# Evolving Teamwork and Role-Allocation with Real Robots

**Matt Quinn[1], Lincoln Smith[1], Giles Mayley[2] and Phil Husbands[1]**
[1]Centre for Computational Neuroscience and Robotics, University of Sussex, Brighton, BN1 9QG, U.K.
[2] Mathematique Appliqué S.A. Brighton Innovation Centre, University of Sussex, Brighton, BN1 9QG, U.K.

**Abstract**

We report on recent work in which we employed artificial evolution to design neural network controllers for small, homogeneous teams of mobile autonomous robots. The robots are evolved to perform a formation movement task from random starting positions, equipped only with infrared sensors. The dual constraints of homogeneity and minimal sensors make this a non-trivial task. We describe the behaviour of a successful evolved team in which robots adopt and maintain functionally distinct roles in order to achieve the task. We believe this to be the first example of the use of artificial evolution to design coordinated, cooperative behaviour for real robots.

## Introduction

In this paper we report on our recent work evolving controllers for robots which are required to work as a team. The word 'team' has been used in a variety of senses in both the multi-robot and the ethology literature, so it is appropriate to start the paper with a definition. We will adopt the definition given by Anderson and Franks (2001) in their recent review of team behaviour in animal societies. They identify three defining features of team behaviour. Firstly, individuals make different contributions to task success, i.e. they must perform different sub-tasks or roles (this does not preclude more than one individual adopting the same role; there may be more individuals than roles). Secondly, individual roles or sub-tasks are interdependent (or "interlocking") requiring structured cooperation; individuals operate concurrently, coordinating their different contributions in order to complete the task. Finally, a team's organisational structure persists over time, although its individuals may be substituted, or swap roles (Anderson & Franks 2001).

The designer of a multi-robot team faces a number of challenges. One of which arises because a team is a structured system. Robots must be designed to behave in such a way that the team will both become and remain appropriately organised. This requires ensuring that all the individual roles or sub-tasks are appropriately allocated. One way to address this problem is to design a team in which each individual's role is predetermined (Balch & Arkin 1998, e.g.). In addition to its organisational advantages, the pre-allocation of roles has the additional advantage of specialisation: Division of labour means that each robot's behavioural and morphological design can be tailored to its particular task. In natural systems, this type to team organisation is often found amongst eusocial insects, where roles may be caste-specific. (Detrain & Pasteels 1992, e.g). Despite the organisational advantages of system heterogeneity and the efficiency benefits of specialisation, we are interested in the design of homogeneous systems. In a homogeneous multi-robot system, each robot is built to the same design, and has an identical controller. Our interest in homogeneous robot teams stems from their potential for system-level robustness and graceful degradation due to the interchangeability of team members (although this is not an issue that we will be addressing in this paper). Since each robot is capable of performing any role or sub-task, homogeneous systems are potentially better than heterogeneous teams at coping with the loss of an individual member. Lack of role specialisation also has potential benefits for organisational flexibility (Stone & Veloso 1999). However, from the perspective of team organisation, the constraint of homogeneity makes the design task more difficult. In a homogeneous team there are no differences between robots' control systems or morphologies which can be exploited for the purposes of team organisation. Other mechanisms must be employed to facilitate the dynamic allocation and coordination of roles.

Dynamic role allocation and closely coordinated cooperation are two areas which have been addressed by a number of researchers in the field of multi-robot systems, resulting in successful implementations of such tasks such as cooperative transport (Chaimowicz *et al.* 2001), robot football (Stone & Veloso 1999), and coordinated group movement (Matarić 1995). Solutions have relied heavily on the use of essentially global information shared by radio communication. For example, in Matarić's (1995) implementation of coordinated movement with homogeneous robots, robots made use of a common coordinate system (through radio beacon triangulation) and exchanged positional information via radio communication in order to remain coordinated (Matarić 1995). Mechanisms for dynamic task or role allocation rely on communication protocols by which robots globally advertise or negotiate their current (or intended) roles (Chaimowicz *et al.* 2001; Matarić & Sukhatme 2001; Stone & Veloso 1999, e.g.).

Our work differs from that of these researchers. We wish to design teams in which system-level organisation arises, and is maintained, solely through local interactions between individuals which are constrained to utilise minimal and ambiguous local information. Systems capable of functioning under such constraints have some interesting potential engineering applications (see, for example, (Hobbs, Husbands, & Harvey 1996) for discussion of the need for minimal systems in the space industry). However, they are also interesting from an adaptive behaviour perspective, providing an example of a phenomenon often referred to as 'self-organising' or 'emergent' behaviour (Camazine *et al.* 2001).

Imposing the joint constraints of homogeneity and minimal sensors leaves us with a complex design task. One which cannot easily decomposed and addressed by conventional 'divide and conquer' design methodologies. Instead, it is a problem exhibiting significant interdependence of its constituent parts. For this reason, we have adopted an evolutionary robotics approach and employed artificial evolution to automate the design process, since such an approach is not constrained by the need for decomposition or modular design (Nolfi 1998).

We believe that the work reported in this paper is the first successful use of evolutionary robotics methodology to develop cooperative, coordinated behaviour for a real multi-robot system. To date, this research field has focussed almost exclusively on single robot systems. (See (Nolfi & Floreano 2000) and (Meyer, Husbands, & Harvey 1998) for a good surveys of evolutionary robotics research). Insofar as we are aware there are only two other published examples of the evolution of controllers instantiated on more than one physical robot. However, neither of these are co-operative systems. The first example is due to Floreano, Nolfi and Mondada. They evolved two populations of neural network controllers for Khepera mini-robots as part of a project investigating the dynamics of predator-prey co-evolution (Floreano, Nolfi, & Mondada 1998; Nolfi & Floreano 1998). One population were evolved to perform a 'predator' role, the other, a 'prey' role. Controllers were downloaded onto real robots and evaluated in pair-wise contests. The behaviour of the controllers that they evolved provides an interesting example of coordination, but in competitive system. The second example is due to Watson, Ficici and Pollack (1999). They evolved minimal neural network controllers for population of eight 'Tupperbot' mini-robots. The robots were evolved to perform photo-taxis—an individual task—and evolution was facilitated by local, probabilistic transfer of genetic material between robots via infrared communication. Their work is interesting as a proof-of-concept example of 'embodied evolution'. However, neither cooperative nor coordinated behaviours were required, nor evident in the behaviour which evolved.

The work which we will describe in this paper represents our first experiments in the evolutionary design of homogeneous multi-robot teams. We used three robots, each minimally equipped with four active infra-red sensors, and two motor-driven wheels. Robot controllers were evolved to perform a formation movement task, in an obstacle-free environment, starting from random initial positions. The robots and their task are introduced in more detail in next two sections. Robots were controlled by neural networks which were evolved in simulation, before being successfully transferred onto real robots. The networks, simulation and evolutionary machinery are covered in the fourth section. The penultimate section describes the successful behaviour of one of the evolved teams in some detail, showing that task success is dependent on the robots performing as team, in accordance with definition given at the beginning of this paper.

## The Robots

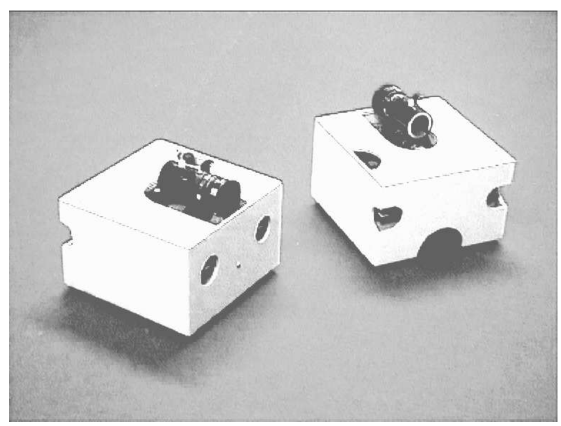

Figure 1: Two members of the three-robot team. The cameras shown are not used for the experiments described in this paper.

For these experiments, we used three robots, each been built to the same specification; two of the robots are shown in figure 1. Each robot's body is 16.75 cm wide by 16.75 cm long by 11 cm high (this excludes the additional height of its unused camera). Two motor-driven wheels, made of foam-rubber, are arranged one on either side of the robot and provide locomotion through differential drive; the robots have an average top speed of 6cm/s. An un-powered castor wheel, placed rear-centre, ensures stability. In the experiments described in this paper, a robot's *only* source of sensory input comes from its four active infrared sensors, each comprising a paired infrared emitter and receiver. Each robot has two infrared sensors at the front and two at rear, as illustrated in figure 2. Although each robots is also equipped with a 64 pixel linear CCD array camera (shown in the diagram), a 360 degree electronic compass, bump sensors, and wheel rotation sensors (i.e. shaft encoders), the controllers we evolved were prevented from making use of any of these devices.

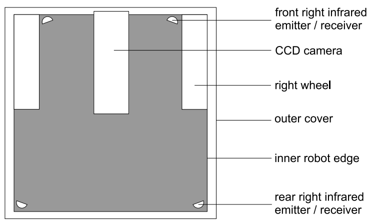


Figure 2: Plan view of a robot.

The robots are controlled by a host computer, with each robot sending its sensor readings to and receiving its effector activations from this machine via a radio link. Each robot uses a 80C537-based micro-controller board for low-level control. The host computer is responsible for running the controller for each robot, updating each controller's inputs with the sensor readings from the appropriate robot, and transmitting the controller's output to the robot. In these experiments, each robot was updated at approximately 5 Hz. It should be noted that although the physical instantiation of the robots has been implemented as a host/slave system, conceptually the robots are to be considered as independent, autonomous agents by virtue of the logical division of control into distinct and self-contained controllers on the host machine.

### Infrared Sensors

The reader may not be familiar with the limitations of active infrared sensors, especially those peculiar to a multi-robot scenario, so we will address them in some detail. An active infrared sensor comprises a paired infrared emitter and receiver. Its normal function is to emit an infrared beam and then measure the amount of infrared which reflects back from nearby objects. In this way our robots can use their sensors to detect other robots up to a maximum of about 18cm (i.e. just over one body length away). The dark grey beams in the left-hand panel of figure 3 approximately indicate the limited areas in which a robot can detect other robots in this manner. IR sensors are sometimes referred to as proximity sensors, however this is somewhat misleading. Whilst the sensor reading due to reflected IR is a non-linear, inverse function of the distance to the object detected, it is also a function of the angle at which the emitted beam strikes the surface of the object, and of the proportion of the beam which strikes that object. It is because an IR sensor reading combines these three factors into a single value that, even in normal function, sensor readings are ambiguous.

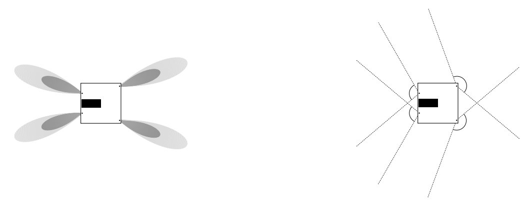

Figure 3: **Left:** The extent to which reflected IR can be sensed (dark grey area), and the extent to which IR beam is perceptible to other robots (light grey area). **Right:** The angles from which a robot can perceive the IR emissions of others

The ambiguity of IR sensors is significantly increased in a multi-robot scenario, because the robots' sensors interfere with one another. Since each robot is constantly emitting IR, a robot's IR emissions can also be directly sensed by other robots. The light grey beams in the left-hand panel of figure 3 indicates the approximate area in which a robot's infrared emissions may be directly detected by other robots. The maximum range at which emissions can be detected is approximately 30cm—almost twice the range at which a robot can detect an object by reflected IR. The right-hand panel of the same diagram illustrates the range of angles at which a robot can receive the IR emissions of other robots. The sensor value due to receiving another's IR emissions is also the combined function of a number of factors: It will depend on the distance between the robots, but also the angle at which the emitted beam strikes the other robot's receiver, and which portion of the beam strikes the receiver (IR is significantly more intense at the centre of the beam than at the edges). Readings due to direct IR are thus ambiguous for the same reasons as reflected IR. However, ambiguity is compounded by the fact that, to the robot, readings due to reflected IR are indistinguishable from those due to the reception of IR emissions of other robots. Moreover, a sensor reading may be the result of a combination of both reflected and direct IR and it may be due to one or both of the other robots.

## The Task

The task with which we present the robots is an extension of that used in previous work which involved two simulated Khepera robots (Quinn 2001). Adapted for three robots, the task is as follows: Initially, the three robots are placed in an obstacle-free environment in some random configuration, such that each robot is within sensor range of the others. Thereafter, the robots are required to move, as a group, a certain distance away from their initial position. The robots are not required to adopt any particular formation, only to remaining within sensor range of one another, and to avoid collisions. During evolution, robots are evaluated on their ability to move the group centroid one metre within the space of three simulated minutes. However, our expectation was that a team capable of this would be able to sustain formation movement of much longer periods. The robots are not required to adopt any particular formation, only to remaining within sensor range of one another, and to avoid collisions. Since the robots start from initial random configurations, we anticipate that successful completion of the task will entail two phases. The first entailing the team organising itself into a formation, and the second entailing the team moving whilst maintaining that formation.

From the characterisation of the robots' sensors in the previous section, it should be clear that these impose significant constraints. They provide very little direct information about a robot's surroundings. Any given set of sensor input can be the result of any one of large number of significantly different circumstances. Furthermore, outside the limited range of their IR sensors, rbots have no indication of each other's position. Any robot straying more than two body-lengths from its teammates will cease to have any indication of their location. Of course, a robot controller may employ strategies to overcome some of the limitations of it sensors. For example, additional information can be gained by strategies which combine sensing and moving, and the integration of sensor input over time. However, it should be clear that the team's situation contrasts strongly with previous work in which robots utilised shared coordinate systems and global communication. It is worth noting that biological models of self-organising coordinated movement assume typically that agents are presented with significantly more information about their local environment than these robots have. For example, in models of flocking and shoaling, agents are typically assumed to have ideal sensors which provide the location, velocity and orientation of their nearest neighbours (see (Camazine *et al.* 2001) for an extensive review of such biological models; see also (Ward, Gobot, & Kendal 2001) for a recent evolutionary simulation model)

The team are also constrained by their homogeneity for the reasons discussed in the introduction. The team will move from their initial random configuration into the formation in which they will maintain whilst moving. In so doing it seems inevitable that different robots will be required to adopt different roles (for example, a leader and two followers). The robots must find some way of appropriately allocating and maintaining these roles despite the lack of any intrinsic differences between them. This is, of course, made all the more challenging by the poverty of the robots' sensory input.

## Implementation

### Evaluating Team Performance

A single genotype was used to generate a team by 'cloning' (i.e. decoding the genotype and then making copies of the resulting controller). Given that different starting positions will present different challenges, it was important that each team (i.e. each evolutionary individual) is evaluated under the same set of initial conditions. To this end, at each generation of the evolutionary algorithm, a set of 75 starting positions was randomly generated, as detailed in figure 4, and used for the evaluation of all the teams. Each evaluation involved multiple trials from the different starting positions. Fitness was measured as the average score over all the trials in the evaluation set.

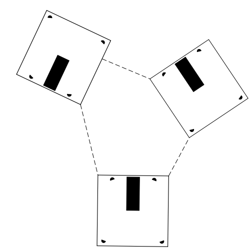

Figure 4: **An example starting positions:** Each robot's orientation is set randomly in the range $[0 : 2\pi]$, and the minimum distance between the edges of each robot and its nearest neighbour is set randomly in the range [10cm:22cm].

Reflecting the task description, the evaluation function seeks to assess the ability of the team to increase its distance from its starting position, whilst avoiding collisions and staying within sensor range. It therefore consists of three main components. First, at each timestep of the trial, the group is rewarded for any gains in distance. Second, this reward is multiplied by a dispersal scalar, reducing the fitness increment when one or more robots are outside of IR sensor range. Third, at the end of a trial, the group's accumulated score is reduced in proportional to the number of collisions which have occurred during that trial. (The maximum number of allowed collisions per trial was 20, if this number was reached, the trial was ended). More specifically, a team's trial score, is:

$$P.\Big(\sum_{t=1}^{T}\Big[f\big(d_t,\, D_{t-1}\big)\,.\big(1\,+\,\tanh\big(s_t/20.0\big)\big)\Big]\Big)$$

Here $P$ is a collision-penalty scalar in the range $[0.5 : 1]$, such that, if $c$ is the number of collisions between robots, and $c_{\max}$ is the maximum number of collisions allowed, then $P = 1 - c/2c_{\max}$. The distance gain component is given by the function $f$. This measures any gain that the team have made on their previous best distance from their initial location. Here a team's location is taken to be the centre-point (or centroid) of the group. If $d_t$ is the Euclidean distance between the group's location at time $t$ and its location at time $t = 0$, $D_{t-1}$ is the largest value that $d_t$ has attained prior to time $t$, and $D_{\max}$ is the required distance (i.e. 100 c.m.), then the function $f$ is defined as:

$$f\left(d_t, D_{t-1}\right) = \begin{cases} d_t - D_{t-1} & \text{if } D_{t-1} < d_t < D_{\max} \\ 0 & \text{otherwise} \end{cases}$$

The final component of a team's trial score, the scalar $s_t$, is a measure of the team's dispersal beyond sensor range at time $t$. If each robot is within sensor range of at least one other, then $s_t = 0$. Otherwise, the two shortest lines that can connect all three robots are found, and $s_t$ is the distance by which the longest of these exceeds sensor range. In other words, the team is penalised for its most wayward member. Note that $s_t$ used in combination with a **tanh** function, this ensures that as the robots begin to disperse, the team's score increment falls away sharply, the gradient of the tanh curve falls off as the distance between the robots increases, ensuring that increases in distance will still receive some minimal reward, even when the robots are far apart.

## Simulation

Controllers were initially evolved in simulation, before being transferred to the real robots. A big problem with evolving in simulation is that robots may become adapted to inaccurate features of the simulation, not present in the real world (Brooks 1992). However, building completely accurate simulation model of the robots and their interactions would be an onerous, and potentially impossible task, moreover, it would be unlikely that such a simulation would have significant speed advantages over evolving in the real world (Matarić & Cliff 1996). To avoid this problem we employed Jakobi's minimal simulation methodology(Jakobi 1997; 1998b). This enabled us to build a relatively crude, fast-running simulation model of the robots and their interactions, based on a relatively small set of measurements. The parameters of this model were systematically varied, within certain ranges, between each evaluation of a team. Parameters included, for example, the orientation of each robots' sensors, the manner in which a robot's position was affected by motor output, and the effects of IR interference. Whilst it was generally either difficult or time-consuming to measure parameters needed for the simulation with great accuracy on the robots, it was relatively easy to specify a range within which each of the parameters lay, even if that range was wide. Varying parameters within these ranges meant that a robot of capable of adapting to the simulation would be adapted to a wide range of possible robot-environment dynamics, including those of the real world. In addition to compensating for inaccuracies in our measurements, variation was used in the same way to compensate for inaccuracies in our modelling, since we were able to estimate the error due to these inaccuracies and adjust parameter ranges to compensate. More importantly, this approach allowed us to sacrifice accuracy for speed and employ cheap, inaccurate modelling where more accurate modelling would have incurred significant computational costs. Space precludes a description of our implementation of this minimal simulation, but full details are available elsewhere (Quinn *et al.* 2002).

## Neural Networks

The robots were controlled by artificial neural networks. Since it was unclear how the task would be solved, we could estimate little about the type of network architecture that would be needed to support the required behaviour. Thus we attempted to place as much of the design of the networks as possible under evolutionary control—specifically, the thresholds, weights and decay parameters, and the size and connectivity of the networks. Each neural network comprised 4 sensor input nodes, 4 motor output nodes, and some number of artificial neurons. These were connected together by some number of directional, excitatory and inhibitory weighted links. The network has no explicit layers, so any neuron may connect to any others, including itself; and may also connect to any of the sensory or motor nodes.

The neurons we use are loosely based on model spiking neurons (see Gerstner and Kistler, (2002) for a comprehensive review of such models). At any time-step, the output, $O_t$, a neuron is given by:

$$O_t = \begin{cases} 1 & \text{if } m_t \geq T \\ 0 & \text{if } m_t < T \end{cases}$$

where $T$ is the neuron's threshold. Here $m_t$ is analogous to membrane potential in a real neuron; it is a function of a neuron's weighted, summed input(s) integrated over time, such that:

$$m_t = \begin{cases} (\gamma_A)m_{t-1} + \sum_{n=0}^{N} w_n i_n & \text{if} \quad O_{t-1} = 0 \\ (\gamma_B)m_{t-1} + \sum_{n=0}^{N} w_n i_n & \text{if} \quad O_{t-1} = 1 \end{cases}$$

where $\gamma_A$ and $\gamma_B$ are decay constants, and $w_n$ designates the weight of the connection from the $n^{th}$ input ($i_n$) that scales that input. $\gamma_A$ and $\gamma_B$ are constrained to the range [0:1], the values of weights and thresholds are unconstrained. For certain parameter settings this neuron will behave like a simple spiking neuron, accumulating membrane potential, firing and then discharging (i.e., with $\gamma_A > \gamma_B$ and $T > 0$). However, the neurons also exhibit a range of other interesting temporal dynamics under different settings.

Each sensor input node outputs a real value in the range [0.0:1.0], which is simple linear scaling of the reading taken from its associated sensor. Motor outputs consist of a 'forward' and a 'reverse' node for each motor. The output, $M_{\text{out}}$, of each motor nodes is a simple threshold function of its summed weighted inputs:

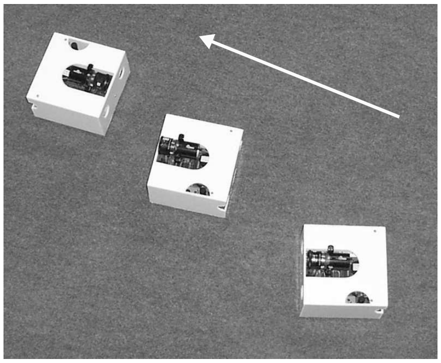

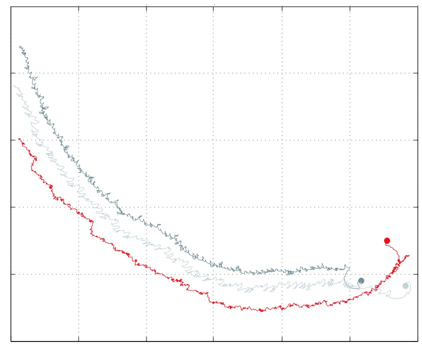

Figure 5: **Left:** Video still of the team travelling in formation. **Right** An example of team trajectory, tracing the position of each robot over a 5 minute period. Grid divisions are at 50cm intervals, robots' initial positions (bottom right) indicated by dots. Data generated in simulation.

$$M_{\text{out}} = \begin{cases} 1 & \text{if} \sum_{n=0}^{N} w_n i_n > 0 \\ 0 & \text{if} \sum_{n=0}^{N} w_n i_n \leq 0 \end{cases}$$

The final output of the each of the two motors is attained by subtracting its reverse node output from its forward node output. This gives three possible values for each motor output: {-1,0,1}.

The network is encoded by a topological encoding scheme, described in (Quinn *et al.* 2002), which is designed to enable the size and connectivity of the network to be placed under evolutionary control with only minimal constraints on network structure. Through macro-mutation operators, described in following section, neurons and connections can be added to or removed from the network, and existing connections can become reconnected.

## The Evolutionary Machinery

A simple, generational, evolutionary algorithm (EA) was employed for these experiments. An evolutionary population contained 50 genotypes. In the initial population, each genotype encoded a randomly generated network with three neurons and an average of 6 connections per gene; weights and thresholds were initially set in the range [-5:5] but were thereafter not constrained. At the end of each generation (i.e. after all individuals had been evaluated), genotypes were ranked by score, the 10 lowest scoring individuals were discarded and the remainder used to generate a new population. The two highest scoring individuals ('the elite') were copied unchanged in the new population, thereafter, genotypes were selected randomly with a probability inversely proportional to their rank. 60% of new genotypes were produced by recombination, and mutation operators were applied to all genotypes except the elite.

Genotypes were subject to both macro-mutation (i.e. structural changes) and micro-mutation (i.e. perturbation of real-valued parameters). Micro-mutation entailed that a random Gaussian offset was applied, with a small probability, to all real-valued parameters encoded in the genotype, such that the expected number of micro-mutations per genotype was 2. The mean of the Gaussian was zero and its standard deviation was 0.33 of that parameter's range (in the case of decay parameters) or its initialisation range (in the case of weights and thresholds). Three types of macro-mutation were employed. Firstly, a new neuron could be added to, or a randomly chosen neuron deleted, from the encoded network. The probability of addition was 0.004, and of deletion was 0.01. Secondly a new connection could be added or a randomly chosen connection deleted with the respective probabilities of 0.02 and 0.04. Finally, a randomly chosen connection could be chosen and reconnected elsewhere, this occurred with a probability of with a probability of 0.04.

## Evolved Behaviour

To date, we have undertaken a total of ten evolutionary runs. Four of these were terminated at early stage because they seemed unpromising. The remaining six runs produced teams capable of a consistently high standard of success after being left to evolve for between two and five thousand generations. There were significant behavioural differences between the successful teams, and we have chosen to focus on a single team rather than attempt to summarise them all. In describing the behaviour of the team, we wish primarily to achieve two objectives. The first is to demonstrate that the robots' behaviour is indeed that of a team, in the sense in which the term was introduced at the beginning of this paper. The second is to illustrate the process by which these roles become allocated in the absence of any intrinsic differences between the robots.

Paper really is too static a format to demonstrate how well the team transferred from simulation to reality, a problem which is lamented in more detail elsewhere (Jakobi 1998b). We can only report that the behaviour observed in simulation was qualitatively reproduced in reality. In simulation, averaged over 5000 trials, this team achieve a mean score 99.7 (out of a possible 100). We have not completed nearly so many trials with real robots, however, we have conducted a sequence of 100 consecutive trials (with random starting positions) with the real robots, the team successfully completed all of them.

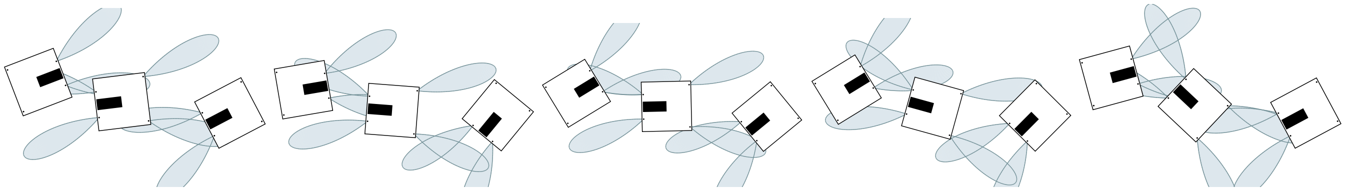

Figure 6: Time sequence illustrating relative positions during formation movement over a short (4 second) period. Robots maintain contact through direct sensing of each other's IR beams.

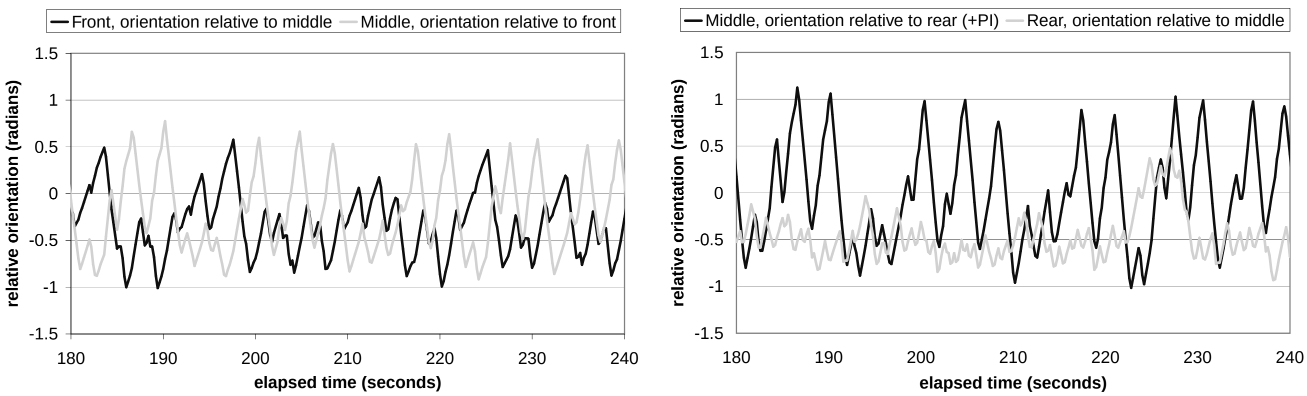


Figure 7: **Relative orientations of robots in formation over a 60 second period.** (Data taken from simulation.) *Left:* The movements of the front and middle robot are closely coordinated, with relative orientations predominantly in anti-phase. *Right:* The coordination of the middle and rear robot is much looser (For ease of presentation, the relative orientation of the middle robot has been offset by $\pi$) in the left-hand graph.

## Formation Movement

The team travel in a line formation, as can be seen from the video still in figure 5. The lead robot travels in reverse, whilst the middle and rear robot travel forwards. When travelling in formation, the team move at just over 1 cm/s, a relatively slow speed compared the 6 cm/s maximum speed that an individual robot is capable. The photograph fails to catch the dynamics of the team's movement which entails each robot swinging clockwise and counterclockwise whilst maintaining its position—watching the video footage sped up, team locomotion appears almost snakelike. The sequence of diagrams in figure 6 is an attempt to illustrates this aspect of the team's locomotion. Note from these diagrams, that the robots rely almost entirely on the direct perception of each other's IR beams (i.e. sensory interference) in order to coordinate their movement.

One illuminating way of illustrating relational movement patterns is changes in individual's orientation relative to the position of the individual with which it is interacting (see Moran *et al.* (1981), where this is used to great effect in their analysis of social interaction between wolves). (Relative orientation is an egocentric measure; the orientation of A relative to the position B is the angle between A's orientation and the line AB). Figure 7 shows the orientation of each robot relative to its neighbours during a period of formation movement. It illustrates the high degree of coordination between the front and middle robot, each responding closely to the other's movements. It also illustrates the much lower degree of coordination between the middle and rear robots, and the difference, with respect to the frequency of angular oscillation, between the movement of the rear robot and the leading pair. Despite the oscillating angular displacement of the robots, their formation is extremely robust. The formation is maintained indefinitely, despite robots only having been evolved for their ability to move the group centroid one metre.

## Roles

It should be clear from the above that robots perform the task we have set them. But are the robots actually operating as a team? In what follows we briefly show that each robot makes some necessary contribution to overall success and that these contributions are different and persist over time. To this end, we are interested in what each individual contributes to the maintenance of the formation and its continued movement. Perhaps the simplest way to assess individual contributions is simply by considering the effects of the removal of individual robot from the formation. To this end, we consider the effects of the removal of either the front or the rear robot (removal of middle robot is unilluminating, merely leaving the remaining two robots out of sensor range).

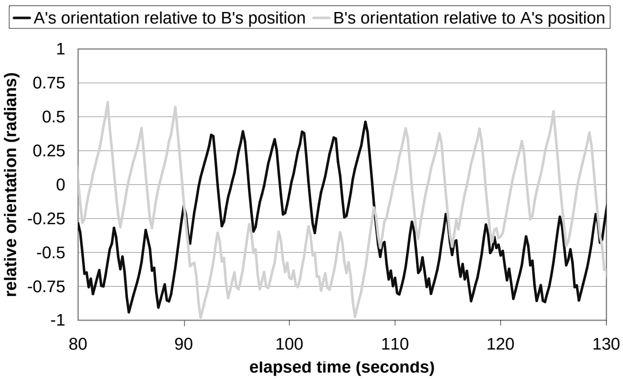


Figure 8: **Relative orientations of two robots, A and B, operating in the absence of a third robot**. Similarly to the front pair in a full formation, orientations are in anti-phase, although here the pattern is more regular. The configuration (and the pattern) is asymmetric, and maintained although robots periodically swapping positions within the configuration (seen at 90 and 110 seconds in the figure). Note that there is no significant displacement of the pairs' position.

If the rear robot is removed from the formation, the locomotion of the remaining pair ceases, there is no further significant displacement of their position. However, this is the only significant effect. The pair maintain the same configuration as when in full formation. Their cycle of angular oscillation relative to one another remains in anti-phase, although the pattern becomes more regular, as illustrated in figure 8. This is a dynamically stable configuration, tightly constrained by sensory feedback, which will persist indefinitely. If the rear robot is replaced, the group will move away once more. Now we consider the front robot. If this is removed from the full formation, the middle robot swings round toward the rear robot, and—after some interaction—the two robots form an opposed pair which maintain the same dynamically stable configuration as was just described.

From the above, we can say the following: Firstly, the rear robot has no significant effect on the other two robots' ability to maintain formation, but it is crucial to sustaining locomotion. Secondly, it is clear that the middle robot responds to the presence of the rear robot by moving forwards, since in the absence of the rear robot, the remaining pair cease to travel. For locomotion to continue, the configuration of the rear and middle robot must persist. That is, the middle robot must continue to sense the rear robot with its back sensors. Finally, in the absence of the front robot, the configuration adopted by the middle and rear robot in the formation is unstable.

This analysis is sufficient to show that these robots are working as team, concurrently performing separate but complementary roles which, in combination, result in coordinated formation movement. A more precise characterisation of each robot's contribution is difficult without presenting detailed analysis of the close sensorimotor coupling between the opposed front pair, and how this coupling is perturbed, but not completely disrupted, by the presence of the rear robot. Nevertheless, it is possible say something further about the team's organisation through investigating the effects of reorganising its formation. Firstly, when the middle robot is quickly picked up and rotated by 180 degrees, the formation is maintained and the team start to move in the opposite direction, with the robots which were previously front and rear adopting the roles appropriate to their new positions in the formation. Secondly, if the rear robot is removed from the formation and appropriately placed behind the front robot, formation again move off in the opposite direction, with each robot performing the role appropriate to its position. Thus, the fact that each robot remains in the same role within the formation is solely by virtue of the spatial organisation of the formation, rather than any long-term differences in internal state. This is not to say that the robots' behaviour is reactive. We know from analysis (not presented here) that the evolved networks rely heavily on temporal dynamics, such as short-term transient states. However, they do not rely on internal state to maintain their roles.

## Role Allocation

How are the roles initially allocated within the team? This is essentially to ask how the robots achieve the formation position from random initial positions, since as has already been noted, that the maintenance of individual roles is a function of the spatial organisation of the team formation. Any discussion of the initial interactions of the robots will be difficult without at least some information about how the robots responds to sensory input, so we will start by giving a very simplified explanation. In the absence of any sensory input, the robots move in a small clockwise forwards circle (the motor output is a cyclic pattern of left motor forward for 3 time-steps, followed by one times-step of right motor forward). A robot is generally 'attracted' to any source of front sensor input. It will rotate anticlockwise in response to any front left input and clockwise in response to front right input. Activation of either (or both) of the rear sensors in the absence of significant front sensor input causes the robot to turn more tightly in a clockwise direction (i.e. the fourth step of the basic motor pattern is removed). This is an incomplete description, but should be sufficient for the purposes of our explanation.

From its initial position, a robot will begin to circle clockwise until its senses another robot. Recall that a robot can sense both IR reflected off the body of another robot and the IR beam of another robot, the latter being perceptible from twice the distance than the former. For this reason a robot will typically first encounter either the front or rear IR beams of another robot (direct IR), or one of its side panels (reflected IR). A robot

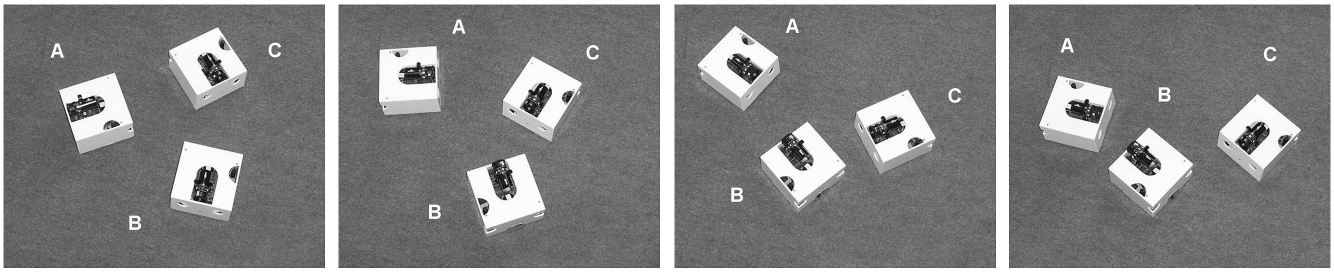


(i) (ii) (iii) (iv)

Figure 9: **An example of the team moving into the formation positions.** (i) The robot's initial positions. Initially, C is attracted B's rear sensors, causing B to turn tightly, A circles away, clockwise (ii) B and C begin to form a pair as A circles round towards them (iii) A disrupts the pair formation of B and C, subsequently pairing with B. (iv) C becomes attracted to B's rear sensors and begins to move into position. Shortly after this, the team achieve their final formation.

'attracted' to the side of another robot will simply be ignored as it cannot be sensed. A robot attracted to the rear IR beams of another robot will in turn activate that robot's rear sensors, causing it to turn sideways on. If however a robot becomes attracted to the front IR beams of another, it will in turn activate the front sensors of that robot as it approaches, both robots will turn to face each other—mutually attracted. The remaining robot will subsequently become attracted to rear sensors of one of the pair, bringing the formation into completion. Prior to the arrival of the third robot, the facing pair maintain the dynamic, stable configuration which was described in the previous section (illustrated in figure 8). In the present context, this serves as 'holding' pattern, in which the pair await arrival of the remaining team member.

The process of achieving formation is not always quite as simple as the above description might imply. The pairing process may have to be resolved between three robots (as for example, in panels ii and iii of figure 9) where one robot may disrupts the pair-forming of the other two. However, the explanation given above should be sufficient to inform the reader of the basic dynamics of the process of team formation. A process which can be seen as a one of progressive differentiation. The robots are initially undifferentiated with respect to their potential roles. The opposed pairing of two robots partially differentiates the team. The excluded robot's role is now determined—it will become the rear robot in the formation. Further differentiation occurs when the unpaired robot approaches the back sensors of one of the waiting pair, thereby determining the final two roles.

## Conclusion

The structured cooperation required for the performance of a team task presents interesting problems for a distributed control system. This is particularly true when individuals are homogeneous, and constrained to make use of limited local information. We have suggested that artificial evolution is a useful tool for automating the design of such systems, and presented an example of an evolved homogeneous multi-robot team. We have shown that the evolved system is capable of organising itself into a team formation, and maintaining this organisation over time.

It is worth noting the novelty of this work within the field of evolutionary robotics. To date, this research field has focussed almost exclusively on single robot systems. Insofar as we are aware, the work reported in this paper represents the first published example of cooperative and coordinated behaviour for a real multi-robot system designed by artificial evolution. By virtue of involving multiple robots, it is also one of the few examples of evolutionary robots research in which controllers must engage with a non-static environment (single-robot exceptions include (Jakobi 1998a; Smith 1998)).

Finally, we suggest that such a system would be extremely difficult to design by hand, given the sensory constraints and the close coupling of the individual robots. Of course, this not an easy claim to prove. However, we will conclude with a quote from someone with a great deal of experience in hand-designing multi-robot systems. Discussing the need for more complex sensors in the design of a following behaviour, Matarić (1995) comments: "If using only IRs, the agents cannot distinguish between other agents heading toward and away from them, and thus are unable to select whom to follow".

### Acknowledgements

We would like to thank Adrian Thompson for investing time and effort in robot modifications. Thanks also to Nick Jakobi and Kyran Dale for useful discussion. This work was funded by the B.N.S.C. Space Foresight project IMAR.

## References

Anderson, C., and Franks, N. 2001. Teams in animal societies. *Behavioural Ecology* 12(5):534–540.

Balch, T., and Arkin, R. 1998. Behavior-based formation control for multiagent robot teams. *IEEE Transactions on Robotics and Automation* 14(6):926–939.

Brooks, R. 1992. Artificial life and real robots. In Varela, F., and Bourgine, P., eds., *Toward a Practice of Autonomous Systems*. MIT Press. 3–10.

Camazine, S.; Denouberg, J.-L.; Franks, N.; Sneyd, J.; Theraulaz, G.; and Bonabeau, E. 2001. *Self-Organization in Biological Systems*. Princeton University Press.

Chaimowicz, L.; Sugar, T.; Kumar, V.; and Campos, M. 2001. An architecture for tightly coupled multi-robot cooperation. In *Proc. IEEE Intl. Conf. Robotics and Automation*, 2292–2297. Seoul, South Korea: IEEE Press.

Detrain, C., and Pasteels, J. 1992. Caste polyethism and collective defense in the ant *pheidole pallidula*. *Behavioural Ecology and Sociobiology* 29:405–412.

Floreano, D.; Nolfi, S.; and Mondada, F. 1998. Competitive co-evolutionary robotics: From theory to practice. In Pfeifer, R.; Blumberg, B.; Meyer, J.-A.; and Wilson, S., eds., *Proc. 4rd Intl. Conf. on Simulation of Adaptive Behavior*, 512–524. MIT Press.

Gerstner, W., and Kistler, W. 2002. *Spiking Neuron Models*. Cambridge University Press.

Hobbs, J.; Husbands, P.; and Harvey, I. 1996. Achieving improved mission robustness. In *Proc. 4th E.S.A. Workshop on Advanced Space Technologies for Robot Applications*.

Husbands, P., and Meyer, J.-A., eds. 1998. *Evolutionary Robotics: Proceedings of the First European Workshop, EvoRobot98*. Springer.

Jakobi, N. 1997. Half-baked, ad-hoc and noisy: Minimal simulation in evolutionary robotics. In Husbands, P., and Harvey, I., eds., *Fourth European Conference on Artificial Life*, 348–357. MIT Press/Bradford Books.

Jakobi, N. 1998a. Evolving motion-tracking behaviour for a panning camera head. In *Proc. 5th Intl Conf. Simulation of Adaptive Behaviour*. MIT Press.

Jakobi, N. 1998b. *Minimal Simulations for Evolutionary Robotics*. Ph.D. Dissertation, University of Sussex, U.K.

Matarić, M., and Cliff, D. 1996. Challenges in evolving controllers for physical robots. *Robots and Autonomous Systems* 19(1):67–83.

Matarić, M., and Sukhatme, S. 2001. Task-allocation and coordination of muliple robots for planetary exploration. In *Proc. 10th Intl Conf. Advanced Robotics*.

Matarić, M. 1995. Designing and understanding adaptive group behaviour. *Adaptive Behaviour* 4(1):51–80.

Meyer, J.-A.; Husbands, P.; and Harvey, I. 1998. Evolutionary robotics: A survey of applications and problems. In Husbands and Meyer (1998), 1–21.

Moran, G.; Fentress, J.; and Golani, I. 1981. A description of relational patterns of movemnet during 'ritualized fighting' in wolves. *Animal Behaviour* 29:1146–1165.

Nolfi, S., and Floreano, D. 1998. Co-evolving predator and prey robots: Do 'arm races' arise in artificial evolution? *Artificial Life* 4(4):311–335.

Nolfi, S., and Floreano, D. 2000. *Evolutionary Robotics: The Biology, Intelligence and Technology of Self-Organizing Machines*. MIT Press.

Nolfi, S. 1998. Evolutionary Robotics: Exploiting the full power of self-organization. *Connection Science* 10(3-4):167–183.

Quinn, M.; Smith, L.; Mayley, G.; and Husbands, P. 2002. Evolving formation movement for a homogeneous multi-robot system: Teamwork and role allocation with real robots. Cognitive Science Research Paper 515, University of Sussex, U.K.

Quinn, M. 2001. A comparison of approaches to the evolution of homogeneous multi-robot teams. In *Proc. Congr. Evolutionary Computation*, 128–135. Seoul, South Korea: IEEE Press.

Smith, T. 1998. Blurred vision: Simulation-reality transfer of a visually guided robot. In Husbands and Meyer (1998), 152–164.

Stone, P., and Veloso, M. 1999. Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork. *Artificial Intelligence* 110(2):241–273.

Ward, C.; Gobot, F.; and Kendal, G. 2001. Evolving collective behaviour in an artificial world. *Artificial Life* 7(2):191–210.

Watson, R. A.; Ficici, S. G.; and Pollack, J. B. 1999. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *Proc. Congr. Evolutionary Computation*, 335–342. Washington D.C., U.S.A.: IEEE Press.