Evolving Communication in Robotic Swarms using On-line, On-board, Distributed Evolutionary Algorithms

No Author Given

No Institute Given

Abstract. Robotic swarms offer flexibility, robustness, and scalability. For successful operation they need appropriate communication strategies that should be dynamically adaptable to possibly changing environmental requirements. In this paper we try to achieve this is through evolving communication on-the-fly. As a test case we use a scenario where robots need to cooperate to gather energy and the necessity to cooperate is scalable. We implement an evolutionary algorithm that works during the actual operation of the robots (on-line), where evolutionary operators are performed by the robots themselves (on-board) and robots exchange genomes with other robots for reproduction (distributed). We perform experiments with different cooperation pressures and observe that communication strategies can be successfully adapted to the particular demands of the environment.

Keywords: swarm robotics, communication, on-line, on-board, distributed

1 Introduction

Swarm robotics has emerged in recent years as an important field of research. Drawing inspiration from the behavior of social insects, the main idea behind swarm robotics is that a group of simple robots, by means of cooperation, are able to perform tasks beyond the capabilities of a single individual. The motivations for this approach are increased robustness, flexibility, and scalability [5].

For robotic swarms to be successful, a key component is the development of appropiate communication strategies, particularly due to the requirement that robots operate in a descentralized manner. Furthermore, robotic swarms are expected to operate in dynamic environments for which a high degree of flexibility and adaptation is required. Thus, instead of using fixed communication policies, it is better to equipt robots with the ability to adapt their communication strategies to environmental requirements.

A promising way to achieve this in through the use of an evolutionary robotics (ER) approach, i.e., using evolutionary algorithms to evolve the robots' controllers [10]. ER techniques have been applied to diverse problems such as gait control for legged robots [14], and navigation for aerial vehicles [2]. The taxonomy offered by Eiben et al. classifies ER tecniques according to *when* does evolution happen (off-line vs. on-line), *where* does it take place (on-board vs. off-board), and *how* does it happen (encapsulated/centralised or distributed or a hybrid of these two) [6].

In this work we study the evolution of communication in robotics swarms using on-line, on-board, and distributed evolutionary algorithms. This means that evolution takes place during the actual operation of the robots (on-line), evolutionary operators are performed exclusively inside each robot (on-board), and robots exchange genomes with other robots instead of maintaining purely local pools of genomes (distributed). In particular, the evolutionary algorithm (EA) used in this work, Hybrid EvAg, is a hybrid between a purely distributed evolutionary algorithm and a purely local one [9]. In Hybrid EvAg, each robot maintains both a local pool of genomes and a cache of robot neighbors for periodical exchange of genomes.

The task we study in this work involves a group of robots that require cooperation to gather energy sources randomly distributed in a rectangular arena. Our experiments draw ideas from the work of Buzing et al. [4], the main one being that communication arises as a mean to facilitate cooperation, and thus no fitness is explicitily given to robots for communicating. We study the effect of different cooperation pressures in the communication preferences evolved and, as in [4], we draw a distinction between talking and listening behaviors.

2 Related Work

Many authors have used computer simulations to study the environmental and evolutionary conditions conducive to communication. According to Perfors [12], work in this area can be divided in two categories: the evolution of syntax [3, 15, 13] and the evolution of comunication and coordination [11, 4, 8, 1].

One key difference between this and other existing work is that we do not intend to establish conclusions about the emergence of communication as an evolutionary construct. Our question is more practical: can we use on-line, on-board, distributed EAs as a tool to allow robotic swarms to develop communication strategies on their own? In this sense, the works of Buzing et al. [4] and Floreano et al. [8] are of particular interest. First, our experimental setting is directly based on that of [4], and we borrowed the ideas of varying degrees of environmental pressure. On the other hand, our work is similar to [8] in that both works studied communication in the context of robotic swarms. A comparison between the present work, and [4] and [8] is shown in Table 1 and a more detailed description of their work is shown next.

Buzing et al. [4] studied the evolution of communication within what they named the VUSCAPE model. This model, based on SUGARSCAPE [7], consists of a discrete landscape in which sugar seeds are periodically redistributed and agents need to collect them in order to survive. In addition, pressure towards cooperation is introduced in the form of a limit to the amount of sugar agents can collect on their own. In order to facilitate cooperation, agents have a hard-wired ability to communicate (using messages with fixed syntax and semantics), but their attitude towards using communication is not fixed and evolves over time. The authors used this model to study how communication evolves under different levels of cooperation pressure, and concluded that higher levels of cooperation pressure translate into increased attitudes towards communication.

On the other hand, Floreano et al. studied the evolutionary conditions that facilitate the emergence of communication. Their setting investigated colonies of robots

Table 1: Comparison between Buzing et al.[4] , Floreano et al.[8], and the present work

	Buzing et al.	Floreano et al.	This work
Dynamic Environment	YES (energy redistributed)	NO	YES (energy redistributed)
Hard-wired semantics	YES	NO	YES
Varying cooperation pressure	YES	NO	YES
Means of communication	Message board. Messages only travel parallel to the axes	Emitting blue light	Broadcasting within a certain circular range
2 agents on 1 location	YES	NO	NO
Agents die	YES	NO	NO
Controller	Rule set	Neural network	Neural network
Actions	2 behavior macros: go to largest sugar seed or random move. Talk / Listen with a certain probability	Spin left/right wheel. Turn on/off blue light	3 behavior macros: random move, avoid obstacle, go to largest energy source. Talk / Listen with a probability
Fitness function	Environmental fitness based on energy	Number of cycles stepping on the energy source minus number of cycles stepping on the poison source	Energy gained
On-line	YES	NO	YES
On-board	YES	NO	YES
Distributed	YES	NO	YES
Selection	No parent selection. Agents mate when at the same location. Environmental survivor selection (agents that run out of sugar die)	Individual and colony-level	Global parent selection, local survivor selection

that could forage in an environment with food and poison sources (one of each), and in which robots could use a blue light to (possibly) signal about the location of the food/poison sources. In contrast with Buzing et al., the semantics of the messages were not hard-wired into the systemn, and they found that different communication strategies evolved depending on the kin structure and selection level of the population (individual- versus colony-level).

3 Experimental Settings

Our experiments were run using the RoboRobo simulator developed by Nicholas Bredeche, a fast and simple 2D robot simulator built in C++. The test scenario proposed is directly based on the VUSCAPE model developed by Buzing et al. [4]. Our scenario consists of a number of robots set in a rectangular arena in which several energy sources (corresponding to sugar in VUSCAPE) are randomly distributed. Each robot's fitness is determined by how much energy it is able to collect over a certain period of time. However, collecting energy is made difficult by the following factors:

- Robots constantly lose energy over time. Whenever a robot's energy counter reaches zero, the robot is immediately switched off for the rest of an evaluation period, thus receiving minimal fitness.
- The environment requires that robots cooperate in order to successfully collect energy. In order to study different levels of cooperation pressure, we add an experimental parameter, the cooperation threshold (CT), specifying how much energy a robot can collect from a single source on its own. Specifically, a source carrying an amount of energy higher than the CT must be collected by two or more robots, in which case the energy is distributed equally among the collecting robots.
- The only way for a robot to gather knowledge (on its own) about the location of an energy source is through a fixed set of sensors of limited range.
- Collected energy sources are randomly relocated, thus increasing the need for robots to have an exploratory behavior.

In order to surmount these difficulties, robots are able to facilitate cooperation and exploration through a hardwired ability to communicate. In particular, robots can use (with a certain probability) information given by other robots about the location and size of energy sources (i.e., listen), and multicast (with a certain probability) the size and location of energy sources they are not able to collect on their own (i.e., talk). Notice that while robots possess an innate ability to communicate, the extent to which they are willing to do so is not fixed; we deliberately leave it subject to adaptation through evolution.

3.1 Controller

Each robot is controlled through a neural network that decides between different preprogrammed control policies. The twelve (12) inputs of the neural network are:

- Measurements from eight (8) distance sensors that detect obstacles and other robots in the vicinity.
- Angle to the largest energy source the robot has knowledge of.
- Distance to the largest energy source the robot has knowledge of.
- Current energy level.
- Bias node.

The five (5) outputs of the neural network are:

- Three (3) outputs corresponding to different actions: Random Walk, Avoid Obstacles, and Go to largest energy source; the highest valued output determines the next action of the robot.

- Talk preference, i.e., the probability that the robot multicasts information about an energy source when it needs to cooperate.
- Listen preference, i.e., the probability that the robot incorporates knowledge about energy sources seen by other robots.

The actions are implemented as follows:

Random Walk The robot chooses a random direction and moves as far as it can in a straight line in the chosen direction.

Avoid Obstacles The robot moves straight in the direction that is currently facing until its sensors detect an obstacle. It then rotates away from the obstacle and moves in a straight line again.

Go to largest energy source The robot rotates so that it faces the largest energy source it is aware of and moves towards this source as fast as it can.

3.2 Evolutionary Algorithm

The controllers in our experiments (i.e., neural networks) were trained using Hybrid EvAg, a variant of the on-line, on-board, distributed evolutionary algorithm for robotics described in [9]. In Hybrid EvAg, in addition to a local cache of neighbors for genome exchange, each robot maintains a local pool of μ+1 genomes (μ stored plus one active controller). Parental selection is performed by selecting two neighbors from the cache and using their current genomes (active controllers) as parents. If, after evaluation, the new genome turns out to be better than the worst one in the local pool of μ genomes, the worst one is replaced by the new. This local pool of genomes is used to randomly choose genomes for reevaluation. Thus, in Hybrid EvAg survival selection is local while parental selection is (approximately) global.

The cache of neighbors in Hybrid EvAg is maintained using the Newscast gossiping protocol as explained in [9]. We compared the performance of the Newscast-based Hybrid EvAg with that of a panmictic variant in which each agent has access to the local pools of all the other agents for parent selection. This allows us to study the effect that the lack of information about the true global genome pool has on gossiping-based distributed evolutionary algorithms.

The genome representation of the neural network was a real-valued vector consisting of the neural networks' weights and a mutation step size for every weight. Mutation was performed using Gaussian perturbation, and the recombination operator was standard two-parent arithmetic crossover. Binary tournament was used for parent selection. Table 2 shows the evolutionary parameters used in our experiments.

3.3 Experimental Details and Performance Measures

We used a group of 20 robots and performed 56 different simulations to account for the stochasticity in the evolutionary algoritms. Each simulation ran for 2,000,000 steps, with a new generation of controllers being evaluated each 1,000 steps. After each evaluation period, the controller's fitness was calculated and the evolutionary algorithm described in Sec. 3.2 was carried to select a new controller. Each robot's energy counter

Table 2: Parameter values for the evolutionary algorithms

μ (size of the local pool of genomes)	10
σ (initial mutation step size)	1
Crossover rate	0.5
Re-evaluation rate	0.2
Mutation rate	1
Newscast cache size	20

was then reset to its initial value and the robot was allowed to move randomly for 250 steps in order to avoid difficult conditions inherited from the previous evaluation.

The performance of the evolutionary algorithms was evaluated in terms of the performance metrics described next. Note that the values reported in Sec. 4 correspond to these measures averaged over the 56 experiments.

– Fitness: the median fitness of the group of robots for each generation.
– Talk/listen preferences: the median average talk/listen preference during 250 controller steps (i.e., not counting the random rellocation steps).
– Frequency of controller actions: the median frequency of controller actions during 250 controller steps.

As we are interested in assessing whether robots can develop successful strategies for different environmental demands, we study the effect of the cooperation threshold (CT), and thus environmental pressure, on the evolved strategies; for this, two values of the CT were considered (CT = 1 and CT = 5). In one case (LOWCT) the CT was set so that robots needed cooperation for all the energy points in the arena. In the other, (HIGHCT) cooperation was not required for any of the energy points.

4 Experimental Results

For both of the CT values considered, both the Newscast-based and panmictic variants of Hybrid EvAg were able to improve the average fitness of the robots over time (as evidenced in Figs. 1a and 1b). Interestingly, although the mating pool for each robot was smaller in the Newscast-based variant, it showed much quicker convergence than the panmictic variant. For the HIGHCT case this resulted in the Newscast-based variant having a somewhat better fitness at the end of the simulation. However, the panmictic variant showed better performance in the LOWCT case.

Evolved talking and listening preferences were very high in the LOWCT case, which indicates that communication evolved as a response to the environmental pressure to cooperate (see Fig. 2). With the panmictic variant of Hybrid EvAg, the average talking and listening probabilites converged to close to 100% after approximately 600,000 controller steps (600 generations). On the other hand, although the Newscast-based variant quickly reached high talking/listening probabilities (aproximately 90% in

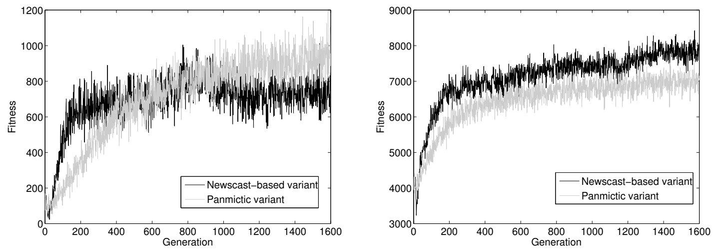

(a) LOWCT case (b) HIGHCT case

Fig. 1: Fitness vs. Number of generations for CT = 1 (LOWCT implying high pressure to cooperate) and CT = 5 (HIGHCT implying low pressure to cooperate). Mind the different scales on the Fitness axes.

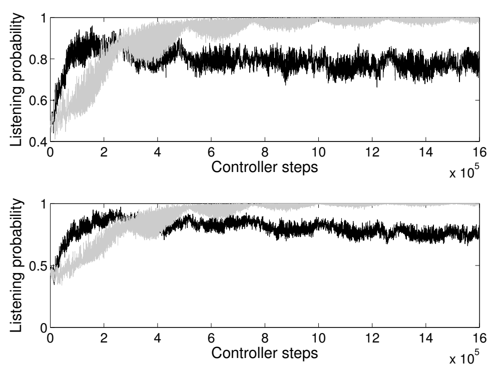

Fig. 2: Talking (upper) and Listening (lower) probabilities vs. Controller steps. Dark line: Newscast-based variant. Light line: Panmictic variant (LOWCT case)

less than 200 generations), the final values were considerable lower than those obtained with the panmictic variant; in fact, talk/listen probabilities show a decreasing tendecy over time. This partially explains why the fitness was lower for the Newscast-based variant in the LOWCT case, as a lower preference for communication was detrimental to the robots' capacity to cooperate.

In the HIGHCT case the talk/listen preferences were considerably lower than in the LOWCT case (see Fig. 3). This is not surprising since cooperation was not required in order for robots to succeed in this arena and, due to cooperation involving a split of the resources among cooperating robots, it would have only resulted in less fitness overall. However, one interesting observation is the significant difference between the talking/listening probabilities obtained with the Newscast-based and the panmictic variant. The Newscast-based variant's evolution history was highly irregular and showed no sign of convergence, in contrast with the typical evolution pattern observed with the panmic-

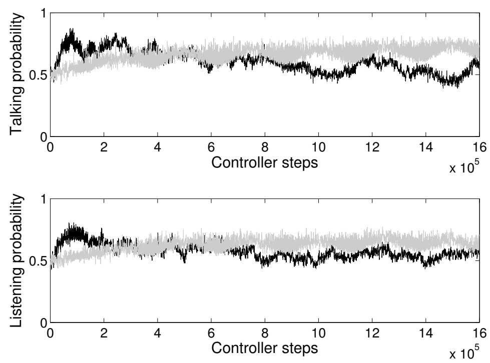

Fig. 3: Talking (upper) and Listening (lower) probabilities vs. Controller steps. Dark line: Newscast-based variant. Light line: Panmictic variant (HIGHCT case)

tic variant. The reason why the strategies evolved by the panmictic and Newscast-based variants were so different requires further investigation.

Finally, regarding the frequency of controller actions, there are significant differences between the strategies evolved using the panmictic and Newscast-based variants of Hybrid EvAg. Both in the LOWCT and HIGHCT cases the controllers evolved using the Newscast-based variant showed a much higher preference for the ”Avoid Obstacles” action than those evolved using the panmictic variant (as shown in Figs. 4 and 5). Significant differences can also be observed in the preferences for the ”Go to Largest Energy Source” action in the LOWCT case (Fig. 4), with the panmictic variant converging to a higher value than the Newscast-based variant.

5 Conclusions and Future Work

In this paper we presented an initial study on the applicability of on-line, on-board, distributed evolutionary algorithms (e.g., Hybrid EvAg) for evolving communication in robotic swarms. For this first study we assumed robots possessed the ability to communicate using messages with fixed semantics, and focused on studying the communication strategies evolved under different degrees of cooperation pressure. We also draw a distinction between the preference for sending messages (i.e., talking) and that for receiving messages (i.e., listen).

The results show that when the environmental pressure to cooperate is large, the controllers evolved using on-line, on-board, distributed evolutionary algorithms develop a high preference for communication, while the opposite behavior is seen in the presence of a low environmental pressure to cooperate. However, we observed a distinction between the communication preferences evolved using a distributed algorithm with full information of the global genome pool (panmictic variant), versus one in which each robot only has a local approximation of the genome pool (Newscast-based variant). The reason for these differences require further investigation, but it is probably related

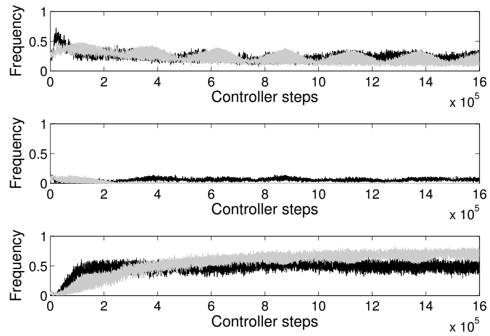

Fig. 4: Frequency of controller actions: Random (upper), Avoid Obstacles (middle), and Go to Largest Energy Source (lower). Dark line: Newscast-based variant. Light line: Panmictic variant (LOWCT case)

to the information loss inherent to the Newscast-based variant. Note that in some cases (e.g, HIGHCT case) the Newscast-based variant can offer a higher performance than the panmictic variant.

In future work we aim to study the evolution of communication on groups of robots having a lesser degree of hard-wired abilities (such as the current fixed controller actions and semantics). Also, we are currently studying larger groups of robots (e.g., 500 robots) since the computational advantages of the Hybrid EvAg algorithm are more relevant in such a context, and different types of communication behavior may emerge.

References

1. C. Ampatzis, E. Tuci, V. Trianni, and M. Dorigo. Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26, 2008.
2. G. J. Barlow. Autonomous controller design for unmanned aerial vehicles using multi-objective genetic programming. In *Proceedings of the Graduate Student Workshop at the 2004 Genetic and Evolutionary Computation Conference (GECCO-2004)*, Seattle, WA, June 2004. Winner of Best Paper award at the Graduate Student Workshop at the 2004 Genetic and Evolutionary Computation Conference (GECCO-2004).
3. E. J. Briscoe. Grammatical acquisition and linguistic selection. In *Linguistic Evolution through Language Acquisition: Formal and Computational Models*, chapter 9. Cambridge University Press, 2002.
4. P. C. Buzing, A. E. Eiben, and M. C. Schut. Emerging communication and cooperation in evolving agent societies. *Journal of Artificial Societies and Social Simulation*, 8(1):1–16, 2005.
5. E. Şahin. Swarm robotics: From sources of inspiration to domains of application. Technical Report METU-CENG-TR-2005-01, Department of Computer Engineering, Middle East Technical University, january 2005.
6. A. E. Eiben, Evert Haasdijk, and Nicolas Bredeche. Embodied, on-line, on-board evolution for autonomous robotics. In P. Levi and S. Kernbach, editors, *Symbiotic Multi-Robot*

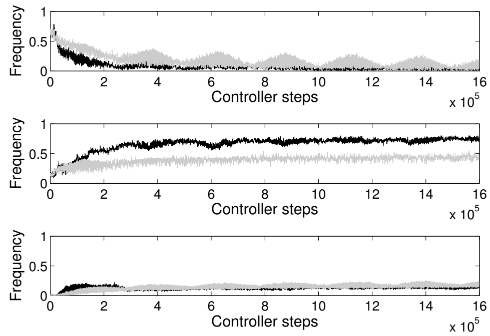

Fig. 5: Frequency of controller actions: Random (upper), Avoid Obstacles (middle), and Go to Largest Energy Source (lower). Dark line: Newscast-based variant. Light line: Panmictic variant (HIGHCT case)

Organisms: Reliability, Adaptability, Evolution, chapter 5.2, pages 361–382. Springer, May 2010.

7. J.M Epstein and R Axtell. *Growing Artificial Societies: Social Sciences from Bottom Up*. Brooking Institution Press and The MIT Press, 1996.
8. D. Floreano, S. Mitri, S. Magnenat, and L. Keller. Evolutionary conditions for the emergence of communication in robots. *Current biology : CB*, 17(6):514–9, March 2007.
9. R. Huijsman, E. Haasdijk, and A.E. Eiben. *An On-line On-board Distributed Algorithm for Evolutionary Robotics*. Angers, France, 24-26 October 2011.
10. Stefano Nolfi and Dario Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press, Cambridge, MA, 2000.
11. M. Oliphant and J. Batali. Learning and the emergence of coordinated communication. *The newsletter of the Center for Research in Language*, 11(1), 1997.
12. A. Perfors. Simulated evolution of language: a review of the field. *Journal of Artificial Societies and Social Simulation*, 5(2):1–62, 2002.
13. L. Steels. Modeling the formation of language: Embodied experiments. In S. Nolfi and M. Mirolli, editors, *Evolution of Communication and Language in Embodied Agents*, pages 235–262. Springer, Berlin, 2010.
14. J. Teo. Darwin + robots=evolutionary robotics: Challenges in automatic robot synthesis. In *2nd International Conference on Artificial Intelligence in Engineering and Technology (ICAIET 2004)*, pages 7–13, Kota Kinabalu, Sabah, Malaysia, august 2004.
15. P. Vogt. The emergence of compositional structures in perceptually grounded language games. *Artificial Intelligence*, 167(1-2):206 – 242, 2005. Connecting Language to the World.