

Synchronization and Gait Adaptation in Evolving Hexapod Robots

Mariagiovanna Mazzapioda, Stefano Nolfi

Institute of Cognitive Sciences and Technologies, National Research Council (CNR)
Via San Martino della Battaglia, 44 , Rome, Italy
mariagiovanna.mazzapioda@istc.cnr.it, stefano.nolfi@istc.cnr.it

Abstract. In this paper we present a distributed control architecture for a simulated hexapod robot with twelve degrees of freedom consisting of six homogeneous neural modules controlling the six corresponding legs that only have access to local sensory information and that coordinate by exchanging signals that diffuse in space like gaseous neuro-transmitters. The free parameters of the neural modules are evolved and are selected on the basis of the distance travelled by the robot. Obtained results indicate how the six neural controllers are able to coordinate so to produce an effective walking behaviour and to adapt on the fly by selecting the gait that is most appropriate to the current robot/environmental circumstances. The analysis of the evolved neural controllers indicates that the six neural controllers synchronize and converge on an appropriate gait on the basis of extremely simple control mechanisms and that the effects of the physical interaction with the environment are exploited to coordinate and to converge on a tripod or tetrapod gait on the basis of the current circumstances.

1 Introduction

In this paper we describe a method for developing the control system for a simulated hexapod robot with twelve degrees of freedom that has to exhibit a walking behaviour. The architecture proposed is fully distributed and consists of six identical neural modules in which each module is located in the corresponding leg and in which neural modules coordinate by producing and detecting signals that diffuse in space like gaseous neuro-transmitters.

This implies that as in other related models [1],[2] leg coordination does not arise from a centralized gait generator, but rather from the interactions between the neural modules controlling the corresponding legs. More precisely, leg coordination is mediated by the physical robot/environment interaction and by the signals produced by neural modules located nearby. However, contrary to the other models referenced above, each neural module influences and is influenced by the neural modules located nearby in the same way (i.e. the control system is constituted by a set of identical neural modules).

In more general terms, we assume that the control system of our robot is composed by a number of homogeneous neural modules that, each separately, exhibit a limit cycle (i.e. a periodic behaviour). Our problem, therefore, is that to define the rules that determine the conditions in which signals are produced and the way in which detected

signals affect nearby neural modules so that the resulting closed loop system exhibits a coordinated limit cycle behaviour that allow the robot to walk effectively (for a related approach see [3]). This problem have been attacked by using a self-organized technique based on artificial evolution [4] in which the free parameters of the neural modules are encoded in a population of evolving genotypes, and variations introduced through genetic operators are retained or discarded on the basis of the overall behaviour exhibited by the neural modules embodied in the robot and tested in the environment.

The goal of this paper is not to understand the biological basis of locomotion control in natural organisms but rather to build real-time walking machines. In particular, we are interested in investigating whether robots that have a modular structure (i.e. that are constituted by repeated homologous body elements) can exhibit coherent and effective behaviour on the basis of modular control systems (i.e. on the basis of distributed control system in which repeated parts of the robots body are controlled by corresponding repeated control units). Progresses toward this objective, in fact, might have a significant impact on robotics with particular reference to self-reconfigurable robotics [5], [6] and evolutionary robotics techniques that allow to co-evolve and co-adapt the robots' control system and body structure [7], [8], [9], [10].

2 The experimental setup

In this section we describe the simulated hexapod robot used in the experiments, its control system, and the evolutionary algorithm used to set the free parameters of the robot's control system. The characteristics of the simulated robot are identical to that described in a previous work [11]. In this paper, however, we present an extended version of the control system and new experimental results that show, in particular, how evolved robots are able to adapt their gait on the fly on the basis of the current circumstances.

2.1 The hexapod robot

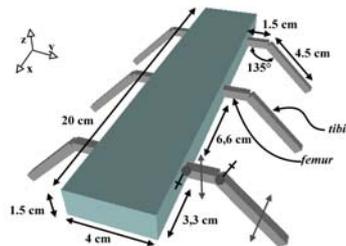


Fig. 1. The simulated hexapod robot. The grey circles shown on the bottom-right side of the picture indicate the position of the joints, while the grey arrows indicate their rotational axis.

The simulated robot (Fig. 1) consists of a main body (with a length of 20 cm, a width of 4 cm, and a height of 1.5 cm) and 6 legs.

Each leg consists of two segments (a “femur” and a “tibia” with a length of 1.5 and 4.5 cm respectively) and has two motors controlling two corresponding joints (the body-femur and the femur-tibia joints). The femur and the body-femur joint allow the robot to raise its central body from the ground and to move the tibia up and down. The body-femur joint is a motorized hinge joint with rotational axis parallel to the x -axis that can rotate from $-\pi/16$ to $+\pi/16$ rad. The femur-tibia joint allows it to move the tibia forward or backward. It is a motorized hinge joint that rotates from $-\pi/8$ to $+\pi/8$ rad with respect to its own axis (i.e. an axis rotated of $\pi/4$ rad with respect to yz -plane). The motors controlling the joints can apply a maximum torque of $0.03Nm$ at maximum speed of 3100 rpm in both directions. For each leg, two simulated position sensors detect the current angular position of the corresponding joint. The total weight of the simulated robot is 387g. Gravity force is -9.8 m/sec^2 . The environment consists of a flat surface. The robot and the robot/environment interaction were simulated by using the VortexTM toolkit (Critical Mass Labs, Canada), that allows to realistically simulate the dynamics and collisions of rigid bodies in 3D.

2.2 The control system

The robot is controlled by a distributed control system consisting of six homogeneous neural modules, located at the junction between the main body and the legs, that control the six corresponding legs (Fig. 2).

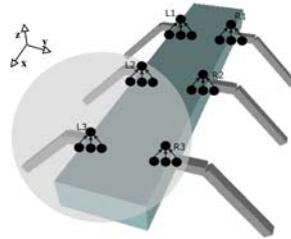


Fig. 2. The robot and its control system consisting of 6 neural modules. L1, L2, and L3 indicate the front, middle and rear leg located on the left side of the robot. R1, R2, and R3 indicate the front, middle and rear leg located on the right side of the robot. The grey circle represent a possible range of diffusion of the signal produced by one neural module (i.e. the neural module controlling the L3 leg)

The six neural modules are identical (i.e. have the same architecture and the same free parameters) and have access to local sensory information only. More specifically, each neural module has access to the current angular position and controls the frequency of oscillation of the two joints of the corresponding leg. Neural modules communicate between themselves by producing signals and by detecting the signals produced by other neural controllers located within a given Euclidean distance. Signals thus are similar to gaseous neuro-transmitters such as nitric oxide that are released by neurons and affect other neurons located nearby in a diffuse manner (see [12], [13], [14]). Signal transmission is instantaneous.

Each of the twelve motors neurons produces a sinusoidal oscillatory movement with a variable frequency of the corresponding joint, within the joint's limits. More specifically, the current desired position of a corresponding joint is computed according to the following equation:

$$\text{pos}(t) = \sin(V(t) \cdot t + \varphi) \quad (1)$$

where $\text{pos}(t)$ indicates the desired angular position of the joint at time t , $V(t)$ (that ranges between 7 and 14 Hz) indicates the current frequency of the oscillator, and φ indicates the starting position of the joint. The output of the neurons is normalized within the range of movement of the corresponding joint and is used to encode the desired position of the corresponding joint. More precisely, motors are activated so as to reach a speed proportional to the difference between the current and the desired position of the joint (maximum motor speed is 3100 rpm, maximum torque is 0.03Nm). We decided to use an high frequency range ([7, 14] Hz) and to update the state of the sensors and motors at an high rate (every 1.5 ms) to avoid instabilities arising from the calculation of the dynamic of the robot/environmental interaction and to reduce the time required to test individuals' behaviour in simulation.

Each neural module has six input neurons directly connected to six output neurons (Fig. 3).

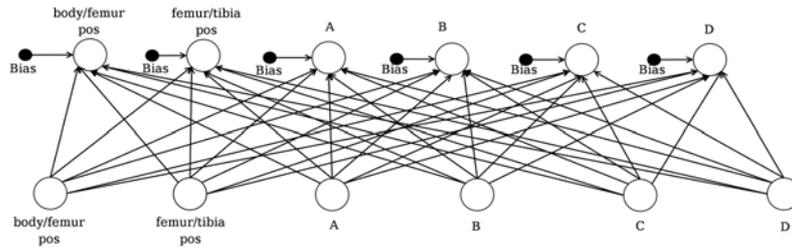


Fig. 3. The topology of each neural module. The six input neurons indicated in the bottom part of the picture encode the current angular position of the two joints of a leg and signal A, B, C and D (see text). The six output neurons are indicated in the top part of the picture. The first two modulate the frequency of oscillation of the two corresponding motorized joints and the others four determine whether or not the signal A, B, C and D are produced.

The input neurons encode the current angular positions of the two joints of the corresponding leg (normalized in the range [0.0, 1.0]) and whether signals, produced by other neural modules, are detected. Each neural module can produce four different signals (A, B, C and D) that diffuse and can be detected up to a certain distance (D_a , D_b , D_c and D_d , in the case of signal A, B, C and D, respectively). The intensity of the detected signal is linearly dependent from the distance of emitting neural module, and vary within 0 (when distance from emitting neuron is D) and 1.0 (when distance from emitting neuron is 0). Furthermore, detected signal is linearly proportional to the number of neural modules that are currently producing the corresponding signal located within the corresponding maximum diffusion distance.

The activation of output neurons is computed by using a standard logistic function. The first two output neurons determine how the frequency of oscillation of the two

corresponding joints varies. More specifically, each time step (i.e. each 1.5ms), the frequency of oscillation of a joint can vary by an amount whose range is [-1.4Hz, +1.4Hz] according to the following equation:

$$\begin{aligned}
 V(0) &= Val & 7 \leq Val \leq 14 \\
 V(t) &= V(t-1) + \begin{cases} (Out-0.75) \cdot 1.4 & Out \geq 0.75 \\ 0 & 0.25 < Out < 0.75 \\ (Out-0.25) \cdot 1.4 & Out \leq 0.25 \end{cases} & (2)
 \end{aligned}$$

Where Val indicates the initial value of frequency of a joint that is randomly set within the range, Out indicates the output of the corresponding motor neuron, and V(t) indicates the current frequency, V(t-1) indicates the frequency at the previous time step. Frequency is bounded in the range [7Hz, 14Hz], i.e. variations that exceed the limits are discarded. This means that each leg oscillates at a given frequency (within a range) and that each neural module can accelerate or decelerate the frequency of oscillation of the corresponding leg by a fixed amount each time step. The other four output neurons determine whether or not signal A, B, C and D are produced. More specifically, signal A, B, C and D are produced when the output of the corresponding output neuron exceeds the corresponding threshold (T_a , T_b , T_c and T_d in the case of signal A, B, C and D, respectively).

2.3 The evolutionary algorithm

The free parameters of the neural modules are evolved through an evolutionary algorithm. Robots were selected for the ability to walk along a straight direction as far as possible. Each robot was allowed to "live" for 5 trials, each lasting 3000 ms (i.e. 2000 time steps of 1.5 ms). The state of the sensor and motor neurons, the torque applied to the motors, and the dynamics of robot/environment interaction are updated each time step (i.e. each 1.5 ms). At the beginning of each trial: the main body of the robot is placed at a height of 3.68 cm with respect to the ground plane (i.e. the whole robot floats in the air at 0.5 cm from the ground). The initial position of the twelve joints and the initial desired velocity of each corresponding motor is set randomly within the corresponding range. The fitness of each robot is computed by measuring the Euclidean distance between the initial and final position of the centre of mass of the robot during each trial. The total fitness is computed by averaging the distance travelled during each trial.

The initial population consisted of 100 randomly generated genotypes that encoded the connection weights and the biases of a neural module, the maximum distance of diffusion of the four signals (D_a , D_b , D_c and D_d), and the thresholds that determine when signals are produced (T_a , T_b , T_c and T_d). Each parameter is encoded as real number. Connection weights and biases, diffusion distances of signals, and thresholds that determine signal emission are normalized within the following ranges: [-15.0, +15.0], [0.0, 10.0], [0, 1.0], respectively. Each genotype is translated into 6 identical

neural modules that are embodied in the robot and evaluated as described above. The 20 best genotypes of each generation were allowed to reproduce by generating five copies each, with 3% of their genotype value replaced with a new randomly selected value (within the corresponding range). The evolutionary process lasted 300 generations (i.e. the process of testing, selecting and reproducing robots is iterated 300 times). The experiment was replicated 10 times starting from different, randomly generated, genotypes.

3. Obtained results

By analysing the results of the evolutionary experiments we observed that evolved robots display an ability to walk effectively, in all replications of the experiment. In particular, evolved robots display an ability to quickly coordinate the phases and the frequency of oscillation of their twelve motorized joints by converging toward a tripod gait independently from the initial position of the joints (see Fig. 4).

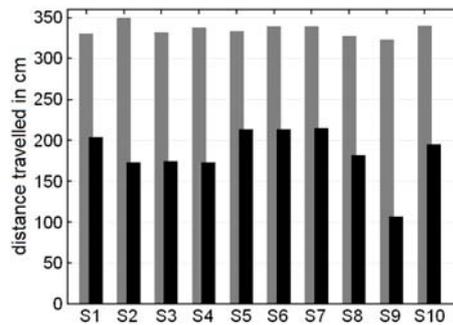


Fig. 4. Average distance travelled by the best robot of each replication in a normal and in a test condition (grey and black histograms respectively) in which the robot is loaded with an additional weight corresponding to 1.5 times the robot's body weight. Average results for 100 trials each lasting 3sec. The robots of all replications display a tripod gait when tested in a normal condition. In the test condition, the robots of replication S1, S5, S6, S7, and S10 display a tetrapod gait. The robots of the other replications, instead, by not being able to select an appropriate gait when loaded with additional weight, display lower performance in this condition.

Surprisingly, we observed that evolved robots generalize their ability to walk in situation in which they have to carry a weight equal to 1.5 of robots' body weight (see Fig. 4). Interestingly, in some of the replications, evolved robots converge on a tripod gait (when they are not loaded with additional weight) and on tetrapod gait (when they are loaded with the additional weight, see Fig. 5).

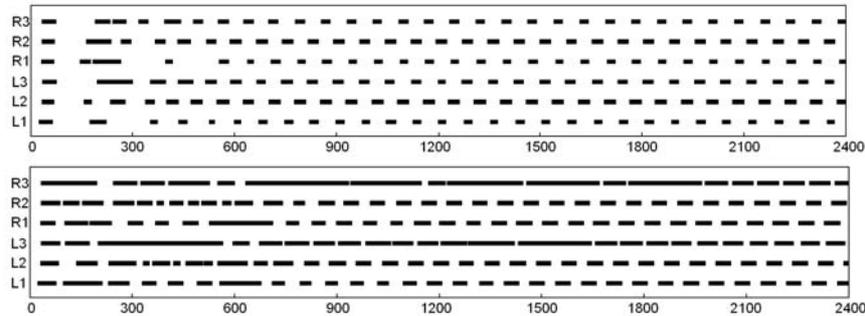


Fig. 5. A typical behaviour exhibited by an evolved robot of one of the best replications during two trials in which the robot is tested in a normal condition or in a test condition in which it is loaded with an additional weight (top and bottom figures, respectively). At the beginning of the trial the position of the joints and frequency of oscillation are randomly initialised within limits. The black lines indicate the phases in which the tibia of the corresponding leg touch the ground. Legs are labelled with L for left and R for right and numbered from 1 to 3 starting from the front of the insect. The horizontal axis indicates time in milliseconds. Gaits remain stable after 2400 ms (results not shown for space reasons). Please notice that these pictures do not indicate the trajectories of the robot in space but only the phase in time during which the legs touch the ground.

This implies that evolved robots, as real insects, select a tripod gait in normal conditions and a tetrapod gait when they are loaded with additional weight.

The ability to converge on a tripod or a tetrapod gait in the two circumstances play a functional role since the tripod gait is faster when the weight of the robot is not too high but is ineffective when the robot is loaded with additional weight. Indeed, as shown in Fig.4, the robots that are not able to switch to a tetrapod gait when they are loaded with additional weight display significantly worse performance. This can be explained by considering that in the tripod and in the tetrapod gait robots are supported by at least three or four legs, respectively. In the tetrapod gait therefore, the robot can exploit the power produced by four legs rather than three legs at the same time. In the tripod gait, in fact, the front and rear leg of one side and the middle leg of the other side perform they swing movement at the same time and the three other legs are in anti-phase. In the tetrapod gait, instead, a “wave” of swing movements passes along the body from rear to front.

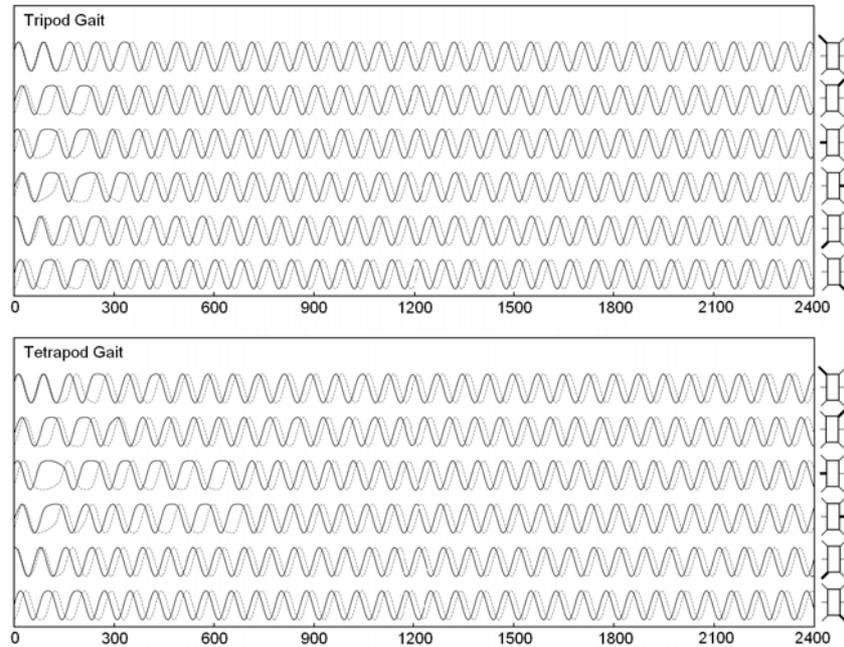


Fig. 6. Desired angular position of the twelve joints during the same trials shown in Figure 5 (top: data for the test in a normal condition, bottom: data for the test with additional weight). Each line indicates the desired angular position of the joints of the leg indicated with a dark line in the right part of the Figure. Full lines and dotted lines indicate the position of the body-femur and femur-tibia joints, respectively. High values indicate positions in which the femur is elevated with respect to the main body and positions in which the tibia is oriented toward the front of the robot.

The dynamical behaviour produced by the walking robots does not only result from the interaction between the six neural modules that control the six corresponding legs but also from the dynamics originating from the interaction between the robot body and the environment. Indeed, the way in which the actual position of the joints varies in time (Fig. 7) is influenced not only from the variation of the desired joint position (Fig. 6) but also from the forces arising from the collision between the legs and the ground. These forces are influenced by several factors such as the actual orientation of the robot with respect to the ground, the total weight of the robot, the current velocity of the robot, the characteristics of the ground, etc.

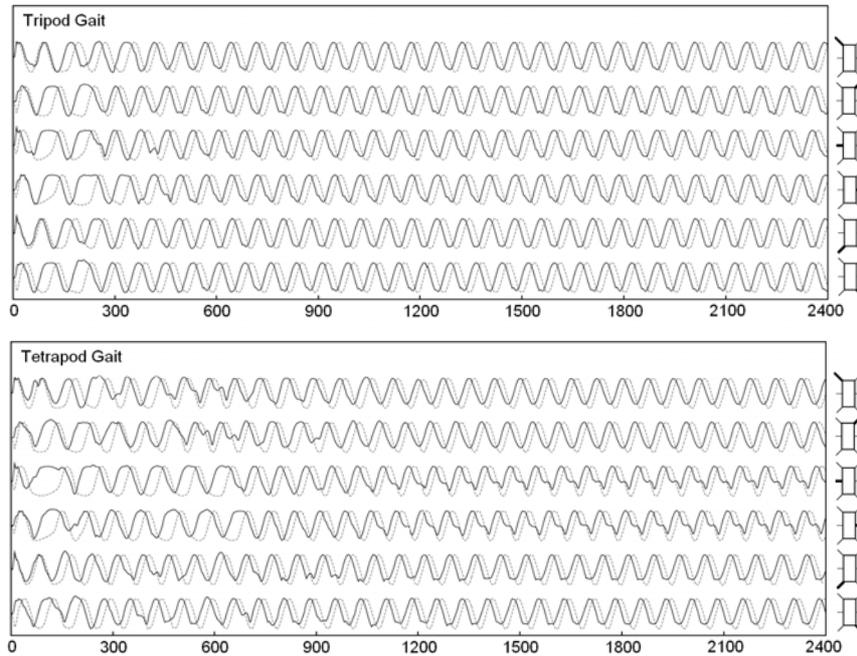


Fig. 7. Actual angular position of the twelve joints during the same trials shown in Fig. 5 and 6 (top: data for the test in a normal condition, bottom: data for the test with additional weight). Each line indicates the actual position of the joints of the leg indicated with a dark line in the right part of the Figure. Full lines and dotted lines indicate the position of the body-femur and femur-tibia joints, respectively. High values indicate positions in which the femur is elevated with respect to the main body and positions in which the tibia is oriented toward the front of the robot.

The results described above refer to robots that have been evolved in a normal condition (in which they were never loaded with additional weight) and have been tested in a normal and over-weight condition (in which in half of the trials they were loaded with additional weight and in half of the trials they were not). By evolving and testing the robots in a normal and over-weighted condition we observed that evolved robots displayed lower performance on the average (with a loss of about 20% with respect to the results shown above). Moreover, in this new experiment evolved robots always displayed tetrapod gaits and were not able to adapt their gait on the fly so to select a tripod gait when tested in normal conditions. This results can be explained by considering that in the latter experiment evolving robot converge on a local minima, i.e. a simple solution that allow them to reach good but sub-optimal performance on the basis of simple control mechanisms. This hypothesis is also supported by the analysis reported in the next sections that suggest how tetrapod gait tend to easily emerge as a result of the effects of the collisions with the ground without necessarily requiring control mechanisms that allow the legs to effectively coordinate and synchronize through signals.

In another replication of the experiment we verified that the role of space in modulating the effect of signals was really necessary to achieve effective results. More precisely, by running a replication of the experiment in which the signals produced by each module affected all other modules in the same way (independently from the distance between modules) we observed that evolved robots displayed significantly lower performance and were never able to converge on stable gaits.

4. Analysis of the mechanisms that lead to leg coordination and gait selection

To understand the mechanisms that lead to the synchronization of the twelve joints, we analysed the behaviour of each neural module and the interaction between different neural modules mediated by signals (i.e. the conditions in which signals are produced and the effects of signals detected). Here we report the analysis conducted in the case of the evolved individuals already described in Figure 5-7. The analysis conducted on individuals of other replications showed qualitatively similar, although in some case slightly more complex, strategies. We will first describe how the six legs converge toward a tripod gait in the normal condition and then how they converge on a tetrapod gait when the robot is loaded with additional weight.

As could be expected, the synchronization between the two joints of each leg is achieved within each single neural controller. More specifically: (a) the body-femur joint decelerates when it is elevated and the tibia is oriented toward the rear, and (b) the femur-tibia joint decelerates when the body-femur joint is lowered and the tibia is oriented toward the rear of the robot. Deceleration results both as the effect of output of the corresponding neural module and as a result of the effect of collisions. The combination of these two mechanisms leads to a stable state, that correspond to stable phase observed after coordination, in which the protraction movement of the tibia is performed when the body-femur joint is elevated the retraction movement is performed when the body-femur joint is lowered.

4.1 Analysis of the mechanisms that lead to a tripod gait

Although neural modules can produce and detect up to four different signals, the evolved individual shown in Fig. 5-7 only produces one of the four signals: signal B. In the other replications of the experiment, evolved robots use 1 or 2 signals. Interesting however, significantly lower performance have been observed in other experiments in which neural modules were allowed to produce and detect only two signals (result not shown). This suggest that the possibility of using many signals plays a crucial role during the first evolutionary phases despite only 1-2 signals are exploited by evolved individuals.

Since the maximum distance of diffusion of signals is 7.81 cm, in the case of the robot shown in Fig. 5-7, the signal produced by each leg affects the contra-lateral leg of the same segment, the previous and succeeding legs of the same segment, and the previous and succeeding legs of the contra-lateral segment (when present). Since the

amount of the signal detected is proportional to the distance (4.0cm, 6.6cm, and 7.71cm respectively) the impact of produced signal is larger on the contra-lateral leg of the same segment, smaller but still significant on the previous and succeeding legs of the same segment, and almost negligible on the previous and succeeding legs of the contra-lateral segment. If we ignore the negligible effect on previous and succeeding legs of the contra-lateral segment, this means that the signal produced by a leg of one group ([L1,L3,R2] or [R1,R3,L2]) affects only the legs of the other group that should be in anti-phase in a tripod gait. The legs that are affected by a signal are 2 out of 3 legs in the case of legs [L1,L3,R1,R3] and 3 out of 3 legs in the case of legs [L2,R2].

To explain how the six legs coordinate we should explain why uncoordinated states are unstable and lead to coordinated phases (through relative acceleration/deceleration of the joints) and why coordinated states are stable.

The latter aspect can be explained by considering that, during coordinated phases, legs belonging to the two groups (A and B) are in phase within the group and in anti-phase between groups. Legs of group A produces a signal when their tibia are oriented toward the rear and their femurs are elevated (i.e. at the starting of the protraction phase) and they reduce their velocity only when their tibia are oriented toward the rear and they detect a signal produced by legs of group B. Since the legs of the two groups are in anti-phase and the signals are produced in an alternate way by the legs of the two groups, in coordinated phases signals do not produce acceleration/deceleration effects.

To explain the former aspect (i.e. why uncoordinated phases are instable) let us consider the case in which the legs of groups A and B, during a retraction phase, have both their tibia oriented toward the rear but the legs of group A are slightly advanced with respect to the legs of group B. Since the interval in which legs emit the signal is larger than the interval in which the legs decelerate when they detect a signal, the deceleration effect of the legs of the group A on the legs of the group B is longer than viceversa. This implies that phase distance between the legs of the two groups tend to increase when they are in phase or almost in phase until the legs of the two groups reach the stable state described in the previous paragraph.

4.2 Analysis of the mechanisms that lead to a tetrapod gait

To understand the mechanisms that allow robots to switch from a tripod to a tetrapod gait when loaded with additional weight we should consider that the additional weight increases the intensity of friction in particular on the femur joints that are no longer able to reach their extreme posterior position (Fig. 7, bottom). Since legs emit the signal when the femur is elevated, this implies that the overall speed of the legs tend to decrease due to the fact that signals are produced for a longer time period.

Another factor to be considered is that when the body of the robot is not perfectly aligned with respect to the ground plane, different legs are subjected to stronger or weaker friction forces. For example, when the body of the robot is inclined toward the front, with respect to the rostro-caudal axis, the frontal legs are subjected to a stronger friction and, as a consequence, these legs produce signals for a longer time period.

The legs that are particularly stressed (due to current inclination of the body and to the fact that are performing a retraction movement) tend to slow down nearby legs

that are also performing a retraction movement. This implies that legs that are under stress tend to recruit to their phase nearby legs. The final result is that the legs of the same side and of adjacent segments do not move in perfect anti-phase, as in the tripod gait, but with a partially overlapping phase.

5. Conclusions

In this paper we present a distributed control architecture for a simulated hexapod robot with twelve degrees of freedom consisting of six homogeneous neural modules controlling the six corresponding legs that only have access to local sensory information and that interact by producing and detecting signals that diffuse in space.

The free parameters of the homogeneous neural modules, that regulate the frequency of oscillation of the corresponding leg and the signals that are emitted on the basis of the current position of the leg and of the signals detected have been set through an evolutionary method in which variation of the free parameters are retained or discarded on the basis of the global behavior exhibited by the robot in the environment. This method allows evolving robots to select solutions that exploit properties emerging from the interaction between the neural modules and between the robot and the environment.

The analysis of the evolved neural controllers indicates that the six homogeneous neural controllers converge on an appropriate gait on the basis of extremely simple control mechanisms. In some of the replications, in particular, coordination and gait selection is achieved on the basis of a single signal. This implies that a single rule (that accelerates or decelerates the frequency of oscillation of nearby legs depending on the state of the leg that detect the signal) is sufficient to converge on an a stable and effective gait.

Finally, we observed that evolved robots generalize their ability to produce an effective walking behaviour also when they are loaded with additional weight by displaying an ability to select a tripod or a tetrapod gait in the normal condition and in test conditions in which they are loaded with an additional weight, respectively. Overall, the obtained results suggest that an hexapod robot can be controlled on the basis of fully homogeneous distributed control system in which the interaction between neural modules is only regulated by the distribution in space.

References

1. Gallagher, J., Beer, D.R., Espenschied, K., Quinn, R.D.: Application of evolved locomotion controllers to a hexapod robot. *Robotics and Autonomous Systems*. (1996) 19: 95-103.
2. Ijspeert, A.J., Crespi, A., and Cabelguen, J.M.: Simulation and robotics studies of salamander locomotion. Applying neurobiological principles to the control of locomotion in robots. *Neuroinformatics*. (2005) 3(3):171-196.
3. Hulse, M., Wischmann, S., Pasemann F.: Structure and function of evolved neurocontrollers for autonomous robots. *Connection Science*. (2004) 16 (4): 249-266.

4. Nolfi, S., Floreano, D.: *Evolutionary Robotics: The Biology, Intelligence, and Technology of Self-Organizing Machines*. Cambridge, MA: MIT Press/Bradford Books (2000)
5. Kamimura, A., Murata, S., Yoshida, E., Kurokawa H., Tomita K., Kokaji S.: Self-reconfigurable modular robot - experiments on reconfiguration and locomotion. In T. J. Tarn et al. (Eds), *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*. (2001) New York: IEEE Press.
6. Yim, M., Zhang, Y., Duff, D.: Modular robots. *IEEE Spectrum*. (2002) 30-34.
7. Sims, K.: Evolving 3D morphology and behavior by competition. In R. Brooks & P. Maes (Eds), *Proceedings of Fourth Conference on Artificial Life*. (1994) Cambridge, MA:MIT Press.
8. Bongard, J., Pfeifer, R.: Evolving complete agents using artificial ontogeny. In Hara F. & R. Pfeifer (Eds), *iMorpho-functional Machines: The New Species: Designing Embodied Intelligence*. Berlin: Springer Verlag (2003)
9. Hornby, G. S., Pollack, J. B.: Creating high-level components with a generative representation for body-brain evolution. *Artificial Life*. (2002) 8(3):223-246
10. Bianco R., Nolfi S. (2004). Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. *Connection Science*, 4: 227-248.
11. Mazzapioda, M., Nolfi, S.: Synchronization within Homogeneous Neural Modules Controlling a Simulated Hexapod Robot. *Proceedings in AlifeX*. (2006) In press.
12. Elphick, M.R., Kemenes, G., Staras, K., O'Shea, M.: Behavioural role for nitric oxide in chemosensory activation of feeding in a mollusc. *Journal of Neuroscience* (1995) 15(11):7653-7664
13. Elphick, M.R., Williams L., O'Shea M.: New Features of the locust optic lobe:evidence of a role for nitric oxide in insect vision. *Journal of Experimental Biology* (1996) 199: 2395-2407
14. Husbands, P., Philippides A., Smith T.M.C., O'Shea M.: Volume Signalling in Real and Robot Nervous Systems. *Theory in Biosciences* (2001), 120: 253-269