Efficiency and Optimization of Explicit and Implicit Communication Schemes in Collaborative Robotics Experiments

Kjerstin I. Easton, Alcherio Martinoli

Collective Robotics Group, California Institute of Technology, Pasadena, U.S.A.
easton, alcherio@caltech.edu

Abstract

This paper presents the investigation of three communication schemes which may be used in a distributed robotic system, two based on implicit forms of communication (mechanical interaction and vision) and one based on an explicit form of communication (infrared signaling). To support the discussion and comparison between the three forms, we have chosen a concrete case study concerned with locating and pulling sticks out of an arena floor, a task successfully achieved only through collaboration between two robots. Communication schemes, among other system features, heavily influence the rate of successful collaborations, the metric adopted in this paper in order to evaluate the performance of the robotic team. Results collected using an embodied simulator show that, as a function of the system constraints (e.g., number of robots, hardware and behavioral parameters) solutions based on more complex individuals do not necessarily lead to an improved team performance. Although the stick pulling is a simple case study without any practical application, it presents all the main difficulties of designing and controlling scalable, distributed robotic systems, characterized by subtle, nested effects between individual and group behavior or hardware and software parameters. We believe that embodied simulations are a key level of implementation in helping us understand these subtle mechanisms, achieve further abstraction, and optimize the system before any real hardware solution is implemented.

1. Introduction

Swarm Intelligence (SI) is a computational and behavioral metaphor for solving distributed problems that takes its inspiration from biological examples provided by social insects [1]. The abilities of such natural 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 is mediated by nothing more than a small set of simple low-level interactions between individuals and between individuals and the environment. The three main advantages of the SI approach to the control of a group of robots are scalability, flexibility, and robustness [2]. Collective systems based on an SI approach are robust not only through unit redundancy but also through the unit minimalistic design [3]. 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 as capable as possible group behavior. Minimizing the individual complexity could in turn help to reduce the costs of the collective solution, an important characteristic particularly for large distributed systems. To minimize the cost of designing a collective system, it would be prudent to optimize and predict the benefit of a hardware or behavioral solution in simulation before implementing it for a large number of units. It has been shown in several cases that an embodied simulator can faithfully reproduce real robot experiments, in particular when higher levels of abstraction may have failed to quantitatively predict the system dynamics without free parameters [4,5,6,7].

2. The Stick-Pulling Experiment

The experiment presented in this article is a follow-up to tests presented in [4]. Although the stick-pulling experiment has no practical application per se, it captures well the class of problems engineers have to face in designing and controlling fully distributed robotic systems. In addition, as shown in [7], the collaborative nature of the stick-pulling task could be easily generalized to applications in which several robots need to coordinate their activity in both space and time to accomplish their mission. For instance, in a distributed sensing problem that requires continuous monitoring of an area as well as additional attention by an array of n sensors if particular events arise, we could use control and communication schemes similar to those discussed in this paper.

All behaviors in the experiment are based on local interactions and communications, according to SI principles. The task is for Khepera robots equipped with grippers and a belt of proximity sensors to locate sticks in a circular arena and to pull them out of the ground. Because of the length of a stick, a single robot is not capable of pulling it out of the ground alone; collaboration between two robots is necessary to complete the extraction.

2.1 Physical Setup

The experiment is carried out in a circular arena 80 cm in diameter, delimited by a white wall. Four holes at the corners of a square with 30 cm edges hold white sticks (15 cm long, diameter of 1.6 cm) which, in their lowest position, protrude 5 cm above the ground. Groups of 2 to 6 Khepera robots pull the sticks out of the ground (see Figure 1). Collaboration between robots is required for success because the stick is too long for one robot to extract in a single pull. After a successful collaboration, the stick taken out of the ground is released by the robot, and replaced in its hole by the experimenter.

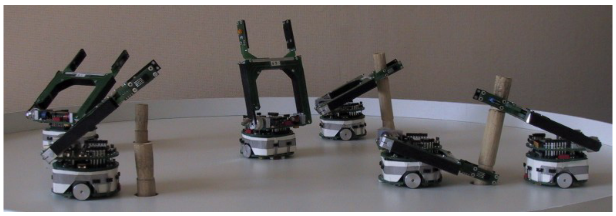

Figure 1: *Physical set-up for the stick-pulling experiment*

2.2 Stick-pulling Robot Controller

The default robot behavior is to wander in the arena in a search mode, moving in a straight line until the frontal proximity sensors detect an object. The robots can distinguish sticks from obstacles (walls, other robots) because of the sticks' thinness. If the object is an obstacle, the robot turns away, performs obstacle avoidance for a few seconds, and returns to the search behavior. If the object is a stick, the robot backtracks a few centimeters, grips the stick and pulls it up.

Basic implicit communication between the robots occurs during the gripping process; the robot determines its role in the collaboration from the effect of the other robot's grip on the stick. While pulling up on a stick, the robot determines whether another robot is already gripping the same stick by measuring the speed of elevation of the gripper arm. If no other robot is holding the stick, we call such a grip a *grip1*. If another robot is already holding that stick and therefore "braking" the elevation, such a grip is called *grip2*. When a robot makes a grip1, it holds the stick partway out of the ground and releases it when either the duration of the grip exceeds a *gripping time parameter* or another robot comes to complete the collaboration, making a grip2. A 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. If a grip2 is made, the grip1 robot will release the stick. To mark the successful collaboration, the robot that made grip2 performs a short "success dance" (moving the arm up and down) and releases the stick, which is replaced in the hole by the experimenter. After releasing a stick, the robots resume searching for sticks.

Because sticks are recognized by their thinness, a stick can only be recognized when approached from the opposite side within a certain angle (approx. 126°), limiting the probability of collaboration, but preventing tangling of the robots' grippers. From such an angle, the approaching robot cannot detect the robot holding the stick. For other angles of approach, both the stick and the robot are detected and the whole is taken to be an obstacle.

2.3 The Embodied Simulator

The experiment has also been implemented in Webots [8], a 3D kinematic, sensor-based simulator of Khepera robots. The simulator computes trajectories and sensory input of the robots in an arena corresponding to the physical set-up. The simulation is sufficiently faithful for the controllers to be transferred to real robots without changes and for the robot behaviors in simulation to be very similar to those of the real robots, as shown in several previous papers [4,5,6].

3. Communication Scheme Implementation

The communication schemes proposed in this paper are each a form of broadcasting; there are no handshaking mechanisms between the emitter and the receiver and the emitter does not target a specific receiver when sending a message. The differences between the communication schemes lie mainly in the physical layer of the communication channel and the range of communication.

3.1 Basic Communication Scheme

The basic communication scheme uses the simplest controller and robot hardware of the three schemes. The controller is described in section 2.2 and the robot is equipped with only the proximity sensors on the robot base and the gripper turret.

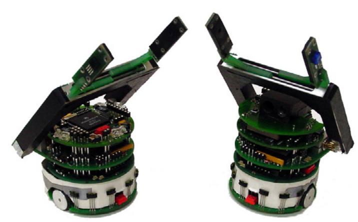

Figure 2: *Khepera robots equipped with gripper and IrDA modules (left),and gripper and camera modules (right)*

3.2 Infrared Signaling Scheme

The experiments exploiting infrared (IR) signaling are based on the IrDA communication turrets developed for the Khepera robots [9] (Figure 2). These turrets allow local communication through four directional IR emitters and receivers, separated by angles of 90 degrees. The robot behavior is modified to exploit the explicit communication capability: when a robot grips a stick, it emits a continuous signal from its frontal emitter in a 60° cone. Any robot in search mode within the cone will receive the signal in one or more of its four IR receivers. Stick-searching robots sensing the signal perform phototaxis towards it until they detect an object, at which point they proceed with the original behavior. Because the emission is directional, robots moving towards the emitter tend to arrive at the calling robot at an acceptable angle for performing a grip2.

Though not enough of the prototype IrDA turrets were available to perform real robot experiments, the scheme was fully implemented in Webots.

3.3 Vision-based Communication Scheme

The experiments using visual communication use the K213 linear camera turret available for the Khepera robot (Figure 2). The K213 captures a horizontal 1-by-64-pixel grayscale image and has a 36° field of view. When installed on a robot, the camera's field of view falls above the top of sticks that have not been lifted. A 15 cm-high black backdrop was installed around the arena and visible parts of the robots were masked with black tape. Because the sticks are white against the black background, they are visible to other robots in the arena when raised.

To exploit the new hardware, stick-searching robots have a modified search behavior with three parameters, the *scan interval*, in centimeters, the *scan speed*, indicated by the time required to complete one scan revolution, and the *scan angle*, in degrees. While searching, if the robot identifies a stick within its field of view, it will drive towards it, centering the stick in the image. In the course of the search behavior, if the robot travels a distance greater than the scan interval, it rotates about its center by a specific amount, defined by the scan angle, in a randomly chosen direction. If a stick becomes visible at any point in the rotation or if the scan angle is reached, the robot will stop scanning and drive forward in search mode. Experiments were performed primarily in Webots, but the vision-enabled behavior was implemented and tested for validation on real robots.

4. Results and Discussion

The basic, IR-, and vision-based communication schemes were implemented in Webots to explore the effects of behavioral and/or hardware parameters for each scheme. 10 runs were performed for each experiment in Webots while only 3 runs were performed with real robots. All the error bars in the plots represent the standard deviation over the multiple runs.

4.1 Explicit Communication using IR Signaling Scheme

The implementation in [4] showed that the IR communication scheme systematically increased the collaboration rate, showing an especially significant increase for small group sizes.

The improvement in performance appears to be most pronounced the less crowded in the environment. For a significantly large number of robots, many more robots than are needed will be summoned by the IR "call for help." Two probabilities governing important dynamics of the system are affected by the introduction of this explicit communication [4]. For robots that receive the signal, the probability they will encounter another robot and enter the obstacle avoidance behavior increases due to the greater density of robots in the neighborhood of the signaling cone and the stick. Alone, this effect would reduce performance because robots performing phototaxis towards the emitting robot are more likely to interfere with one another, causing them to enter obstacle avoidance mode, therefore reducing the amount of time they spend searching for sticks. This effect is overpowered, however, by the increase in the probability a robot will grip the signaling robot's stick, especially in small groups of robots.

The original Webots IR implementation, however, did not take into account occlusion of the signal by robots in the arena; the robots are opaque to the IR signal. After introducing occlusion to the simulation, robots in the signaling cone but not in line-of-sight of the transmitter do not receive the signal. This lowers the probability of interference between robots in the signaling cone, leading to a performance increase over the original implementation of IrDA signaling (Table 1).

Table 1: *Comparison of IR communication in [4] to IR communication with occlusion. C_o is the optimal collaboration rate (1/min). "%" indicates the percent improvement over stick pulling without IR communication. C_o was obtained by systematic search over all possible gripping time parameters*

	IR [4]			IR with occlusion		
Group size	$\mu(C_o)$	$\sigma(C_o)$	%	$\mu(C_o)$	$\sigma(C_o)$	%
2	0.250	0.095	96.8	0.240	0.078	89.0
4	1.060	0.105	58.9	1.063	0.144	59.4
6	1.890	0.249	24.3	2.043	0.108	34.4

4.2 Significance of IR Emission Angle

Though the signaling robot's message is a simple 1-bit signal (ON or OFF), the angle of emission is an integral part of that message. It limits and identifies which robots are eligible to receive the message. The angle of emission directly affects how many robots are summoned to help collaborate and at which angle relative to the signaling robot they will arrive.

Changing the emission angle to 126°, the angle of acceptable approach for collaboration, yielded performance similar to that achieved with an emission angle of 60° for all group sizes tested (an example, with group size of 6 is in Figure 3). It appears the increased probability of a robot making a grip2 (due to the larger area covered by the signal) is counterbalanced by the increased probability of interference between robots. It was shown in [4] that these two probabilities have a nonlinear relationship, so we do not expect to see this counterbalancing effect for all choices of emission angle.

Increasing the angle to 270°, for example, showed a decrease in performance relative to the other emission angles. The most dramatic decrease was seen in the largest group size simulated, 6 robots, which has a maximal collaboration 32% lower than the maximum achieved by robots with no IR communication. This indicates that the larger emission angle attracts robots that are not capable of making a grip2 and their presence in the neighborhood of the stick, as result, prevents robots approaching from an acceptable angle from making a grip2.

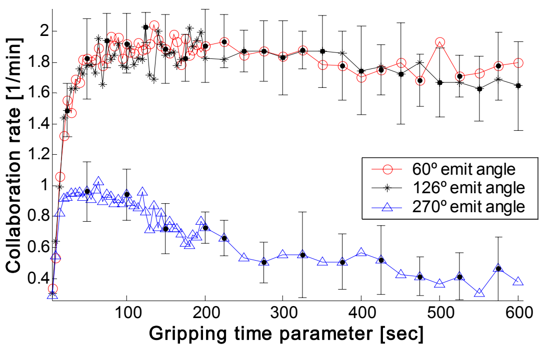

Figure 3: *Performance of a group of 6 robots for different IR emission angles.*

4.3 Implicit Communication using Vision

To compensate for the limited 36° angle of view of the K213 vision turret, we introduced a scanning behavior, characterized by the three parameters mentioned above, the scan interval, the scan speed, and the scan angle. In the following, we describe systematic experiments varying a single parameter at a time (linear search). We note that, when raised, the grip1 robot's camera and gripper block approximately 90° of the stick's visibility, i.e., robots approaching the grip1 robot from behind will not see the stick, but those approaching from the side will.

Scan Interval. For the experiments exploring the significance of the scan interval, the scanning speed was set such that a complete revolution about the robot's axis would require approximately 7 seconds, the fastest rotation possible that would allow a frame for every 10° of rotation given the camera's default frame rate of 5 fps. Furthermore, the scan angle was set to 360° and simulations run for group sizes of 2, 4, and 6 for scan intervals every 10 cm from 10 cm to 70 cm.

Since the arena is 80 cm in diameter, the robot will encounter an obstacle and reset its scan interval counter before executing a scan; as scan intervals approach 80 cm, the behavior becomes equivalent to the original stick-pulling behavior.

Groups of 4 and 6 robots reach the maximal mean collaboration rate as the scan interval approaches the diameter of the arena. For smaller scan intervals, performance is much lower than that achieved by the original stick-pulling behavior (Figure 4). This is due to the time penalties introduced by the scanning behavior.

For teams of two robots, however, the time penalty associated with the scanning behavior is counterbalanced by collaborations stimulated with help from the vision system. Because the only other robot present is gripping the stick, if a robot moves towards a raised stick, it will not interfere with any other robots. There is also a zero probability of the view of an accessible stick being obstructed by another robot.

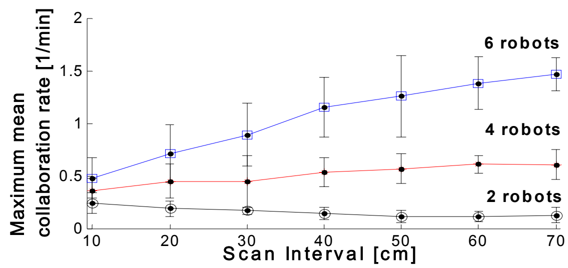

Figure 4: *Maximal performance for each scan interval, with scan speed of 1 revolution per 7 seconds.*

Scan Speed. To explore the effect of scanning speed, the experiments were repeated with scanning performed at the robot's minimum speed, with a full revolution requiring approximately 21 seconds, a decrease in speed by a factor of three over the previous experiments. The result is counterintuitive: despite increased time spent scanning, slow-scanning behavior outperforms the fast-scanning behavior (Figure 5). This implies that slower scanning stimulates more collaborations than fast scanning, though we note that the benefit of scanning is still outweighed by the time cost for the larger team sizes. We suggest that, due to the dynamic nature caused by crowding of the environment in the 4-and 6-robot cases, slower scanning significantly increases the probability of the scanning robot observing a raised stick, whether it was raised before the scan began or is lifted while in the robot's field of view.

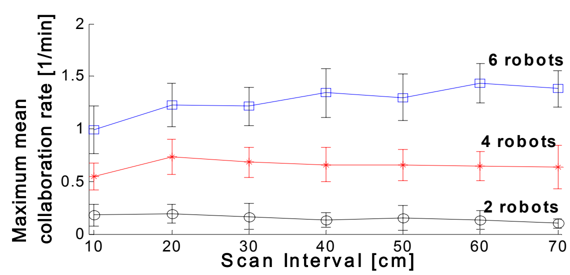

Figure 5: *Maximal performance for each scan interval, with scan speed of 1 revolution per 21 seconds.*

Scan Angle. For the experiments exploring the significance of the scan angle, the scan interval was tested at 20cm and 70cm and simulations were run for group sizes of 2, 4, and 6 for scan angles every 60° from 0° to 360°.

The scan speed was set to one revolution per 21 seconds. The smaller the scan angle, and thus, the effective field of view, the less time is spent scanning. For small scan angles, any benefit of vision is counterbalanced by the time penalty incurred by the scanning behavior (Figure 6). At best, for this system, the vision-based scheme achieves approximately the same performance as the non-IR, non-vision stick-pulling behavior.

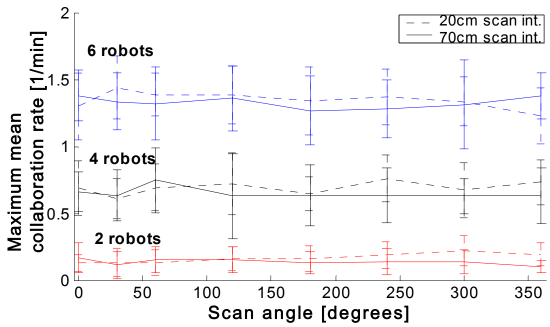

Figure 6: *Maximal performance for each scan angle.*

4.4 Sensor Position: IR vs. Vision

One reason the vision scheme does not compete with the IR scheme in these cases is the sensor position. While robots in the IR system can receive a signal in any of 4 IR receivers positioned to receive a call from any angle, the vision system has one sensor, facing forward, with an angle of view of 36°. The robot with a camera can only "receive" the message "sent" by a raised stick when it scans, in contrast to the immediate transmission and receipt of the signals in the IR communication scheme. Furthermore, when the robot scans, it will only receive the raised stick signal if it is in line of sight of the stick, with no robots in the way, and is within the 270° visibility angle of the stick, more than half the range of which is unacceptable as an angle of approach.

To explore further the nature of the difference between the IR scheme and the vision-based scheme, in simulation, teams of IrDA-equipped robots were endowed with the same behavioral controller as the camera-equipped team. The IR team uses only the frontal IR receivers, which have an effective field of view of 120°, and emission angles set to 270°. Both teams had scan speeds of 1 revolution per 7 seconds.

For experiments with a scan angle of 0° or a large scan interval relative to the arena diameter, the two teams gave equivalent performances. Figure 7 shows one example.

For smaller scan intervals, however, there is a large discrepancy in performance (Figure 8). This is due to the difference in the field of view between the two teams. When the IR team scans, "looking" for sticks, a robot can receive a signal even when the stick would not be in view of a camera-equipped robot in the same position. Because the two teams have the same controller, this means the IR-equipped robot will exit the scan much sooner than a camera-equipped robot in the same scenario. The IR-equipped robot will not necessarily move in the direction of the signal's source because the IR field of view is so large and no IR-specific centering behavior was introduced, but the robot will spend much less time scanning than the robots in vision-based team and, because the direction is not as closely determined, will be less likely to interfere with robots moving toward the stick. The IR-equipped team thus achieves a collaboration rate near the vision-based optimum for scan intervals and angles that are far from optimal for the vision-based team.

We note a major difference between the front-sensing IR team and the IR team in section 4.2 with a 270° emission angle. Because the front-sensing IR team does not perform phototaxis towards the signal, no significant increase in interference is observed. The same limitation in front-only sensing that prevents improved performance thus also prevents decreased performance due to interference.

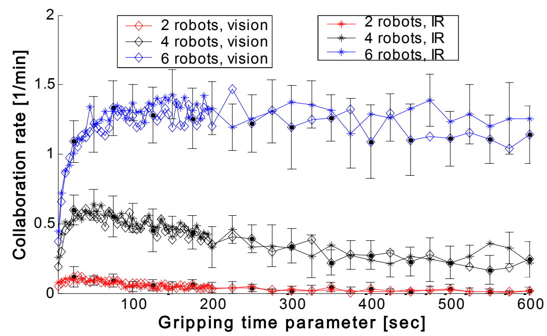

Figure 7: *For a scan interval of 70 cm and scan angle of 360°, the performances of both teams are equivalent.*

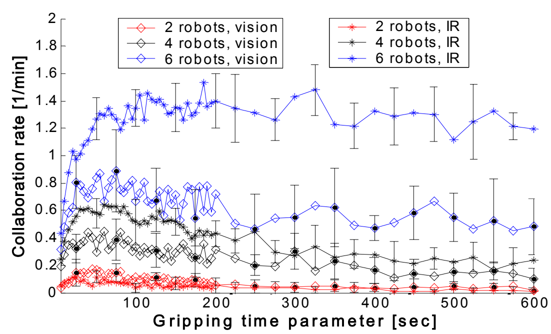

Figure 8: *For a scan interval of 30 cm and scan angle of 360°, the advantage of the larger field of view of the IR team is apparent.*

4.5 Validation of Vision Behavior with Real Robots

Real robots were equipped with cameras and programmed with the optimum parameters found in section 4.3: no scanning is performed, but the robots do center on any sticks appearing in the field of view while searching for sticks. Three experiments were performed for each gripping time parameter of 30 seconds, 100 seconds, and 500 seconds for groups of 2, 4, and 6 robots (Figure 9).

We observe the real robot performance is consistently lower than the performance predicted using Webots. We believe this is due to two effects not taken into account in the Webots simulation. The first is the observation that occasionally, immediately after gripping and lifting, the robots will drop the stick, execute obstacle avoidance, and resume the stick-searching behavior. We hypothesize this is due to miscalibration of the arm position caused by slippage in the arm-raising mechanism. We also observed occasional entangling of grippers when robots approached at the extreme edge of the acceptable approach angle, an interaction we cannot capture directly using Webots. We note that though the sensor-based embodied simulator is an extremely useful predictive tool, effects such as these must be observed, then integrated into the simulation probabilistically to make a quantitatively correct prediction of non-ideal real robot performance.

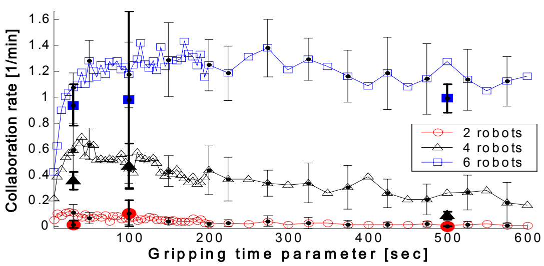

Figure 9: *Real robot experiment results and the Webots data for the same parameters. Solid color markers represent the results for real robot experiments.*

5. Conclusion

Extensive simulation of the three communication schemes presented with a sensor-based simulator allows us to explore the subtle, nested effects between individual and group behavior or hardware and software parameters, allowing us to conclude that while the IR-based communication scheme offered significant performance advantages over the basic scheme in this environment, the vision-based scheme presented here is not well suited to this distributed robotic system. This conclusion is limited, however, to this robotic system with the performance metric based solely on collaboration rate; redefining the performance metric will change the effectiveness of a proposed solution.

The cluttered nature of the experimental setup combined with the very limited field of view of the camera justifies the difference in performance as compared to the omnidirectional view and the instantaneous communication in the IR scheme. Though the implication that the benefits of the vision sensor, especially in a task where identification of a collaboration opportunity is key, are outweighed by the behavioral cost of implementation is counterintuitive, we observe that additional sensing and signaling capabilities may not offer any advantage over a simpler system. Extensive simulation prior to large-scale implementation can help keep costs low in the evaluation of a proposed hardware solution for a collective system. This supports the idea of minimalism in unit design for large distributed systems.

Acknowledgements

We would like to thank Tom Vanderslice for contributions on robot programming aspects of the project and Adam Hayes for his constructive comments in reviewing this paper. This work is supported in part by the Center for Neuromorphic Systems Engineering under grant EEC-9402726. Kjerstin Easton is supported by a NSF Graduate Research Fellowship.

References

[1] Bonabeau E., Dorigo M., and Theraulaz G., “Swarm Intelligence: From Natural to Artificial Systems”, SFI Studies in the Science of Complexity, Oxford University Press, New York, NY, 1999.

[2] A. Martinoli. Collective Complexity out of Individual Simplicity. Invited book review on "Swarm Intelligence: From Natural to Artificial Systems", by E. Bonabeau, M. Dorigo, and G. Theraulaz. *Artificial Life*, Vol. 7, No. 3, pp. 315-319, 2001. © MIT Press

[3] Böhringer K., Brown R., Donald B., Jennings J., and Rus D., “Distributed Robotic Manipulation: Experiments in Minimalism”. In Khatib O. and Salisbury J. K., editors, Proc. of the Fourth Int. Symp. on Experimental Robotics, Stanford, June, 1995, Lecture Notes in Control and Information Sciences, Springer Verlag, pp. 11-25.

[4] Ijspeert A. J., Martinoli A., Billard A., and Gambardella L.M., “Collaboration through the Exploitation of Local Interactions in Autonomous Collective Robotics: The Stick Pulling Experiment”. *Autonomous Robots*, Vol. 11, No. 2, pp. 149-171, 2001.

[5] Hayes A. T., Martinoli, A. and Goodman R. G., “Swarm Robotic Odor Localization”. *Proc. of the IEEE Conf. on Intelligent Robots and Systems IROS-01,* October 2001, Wailea, Hawaii, USA, pp. 1073-1078.

[6] Martinoli A., Ijspeert A. J., and Mondada F., “Understanding Collective Aggregation Mechanisms: From Probabilistic Modelling to Experiments with Real Robots”. *Robotics and Autonomous Systems*, 1999, Vol. 29, pp. 51-63. Invited paper.

[7] Lerman K., Galstyan, Martinoli A., and Ijspeert A. J. A Macroscopic Analytical Model of Collaboration in Distributed Robotic Systems. *Artificial Life*, Vol. 7, No. 4, pp. 375-393, 2001.

[8] Michel O., “Webots: Symbiosis Between Virtual and Real Mobile Robots”. In Heuding J.-C., editor, Proc. of the First Int. Conf. on Virtual Worlds, Paris, France, July, 1998, Springer Verlag, pp. 254-263. See also http://www.cyberbotics.com/webots/.

[9] 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.