

Self-Organising Sync in a Robotic Swarm

Vito Trianni and Stefano Nolfi

Institute of Cognitive Sciences and Technology - CNR
via S. Martino della Battaglia, 44, 00185 Rome, Italy
Email: {vito.trianni,stefano.nolfi}@istc.cnr.it

Abstract—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. Owing to these properties, synchronisation appears particularly interesting for swarm robotic systems, as it allows to robustly coordinate through time the activities of the group while keeping a minimal complexity of the individual controllers. The goal of the experiments presented in this paper is the study of self-organising synchronisation for robots that present an individual periodic behaviour. In order to design the robot controllers, we make use of artificial evolution, which proves capable of synthesising minimal synchronisation strategies based on the dynamical coupling between robots and environment. The obtained results are analysed under a dynamical systems perspective, which allows us to uncover the evolved mechanisms and to predict the scalability properties of the self-organising synchronisation with respect to varying group size.

I. INTRODUCTION

Synchrony is a pervasive phenomenon: examples of synchronous behaviours can be found in the inanimate world as well as among living organisms [1], [2]. The discovery of the basic mechanisms behind self-organised synchronisation aroused research for many years, until the appropriate analytical methods were developed [3], [4]. Self-organising synchronisation phenomena can be modelled as systems of multiple coupled oscillators. Consider for example the synchronous flashing of fireflies [5]. Fireflies can be modelled 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 details, see [5]–[7]).

The synchronisation behaviours observed in Nature can be a powerful source of inspiration for the design of robotic systems. Synchronisation is an important mean to achieve coordination. This holds true particularly for swarm robotic systems [8], where emphasis is given to the emergence of coherent group behaviours from simple individual rules. Much work takes inspiration from the self-organised behaviour of fireflies or similar synchronisation behaviours observed in Nature [9]–[13]. The goal of the experiments presented in this paper is the study of self-organising synchronisation in a group of robots based on minimal behavioural and communication strategies. We follow the basic idea that if an individual displays a periodic behaviour, it can synchronise with other (nearly) identical individuals by temporarily modifying its

behaviour in order to reduce the phase difference with the rest of the group. In other robotic studies, synchronisation is based on the entrainment of the individual internal dynamics through some form of communication. In this paper, instead, we do not postulate the need of internal dynamics. Rather, 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 synchronisation, allowing to keep a minimal complexity of both the behavioural and the communication level. In order to define a robot controller able to exploit such dynamical agent-environment interactions, we use artificial evolution [14], [15]. The obtained results are analysed under a self-organising perspective, evaluating their scalability to large groups of robots.

The main contribution of this paper consists in the analysis of the evolved behaviours, which is brought forth exploiting a dynamical systems approach [16]. In this paper, we introduce a dynamical system model of the robots interacting with the environment and among each other. This model offers us the possibility to deeply understand the evolved behaviours, both at the individual and collective level, by uncovering the mechanisms that artificial evolution synthesised to maximise the user-defined utility function. Moreover, we show how the developed model can be used to predict the ability of the evolved behaviour to efficiently scale with the group size. We believe that such predictions are of fundamental importance to quickly select or discard obtained solutions without performing a time-demanding scalability analysis, as well as to engineer swarm robotic systems that present the desired properties.

II. EVOLUTION OF SELF-ORGANISING SYNCHRONISATION

In this section, we present the experimental scenario defined for the evolution of synchronisation behaviours. The task 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 the rectangular arena (see Figure 1). Oscillations are possible through the exploitation of a symmetric gradient in shades of grey painted on the ground, which can be perceived by the robots through the infrared sensors placed under their chassis. The gradient presents a black stripe for $|y| > 1$, in which the robots are not supposed to enter. Collisions with

walls or other robots are avoided using the infrared proximity sensors placed around the cylindrical body of the robots. Finally, synchronisation of the movements can be achieved by exploiting a binary communication system: each robot can produce a continuous tone with fixed frequency and intensity. When a tone is emitted, it is perceived by every robot in the arena, including the signalling one. The tone is perceived in a binary way, that is, either there is someone signalling in the arena, or there is no one.

The robots used in this experiments are the *s-bots*, which are small autonomous robots with a differential drive system [17]. The evolutionary experiments presented in this paper are performed in simulation, using a simple kinematic model of the *s-bots*, and the results are afterwards validated on the physical platform. Artificial evolution is used to set the connection weights and the bias terms of a fully connected, feed forward neural network—a perceptron network. 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 4 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 *s-bots* taking part in the experiment (i.e., we make use of a homogeneous group of *s-bots*). The performance of a genotype is evaluated by a 2-component 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. 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. In this way, synchronous oscillations are rewarded also when robots are in perfect anti-phase. In addition to the fitness computation described above, two indirect selective pressures are present. First of all, a trial is stopped when an *s-bot* 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 an *s-bot* 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.

III. EVOLUTIONARY RESULTS

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 fitness function. 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. In order to produce self-sustained oscillations, signalling is exploited. The main role of the evolved communication strategy is to provide a coupling between the oscillating *s-bots*, in order to achieve synchronisation. In fact, each evolved controller produces a signalling behaviour that varies while the robots oscillate. In this way, the signal emitted by a robot carries information about its position (or *phase*), which can be exploited by other robots for synchronisation. In summary, the evolved synchronisation behaviours are the results of the dynamical relationship between the robot and the environment, modulated through the communicative interactions among robots. No further complexity is required at the level of the neural controller: simple and reactive behavioural and communication strategies are sufficient to implement effective synchronisation mechanisms.

A qualitative analysis of the obtained controllers reveals that the behaviours produced are quite similar one to the other. 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 *s-bots* in the environment, so that further interferences to the individual oscillations are limited and synchronisation can be achieved. During the synchronous phase, collision avoidance is therefore less probable, but still possible due to the environmental noise, which may let robots deviate from their normal movements and approach other robots. Otherwise, this phase is characterised by stable synchronous oscillations of all *s-bots*, and small deviation from synchrony are immediately compensated.

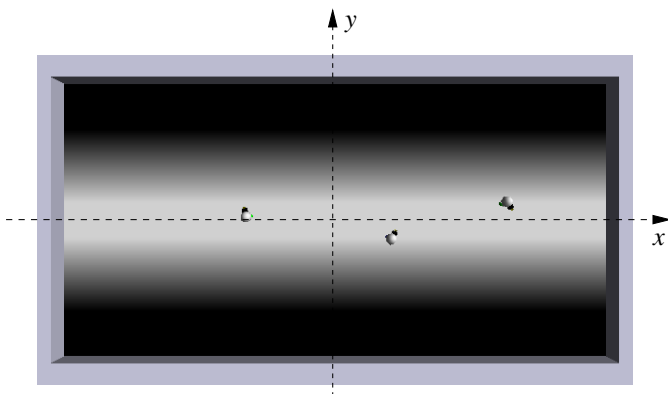


Fig. 1. Snapshot of a simulation showing three robots in the experimental arena. The dashed lines indicate the reference frame used in the experiments.

IV. DYNAMICAL SYSTEM MODELLING

We want to analyse the behaviour of a group of robots that synchronise their periodic oscillations. Our main interest is the understanding of both the individual behaviour and the synchronisation mechanism. Such understanding may be useful to predict some features of the evolved behaviour, e.g., scalability. To do so, we model the behaviour of the single robot looking only at the relevant features of the agent-environment dynamics. In particular, we ignore physical interactions among robots and between robots and walls. Moreover, we neglect the environmental noise and second order dynamics in the robot motion. As a consequence of such simplifications, the oscillatory behaviour of the robot r can be modelled as follows:

$$\langle y_r, \theta_r, S_r \rangle|_{t+1} = \mathcal{B}_c(y_r, \theta_r, s)|_t. \quad (1)$$

where y_r is the y coordinate of robot r at time t , θ_r its orientation, s is the binary communication signal perceived at time t and S_r is the signal emitted by robot r at time t . The function \mathcal{B}_c encodes the fundamental features of the individual behaviour, as it is produced by the parameters c of the evolved controller. In other words, given the above simplification and considering the features of the gradient painted on the arena floor, it is possible to neglect the x coordinate of a robot, as it does not influence the individual behaviour. The latter can be described as a trajectory in the 3D space $\langle y, \theta, s \rangle$, which is determined by \mathcal{B}_c . Notice that when only one robot is present, the perceived sound s corresponds to the self-emitted signal S_r . With R interacting robots, the communication channel determines the following coupling rule:

$$s(t) = \max_r S_r(t) \in \{0, 1\}, \quad (2)$$

which specifies that a binary signal is perceived if and only if it exists at least one s -bot r that is signalling. Notice that the sound perception s is equal for all robots in the environment, because communication is global and binary. What happens with R robots? The only interaction among s -bots is a communicative one, given by the coupling introduced in equation (2). It is therefore possible to define the following discrete-time dynamical system of $3R + 1$ equations:

$$\begin{cases} \langle y_1, \theta_1, S_1 \rangle|_{t+1} = \mathcal{B}_c(y_1, \theta_1, s)|_t \\ \vdots \\ \langle y_R, \theta_R, S_R \rangle|_{t+1} = \mathcal{B}_c(y_R, \theta_R, s)|_t \\ s|_{t+1} = \max_r S_r|_{t+1} \end{cases}. \quad (3)$$

In the following, we make use of this model to discuss about the behaviour of a single s -bot and the evolved synchronisation mechanism.

V. BEHAVIOURAL ANALYSIS

The behaviour of the individual s -bot can be studied looking at how position y , orientation θ and perceived sound s vary through time. We analyse the behaviour produced by the best evolved controller among the 20 evolutionary replications, namely the controller evolved in the 8th replication, which will be referred to as c_8 . To do so, we numerically integrate equation (3) for $R = 1$ to compute a *vector field* showing

the instantaneous direction and magnitude of change for each point in the state space $\langle y, \theta, s \rangle$ (see the top-left plot in Figure 2). This is a 3-dimensional space where y and θ are continuous variables that vary respectively in the range $[-1, 1]$ and $[0, 2\pi]$, while s is a binary variable. The plot suggests how the state of an s -bot starting at any point in its space evolves through time. Together with the vector field, the continuous line indicates the limit cycle attractor to which every trajectory converges. Notice that the continuous line is actually a closed trajectory, due to the 2π -periodic boundary conditions of θ . The existence of such a limit cycle attractor indicates that the individual behaviour is actually periodic, and defines the dynamics of convergence toward a stable motion of the robot.

Another important information can be extracted from the vector field: the signalling behaviour. For each point in the plane $\langle y, \theta \rangle$, it is possible to distinguish 4 different signalling behaviours:

- **no signalling:** the robot never emits a signal when placed at position $\langle y, \theta \rangle$.
- **environment-driven signalling:** the robot always emits a continuous signal when placed at position $\langle y, \theta \rangle$, no matter what signal is perceived. Signalling depends entirely on the position of the s -bot in the environment.
- **signal-driven signalling:** the robot emits a continuous signal when placed at position $\langle y, \theta \rangle$, but only in response to a perceived signal. Otherwise, no signal production is observed.
- **alternate signalling:** the robot emits a signal when placed at position $\langle y, \theta \rangle$ if no signal is perceived, and signalling is stopped in response to a perceived signal. As a consequence, the s -bot continuously switches on and off its loudspeaker.

We show the signalling behaviour of the best evolved controller in the top-right plot of Figure 2. Different signalling behaviours are indicated by circles filled with varying grey-level. It is possible to notice that the limit cycle traverses areas of the state space characterised by varying signalling behaviour. A signal is produced when the s -bot enters the “environment-driven” area, and it is stopped when the s -bot exits from the “signal-driven” area. Notice that entering in the signal-driven area having $s = 0$ does not lead to the production of a signal, while entering with $s = 1$ maintains the previous signalling status.

In order to describe the individual behaviour, notice that the limit cycle attractor jumps between the planes characterised by $s = 0$ and $s = 1$. In other words, the system switches between two different dynamics. The vector fields for these two conditions determine the quality of the individual oscillations, as shown in the bottom plots of Figure 2. When $s = 0$, the robot follows the left vector field, moving straight until it enters in the environment-driven signalling area. At this point, the production of a signal corresponds to a switch to the dynamics described by the right vector field, which presents a limit cycle attractor displayed by a dotted line. It is possible to notice how the normal limit cycle approaches this attractor when $s = 1$ (see the grey segments of the limit cycle in Figure 3). However, before converging onto this attractor, the limit cycle enters the

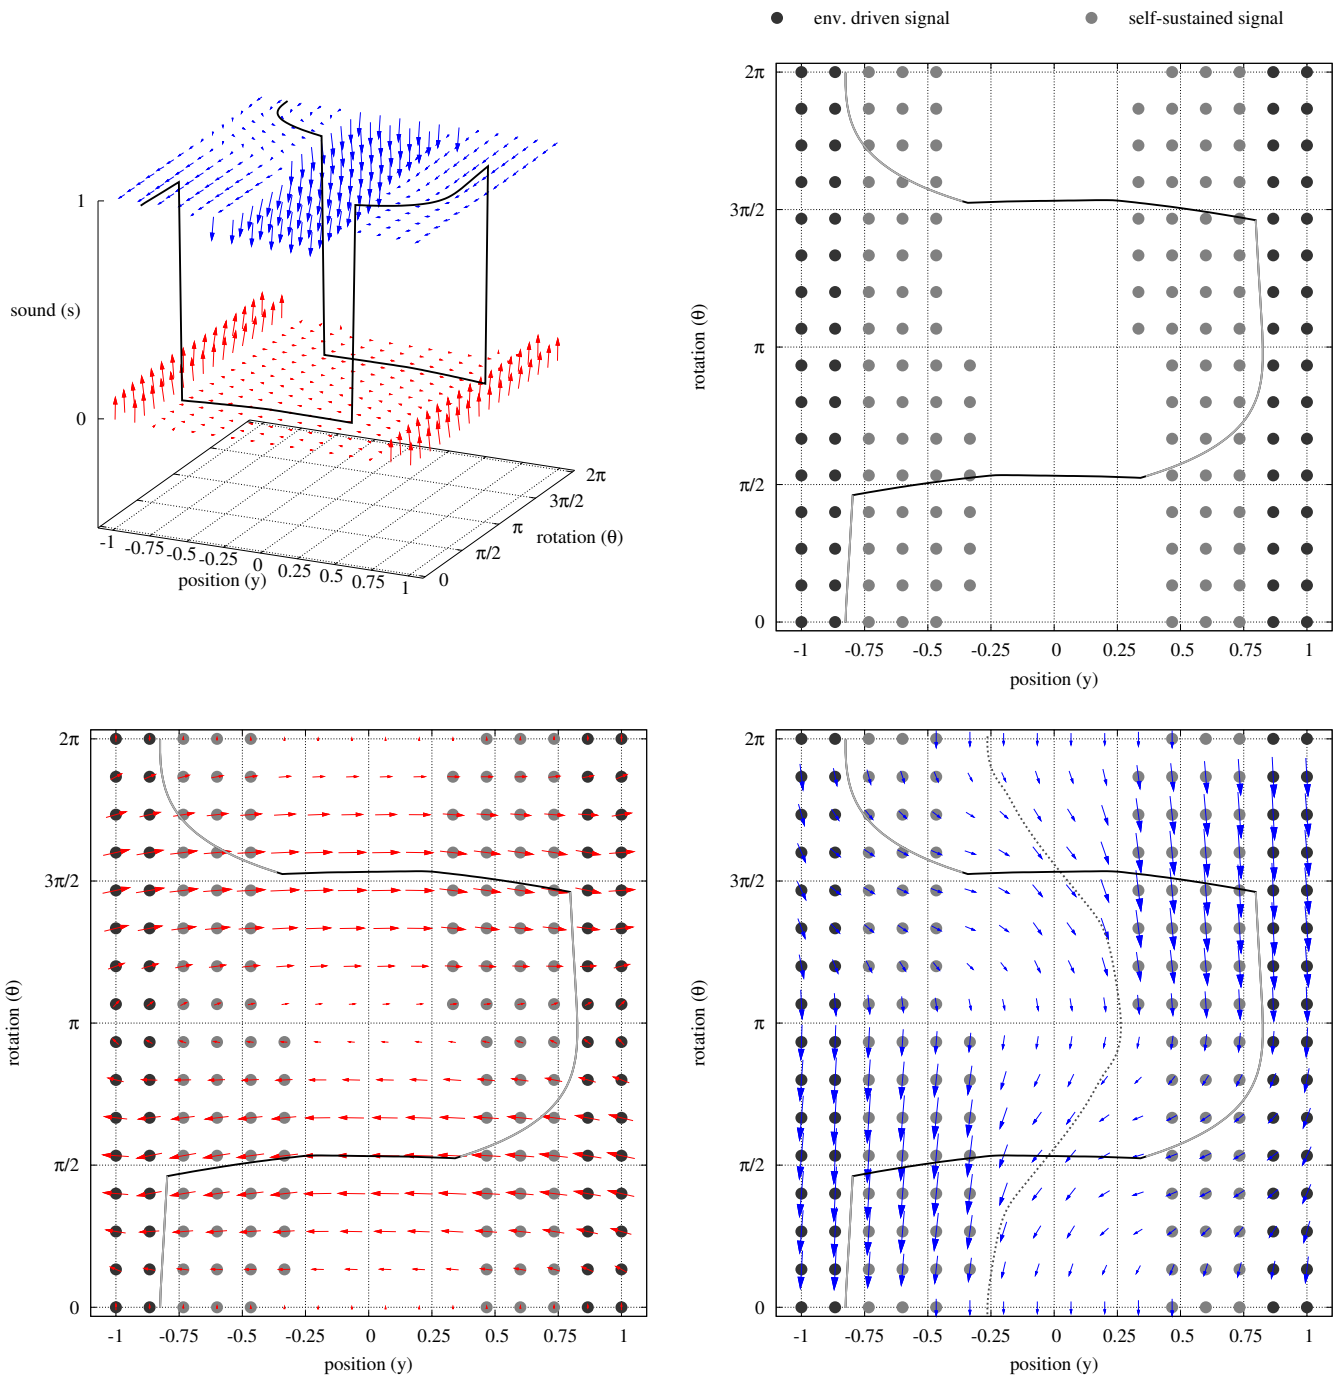


Fig. 2. Individual behaviour produced by controller c_8 . Top-Left: 3D vector field showing for each point in the state space the direction of variation and its magnitude. The θ dimensions is characterised by 2π -periodic boundary conditions. The continuous line represents the limit cycle attractor. Top-Right: signalling behaviour of the controller for each position and orientation (see text for details). The continuous line represents a projection of the limit cycle on the $y\theta$ plane: a black line colour indicates that the trajectory belongs to the plane $s = 0$, while the grey colour corresponds to the portion of trajectory that belongs to the plane $s = 1$. Bottom-Left/Right: projection on the $y\theta$ plane of the vector fields for a perceived signal $s = 0$ and $s = 1$. The dotted line in the bottom-right vector field represents the limit cycle for a constant perceived signal forced to 1, despite the individual behaviour.

“no signalling” area, and therefore the s -bot switches back to movements dictated by the vector field for $s = 0$.

Once decoded the individual behaviour, we analyse the system (3) with $R = 2$ robots. In this case, the dimensionality of the system does not allow an easy visualisation of the trajectories. However, we observed that the s -bots’ movements are governed solely by the individual behaviour \mathcal{B}_c and by

the coupling rule (2), which states that a signal is perceived whenever some s -bot emits a signal. As a consequence, it is possible to describe the behaviour of synchronising s -bots by looking at how the individual movements change with respect to incoming signals. Figure 3 presents various plots that represent different phases of the synchronisation. In the upper part, the position y for the two robots is plotted

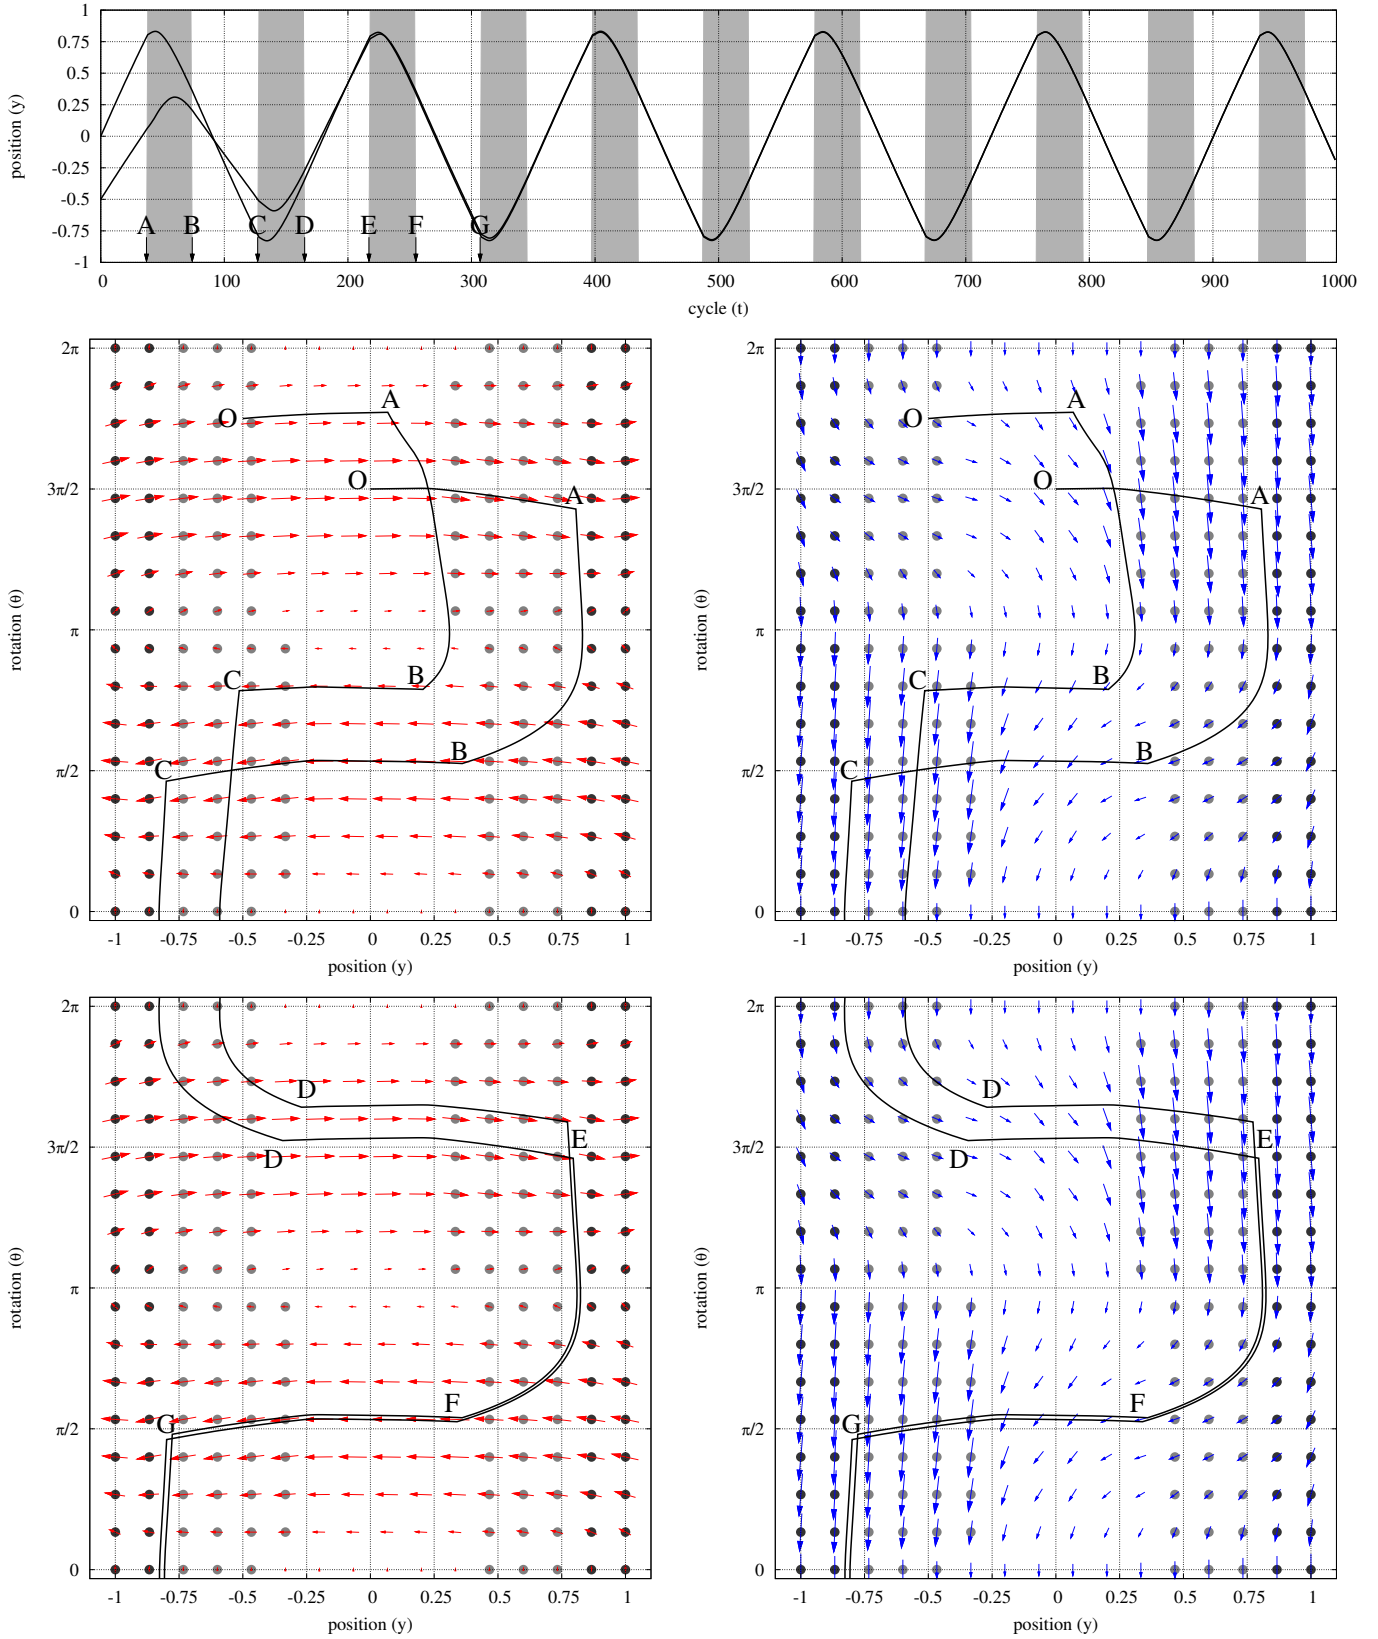


Fig. 3. Synchronisation behaviour of controller c_8 . Top: the position y of two s -bots that synchronise is plotted through time. The gray band in the background indicate that a signal is being perceived. Centre and bottom: vector fields for the condition $s = 0$ (left) and $s = 1$ (right). For each point, the individual signalling behaviour is displayed as a dot with varying grey level (see also Figure 2). The trajectories of the two synchronising robots are shown, and relevant events are marked with capital letters. The same letters indicate the time of the corresponding events in the top graph.

with respect to time. It is possible to observe that after an initial transitory phase, the robots converge towards coordinated movements. In particular, the position y is “modulated” through communicative interactions: the robot that signals first influences the behaviour of the other robot, which anticipates the turnabout in response to the perceived signal (see the gray bands in the background that indicate a continuous perceived signal). A better idea on how synchronisation is achieved is given by plotting the trajectories of the two robots over the vector fields for $s = 0$ and $s = 1$ (see the central and bottom plots of Figure 3). The two s -bots start in the points indicated by ‘O’, and none of them is signalling. As a consequence, the s -bots follow the top-left vector field, until they reach the point indicated by an ‘A’. Here, one of the robot enters the environment-driven signalling area, and therefore emits a signal, that triggers a behavioural change in both robots. The robots now follow the top-right vector field and both perform a clockwise turn, as indicated by the arrows. However, this turn is not performed at the same speed by the two robots: the one at larger y moves faster than the other, as indicated by the size of the arrows of the vector field. Consequently, the difference in distance among the two robots is consistently reduced in this phase, which ends with the robots reaching the points indicated with ‘B’. In the interval from points ‘B’ to points ‘C’ no robot is signalling and no interaction is present. The same interaction characterises the phases between points ‘C’ and ‘D’ and between ‘E’ and ‘F’, until synchronisation is achieved. This synchronisation mechanism is therefore based on the modulation of the position y during the oscillation: the first robot that reaches the environment-driven signalling area triggers a U-turn in the other robot, which is however performed at a lower speed, allowing the trajectories to approach and eventually converge into synchronous oscillations.

VI. SCALABILITY ANALYSIS

The analysis of the synchronisation behaviour for two s -bots is accompanied by a scalability analysis in which we test all evolved behaviours with groups of 3, 6, 12, 24, 48 and 96 s -bots. We first test the evolved behaviour in simulation, and we found that physical interactions may prevent the system from scaling to very large number of robots (data not shown). In fact, physical interactions occur with a higher probability per time step as the group size increases. Every collision avoidance provokes a temporary de-synchronisation of at least two robots, which have to adjust their movements in order to re-gain synchronous oscillations with other robots. As a consequence, the performance of the group as a whole is negatively affected. Still, the evolved synchronisation mechanism may scale well if there are no physical interactions. To prove so, we performed a further scalability analysis by ignoring collisions among robots (see Figure 4). We found that many controllers present perfect scalability, with only a slight decrease in performance due to the longer time required by larger groups to perfectly synchronise. However, other controllers present poor performance for large groups. By observing the actual behaviour produced by these controllers, we realised that the absence of scalability is caused by a

communicative interference problem: the signals emitted by different s -bots overlap in time and are perceived as a constant signal (recall that the sound signals are global and that they are perceived in a binary way, preventing an s -bot 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. This problem is the result of the fact that we used a “global” communication form in which the signal emitted by an s -bot is perceived by any other s -bot everywhere in the arena. Moreover, from the perception point of view, there is no difference between a single s -bot and a thousand signalling at the same time. The lack of locality and of additivity is the main cause of failure for the scalability of the evolved synchronisation mechanisms. However, as we have seen, this problem affects only some of the analysed controllers. In the remaining ones, the evolved communication strategies present an optimal scalability that is only weakly influenced by the group size.

Is it possible to predict whether a given evolved behaviour will scale or not with increasing group size? We try to give an answer by exploiting the mathematical model introduced in Section IV. We start from the observation that, if a synchronisation mechanism does not scale with the group size R , there exist an alternative attractor to the synchronous one, in which robots move incoherently. In other words, the dynamical system (3) undergoes a bifurcation with varying parameter R , so that two attractors are observable for large R : the coherent, synchronous one, and the incoherent one. In order to predict from the individual behaviour whether such a bifurcation exists, it is necessary to understand which are the conditions for the existence of an incoherent attractor. Recall that, whenever an evolved synchronisation mechanism does not scale, the perceived signal does not vary in time. In such a situation, in fact, the s -bots do not receive information about the position and orientation of other robots. If an s -bot r perceives a constant signal, its behaviour can be predicted as follows:

$$\langle y_r, \theta_r, S_r \rangle|_{t+1} = \mathcal{B}_c(y_r, \theta_r, f(s))|_t, \quad (4)$$

where $f(s)|_t$ indicates the constant perceived signal. It is therefore possible to plot the vector field for the above behaviour, and analyse possible attractors—be they fixed points or limit cycles—towards which all trajectories of the s -bot converge. We claim that, if such attractors exist and if they entirely lay out of the portion of state space in which $S_r(y_r, \theta_r, f(s))|_{t+1} = f(s)|_{t+1}$ —which we refer to as the *non-interaction area*—then the evolved synchronisation mechanism is scalable, no matter the group size R .

To prove the above claim, simply observe that the incoherent attractor exists contextually to a perceived signal that does not vary in time. Given that the s -bots themselves are responsible for signal production, the existence of the incoherent attractor requires that all s -bots participate in the signal production, so that:

$$\forall t \exists r \in \{1, \dots, R\} : S_r(y_r, \theta_r, f(s))|_{t+1} = f(s)|_{t+1}. \quad (5)$$

However, this requires that the attractor for the system (4) is contained at least partially within the *non-interaction area*,

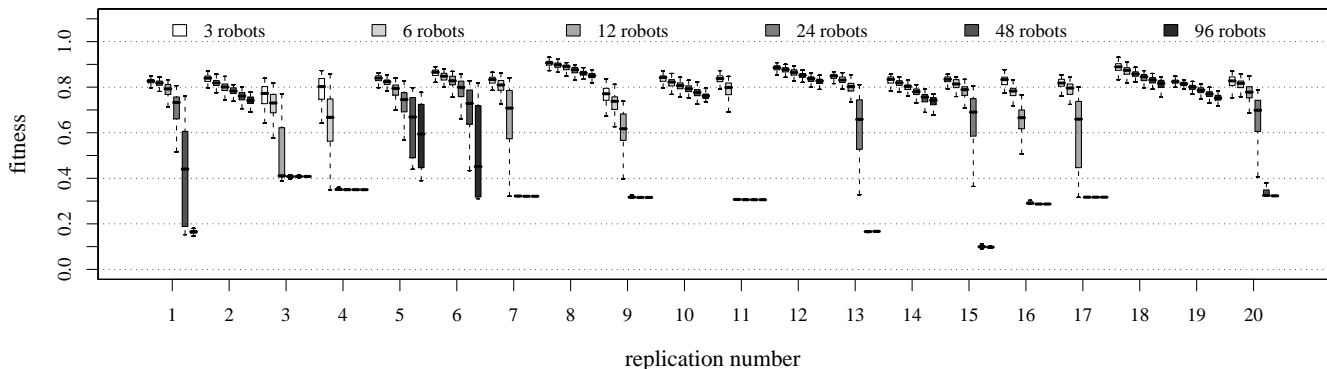


Fig. 4. Scalability of the synchronisation mechanism. The boxplot shows, for each evolved controller, the performance obtained in tests with 3, 6, 12, 24, 48, and 96 *s-bots*. Each box represents the inter-quartile range of the data, while the black horizontal line inside the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. Outliers are not shown.

which contradicts our hypothesis.

The controller c_8 analysed in Section V produces a scalable behaviour, as shown in Figure 4. In fact, it presents a limit cycle attractor shown as a dotted line in the bottom-right vector field of Figure 2, which is completely contained within the no-signalling area. On the contrary, controller c_{13} does not present scalability (see Figure 4). The evolved behaviour can be appreciated and analysed with the 3D vector field of Figure 5 that shows the individual behaviour under normal conditions. The right vector field in Figure 5 corresponds to the behaviour of the *s-bot* when a continuous signal $s = 1$ is constantly perceived. It is possible to notice that the limit cycle attractor for this condition traverses the environment-driven signalling area. As a consequence, with a sufficiently large number of *s-bots* the evolved synchronisation mechanism does not scale, as can be appreciated in Figure 4.

A further prediction from the mathematical model consists in the minimum group size R_m for which the incoherent attractor exists (i.e., the bifurcation point). This group size depends on the time each robot spends in the *non-interaction area* while moving over the limit cycle. In fact, in order to satisfy condition (5), it is necessary that while a robot moves within the *non-interaction area*, another robot prepares to enter in it. In other words, R_m robots should be evenly spaced over the limit cycle so that, when one *s-bot* exits the *non-interaction area*, another one enters in it, therefore sustaining the production of the constant signal. As a consequence, the minimum group size R_m is given by:

$$R_m = \left\lceil \frac{T}{T_n} \right\rceil, \quad (6)$$

where T is the period of a single oscillations, and T_n is the fraction of this period spent within the *non-interaction area*. For controller c_{13} , we experimentally obtained $R_m = 6$, which is to be considered a theoretical lower bound for the minimum group size. We actually observed the appearance of the incoherent attractor for a minimum group size of 9 (data not shown).

VII. CONCLUSION

Much as natural evolution produced swarms of fireflies able to self-organise to achieve coherent group behaviour, artificial evolution can synthesise self-organising swarms of robots that accomplish complex tasks. In this respect, swarm intelligence can benefit from the study and analysis of natural as well as artificial systems: in both cases, a deep understanding of the dynamics that govern the individual behaviour and the social interactions can underpin novel developments in the engineering of swarm intelligent systems. In this paper, we have presented an artificial evolutionary process that has shaped the behaviour of a robotic system to display self-organised synchronisation. We have also shown how the dynamical system analysis can explain the evolved mechanisms and predict the behaviour of the robotic system for varying group size. We believe that this analysis can bring useful insights on how to build—through automatic techniques or hand-design—swarm robotics systems that are capable of self-organised synchronisation and that scale to large number of robots. In fact, we have given a clear description of the building blocks necessary to produce synchronised behaviours, and, most importantly, we have decoded the individual behaviour to find the conditions that allow the system as a whole to synchronise, no matter the group size. In conclusion, we believe that studies about synchronisation such as the one presented in this paper, notwithstanding the explicitly simplified experimental setup, can have a strong impact on future studies in swarm robotics.

REFERENCES

- [1] S. H. Strogatz, *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press, New York, NY, 2003.
- [2] A. Pikovsky, M. Rosenblum, and J. Kurths, *Synchronization: A Universal Concept in Nonlinear Sciences*. Cambridge University Press, Cambridge, UK, 2001.
- [3] Y. Kuramoto, “Phase dynamics of weakly unstable periodic structures,” *Progress of Theoretical Physics*, vol. 71, no. 6, pp. 1182–1196, 1984.
- [4] S. H. Strogatz, “From kuramoto to crawford: exploring the onset of synchronization in populations of coupled oscillators,” *Physica D: Nonlinear Phenomena*, vol. 143, no. 14, pp. 1–20, 2000.
- [5] J. Buck and E. Buck, “Mechanism of rhythmic synchronous flashing of fireflies,” *Science*, vol. 159, no. 3821, pp. 1319–1327, 1968.
- [6] R. E. Mirollo and S. H. Strogatz, “Synchronization of pulse-coupled biological oscillators,” *SIAM Journal on Applied Mathematics*, vol. 50, no. 6, pp. 1645–1662, 1990.

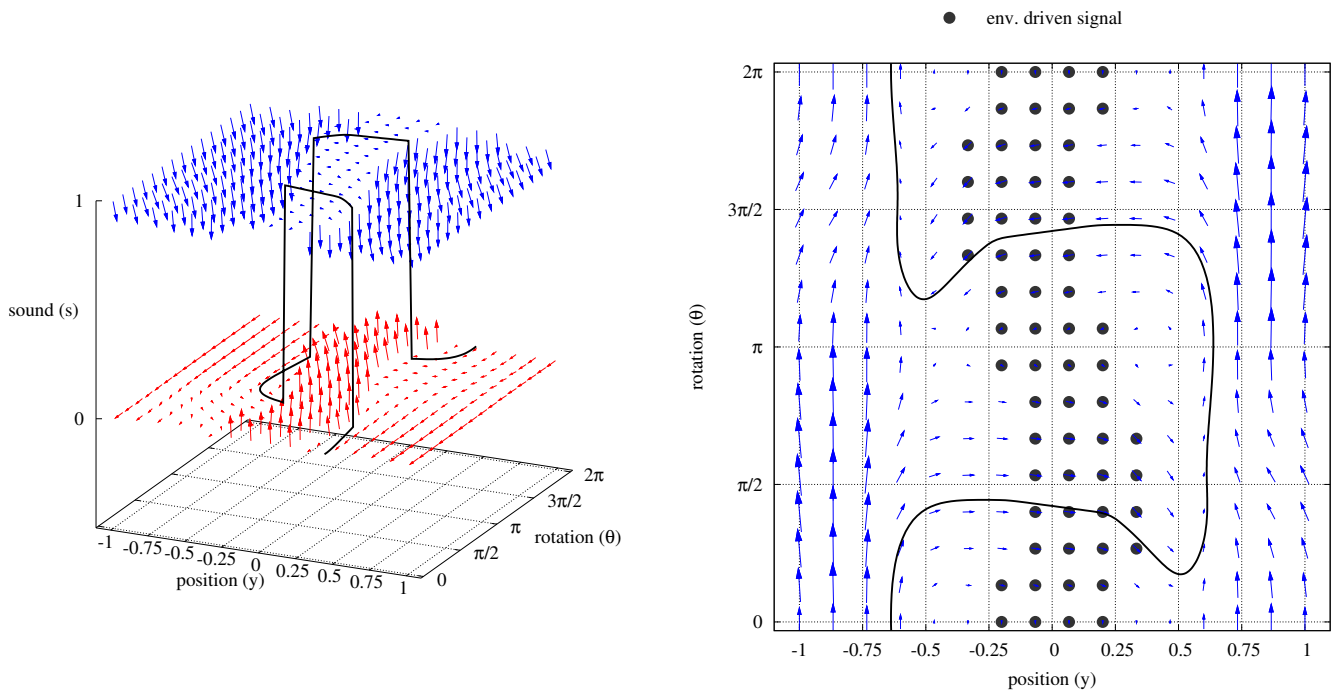


Fig. 5. Individual behaviour of controller c_{13} . Left: the 3D vector field shows for each point in the state space the direction of variation and its magnitude. Right: projection on the $y\theta$ plane of the vector field for a constant perceived signal $s = 1$. The black line represent the limit cycle for this condition. The black dots represent the non-interaction area.

- [7] S. H. Strogatz and I. Stewart, "Coupled oscillators and biological synchronization," *Scientific American*, vol. 269, no. 6, pp. 102–109, 1993.
- [8] M. Dorigo and E. Şahin, "Swarm robotics — special issue editorial," *Autonomous Robots*, vol. 17, no. 2–3, pp. 111–113, 2004.
- [9] O. Holland and C. Melhuish, "An interactive method for controlling group size in multiple mobile robot systems," in *Proceedings of the 8th International Conference on Advanced Robotics (ICAR '97)*. IEEE Press, Piscataway, NJ, 1997, pp. 201–206.
- [10] C. Melhuish, O. Holland, and S. Hodde, "Convoying: using choring to form travelling groups of minimal agents," *Robotics and Autonomous Systems*, vol. 28, pp. 207–216, 1999.
- [11] S. Wischmann, M. Huelse, J. F. Knabe, and F. Pasemann, "Synchronization of internal neural rhythms in multi-robotic systems," *Adaptive Behavior*, vol. 14, no. 2, pp. 117–127, 2006.
- [12] M. Hartbauer and H. Römer, "A novel distributed swarm control strategy based on coupled signal oscillators," *Bioinspiration and Biomimetics*, vol. 2, pp. 42–56, 2007.
- [13] S. Wischmann and F. Pasemann, "The emergence of communication by evolving dynamical systems," in *From animals to animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*, ser. LNAI, S. N. et al., Ed., vol. 4095. Springer Verlag, Berlin, Germany, 2006, pp. 777–788.
- [14] S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
- [15] I. Harvey, E. A. Di Paolo, R. Wood, M. Quinn, and E. Tuci, "Evolutionary robotics: A new scientific tool for studying cognition," *Artificial Life*, vol. 11, no. 1–2, pp. 79–98, 2005.
- [16] S. H. Strogatz, *Nonlinear Dynamics and Chaos: With Applications to Physics, Biology, Chemistry, and Engineering*. Westview Press, Cambridge, MA, 1994.
- [17] F. Mondada, G. C. Pettinaro, A. Guignard, I. V. Kwee, D. Floreano, J.-L. Deneubourg, S. Nolfi, L. M. Gambardella, and M. Dorigo, "SWARM-BOT: A new distributed robotic concept," *Autonomous Robots*, vol. 17, no. 2–3, pp. 193–221, 2004.

isti



Dr. Vito Trianni received the Laurea (Master of Technology) degree in computer science engineering from Politecnico di Milano, Milan, Italy, in 2000, the Master in Information Technology from the ICT Center of Excellence For Research, Innovation, Education and Industrial Labs Partnership (CEFRIEL), Milan, Italy, in 2001, and the Diplôme d'Études Approfondies and the Ph.D. in Applied Sciences from the Université Libre de Bruxelles, Brussels, Belgium, in 2003 and in 2006, respectively. Currently, he is a Research Fellow at the Laboratory of Autonomous Robotics and Artificial Life, Institute of Cognitive Sciences and Technologies of the Italian National Research Council (LARAL-ISTC-CNR) in Rome, Italy. His research interests span over the field of evolutionary robotics, swarm intelligence, and self-organisation.



Dr. Stefano Nolfi is a research director at the Institute of Cognitive Sciences and Technologies of the Italian National Research Council (ISTC-CNR), where he heads the Laboratory of Autonomous Robotics and Artificial Life (LARAL). He is also coordinator of the RTD European Integrated Project ECAgents: Embodied and Communicating Agents. His research interests are in the field of neuro-ethological studies of adaptive behaviour in natural and artificial agents and include: evolutionary robotics, artificial life, complex systems, neural networks, and genetic algorithms. He is the author of the book S. Nolfi and D. Floreano, *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines* (MIT Press/Bradford Books, 2000) and of more than 100 peer-reviewed articles.