

Re-Engineering Evolution: A Study In Self-Organising Synchronisation

Vito Trianni and Stefano Nolfi

Institute of Cognitive Sciences and Technologies,
National Research Council, Rome, Italy
vito.trianni@istc.cnr.it

Abstract

Evolutionary Robotics (ER) is a powerful approach for the automatic synthesis of robot controllers, as it requires little a priori knowledge about the problem to be solved in order to obtain good solutions. This is particularly true for collective and swarm robotics, in which the desired behaviour of the group is an indirect result of the control and communication rules followed by each individual. However, the experimenter must make several arbitrary choices in setting up the evolutionary process, in order to define the correct selective pressures that can lead to the desired results. In some cases, only a deep understanding of the obtained results can point to the critical aspects that constrain the system, which can be later modified in order to re-engineer the evolutionary process towards better solutions. In this paper, we present a case study about self-organising synchronisation in a group of robots, in which some arbitrarily chosen properties of the communication system hinder the scalability of the behaviour to large groups. We show that by modifying the communication system, artificial evolution can synthesise behaviours that properly scale with the group size.

Introduction

The synthesis of controllers for autonomous robots is a complex problem that has been faced with a large number of different techniques (Siciliano and Khatib, 2008). Among the various possibilities, Evolutionary Robotics (ER) represents a viable approach for the automatic synthesis of robot controllers requiring little a priori knowledge about the solution of a given problem (see Nolfi and Floreano, 2000). In fact, the evolutionary process proceeds in the bottom-up direction, directly evaluating controllers for their suitability to the requirements defined by the designer. When dealing with collective or swarm robotics systems, the usage of automatic techniques like ER is even more compelling, in particular when the group behaviour should be the result of a self-organising process arising from numerous interactions among robots. In such conditions, in fact, there is an indirect relationship between the desired group behaviour and the individual control rules. By evaluating the robotic system as a whole (i.e., by testing the global behaviour that results from the individual rules encoded into the individual genotype),

ER provides an automatic process for identifying the mechanisms that produce and support the collective behaviour, and for implementing those mechanisms into the individual controller rules that regulate the robot/environment interactions (Trianni et al., 2008).

However, the advantages offered by Artificial Evolution are not costless, as pointed out by Matarić and Cliff (1996). In particular, it is necessary to identify the conditions that assure the *evolvability* of the system, i.e., the possibility to progressively synthesise better solutions starting from scratch. To do so, the experimenter has to make several choices in setting up the evolutionary process. Some of these choices are arbitrary if performed without any *a priori* knowledge of the system features, and may have a strong impact on the solutions found. This is often the case for the communication abilities provided to a collective robotics system. In fact, communication regulates the interactions among robots, and should be rich enough to support the emergence of the desired group behaviour. On the other hand, ER privileges simple sub-symbolic communication forms, as it contextually develops the behavioural and communication strategies, which co-evolve as a single whole. The selection of the best communication protocol should therefore face this tradeoff, and often only the experimenter intuition makes the difference between a valuable or an unfortunate choice.

Negative results should however be exploited to acquire information on the system dynamics and re-engineer evolution accordingly. In fact, by understanding the properties of unsuccessful systems it may be possible to recognise which are the critical aspects that constrain the system in sub-optimal solutions. In this paper, we present a case study of such an approach. We have studied self-organising synchronisation, in order to understand which are the minimal behavioural and communication strategies that would allow a group of robots to synchronise their periodic behaviour (Trianni and Nolfi, 2009). In particular, we are interested in the scalability property of the evolved behaviours to large groups. By analysing the evolved behaviours, we discovered that the arbitrary choice made in the communication protocol was hindering the evolved behaviour to suitably scale

to large groups. This finding allowed us to re-engineer the characteristics of the robots by identifying a new communication protocol, and to run further evolutionary experiments that resulted in properly scalable behaviours.

Evolution of Self-Organised Synchronisation

Self-organised synchronisation is a common phenomenon observed in many natural and artificial systems: simple coupling rules at the level of the individual components of the system result in an overall coherent behaviour (Strogatz, 2003). Probably, the most common synchronisation phenomenon is related to the flashing behaviour of some firefly species in South-East Asia, which aggregate at dusk and engage in massively synchronous displays (Buck, 1988). Models of this behaviour describe fireflies as a population of pulse-coupled oscillators with equal or very similar frequencies. These oscillators can influence each other by emitting a pulse that shifts or resets their oscillation phase. The numerous interactions among the individual oscillator-fireflies are sufficient to explain the synchronisation of the whole population (for more detail, see Buck (1988); Mirollo and Strogatz (1990); Strogatz and Stewart (1993)). This model has been often exploited to engineer systems capable of synchronous behaviour, also in collective and swarm robotics (Wischmann et al., 2006; Christensen et al., 2009). In this study, we have investigated which are the minimal behavioural and communicative conditions that can lead to synchronisation in a group of robots, in which each individual presents a periodic behaviour. For this purpose, we chose to provide robots with simple reactive controllers and basic communication abilities. The period and the phase of the individual behaviour are defined by the sensory-motor coordination of the robot, that is, by the dynamical interactions with the environment that result from the robot embodiment. We show that such dynamical interactions can be exploited for self-organised synchronisation, allowing to keep a minimal complexity of both the behavioural and the communication level (for more details, see Trianni and Nolfi, 2009).

Experimental setup

The evolutionary experiments are performed in simulation, using a simple kinematic model of the *s-bot* robot (see Fig. 1 and refer to Mondada et al., 2004, for details), and the results are afterwards validated on the physical platform. The experimental scenario for the evolution of self-organising synchronisation requires that each robot in the group displays a simple periodic behaviour, which should be entrained with the periodic behaviour of the other robots present in the arena. The individual periodic behaviour consists in oscillations along the y direction of a rectangular arena (see Figure 2). Oscillations are possible through the exploitation of a symmetric gradient in shades of grey painted on the ground. The gradient presents a white stripe for $|y| < 0.2 m$, and black stripe for $|y| > 1 m$.

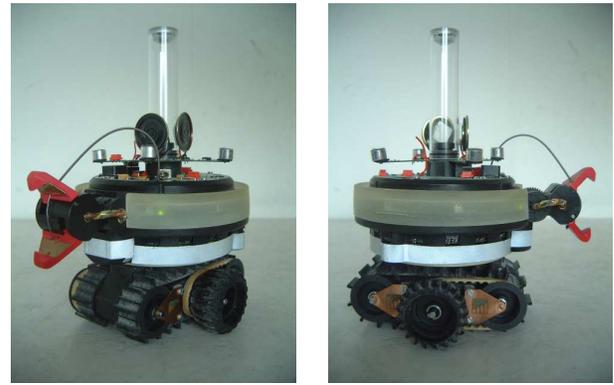


Figure 1: The *s-bot*, the robot used in the experiments.

For the purpose of engineering the evolutionary system, both the characteristics of the arena and the capabilities of the robots give several constraints to the experimental setup. According to these constraints, we select among the various possibilities the minimal set of sensors and actuators that are required to accomplish the task, that is, individual periodic oscillations over the grey gradient and synchronisation of the oscillation phase. Certainly, the controller needs access to the wheels' motors, and we set $\omega_M \approx 4.5 s^{-1}$ as the maximum angular speed of the wheels. The grey gradient of the arena can be perceived by the robots through four infrared sensors placed under their chassis (ground sensors), which are appropriately scaled to encode the grey-level in the range $[0, 1]$, where 0 corresponds to black and 1 to white. The perception of the gradient through these sensors provides the robot with enough information to perform oscillations along the y axis. Additionally, robots need to use the infrared proximity sensors placed around their cylindrical body, in order to avoid collisions with walls or with other robots. These choices, which are mainly constrained by the arena setup and by the features of the physical robot, are sufficient for

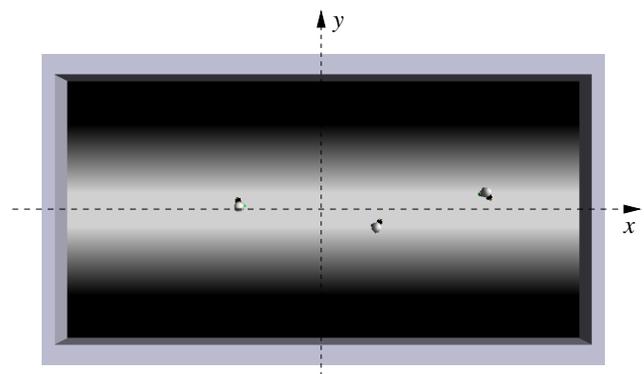


Figure 2: Snapshot of a simulation showing three robots in the experimental arena. The dashed lines indicate the reference frame used in the experiments.

the individual behaviour.

For what concerns the group behaviour, instead, we need to provide the robots with suitable interaction modalities that can lead to synchronisation of their movements. The choice of the communication system is the aspect we focus on in this paper. In fact, the *s-bot* platform features various communication devices, and we need to select among them the one that fits our experimental scenario. Robots are provided with speakers and microphones for sound communication. Moreover, robots can exploit coloured LEDs positioned around their turret to display a colour pattern that can be perceived through the omni-directional camera. Finally, robots have wireless communication abilities. Therefore, there is a large freedom in choosing the communication system. In order to maintain a minimal configuration, we decided to provide the robots with a *global* and *binary* communication system:

$$s(t) = \max_r S_r(t), \quad (1)$$

where $S_r(t) \in \{0, 1\}$ is the binary signal emitted by robot r at time t , and $s(t) \in \{0, 1\}$ is the binary signal perceived by all robots. In other words, each robot r can produce a signal $S_r(t)$. Signals produced by different robots cannot be distinguished, and result in a single signal $s(t)$ perceived by every robot in the arena, including the signalling one. Signals are perceived in a binary way: either there is someone signalling in the arena, or there is no one. This communication protocol is probably the poorest one in terms of the amount of information that can be conveyed. However, this is sufficient for our purposes, as we will see in the following. Note that this communication protocol can be easily implemented with sound signals: a robot can emit a single frequency tone with an intensity high enough to be perceived everywhere in the arena. Note that, differently from the other sensors and actuators, the choice of the communication system is not constrained by the robotic hardware or by other aspects of the experimental setup, but is only dictated by the communication protocol we have chosen.

Evolutionary Setup

Evolution was carried out using homogeneous groups of three robots, each controlled by a fully connected, feed forward neural network—a perceptron network. The neural controller takes as input the information coming from ground sensors, proximity sensors and perceived signals, and it controls the two wheels of the robot's differential drive system and the emission of binary signals. Connection weights and bias terms are genetically encoded parameters. The evolutionary algorithm is based on a population of 100 genotypes, which are randomly generated. This population of genotypes encodes the connection weights of 100 neural controllers. Each connection weight is represented with a 8-bit binary code mapped onto a real number ranging in $[-10, +10]$. Subsequent generations are produced by

a combination of selection with elitism and mutation. Recombination is not used. At each generation, the four best individuals—i.e., the *elite*—are retained in the subsequent generation. The remainder of the population is generated by mutation of the 20 best individuals. Each genotype reproduces at most 5 times by applying mutation with 3% probability of flipping a bit. The evolutionary process runs for 500 generations.

The evolved genotype is mapped into a control structure that is cloned and downloaded onto all the robots taking part in the experiment, therefore obtaining a homogeneous group of robots. During evolution, we use groups composed of three robots only in order to obtain fast simulations. The performance of a genotype is evaluated by a 2-components function: $F = 0.5 \cdot F_M + 0.5 \cdot F_S \in [0, 1]$. The movement component F_M simply rewards robots that move along the y direction within the arena at maximum speed. This component rewards the movements of the robot from the observer perspective, without explicitly indicating how to perform a periodic behaviour: the oscillatory behaviour derives from the fact that the arena is surrounded by walls, so that oscillations during the whole trial are necessary to maximise F_M . The second fitness component F_S rewards synchrony among the robots as the cross-correlation coefficient between the distance of the robots from the x axis. This component is therefore maximised by robots performing synchronous oscillations (either in-phase or anti-phase), and it is null when robots are maximally desynchronised. In addition to the fitness computation described above, two indirect selective pressures are present. First of all, a trial is stopped when a robot moves over the black-painted area, and we assign to the trial a performance $F = 0$. In this way, robots are rewarded to exploit the information coming from the ground sensors to perform the individual oscillatory movements. Secondly, a trial is stopped when a robot collides with the walls or with another robot, and also in this case we set $F = 0$. In this way, robots are evolved to efficiently avoid collisions. For more details on the fitness computation, refer to Trianni and Nolfi (2009).

Design and Evolution

Before presenting the obtained results, it is useful to discuss which are the features that are fixed by the experimenter, and those that are adaptively set by the evolutionary process. We have defined an experimental scenario that is intrinsically cooperative, because robots are homogeneous and are explicitly rewarded to display a desired group behaviour. We have also fixed the sensory-motor configuration and the controller architecture. In particular, we have fixed the interaction modality between different robots, which mainly happens through the binary and global communication signal. Notwithstanding this, the motor and communicative behaviour is not at all pre-determined, but it is the result of the evolutionary process. The individual behaviour and the syn-

chronisation mechanisms are completely determined by the parameters of the neural controller (i.e., connection weights and biases). Individual behaviour and communication signals co-evolve and mutually influence: the individual behaviour determines how the robot moves and experience the environment, which influences the signals emitted. In turns, perceived signals change the way in which the robot reacts to the environment. During evolution, the group behaviour is shaped in order to maximise the user-defined utility metric, within the constraints imposed by pre-determined features. In the following, we will see how the communication protocol we have chosen influences the obtained results.

Behavioural and scalability analyses

We performed 20 evolutionary replications, each starting with a different population of randomly generated genotypes. Each replication produced a successful synchronisation behaviour, in which robots display oscillatory movements along the y direction and synchronise with each other, according to the requirements of the devised fitness function. In general, it is possible to distinguish two phases in the evolved behaviours: an initial transitory phase during which robots achieve synchronisation, and a subsequent synchronised phase. The transitory phase may be characterised by physical interferences between robots due to collision avoidance, if robots are initialised close to each other. The collision avoidance behaviour performed in this condition eventually leads to a separation of the robots in the environment, so that further interferences to the individual oscillations are limited and synchronisation can be achieved. The synchronous phase is characterised by a stable synchronous oscillations of all robots, and small deviations from synchrony are immediately compensated.

The individual ability to perform oscillatory movements is based on the perception of the gradient painted on the arena floor, which gives information about the direction parallel to the y axis and about the point where to perform a U-turn and move back towards the x axis, therefore avoiding to end up into the black painted area. Each evolved controller produces a signalling behaviour that varies while the robots oscillate. The main role of the evolved signalling behaviour is to provide a coupling between the oscillating robots, in order to achieve synchronisation. In response to a perceived signal, robots react by moving in the environment, changing the trajectory of their oscillations. This results in a modulation of the oscillation amplitude and frequency, which allows the robots to reduce the phase difference among each other, and eventually synchronise. In a previous work (Trianni and Nolfi, 2009), we developed a mathematical model and exploited dynamical systems theory to thoroughly analyse the synchronisation behaviour. We invite the reader to refer to that work for further details on the synchronisation mechanisms, which are out of the scope of the present paper.

Once analysed the synchronisation behaviours evolved

using three robots only, we tested their ability to scale up with the group size. To do so, we compared the performance of the evolved behaviour varying the group size. To avoid overcrowding, we performed the scalability analysis in larger arenas, ensuring a constant density of robots across the different settings. By ensuring a constant initial density we limit the negative effects of overcrowding and we are able to compare the performance of robotic systems with varying group size. In order to keep a constant robot density equal to the one used in the evolutionary experiments, we lengthened the arena in the x direction, trying to keep an initial density of 0.25 robots per square meter. Despite the increased arena length, we still keep the same communication protocol, that is, communication continues to be binary and global, with all robots affecting each other. This choice allows us to evaluate the scalability of a behaviour as it was evolved, without modifying the features of the communication channel. We evaluated all best evolved controllers 100 times using six different group sizes (3, 6, 12, 24, 48 and 96 robots). The obtained results are presented in the top part of Figure 3. It is possible to notice that most of the best evolved controllers have a good performance for groups composed of 6 robots. Performance degrades for larger group sizes and only few controllers produce scalable behaviours up to groups formed by 96 robots. The main problem that reduces the scalability of the evolved controllers is given by the physical interactions among robots. Despite the constant initial density we introduced in order to limit the disruptive effect of collision avoidance, physical interactions nevertheless occur with a higher probability per time step, as the group size increases. Every collision avoidance action provokes a temporary desynchronisation of at least two robots, which have to adjust their movements in order to re-gain synchronous oscillations with other robots. In such cases, the whole group is influenced by the attempt of few robots to re-gain synchronisation, due to the global and binary communication.

To summarise, the above analysis showed that physical interactions and collision avoidance have a disruptive effect on the synchronisation ability of the robots, and this effect is more and more visible as the group size increases. However, the synchronisation mechanism evolved may scale with the group size if we ignore physical interactions. To test this hypothesis, we performed an identical scalability analysis, but in this case we ignore the physical interactions among the robots, as if each robot was placed in a different arena and perceived the other robots only through communication signals. The obtained results are plotted in the bottom part of Figure 3. Differently from what was observed above, in this case many controllers present good scalability, with only a slight decrease in performance due to the longer time required by larger groups to perfectly synchronise (namely, controllers evolved in replication number 2, 8, 10, 12, 14, 18 and 19). This result confirms the analysis about the neg-

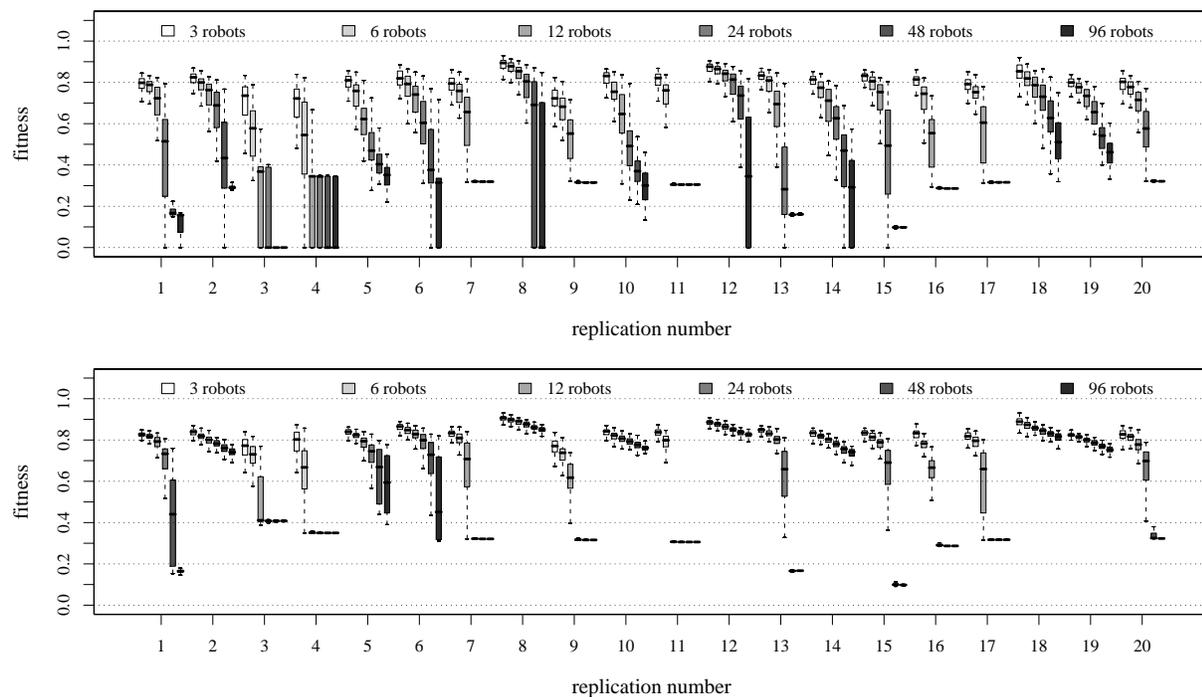


Figure 3: Scalability analysis. The boxplot shows, for each evolved controller, the performance obtained in tests with 3, 6, 12, 24, 48, and 96 robots. Each box represents the inter-quartile range of the data, while the black horizontal line inside the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. Outliers are not shown. Top: scalability of the evolved controllers under normal conditions. Bottom: scalability of the synchronisation mechanism.

ative impact of physical interferences and collisions among robots. In fact, removing the necessity to avoid collisions leads to scalable self-organising behaviours.

Nevertheless, many other controllers present a strange behaviour (namely, controllers evolved in replication number 3, 4, 7, 9, 11, 13, 15, 16, 17, 20). It is possible to notice that the performance presents a high variability up to a certain group size. The variable performance indicates that in some cases the robots are able to synchronise, and in other cases not. With larger group sizes, the performance stabilises to a low, constant value, independent from the initial conditions and the number of robots used. This value, which is characteristic of each non-scaling controller, represents the performance of the robotic system trapped into the basin of an *incoherent attractor*. In other words, the robotic system always converges into a dynamical condition in which no robot can synchronise with any other. By observing the actual behaviour produced by these controllers, we realised that the incoherent condition is caused by a communicative interference problem: the signals emitted by different robots overlap in time and are perceived as a constant signal (signals are global and are perceived in a binary way, preventing a robot from recognising different signal sources). If the perceived signal does not vary in time, it does not bring

enough information to be exploited for synchronisation, and the system remains desynchronised. This result is confirmed by the dynamical system analysis that we performed, which revealed how the individual signalling behaviour is responsible for producing such communicative interference, allowing also to predict which controllers present scalability just looking at the individual behaviour (see Trianni and Nolfi, 2009, for more details).

Re-engineering for scalability

The analysis of the unsuccessful controllers revealed that scalability cannot be always obtained, due to the physical and communicative interferences among robots. In particular, the communication protocol we selected has a strong impact on the scalability of the system. In fact, communication is global and binary, that is, the signal emitted by a robot is perceived by any other robot everywhere in the arena. Moreover, from the robot point of view, there is no difference between a single robot and a thousand signalling at the same time. Therefore, a single robot can influence the whole group. This has no negative effect as long as robots are synchronous, but can have severe consequences when a robot modifies its behaviour due to collision avoidance following some physical interaction with other robots. Further-

more, the binary communication channel generates the communicative interference we described above, which prevent the group from synchronising in certain conditions.

The main problems are therefore related to the absence of *locality*—i.e., signals are perceived everywhere in the arena—and of *additivity*—i.e., signals overlap without adding, preventing to recognise how many robots are contemporaneously signalling. The lack of locality and additivity is the main cause of failure for the scalability of the evolved synchronisation mechanisms.¹ We therefore decided to re-engineer our evolutionary experiments changing the communication protocol, which was arbitrarily chosen in the first place. Given that we are interested in studying global synchronisation, we decided to re-engineer our experiments focusing only on the additivity of the communication system. This allows us to make only minor changes to the experimental setup and directly compare the effects of the re-engineering approach.

Modified Experimental Setup

We evolved self-organising synchronisation behaviours exploiting exactly the same setup as above, but changing the way robots signal and perceive emitted signals. Specifically, we change the binary communication system with a continuous one:

$$\tilde{s}(t) = \frac{1}{N} \sum_{r=1}^N \tilde{S}_r(t), \quad (2)$$

Now, robots always emit a signal $\tilde{S}_r(t) \in [0, 1]$, encoding a number in a continuous range. The emitted signals are perceived as the average $\tilde{s}(t)$ among all the perceived signals. By doing so, the influence of an individual robot on the global perceived signal—which is equal for all robots in the arena—depends on the signalling behaviour of the whole group: the bigger the group, the smaller the influence of the single individual. This communication protocol can be easily implemented on the *s-bots*. For instance, signals could be sent as messages over the wireless network containing a real number in $[0,1]$. On the basis of the analysis performed so far, we expect that self-organising synchronisation behaviour can be evolved with such a communication system, and that they are more scalable.

Analysis of the Obtained Results

Also in this case, we performed 20 evolutionary runs for groups of three robots. All evolutionary runs were successful, and produced synchronisation behaviours that are qualitatively similar to those obtained with the binary communication system: robots perform oscillations over the painted gradient and react to the perceived signal by modifying the individual behaviour, in order to synchronise with other robots. The scalability analysis was performed with

¹However, as we have seen, this problem affects only some of the analysed controllers.

the same modalities as described above, and the obtained results are presented in Figure 4.

In the upper plot, scalability is tested including physical interactions. Also in this case, we notice that collisions prevent the scalability of some controllers, in which a good avoidance behaviour was not evolved. Recall that when a collision is detected, the group scores a null performance. However, it is possible to notice that the usage of an additive communication system leads to better performance even with large groups. Most controllers present good scalability for every tested group size, and only collisions substantially reduce the performance. Here, differently from what was observed before, physical interactions and collision avoidance do not have a severe impact on the performance of the whole group. In fact, the signals of few non-synchronous robots are averaged with those emitted by the rest of the group. As a consequence, the influence on the group of a robot attempting to synchronise decreases with increasing group size. This leads to a quick convergence to synchrony and to an improved group performance.

To better understand the effects of the re-engineering approach, we also performed a scalability analysis for the evolved synchronisation mechanisms, again removing the physical interactions among robots. The results plotted in the lower part of Figure 4 show that all evolved synchronisation mechanisms perfectly scale, and they do not suffer from the communicative interference observed with binary signals. In fact, the perceived signal brings information about the average signalling behaviour of all robots. As a consequence, synchronisation is always achieved, no matter the group size. Notice also that all controllers present a linear decrease in performance in correspondence to an exponential growth of the group size. This observation suggests that the self-organising synchronisation mechanism is very efficient, and is only slightly affected by the group size.

Discussion and Conclusions

In this paper, we have presented a case study about the evolution of self-organising synchronisation in a robotic system. In setting up the experiments, some characteristics of the system were chosen arbitrarily, given that no *a priori* knowledge was available about the possible solutions to the given problem. The results obtained with the initial approach proved that self-organising synchronisation can be actually achieved with a minimal complexity at the level of the control and communication strategy. However, the analysis of the scalability results also pointed to some characteristics of the system that hindered the group from scoring a good performance. We identified the problem in the communication system being global and binary, and to the effects of physical and communicative interferences. To solve this problem, we re-engineered the arbitrarily-chosen communication protocol exploiting the knowledge acquired by analysing the evolved behaviours. The newly devised continuous signals

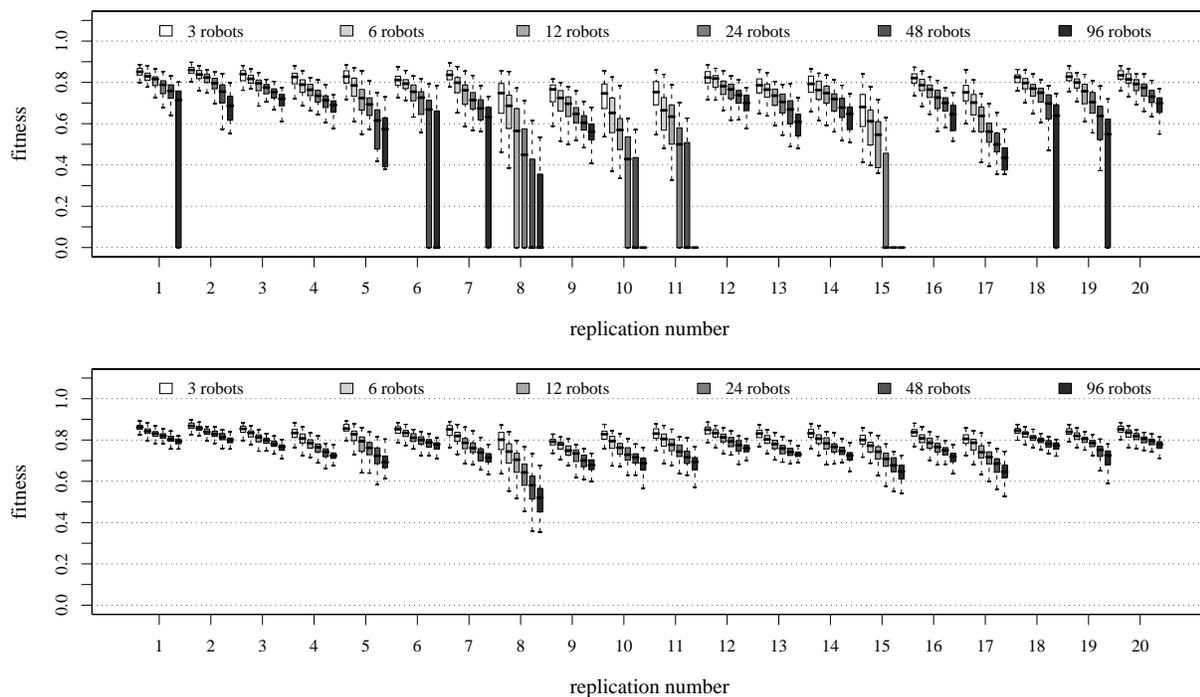


Figure 4: Scalability analysis for the continuous communication system. Top: scalability of the evolved controllers under normal conditions. Bottom: scalability of the synchronisation mechanism.

resulted in better synchronisation behaviours, and in an optimally scaling communication system.

The methodology described here may be generalised. Evolutionary Robotics is actually very useful for the automatic synthesis of controllers for robotic systems. However, it does not exclude arbitrary choices. The advantage given by ER is that, despite such arbitrary choices, it can find good solutions to a given problem. However, much as in conventional engineering methods, multiple design loops may be needed to find optimal results. This paper demonstrates that it is possible to engineer some features of a system undergoing artificial evolution on the basis of the outcomes of the evolutionary process itself. Contrary to trial and error methods without any guidance, we showed that an attentive analysis of negative results conveys knowledge on how to modify the system for evolving better solutions. Note that this is not in contradiction with respect to the need of little *a priori* knowledge in the design of the evolutionary experiment, as mentioned in the introduction. The knowledge we put into the system should not be related to the design of the solution, which is left to the evolutionary process, but rather to the preconditions required for obtaining good solutions.

We believe that it is necessary to formalise an engineering approach to Evolutionary Robotics, which can guide the design of evolutionary experiments. This is particularly true for collective and swarm robotics, in which the desired behaviour of the group is an indirect result of the control

and communication rules followed by each individual. Let's consider here the case in which the robotic hardware available is fixed, and the problem to be solved is well defined, as in any engineering application. In these conditions, it is possible to identify four major issues in the design of the evolutionary system: (i) the definition of the robot sensory-motor configuration (ii) the definition of the genotype-to-phenotype mapping, (iii) the definition of the fitness function, and (iv) the definition of the ecological selective pressures. In this paper, we have just dealt with the robot configuration, and in particular with the communication protocol. In the following, we briefly discuss the other issues.

With respect to the genotype-to-phenotype mapping, the design choices concern mainly the type of controller to be used, and the way in which the genotype is translated into such controller. A widely used approach in the literature consists in encoding into the genotype a fixed number of parameters of the robot controller (typically realized through an artificial neural network), while keeping constant the controller structure. Other approaches are possible, such as evolving the controller architecture (Stanley and Mikulainen, 2002), or evolving controller programs instead of neural networks (Koza, 1992). In collective robotics, another characteristic that has to be determined concerns the genetic relatedness between the individuals forming the group, that is, whether they are *genetically homogeneous* (i.e., they are clones) or *heterogeneous* (i.e., they differ from

each other). The advantage of homogeneous groups are given by a very compact encoding for the parameters of the controllers of the whole group, independently of its size. This advantage comes at the cost of a higher difficulty in obtaining roles that are well defined and differentiated. If this is a requirement, then heterogeneous groups might be more indicated. On the other hand, heterogeneous groups lead to a larger search space, require to estimate each individual contribution to the group performance, or need to identify in advance the role played by different individuals.

For what concerns the fitness function, it is difficult to suggest general principles for properly engineering it, because it strongly depends on the particular experimental conditions. Floreano and Urzelai (2000) propose the usage of a three-dimensional *fitness space*, in which the different dimensions refer to important features of a fitness function. In a collective robotics setup, the definition of a fitness function is more complex, due to the indirect relationship between individual actions and group organisation. A viable approach is given by functions that reward the final outcome of the collective behaviour, rather than the way in which the goal is achieved. This can be done, whenever possible, by measuring group variables that are available to the observer.

Finally, a typical problem of ER is the correct estimation of the performance of a genotype. The fitness function should evaluate the quality of the robot behaviour with respect to some variability of the environment. Typically, the behaviour must be robust with respect to varying initial position and orientation of the robot, and with respect to other parameters that contribute to define the *ecological niche* in which the behaviour is evolved. In order to obtain a reasonable fitness estimate, it is necessary to sample the space of the possible ecological conditions in an appropriate way. In a collective robotics setup, the problem is worsened by the presence of multiple robots, which increase the variability of the ecological niche. It is important to notice that indirect selective pressures may be created through the definition of the ecological niche and through the sampling employed to estimate the fitness. Given that the group is evaluated for presenting a robust behaviour within the parameter space of the ecological niche, the choice of the sampling may influence the evolutionary path. For these reasons, a careful design is required.

In our view, these are the main methodological choices that need to be performed when setting up an evolutionary experiment. In future work, we plan to carefully analyse these issues with both a theoretical and experimental work, in order to better formalise an engineering approach to Evolutionary Robotics

Acknowledgements

This work was supported by the Swarmanoid project, funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission under grant IST-022888.

The information provided is the sole responsibility of the authors and does not reflect the European Commission's opinion. The European Commission is not responsible for any use that might be made of data appearing in this publication.

References

- Buck, J. (1988). Synchronous rhythmic flashing of fireflies. II. *The Quarterly Review of Biology*, 63(3):256–289.
- Christensen, A. L., O'Grady, R., and Dorigo, M. (2009). From fireflies to fault tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766.
- Floreano, D. and Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13(4–5):431–443.
- Koza, J. (1992). *Genetic Programming*. MIT Press, Cambridge, MA.
- Matarić, M. J. and Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1):67–83.
- Mirrollo, R. E. and Strogatz, S. H. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662.
- Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA.
- Siciliano, B. and Khatib, O., editors (2008). *Springer Handbook of Robotics*. Springer Verlag, Berlin, Germany.
- Stanley, K. and Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127.
- Strogatz, S. H. (2003). *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press, New York, NY.
- Strogatz, S. H. and Stewart, I. (1993). Coupled oscillators and biological synchronization. *Scientific American*, 269(6):102–109.
- Trianni, V. and Nolfi, S. (2009). Self-organising sync in a robotic swarm. A dynamical system view. *IEEE Transactions on Evolutionary Computation*, 13(4):722–741.
- Trianni, V., Nolfi, S., and Dorigo, M. (2008). Evolution, self-organisation and swarm robotics. In Blum, C. and Merkle, D., editors, *Swarm Intelligence. Introduction and Applications*, Natural Computing Series, pages 163–192. Springer Verlag, Berlin, Germany.
- Wischnann, S., Huelse, M., Knabe, J. F., and Pasemann, F. (2006). Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14(2):117–127.