Evolving Embodied Genetic Regulatory Network-driven Control Systems

Tom Quick[1], Chrystopher L. Nehaniv[2], Kerstin Dautenhahn[2], and Graham Roberts[1]

[1] Department of Computer Science
University College London
Gower Street
London WC1E 6BT, U.K.
[2] Adaptive Systems Research Group
University of Hertfordshire
Hatfield, Herts AL10 9AB, U.K.

{t.quick,g.roberts}@cs.ucl.ac.uk,{c.l.nehaniv,k.dautenhahn}@herts.ac.uk

Abstract. We demonstrate the evolution of simple embodied Genetic Regulatory Networks (GRNs) as real-time control systems for robotic and software-based embodied Artificial Organisms, and present results from two experimental test-beds: homeostatic temperature regulation in an abstract software environment, and phototactic robot behaviour maximising exposure to light. The GRN controllers are continually coupled to the organisms' environments throughout their lifetimes, and constitute the primary basis for the organisms' behaviour from moment to moment. The environment in which the organisms are embodied is shown to play a significant role in the dynamics of the GRNs, and the behaviour of the organisms.

1 Introduction: Genetic Regulatory Networks and Environmental Coupling

In this paper we present initial results from a novel experimental system, *Biosys*, in which Genetic Regulatory Networks (GRNs) play a critical and active role throughout the lifetimes of evolved artificial organisms, continually driving the interplay between organism and environment, giving rise to coherent observable emergent behaviours.

The role of GRNs in the development of living organisms is well-known [1, 2]. Moreover, they play a key role in cellular metabolism throughout organisms' lifetimes [3]. For both prokaryotes and eukaryotes, there is continual, tightly woven interaction between an organism's genes, the protein machinery of the cell or cells in which the genes exist, and the organism's environment. Our model can be applied to achieving a basis for such control in artificial organisms. In this work we focus on the evolution of heritable regulatory mechanisms based on model GRNs in populations of artificial organisms, and show how these mechanisms can drive behaviour via continual coupling to an environment.

In Artificial Life (ALife) research, GRNs typically appear in one of two broad research contexts. The first is to study them as dynamical systems in isolation from environmental stimulus, for example, identifying basins of attraction and other dynamical systems features [4–6]. The second is to recognise them as an emulatable feature of biological developmental processes, contributing to an at least partially realistic mapping between genotype and phenotype. This approach is outlined by Dellaert and Beer [7], and has been utilised to grow neural networks [8], as well as body plans and nervous systems for multi-cellular artificial agents [9]. However, a common feature of these models that stands in stark contrast to biological organisms, and which we seek to address, is that once a phenotype has been produced from a genotype, the genome becomes totally redundant, having no further role to play in the structure or behaviour of the resultant artificial organism.

A similar developmental approach, though with greater attention to biological detail is employed by Eggenberger [10] and Kumar [11], who focus on growth and morphogenesis rather than agent control. In this context, issues regarding the post-development role of the genome are not directly relevant; however even during development, the role of the environment is minimal in these models, in contrast to the 'eco-devo' approach to development [12, 13].

Artificial chemistry-based control systems modelled on cellular signal transduction have been developed that do exploit moment to moment interaction between an artificial organism and an environment as a basis for producing emergent behaviours [14]. However, genes only factor as part of the Evolutionary Algorithms used to produce the chemistries involved. Cellular components such as proteins and regulatory networks do not feature at all.

In contrast, real-time GRN-driven control does feature in wet-ware synthetic gene network engineering research [15]. Cross-fertilisation of principles and theory with ALife seems likely, particularly as the level of biological detail in ALife models increases to a level that makes the two programmes commensurate.

In section 2, we discuss our experimental methods, focusing on how *Biosys* works, how it is embodied, and describe the Genetic Algorithm (GA) used. We then present results from experiments evolving GRN-driven controllers embodied in two different environments, as follows. First, section 3 presents homeostatic behaviour in an abstract thermal environment, illustrating the workings of our model and some key principles. Second, section 4 presents phototactic, light-maximising robot behaviour in a simulated physical environment. Finally, in section 5 we present our conclusions.

2 Methodology: Embodiment, GRNs and evolution

2.1 Embodied systems

In previous papers we have defined 'embodiment' in terms of the conditions necessary for structural coupling to occur between a system (for example an organism or a software agent) and an environment:

> A system X is embodied in an environment E if perturbatory channels exist between the two. That is, X is embodied in E if for every time t at which both X and E exist, some subset of E's possible states have the capacity to perturb X's state, and some subset of X's possible states have the capacity to perturb E's state. [16]

We constructed the model presented here with this definition in mind. One of the roles played by proteins in our model is to provide the basis for cellular embodiment, by conveying 'bridging' input and output signals between the GRN and the environment (see fig. 1). Sensory input from the environment is manifest as an increase in protein concentrations in the cell, and effector output is determined by protein concentrations. Hence the GRN is *continually coupled* to the organism's environment, perturbing and being perturbed, acting as a real-time control system. The concept of behaviour emerging from such coupling is central to autopoietic theory [17], and is analysed from a dynamical systems perspective by Beer [18], in the context of coupling between a neural network and an environment.

2.2 The *Biosys* model

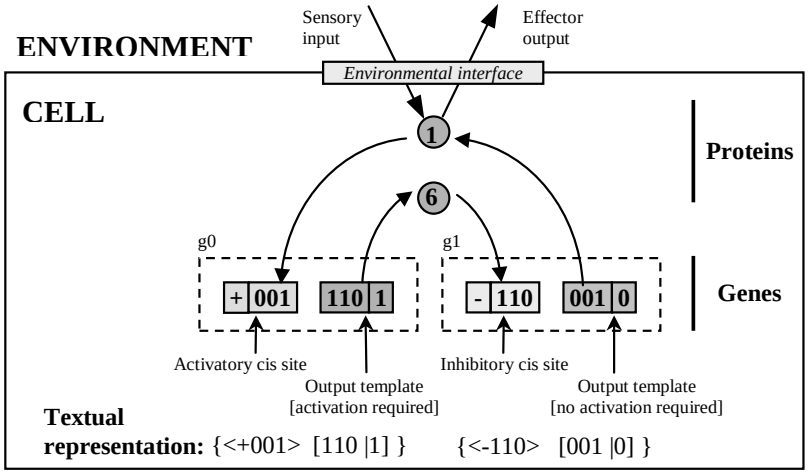

Fig. 1. Schematic gene–protein–environment interaction, with a textual representation of genes illustrated

In the simple single-celled non-spatial biological model currently used in *Biosys*, a cell consists of a genome and proteins (fig. 1). Genes produce proteins, which in turn both regulate the activity of other genes and control effector output. Sensor input is manifest via proteins introduced into the cell.

We use the term 'cell' loosely, to refer to this collection of genes and proteins. When we later refer to a cell responding to an environment in some way, this

is very close to saying that the cell's genome's regulatory dynamics respond in some way, as the only other constituents of the cell are the proteins, which do relatively little other than convey signals between genes, and to and from the environment.

The number of genes in the model does not vary between evolutionary runs. There is also a fixed alphabet of *protein types*—eight in each experiment, represented as a binary integer between 0 and 7 (000–111). Each protein type has an evolvable *decay rate*, and is present at a concentration ranged between 0 and an evolvable *saturation value*. Proteins are produced at a variable rate within the cell via gene expression, and do not directly interact with one another. We refer to protein types with the notation Pn, where n is an integer between 0 and 7.

In a simplification of biological GRN control (cf. [1, 19]), each gene on the genome has the same number of 'cis' *regulatory sites*, which act as *activators* and *inhibitors*, illustrated in fig. 1. Each site has a binary value between 0 and 7, representing a protein type. Identity-based template matching is used to *bind* proteins to these sites. For example, if a regulatory site has the value 5 (101), then instances of P5 (101) will bind to it. An evolvable global *binding proportion* parameter determines how many protein instances out of all available matching instances bind. Where a given value appears on more than one regulatory site on the genome, available proteins are distributed evenly amongst them.

The level of protein binding at each gene's cis regulatory sites is used to calculate the gene's *activation level*. Starting from zero, this is derived by adding the number of protein instances bound to activatory sites and subtracting the number bound to inhibitory sites.

Each gene produces a single protein type as its *output*. The quantity output depends on the gene's activation level, which is passed through a sigmoid *output function* (derived from tanh, the hyperbolic tangent), to calculate an output value between zero and a ceiling. Each gene has one of two *activity types* (also under evolutionary control), corresponding to one of two possible output functions. One type produces output as the activation level rises above zero, and hence is 'off' by default. The other produces output from just below zero, and hence is 'on' by default.

This model bears formal similarities to those used by Kauffman [4] and Reil [5]. Of the differences between those and the research presented here, the most significant are environmental coupling—which we focus on in the experimental sections below, and variable gene activation and output. Variable activation and output are illustrated in fig. 2, which shows output from the simple regulatory network illustrated in fig. 1, but isolated from environmental coupling. The first gene (g0), which is 'off' by default, produces P6 (110), and is up-regulated by the product of the second gene (g1), P1 (001), which is 'on' by default. g1 is in turn down-regulated by g0's P6 output. This creates an oscillating negative feedback loop controlling g1—as its P1 output increases, this increases g0's P6 output, which in turn inhibits g1. As g1's P1 output subsequently falls, so does g0's P6 output, reducing the inhibitory pressure on g1, which becomes active once more.

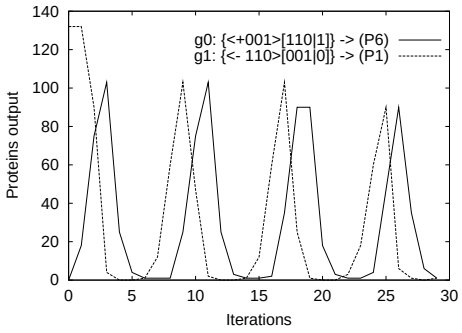

Fig. 2. Oscillating output in a simple two-gene network. *g0*, which is off unless activated, inhibits *g1*; *g1*, which is on by default, activates *g0*

2.3 Genetic Regulatory Control during the life-cycle

The model is run as a series of iterations, where an iteration is one step in an individual cell's 'lifetime'. Each iteration comprises the following sequence. First, unbound proteins are bound to matching regulatory sites. Second, transcription and translation occurs. For each gene, an activation level is calculated based on the level of binding at each cis regulatory site, then passed through the gene's output function to determine how many protein instances to produce. Third, proteins are decayed using each type's decay rate. Finally, an environmental interface performs effector output based on unbound protein concentrations, and adds proteins to the cell based on sensory readings (cf. fig. 1).

2.4 Evolution

We use a binary genome to encode individuals. Each gene is encoded with 8 bits (see fig. 1). There are four bits per cis regulatory site: one bit to indicate whether the site has an inhibitory or activatory effect, and three to encode the integer value of a matching protein type. The output region uses three bits to encode an output integer value, and one bit to indicate whether an activation level greater than zero is needed to produce an output.

The genome also encodes a number of global parameters not illustrated in fig. 1. There are nine four-bit parameters that each represent a proportion value in the range 0.0 to 1.0, via a lookup table. Eight of these parameters specify a rate of decay for each of the eight protein types (*decay rate*). The ninth describes what proportion of available proteins to bind to matching cis regulatory sites (*binding proportion*). An additional three-bit parameter specifies the maximum number of proteins of each type that can be present in the cell (*saturation value*), again via a look-up table.

Individuals were evolved using a basic fixed-length haploid GA, with single point cross-over, tournament selection and weak elitism (a single fittest individual is retained from each population into the next). Cross-over occurred with a probability of 0.9, and mutation with a probability of 0.01 per bit. Fitness functions were specified in terms of high-level phenotype behaviour in relation to an environment.

3 Homeostatic behaviour

In this section we describe results from evolving GRN-controlled cells in an abstract software environment. The environment consisted of a single variable, which we arbitrarily termed 'temperature' (this makes the target behaviour easy to visualise, as being like a heating unit attached to a thermostat). Individual cells' fitnesses were assessed in terms of their homeostatic ability with regard to this variable over 100 iterations. Our intention was to observe coupling between regulatory dynamics and environment in as simple a context as possible.

3.1 Experimental setup

Two protein types were used to couple cells to the environment. For every 'degree' the temperature was below the target on a given iteration, the same number of instances of P0 were introduced into the cell, representing input from sensors. For effector output, the number of unbound instances of P7 in the cell was summed, and the temperature increased by that amount (always greater than or equal to 0).

For each cell, the temperature was reset to a target figure of 50, with a possible range of 0–100. The temperature was then altered by 5 'degrees' at each iteration, initially falling, then switching direction at iterations 25, 75 and 90.

The GA was configured to maximise fitness, normalised to the range 0.0–1.0, calculated by summing deviations from the target, not including environmental increases. Individual fitness scores were inversely proportionate to this deviation measure. Individuals were immediately terminated with a fitness of 0.0 if the temperature fell to zero. Penalty points were accumulated and used to scale fitness downwards for failing to output heat when the temperature was below the target, and for outputting heat when above the target.

Each evolutionary run used a population size of 200 over 250 generations. We performed three evolutionary runs for each of five different specification cells—increasing the number of genes (N) from one to five. In all cases, each gene had a single cis regulatory site ($K = 1$). Note that unlike Kauffman's models, if $N > 1$, $K = 1$ does not imply that each gene is only affected by one other gene, as many genes may output proteins of a specific type, and hence exert a regulatory effect on a single gene.

3.2 Results

Independently of the value of N, individuals' fitnesses fall into three broad bands, most visible in the $N = 1$ runs (fig. 3). At the lowest level of fitness are cells whose genomes do not code for the output protein P7: they do not output any heat, and quickly 'die'. At a low to mid range are cells that do produce P7, but whose genomes have no cis regulatory sites matching P0, which effectively provides information about the environment. Their heat output is therefore unresponsive

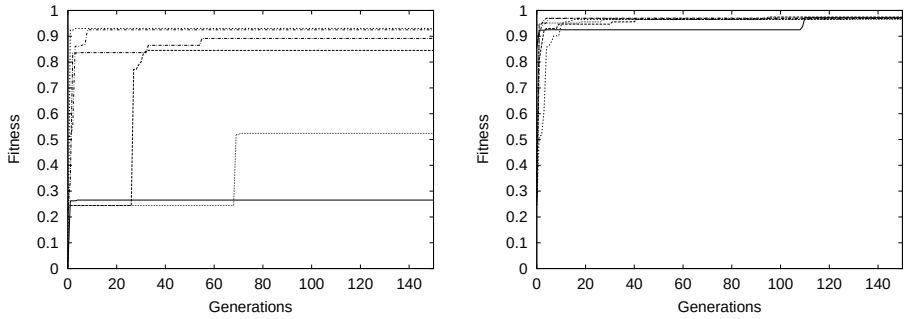

Fig. 3. Fitness over 6 evolutionary runs, showing the first 150 of 250 generations (pre-convergence), population size 200, for temperature controller cells with $N = 1, K = 1$ (*left*) and $N = 5, K = 1$ (*right*)

to the environment's ambient temperature. Finally, with fitnesses around 0.8–1.0 are cells that are fully coupled to the environment, sensitive to P0 and outputting P7 in response. For $N = 1$ genomes, this requires a gene of the form: $\{\langle +000\rangle[111|1]\}$ (see textual gene representation in fig. 1). This gene is inactive until activated by P0, producing P7. Hence the gene is activated when the cell senses a fall in temperature below the target, causing its effector surface to output heat in response. Additional minor increases in fitness can be achieved through alterations to the global parameters that are also encoded on the cells' genomes, such as protein decay rates.

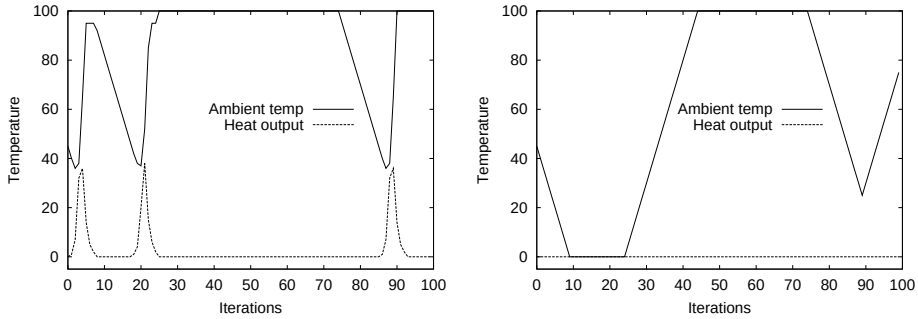

Fig. 4. Ambient temperature and cellular heat output with (*left*) and without (*right*) environmental coupling. Cell and environment are otherwise identical. Temperature changes visible in non-coupled environment (*right*) are autonomous

The left plot in fig. 4 shows the behaviour of the fittest cell from generation 34 of the second $N = 1$ run (fitness: 0.84). As the temperature falls below 50 degrees, the cell's output quickly rises, raising the temperature, causing in turn a drop in output, due to falling levels of P0, which reduces the gene's activation level. For contrast, the right plot shows data for the same environment and cell with no interaction between them—the cell outputs nothing, whilst the

environment's temperature drifts up and down. The marked difference between the two illustrates clearly the impact of the coupling relationship between the cell's genome and the environment. It is from this tightly coupled relationship that the emergent behaviour exhibited by the organism arises.

This is an important point. Simulated in isolation, a GRN with $N = 1, K = 1$ cannot hope to display much variety of behaviour. However, coupling to an environment creates a situation functionally equivalent to the GRN being a sub-set of a larger, effectively invisible regulatory network—with a potentially infinite non-repeating sequence of states therefore available to the sub-network, depending on the nature of the environment. The relevance of this phenomenon becomes apparent in the next section, in which we show phototactic, light-maximising robot controller behaviour emerging from the interplay between an environment and a rudimentary two-gene network.

4 Phototactic, light-maximising behaviour

This section describes results from evolving GRNs in a slightly more complex and realistic environment—a bounded flat physical surface with a light-source suspended in the middle, simulated using Cyberbotics' 'Webots' version 2. Individual cells' fitnesses were assessed in terms of their ability to maximise the amount of light encountered during their lifetimes.

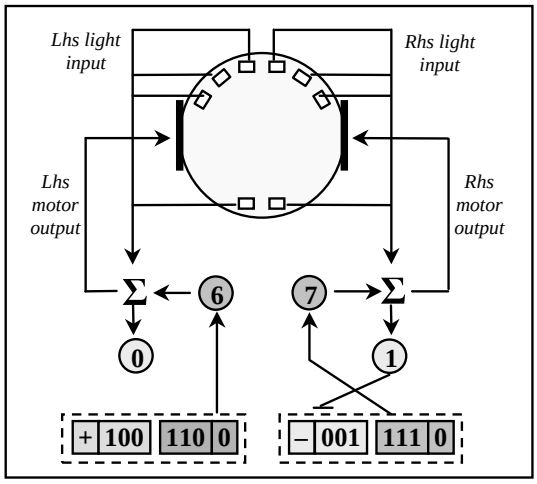

Fig. 5. Interface between Khepera robot and cell. Lhs light input generates instances of P0; rhs light instances of P1. Instances of P6 set the left wheel's speed; P7 the right wheel's speed. The genome shown has evolved to constantly drive both wheels, with rhs light input inhibiting right wheel speed. This produces the behaviour shown on the third plot in fig. 7

4.1 Experimental setup

In this experiment we created an environmental interface for the cells, illustrated in fig. 5, inspired by Braitenberg's 'vehicles' [20]. The Khepera (K-Team) robot has eight infrared sensors capable of detecting visible light, spaced at intervals around a cylindrical body, with a pair of motor-driven wheels on opposite sides. At each iteration the environmental interface sums light readings from the left hand side (lhs) sensors, scales this figure by a hard-coded gain value, and inputs the equivalent number of instances of P0 into the cell. The same process is repeated on the right hand side (rhs), and input as instances of P1. For motor output, the interface sets the speed of the left wheel to the number of instances of P6, and the speed of the right wheel to the number of instances of P7. No other pre- or post-processing was performed.

The mapping from sensor readings to proteins, and proteins to motor output is arbitrary, and remained fixed throughout evolutionary runs. It is the task of the evolutionary process to produce genomes that respond to and produce these proteins in such a way that behaviours with high fitness values emerge.

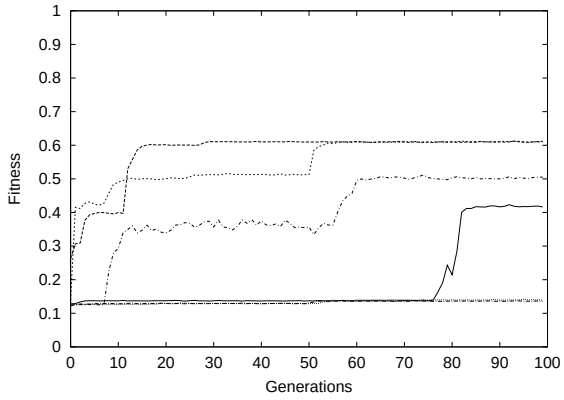

Fig. 6. Fitness over 6 evolutionary runs of 100 generations, population size 150, for phototactic robot controller cells with $N = 2, K = 1$

The GA was configured to maximise fitness, measured as: the amount of light encountered during the cell's lifetime as a proportion of a maximum value. Optimal behaviour would be to move rapidly to the region where light-intensity is greatest (a circle of some radius r from the light-source), and then to remain in that region with as many sensors as possible facing the light.

Individuals that exhibited no movement in the first 10 iterations of their lives were terminated, retaining their fitness scores. Six evolutionary runs were performed, each using a population of 150 over 100 generations. Each individual had two genes with one cis regulatory site per gene ($N = 2, K = 1$). Each individual began its lifetime in the same location.

4.2 Results

Over evolutionary time (fig. 6), the populations took a similar strategic route to that seen in the single variable homeostatic experiments: exploiting increased coupling to the environment to achieve greater fitness. Fig. 7 illustrates typical behaviours.

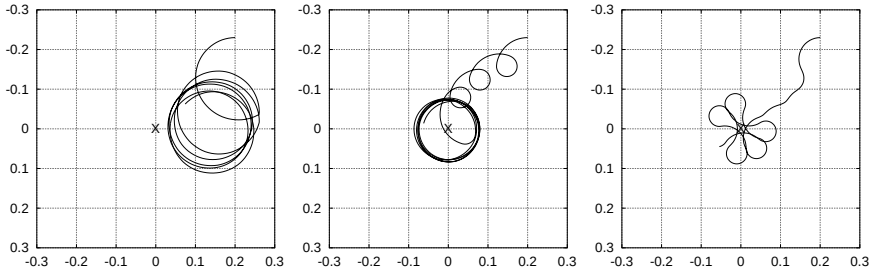

Fig. 7. GRN-driven phototactic robot controller trajectories (N=1,K=1), on a 2D surface. Axes show distance in metres, central 'X' marks light-source. From left to right, plots show fittest individuals (*looper*, *biased-looper* and *weaver*) from run 3 at generations 8, 41 and 83

Initial fitness improvements were quickly achieved through outbound coupling: producing P6 and P7, and hence motor output and movement. The first trajectory plot in fig. 7, generated by individual we will call *looper*, from generation 8 of run 3 (fitness: 0.42), is an example of this. Its genome (cf. fig. 1):

$$\{\langle +100\rangle[111|0]\}, \{\langle -101\rangle[110|0]\}$$

is not sensitive to environmental input (P0 or P1), but continually produces P6 and P7. Its genome also encodes a lower decay rate parameter (not shown) for P7 than P6, so the cell outputs more of the former than the latter, and hence the right wheel turns more quickly than the left, resulting in a continual anti-clockwise looping motion.

Inbound coupling appeared next, 'reading' the changes in input arising from motor output, with cis regulatory sites matching P0 or P1. The second trajectory plot in fig. 7, from generation 41 of run 3, is from a cell (*biased-looper*; fitness: 0.51) with the same genes as *looper*, but with an inhibitory cis site on the second gene matching P1 ($\langle -001\rangle$), corresponding to light from the right hand side. This produces a biased looping behaviour. As the robot loops away from the light source, light levels increase on the rhs, resulting in increased levels of P1 in the cell. This inhibits the production of P6 by the second gene, slowing the left hand wheel, tightening the loop. As the robot turns to face the light again, the rhs faces away from the light, and P6 production increases. This results in a net movement towards the light source, until the robot is circling it. At this point, interaction between cell and environment is stabalised, with the rhs continually facing away from the light.

In the transition between the second and third plots in fig. 7 (from *biased-looper* to *weaver*), the relationship between inputs and outputs were re-arranged, producing a genome as shown, from generation 83 of run 3 (fitness: 0.61):

$$\{\langle +100\rangle[110|0]\}, \{\langle -001\rangle[111|0]\}$$

This causes rhs sensory input (P1) to inhibit rhs effector output (P7). The first gene's output is changed to P6, which it produces constantly. This relationship is shown in fig. 5. The net effect is that whilst the controller maintains a tendency to loop to the left seen in *looper* and *biased-looper*, as the robot turns away from the light, increasing rhs input (P1), production of P7 is inhibited. This slows the right wheel, causing a loop to the right. This control strategy allows for a self-correcting, more or less direct route towards the light-source, around which the robot then orbits, weaving a series of broad left and tight right turns.

It is important to recognise the role the environment plays, through the cells' embodiment, in generating the behaviour observed, selecting from the cell's range of available 'next' states. As the space of possible states of system and environment increases and the scope for mutual perturbation between them becomes greater, so the space of potential behaviour (observable interaction) increases.

5 Conclusion

A key conclusion from a biological perspective, and certainly for other ALife models, is that the significance of the environment as a source of input external to the model should not be overlooked. Particularly from a dynamical systems perspective, the consequences of unpredictable input constantly perturbing the system are profound.

We have demonstrated the potential of environmentally coupled systems based on a GRN model to act as control systems, producing simple coherent behaviours. The simplicity of the GRN controllers is surprising given the behaviours exhibited during their evolution (see fig. 7). In a sense this illustrates the power of adaptive structural coupling that resides in even the most 'basic' of real biological systems—which combine GRN dynamics, sophisticated protein-based signal transduction pathways and many other features besides, coupled to the rich environment of the physical world, all under the aegis of the laws of physics.

Many interesting questions remain to be addressed. Given a more feature-rich, structurally plastic environment and a larger space of possible behaviours, what level of control system sophistication can be achieved with GRNs with higher values of N and K? How would this impact on evolvability? What consequences follow from increasing the sophistication and realism of the model, for example by adding protein–protein interactions and support for development involving differentiated multi-cellularity? Given a definition of embodiment in terms of structural coupling, and given the apparent relationship between system–environment coupling and potential behavioural complexity, what is the precise relationship between embodiment and adaptive behaviour, across different types of environment?

Acknowledgments

The first author wishes to thank Sanjeev Kumar for his advice and suggestions regarding biological processes and the binary representation and implementation of cis regulatory sites.

References

1. Davidson, E.H.: Genomic Regulatory Systems: Development and Evolution. Academic Press (2001)
2. Arthur, W.: The Origin of Animal Body Plans: A Study in Evolutionary Developmental Biology. Cambridge (2000)
3. Alberts, B., et al.: Molecular Biology of the Cell. 4th edn. Garland (2002)
4. Kauffman, S.: Metabolic stability and epigenesis in randomly constructed genetic nets. Journal of Theoretical Biology **22** (1969) 437–467
5. Reil, T.: Dynamics of gene expression in an artificial genome—implications for biological and artificial ontogeny. [21] 457–466
6. de Jong, H., et al.: Hybrid modeling and simulation of genetic regulatory networks: A qualitative approach. In Pnueli, A., Maler, O., eds.: Hybrid Systems: Computation and Control. Volume 2623 of LNCS., Springer-Verlag (2003) 267–282
7. Dellaert, F., Beer, R.: A developmental model for the evolution of complete autonomous agents. In Maes, P., et al., eds.: From Animals to Animats 4: Proc. 4th Int. Conf. on Simulation of Adaptive Behaviour, MIT Press (1996) 393–401
8. Astor, J.C., Adami, C.: A developmental model for the evolution of artificial neural networks. Artificial Life **6** (2000) 189–218
9. Bongard, J.C.: Evolving modular genetic regulatory networks. In: Proc. IEEE 2002 Congress on Evo. Comp., IEEE Press (2002) 1872–1877
10. Eggenberger, P.: Evolving morphologies of simulated 3D organisms based on differential gene expression. In Husbands, P., Harvey, I., eds.: Proc. 4th European Conf. on Artificial Life, MIT Press (1997) 205–213
11. Kumar, S., Bentley, P.J.: Biologically inspired evolutionary development. In: Proc. 5th Int. Conf. on Evolvable Systems: From Biology to Hardware. LNAI, Springer (To Appear)
12. Gilbert, S.: Ecological developmental biology: Developmental biology meets the real world. Developmental Biology **233** (2001) 1–12
13. Bolker, J.A.: From genotype to phenotype: Looking into the black box. In Kumar, S., Bentley, P., eds.: On growth, form and computers. Academic Press (To appear)
14. Ziegler, J., Banzhaf, W.: Evolving control metabolisms for a robot. Artificial Life **7** (2001) 171–190
15. Weiss, R., et al.: Genetic circuit building blocks for cellular computation, communications, and signal processing. Natural Computing: An Int. Journal **2** (2003)
16. Quick, T., Dautenhahn, K., Nehaniv, C.L., Roberts, G.: On bots and bacteria: Ontology independent embodiment. [21] 339–343
17. Maturana, H.R., Varela, F.J.: Autopoiesis and Cognition. D. Reidel (1980)
18. Beer, R.D.: A dynamical systems perspective on agent-environment interaction. Artificial Intelligence **72** (1995) 173–215
19. Ptashne, M.: A Genetic Switch: Phage Lambda and Higher Organisms. Blackwell Science (1992)
20. Braitenberg, V.: Vehicles: Experiments in Synthetic Psychology. MIT Press (1984)
21. Floreano, D., Nicoud, J.D., Mondada, F., eds.: Advances in Artificial Life: 5th Euro. Conf. on Artificial Life. Volume 1674 of LNAI., Springer (1999)