

Minimal Communication Strategies for Self-Organising Synchronisation Behaviours

Vito Trianni and Stefano Nolfi

LARAL-ISTC-CNR, Rome, Italy

Email: vito.trianni@istc.cnr.it, stefano.nolfi@istc.cnr.it

Abstract—The ability to synchronise the individual actions within large groups is an adaptive response observed in many biological systems. Indeed, synchrony can increase the efficiency of a group by maximising the global outcome or by minimising the interference among individuals. In any case, synchronisation appears desirable for a robotic system as it allows to coordinate through time the activities of the group. The main goal of the experiments presented in this paper is the study of self-organising synchronisation behaviours for a group of robots. To do so, we do not postulate the need of internal dynamics. Instead, we stress the importance of the dynamical coupling between robots and environment, which can be exploited for synchronisation, allowing to keep a minimal complexity of both the behavioural and the communication level. We use artificial evolution to synthesise the robot controllers, and we show how very simple communication strategies can produce self-organising synchronisation behaviours that scale to very large groups and that can be transferred to physical robots.

I. INTRODUCTION

Synchrony pervades the world: examples of synchronous behaviours can be found in the inanimate world as well as among living organisms. One of the most commonly cited synchronous behaviour is the one of fireflies from Southeast Asia: thousands of insects have the ability to flash in unison, perfectly synchronising their individual behaviour (see [1], [2]). This phenomenon, reported by amazed travellers since the seventeenth century, has been thoroughly studied and a self-organising explanation has been proposed in order to account for the emergence of synchrony [3], [4]. Fireflies are modelled as a population of pulse-coupled oscillators with equal or very similar frequency. These oscillators can influence each other by emitting a pulse that shifts or resets the oscillation phase.¹ The numerous interactions among the individual oscillator-fireflies are sufficient to explain the synchronisation of the whole population (for more details, see [3], [4]). Despite the effort in understanding the above synchronising mechanism, the adaptive significance of synchronous flashing is not clear yet. Some tend to support a cooperative explanation, for which a cluster of synchronous flashing would result in a very attractive mating signal for faraway female insects. Others support a competitive explanation, for which synchronous flashing is a by-product of the individual attempt to anticipate any other flash (see [2], p. 35). Similar explanations hold for other synchronous behaviours observed in nature, such

¹In some firefly species, it is rather the oscillation frequency that is temporarily altered, having an effect comparable to a phase shift.

as frogs chorusing or crickets chirping. Despite the particular evolutionary pressures leading to synchrony, the latter appears to be a powerful mean for maximising the outcome while minimising the collective effort and/or the interferences among individuals. In this perspective, obtaining synchrony in a robotic system is highly desirable, as it offers the possibility to regulate through time the coordinated effort of a group.

The synchronisation behaviours observed in nature can be a powerful source of inspiration for the design of distributed robotic systems. For example, the self-organising synchronisation mechanism exploited by fireflies was successfully replicated in a group of robots [5]. In this study, the authors designed a specialised neural module for the synchronisation of the group foraging/homing activities, in order to maximise the overall performance. Much as fireflies that emit light pulses, robots communicate through sound pulses that directly reset the internal oscillator designed to control the individual switch from homing to foraging and vice versa. In a follow-up research, it was shown how similar synchronisation mechanisms can be synthesised by artificial evolution [6]. The authors evolved a cooperative foraging behaviour incrementally. Initially, they rewarded the individual ability to explore the environment and find the food source. Then, evolution was continued in a social scenario, and the emergence of communicative behaviours was observed. A further evolutionary refinement led to the emergence of a self-organising synchronisation behaviour based on exactly the same mechanism that was hand-crafted in [5]. Other works study synchrony as the coordination between the body parts of a single robot: for instance, synchronisation issues were considered while studying the gait of an hexapod robot [7].

The main goal of the experiments presented in this paper is the study of self-organising synchronisation based on minimal behavioural and communication strategies. Similarly to the studies presented above, 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 [5], [6], synchronisation is based on the entrainment of the individual internal dynamics through communication: the internal oscillator defines the period and the phase of the individual behaviour, and it is also responsible for communication and synchronisation. 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 [8], 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.

This result is a direct consequence of the attempt to obtain a *complete synchronisation* of the robots movements. Much as ballet dancers feature choreographed gestures, complete synchronisation requires that robots perform perfectly synchronous actions. Additionally, the robots' perceptual flows should be synchronised as well, resulting in a perfect entrainment of the dynamical relationship that each robot has with the environment. This is a stricter requirement than simply synchronising the robots' activities, such as foraging or homing in [5]. It however opens the way to the exploitation of agent-environment interactions rather than internal dynamics. In fact, a sequence of activities defines the phase of a periodic behaviour with a coarse-grained time scale—e.g., the switch from foraging to homing—while a sequence of movements offers a much finer way to recognise the behaviour's phase.

Now, the main problem is defining a robot controller able to exploit the dynamical agent-environment interactions. We use artificial evolution to search the space of the possible behavioural and communication strategies for the synchronisation problem [9], [10]. In particular, we avoid to explicitly reward the use of communication, in order to leave evolution free to explore the space of the possible solutions that lead to a synchronous behaviour. This, however, makes the evolution of communication particularly challenging. In fact, for communication to be in place, it is necessary to contemporary have both the ability to produce a signal and the ability to properly react to the perceived signal [11].

If the evolution of communication is not trivial, it is even less trivial the exploitation of communicative interactions for self-organisation. In this paper, we analyse the properties of the evolved behaviours under a self-organising perspective, evaluating their scalability to large groups of robots. Moreover, we investigate the scalability of communication *per se*, in order to evaluate the efficiency of the evolved strategy when not constrained by the physical interactions among the robots. Finally, we analyse the robustness of the evolved behaviours by testing them on physical robots.

This paper is organised as follows. In Section II we present the experimental setup devised to evolve in simulation the self-organising synchronisation behaviours. Section III presents the obtained results, analysing the communication strategies and the scalability properties of the evolved controllers. Section IV discusses the results obtained by testing the controllers with physical robots. Section V concludes the paper.

II. EXPERIMENTAL SETUP

As mentioned above, in this work we aim at studying the evolution of behavioural and communication strategies for synchronisation. For this purpose, we define a simple, idealised scenario that anyway contains all the ingredients needed for




Fig. 1. The *s-bot*. The traction system is composed by both tracks and wheels, which ensure a good mobility over moderately rough terrains. The rotating turret holds the microphones and the loudspeakers used for communication.

our study. The task requires that each robot in the group displays a simple periodic behaviour, that is, moving back and forth from a light bulb positioned in the centre of the arena. Moreover, robots have to synchronise their movements, so that their oscillations are in phase with each other.

The robots used in this experiments are the *s-bots* (see Fig. 1), which are small autonomous robots having the ability to connect one to the other and to self-assemble [12], [13]. The *s-bot* has a differential drive traction system composed of tracks and wheels. Above the traction system, the turret holds various sensory systems and the gripper for making connections with other *s-bots*.² The evolutionary experiments presented in this paper are performed in simulation, using a simple kinematic model of the *s-bots*. Each *s-bot* is provided with infrared sensors and ambient light sensors, which are simulated using a sampling technique [15]. In order to communicate with each other, *s-bots* are provided with a very simple signalling system, which 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 *s-bot*. The tone is perceived in a binary way, that is, either there is someone signalling in the arena, or there is no one.

The arena is a square of 6×6 meters. In the centre, a cylindrical object supports the light bulb, which is always switched on. The light intensity perceived by the *s-bot*'s sensors decreases quadratically with the distance from the light bulb. However, light can be perceived from every position in the arena. At the beginning of every trial, three *s-bots* are initially positioned in a circular band ranging from 0.2 to 2.2 meters from the centre of the arena. The robots have to move back and forth from the light, making oscillations with an optimal amplitude of 2 meters.

A. The Controller and the Evolutionary Algorithm

In the experiments reported here, artificial evolution is used to synthesise the connection weights of simple neural controllers with fixed architecture. The controller of each *s-bot*

²The assembling capability of the *s-bots* is not the focus of these experiments. For more details on self-assembling *s-bots*, see [14].

consists in a fully connected, feed forward neural network—a perceptron network. The neural network has 11 sensory neurons directly connected to 3 motor neurons. The sensory neurons are simple relay units while the output neurons are sigmoid units whose activation is computed as follows:

$$y_j = \sigma \left(\sum_i w_{ij} I_i + \beta_j \right), \quad \sigma(z) = \frac{1}{1 + e^{-z}}, \quad (1)$$

where I_i is the activation of the i^{th} input unit, β_j is the bias term, y_j is the activation of the j^{th} output unit, w_{ij} is the weight of the connection between the input neuron i and the output neuron j , and $\sigma(z)$ is the sigmoid function.

Four sensory neurons— I_1 to I_4 —are dedicated to the readings of four ambient light sensors, positioned in the front and in the back of the *s-bot*. Six sensory neurons— I_5 to I_{10} —receive input from a subset of the infrared proximity sensors evenly distributed around the *s-bot*'s turret. The last sensory neuron I_{11} receives a binary input corresponding to the perception of sound signals. The activation states of the first two motor neurons— y_1 and y_2 —is scaled onto the range $[-\omega_M, +\omega_M]$, where ω_M is the maximum angular speed of the wheels ($\omega_M \approx 4.5$ rad/s). The third motor neuron controls the speaker in such a way that a sound signal is emitted whenever the activation state y_3 is greater than 0.5.

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 20 best individuals are selected for reproduction and retained in the subsequent generation. Each genotype reproduces four times, applying mutation with 5% probability of flipping a bit. The evolutionary process is run for 500 generations.

B. The Fitness Computation

During the evolution, a genotype is mapped into a control structure that is cloned and downloaded in all the *s-bots* taking part in the experiment (i.e., we make use of a homogeneous group of *s-bots*). Each genotype is evaluated 5 times—i.e., 5 trials. Each trial differs from the others in the initialisation of the random number generator, which influences both the initial position and orientation of the *s-bots* within the arena. Each trial lasts $T = 900$ simulation cycles, which correspond to 90 seconds of real time.

The fitness of a genotype is the average performance computed over the 5 trials in which the corresponding neural controller is tested. During a single trial, the behaviour produced by the evolved controller is evaluated by a 2-component fitness function: $F = 0.6 \cdot F_M + 0.4 \cdot F_S$. The movement component F_M rewards robots that oscillate back and forth from the light bulb. For each *s-bot* s , we look at the closest and farthest distances from the centre of the arena (see the




Fig. 2. Computation of the individual movement fitness. The continuous line represents the trajectory of a robot, which oscillates while circuiting around the centre of the arena. The black dots represent the farthest and closest distances reached. The arrow $A_{s,i}$ indicates the oscillation amplitude between the second and third point on the trajectory.

black dots in Fig. 2). We compute the oscillation amplitudes $A_{s,i}$, with $i = [1, M]$, corresponding to the radial distance covered between two consecutive points, as shown in Fig. 2. Then, the individual movement fitness $F_{M,s}$ is computed as follows:

$$F_{M,s} = \frac{1}{M - m + 1} \sum_{i=m}^M \Theta(A_{s,i}/A_o), \quad (2)$$

where A_o is the optimal amplitude (2 meters). The function $\Theta(x) = 1 - |1 - x|$ simply rewards those oscillations that better approximate the optimum. Given the maximum speed of the *s-bots*, it is possible to compute the maximum number of oscillatory movements having amplitude A_o that can be performed in T control cycles, referred to as $M_o(T)$. In computing $F_{M,s}$, we consider only the last $M_o(T)$ oscillatory movements performed, which corresponds to set $m = \max(1, M - M_o(T))$ as the first oscillatory movement to be considered. The overall movement fitness F_M is computed as the minimum among the individual values of the single *s-bots*.

The second fitness component F_S rewards synchrony among the robots. Synchrony among two *s-bots* can be evaluated as the cross-correlation coefficient between the sequences of the distances from the light bulb. The cross-correlation coefficient ϕ_{xy} of two distance sequences $d_x(t)$ and $d_y(t)$ can be defined as:

$$\phi_{xy} = \frac{\Phi_{xy}}{\sqrt{\Phi_{xx}\Phi_{yy}}}, \quad \Phi_{xy} = \frac{1}{T} \sum_{t=1}^T d_x(t)d_y(t). \quad (3)$$

The coefficient ϕ_{xy} can take values in $[-1,1]$, where 1 indicate perfect synchrony and -1 perfect asymmetry. The synchrony component F_S is computed as the product among the cross-correlation coefficients of all possible pairs $\langle x, y \rangle$ among the *s-bots*:

$$F_S = \prod_{x,y} \max(0, \phi_{xy}), \quad (4)$$

TABLE I

POST-EVALUATION RESULTS FOR THE BEST CONTROLLERS OF THE 20 EVOLUTIONARY RUNS. THE AVERAGE FITNESS AND THE STANDARD DEVIATION COMPUTED OVER 500 TRIALS ARE SHOWN. BOLD VALUES REFERS TO EVOLUTIONARY RUNS THAT PRODUCED A COMMUNICATION STRATEGY IN WHICH SIGNALLING WAS EXPLOITED FOR SYNCHRONISATION.

c_i	c_1	c_2	c_3	c_4	c_5
F	0.63 ± 0.13	0.49 ± 0.16	0.58 ± 0.06	0.67 ± 0.07	0.65 ± 0.16
c_i	c_6	c_7	c_8	c_9	c_{10}
F	0.56 ± 0.19	0.66 ± 0.02	0.51 ± 0.22	0.65 ± 0.12	0.55 ± 0.21
c_i	c_{11}	c_{12}	c_{13}	c_{14}	c_{15}
F	0.60 ± 0.10	0.60 ± 0.11	0.73 ± 0.09	0.74 ± 0.07	0.68 ± 0.13
c_i	c_{16}	c_{17}	c_{18}	c_{19}	c_{20}
F	0.65 ± 0.11	0.52 ± 0.14	0.56 ± 0.09	0.66 ± 0.16	0.61 ± 0.14

Notice that F_S is bounded into the interval $[0,1]$.

In addition to the fitness computation described above, an indirect selective pressure for the evolution of obstacle avoidance is given by blocking the motion of robots that collide. When this happens, the performance is negatively influenced. Additionally, a trial is normally terminated after $T = 900$ simulation cycles. However, a trial is also terminated if any of the *s-bots* crosses the borders of the arena.

III. RESULTS

We performed 20 evolutionary runs, each starting with a different population of randomly generated genotypes. After the evolutionary phase, we selected a single genotype per evolutionary run, chosen as the best individual of the final generation. We refer to the corresponding controllers as $c_i, i = 1, \dots, 20$. In order to evaluate their performance, these controllers have been evaluated in 500 different trials. The results are summarised in Table I, showing the average performance and the standard deviation. It is possible to notice that some controllers do not achieve a good performance, while c_{13} and c_{14} outperform all the other controllers.

Direct observation of the evolved behaviours showed that in some evolutionary runs—9 out of 20—communication was not evolved. In fact, either robots always signal or they never do: in any case, there is no information transfer to be exploited for synchronisation. This justifies the lower performance obtained by the corresponding controllers, as shown in Table I. These results confirm the difficulty of evolving suitable communication strategies for synchronisation. In fact, as mentioned in Section I, the evolution of signalling must be accompanied by a suitable reaction to the signal. If this is not the case, signalling may just interfere with the sensory-motor coordination of partially evolved solutions. Therefore, a certain number of fruitful mutations are required to obtain a successful communication, which concerns both the signalling behaviour and the reaction to the perceived signal.

Despite the above difficulties, 11 out of 20 evolutionary runs were successful, resulting in simple communication strategies in which signalling was exploited for synchronisation (the performance of the corresponding controllers is indicated in bold in Table I). All evolved solutions result in similar behaviours, characterised by two stages, that is, phototaxis when the *s-bots* approach the light bulb, and antiphototaxis when the *s-bots*

move away from it. Signalling is generally performed only during one of the two stages. We can classify the evolved controllers in three classes, according to the individual reaction to the perception of a sound signal.

The first class—composed of c_1, c_4, c_7, c_{19} and c_{20} —involves behaviours in which signalling strongly correlates with antiphototaxis. This can be appreciated looking at the top part of Fig. 3, in which the *s-bots*' distances from the centre and the group signalling behaviour are plotted through time. It is possible to notice that whenever a robot signals, its distance from the light increases and, vice versa, when no signal is perceived the distance decreases. Synchronisation is normally achieved after one oscillation and it is maintained for the rest of the trial, the robots moving in complete synchrony with each other. This is possible thanks to the evolved behavioural and communication strategy, for which a robot emits a signal while performing antiphototaxis and reacts to the perceived signal by reaching and keeping a specific distance away




Fig. 3. The synchronisation behaviour of c_7 . Top: *s-bots*' distances from the light bulb are plotted against the simulation cycles, in order to appreciate the synchronisation of the individual movements. The coloured areas indicate when a signal is emitted by any of the *s-bots* in the arena. Such a signal is perceived by the robots and exploited for synchronisation (see text for details). Bottom: the distance and signalling behaviour of a single *s-bot* are plotted against the simulation cycles. From cycle 500 to 1000, a signal is artificially created, which simulates the behaviour of an *s-bot*. This allows to visualise the reaction of an *s-bot* to the perception of a sound signal.




Fig. 4. The synchronisation behaviour of c_{13} . See Fig. 3 for details.

from the centre of the arena. As shown in the bottom part of Fig. 3, in presence of a continuous signal—artificially created from cycle 500 to cycle 1000—an *s-bot* suspends its normal oscillatory movement to maintain a constant distance from the centre. As soon as the sound signal is stopped, the oscillatory movement starts again. Synchronisation is possible because robots are homogeneous, therefore they all present an identical response to the sound signal that makes them move to the outer part of the arena. As soon as all robots reach the same distance from the centre, signalling ceases and synchronous oscillations can start. To better understand the synchronisation mechanism that characterises this class of controllers, it is worth considering an *s-bot* as an “embodied oscillator”, its individual behaviour being characterised by a period and a phase. The latter is defined by the *s-bot*’s position in its configuration space, which can be defined as the set of possible distances and orientations with respect to the light bulb. Whenever a robot perceives a sound signal, it *resets* the oscillation phase by attaining a particular configuration—i.e., reaching and maintaining a specific distance and orientation with respect to the light bulb. This *reset mechanism* results in the complete synchronisation of the robots movements, as it is exploited by the robots to reduce and cancel the phase difference of their oscillations. Clearly, the phase reset is not instantaneous, because *s-bots* need time to reach the target configuration, due to their embodiment. It follows that the reset configuration must be maintained for enough time to let all *s-bots* converge. In conclusion, the evolved behavioural and communication strategies allow a fast synchronisation of the robots activities, because they force all robots to perform synchronously phototaxis or antiphototaxis since the beginning of a trial, as a reaction to the presence or absence of a sound signal respectively. It also allows a fast synchronisation of the movements thanks to the reset of the oscillation phase. Finally, it provides a mean to fine-tune and maintain through time a complete synchronisation, because the reset mechanism allows to continuously correct even the slightest phase difference.

The second class—composed of c_5 , c_9 , c_{13} , c_{15} and c_{16} —




Fig. 5. The synchronisation behaviour of c_{14} . See Fig. 3 for details.

features the same synchronisation strategy described above, but with inverse movements. Here, signalling strongly correlates with phototaxis, while robots move away from the centre when no signal is perceived (see the top part of Fig. 4). Also in this case synchronisation is the result of a reset mechanism: whenever a signal is perceived, *s-bots* perform phototaxis and keep a constant distance close to the light bulb (see the bottom part of Fig. 4). As soon as the whole group reaches similar distances from the centre of the arena, signalling ceases and the oscillatory movement starts again. We observed here a better precision in synchronisation, probably due to the steeper intensity gradient perceived in proximity of the light bulb. This allows to precisely maintain the reset configuration and therefore to rapidly achieve a complete synchronisation.

The last controller— c_{14} —makes a class on its own, producing a peculiar behaviour. In this case, it is rather the *absence* of a signal that strongly correlates with phototaxis. The individual reaction to the perceived signal can be appreciated looking at the bottom part of Fig. 5. When the continuous signal is artificially created (see simulation cycles 500 to 1000), the *s-bot* performs both phototaxis and antiphototaxis. However, as soon as the signal is removed, the *s-bot* approaches the light bulb. This behaviour results from the evolved communication strategy, which enables the synchronisation of the *s-bots*’ activities, that is, phototaxis and antiphototaxis. Differently from the mechanism presented above, here there is only a coarse-grained phase reset, which concerns the activities rather than the very movements: *s-bots* initially synchronise only the movement direction but not the distance at which the oscillatory movements are performed (see the top part of Fig. 5).³ Despite this limitation, this mechanism allows a very fast and precise synchronisation of the *s-bots*’ phototaxis and antiphototaxis, which is probably the reason why it was evolved in the first place. In order to achieve a complete synchronisation, an additional mechanism was synthesised, which

³Notice that this is a reset mechanism that works on a subset of the variables describing the configuration space.




Fig. 6. Scalability of the successful controllers. 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. The empty circles mark the outliers. The horizontal grey line shows the mean value over 500 trials measured in the evolutionary conditions, in order to better evaluate the scalability property (see also Table I).

allows to precisely entrain the movements of the robots on a fine-grained scale. This mechanism influences the distance covered by an *s-bot* during antiphototaxis: *s-bots* that are farther away from the light bulb slightly bend their trajectory and therefore cover a distance range shorter than the one covered by the other robots in the same time. In this way, the phase difference among *s-bots* is progressively reduced, until all *s-bots* are completely synchronised (see Fig. 5 top).

A. Scalability of the Evolved Behaviours

The above analysis clarified the role of communication in determining the synchronisation among the different robots. In this section, we analyse the scalability of the evolved neural controllers when tested in larger groups of robots. For this purpose, we evaluated the behaviour of the successful controllers using 3, 6, 9 and 12 *s-bots*. In order to evaluate the performance, we use the fitness function defined in Section II-B, but here we measure synchrony in a slightly different way,⁴ computing the minimum cross-correlation coefficient ϕ_{xy} among all possible pairs $\langle x, y \rangle$, as follows:

$$\hat{F}_S = \min_{x,y} \phi_{xy} \quad (5)$$

The obtained results are plotted in Fig. 6. It is possible to notice that most of the best evolved controllers have a good performance for groups composed of 6 *s-bots*. In such condition, in fact, *s-bots* are able to distribute in the arena without interfering with each other. An exception is given by c_1 : in this case, the initial coordination takes longer time and some robots move too distant from the centre of the arena, either exceeding the arena bounds or being not able to perceive the light bulb. In both cases, the performance is close to 0.

⁴We do not use F_S (see (4)) because it is based on a product and it does not scale well with the group size: for example, 10 robots form 45 pairs. If $\phi_{xy} = 0.99$ for all pairs—a very good behaviour—then $F_S \approx 0.64$.

Many controllers present a good behaviour also when groups are composed of 9 *s-bots*. However, we also observe various failures due to interferences among robots and collisions. The situation gets worse when using 12 *s-bots*: the higher the density of robots, the higher the number of interferences that lead to failure. In this case, most controllers achieve a good performance only sporadically. Only c_4 and c_7 makes exception, being able to systematically achieve synchronisation despite the increased difficulty of the task.

B. Scalability of the Synchronisation Mechanism

The scalability analysis performed in the previous section takes into account the complete behaviour, therefore including collision avoidance. When the density of robots is too high, the spatial constraints limit the synchronisation ability, because avoiding collision interferes with the ability to maintain a periodic behaviour and to synchronise with the rest of the group. Moreover, possible collisions among robots prevent the group from synchronising. In order to analyse the scalability property of the synchronisation mechanism only, we evaluate the evolved controllers removing the physical interactions among the robots, as if each *s-bot* is placed in a different arena and perceives the other *s-bots* only through sound signals.

Removing the robot-robot interactions allows us to test large groups of robots—we used 12, 24, 48 and 96 *s-bots*. The obtained results are summarised in Fig. 7. Not surprisingly, c_1 does not scale, being affected by the same problems described above. A similar problem affects the behaviour produced by c_{19} . However, many controllers—namely c_4 , c_7 , c_{13} , c_{14} , c_{15} and c_{20} —perfectly scale, having a performance very close to the mean performance measured with 3 *s-bots*. A slight decrease of performance is justified by the longer time required by larger groups to converge to perfectly synchronised movements (see for example c_7 and c_{20}).

Some controllers—namely c_4 , c_5 , c_9 , c_{14} and c_{16} —present




Fig. 7. Scalability of the synchronisation mechanism. See the caption of Fig. 6 for details.

an interference problem that prevents the group from synchronising when a sufficiently large number of robots is used. In such condition, the signals emitted by different *s-bots* at different times may overlap and may be perceived as a single, continuous tone (recall that the sound signals 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. Such interference can be observed only sporadically for c_4 and c_{14} , but it strongly affects the performance of the other controllers—namely c_5 , c_9 and c_{16} . 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 mechanism. 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.




Fig. 8. Distances from the light bulb and collective signalling behaviour of the real *s-bots*.

IV. TESTS WITH PHYSICAL ROBOTS

So far, we have shown how artificial evolution can synthesise efficient and scalable synchronisation mechanisms which are based on minimal communication strategies. In this section, we test the robustness with respect to the transfer to physical robots. Among the best evolved controllers, we chose c_{13} as it presented a high performance and good scalability properties. The neural network controller is used on the physical *s-bots* exactly in the same way as in simulation. The sensor readings are taken every 100 ms, they are scaled in the range $[0,1]$ and finally fed to the neural network. The outputs of the network are used to control the wheels and the loudspeaker. The only differences with the simulation experiments are in the experimental arena, which is four times smaller in reality (1.5×1.5 meters), and accordingly the light bulb is approximately four times less intense. An overhead camera was used to record the movements of the *s-bots*, and their trajectories were extracted using a tracking software [16].

The behaviour of the physical robots presents a good correspondence with the results obtained in simulation. Fig. 8 shows the *s-bots*’ distance from the light bulb recorded during a successful trial.⁵ It is possible to notice how synchrony is quickly achieved and maintained throughout the whole trial, notwithstanding the high sensors and actuator noise and the differences among the three robots. The latter deeply influence the group behaviour: *s-bot* happened to have a different maximum speed which let them cover different distances in the same time interval. Therefore, if phototaxis and antiphototaxis are very well synchronised, as a result of the communication strategy exploited by the robots, it is possible to notice some differences in the maximum distance reached.

V. CONCLUSIONS

In this paper, we studied the evolution of behavioural and communication strategies for the synchronisation of the

⁵See <http://lara1.istc.cnr.it/trianni/sync.html> for some movies.

movements in a group of robots. The obtained results indicate that self-organising synchronisation behaviours can be successfully evolved, which also scale to large groups. We adopted a minimal approach that does not postulate the need of internal dynamics for the robots to be able to synchronise. Instead, we stress the importance of the dynamical coupling between robots and environment. Robots can be described as embodied oscillators, their behaviour being characterised by a period and a phase. In this perspective, the movements of the robot correspond to an advancement of the oscillation phase. Therefore, *s-bots* can control and modify their phase simply by moving in the environment and by modifying their dynamical relationship with it. In this way, simple and reactive behavioural and communication strategies are sufficient to implement efficient synchronisation mechanisms. Most of the evolved solutions rely on a particular *reset mechanism*, that allows robots to quickly achieve a complete synchronisation. This mechanism is based on the identification of a particular position in the configuration space in which *s-bots* do not signal. This position corresponds to a particular phase of the periodic behaviour, which is held until a signal emitted by any other robot is perceived. When all *s-bots* reach the reset position, the phase differences are cancelled and synchronisation is achieved. This is not the only mechanism observed in the evolved solutions. One of the evolved controllers present two synchronisation mechanisms, which allow a very quick activity synchronisation—i.e., phototaxis and antiphototaxis are immediately performed in perfect sync—followed by a slower convergence to a complete synchronisation—i.e., *s-bots* gradually entrain their movements.

We have also analysed the scalability of the evolved controllers, showing that synchronisation can be obtained also in large groups, despite controllers were evolved using three *s-bots* only. However, when many *s-bots* are placed in the same arena, the avoidance behaviour and possible collisions strongly interfere with the ability to synchronise. This is the principal limitation of the evolved controllers, which does not allow to perform tests with large groups. However, we could appreciate the scalability of the synchronisation mechanism to very large groups by neglecting the physical interactions among the robots. We tested the evolved controllers with up to 96 *s-bots*, and we found that many evolved solutions have a very good scalability. Few controllers presented an interference problem at the level of the signalling behaviour, that prevented robots from extracting a relevant information to be used for synchronisation. Finally, we tested one of the evolved controllers with the real *s-bots*, and also in this case we could observe synchronisation, further proving the robustness of the evolved controller.

In future work, we intend to study synchronisation under two different perspectives. From one hand, we aim at investigating the adaptive significance of synchronous behaviours. Here, synchrony should result only from the need to coordinate the activities of a group. On the other hand, we are interested in exploiting synchrony as a general principle—among others, see for example [17]—to evolve cooperative behaviours.

ACKNOWLEDGEMENT

This work was supported by the ECAgents project, a project funded by the Future and Emerging Technologies programme (IST-FET) of the European Commission, under grant 001940. The information provided is the sole responsibility of the authors and does not reflect the Community's opinion. The Community is not responsible for any use that might be made of data appearing in this publication. The authors thank Nikolaus Correll and Alcherio Martinoli for providing the tracking software used in the experiments presented in this paper.

REFERENCES

- [1] S. Camazine, J.-L. Deneubourg, N. Franks, J. Sneyd, G. Theraulaz, and E. Bonabeau. *Self-Organization in Biological Systems*. Princeton University Press, Princeton, NJ, 2001.
- [2] S. H. Strogatz. *Sync: The Emerging Science of Spontaneous Order*. Hyperion Press, New York, NY, 2003.
- [3] R. E. Mirolo and S. H. Strogatz. Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662, 1990.
- [4] S. H. Strogatz and I. Stewart. Coupled oscillators and biological synchronization. *Scientific American*, 269(6):102–109, 1993.
- [5] S. Wischmann, M. Huelse, J. F. Knabe, and F. Pasemann. Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14(2):117–127, 2006.
- [6] S. Wischmann and F. Pasemann. Neural dynamics of emergent unstructured communication systems in cooperating robots. *Artificial Life*, 2006. submitted.
- [7] M. Mazzapioda and S. Nolfi. Synchronization and gait adaptation in evolving hexapod robots. In S. Nolfi et al., editor, *From animals to animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*, volume 4095 of *LNAI*, pages 113–125. Springer Verlag, Berlin, Germany, 2006.
- [8] R. Pfeifer and C. Scheier. Sensory-motor coordination: The metaphor and beyond. *Robotics and Autonomous Systems*, 20(2–4):157–178, 1997.
- [9] S. Nolfi and D. Floreano. *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. MIT Press/Bradford Books, Cambridge, MA, 2000.
- [10] I. Harvey, E. A. Di Paolo, R. Wood, M. Quinn, and E. Tuci. Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1–2):79–98, 2005.
- [11] S. Nolfi. Emergence of communication in embodied agents: Co-adapting communicative and non-communicative behaviours. *Connection Science*, 17(3–4):231–248, 2005.
- [12] M. Dorigo, V. Trianni, E. Şahin, R. Groß, T. H. Labelle, G. Baldassarre, S. Nolfi, J.-L. Deneubourg, F. Mondada, D. Floreano, and L. M. Gambardella. Evolving self-organizing behaviors for a *swarm-bot*. *Autonomous Robots*, 17(2–3):223–245, 2004.
- [13] 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*, 17(2–3):193–221, 2004.
- [14] E. Tuci, R. Groß, V. Trianni, F. Mondada, M. Bonani, and M. Dorigo. Cooperation through self-assembling in multi-robot systems. *ACM Transactions on Autonomous and Adaptive Systems*, 1(2):115–150, 2006.
- [15] O. Miglino, H. H. Lund, and S. Nolfi. Evolving mobile robots in simulated and real environments. *Artificial Life*, 4:417–434, 1995.
- [16] N. Correll, G. Sempo, Y. Lopez de Meneses, J. Halloy, J.-L. Deneubourg, and A. Martinoli. SwisTrack: A tracking tool for multi-unit robotic and biological systems. In *Proceedings of the 2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS '06)*, pages 2185–2191. IEEE Press, Piscataway, NJ, 2006.
- [17] M. Prokopenko, V. Gerasimov, and I. Tanev. Evolving spatiotemporal coordination in a modular robotic system. In S. Nolfi et al., editor, *From animals to animats 9: Proceedings of the Ninth International Conference on Simulation of Adaptive Behaviour*, volume 4095 of *LNAI*, pages 558–569. Springer Verlag, Berlin, Germany, 2006.