L. (2000). Lifelong adaptation in heterogeneous multi-robots teams: reponse to continual variation. *Autonomous Robots*, 8(3), 239-267.
- Touzet, C. (2000). Robot awareness in cooperative mobile robot learning. *Autonomous Robots*, 8(1), 87-97.
- Wang, J., & Premvuti, S. (1995). Distributed traffic regulation and control for multiple autonomous mobile robots operating in discrete space. In *Proc. of the IEEE Int. Conf. on Robotics and Automation* (p. 1619-1624). IEEE Computer Science.
- Wang, Z. D., Kimura, Y., Takahashi, T., & Nakano, E. (2000). A control method of a multiple non-holonomic robot system for cooperative transportation. In L. E. Parker, G. Bekey, & J. Barhen (Eds.), *Proc. of the Fifth Int. Symp. on Distributed Autonomous Robotic Systems, DARS-00, October, 2000, Knoxville, TN* (p. 447-456). Springer Verlag.
- Yamauchi, B. (1999). Decentralized coordination for multirobot exploration. *Robotics and Autonomous Systems*, 29(1), 111-118.

- Yoshida, E., Arai, T., Yamamoto, M., & Ota, J. (1998). Local communication of multiple mobile robots: Design of optimal communication area for cooperative tasks. *J. of Robotic Systems*, 15(7), 407-419.
- Yoshida, E., Murata, S., Tomita, K., Kurokawa, H., & Kokaji, S. (1999). An experimental study on a self-repairing modular machine. *Robotics and Autonomous Systems*, 29(1), 79-89.