

Methods for Open-box Analysis in Artificial Development

Adrian Grajdeanu
George Mason University
Fairfax, VA 22030
agrajdea@cs.gmu.edu

ABSTRACT

Even a developmental system with very simple building blocks can evolve significantly complex artifacts. Understanding such complexity poses a significant challenge. This paper shows how various investigative methods that are typically used in biology, can be transferred and used in an artificial development context. As an instance of evolved complexity, a self-repairing artifact is analyzed using the following methods: ablation of environmental features, chemical concentrations monitors, *in silico* subsystem simulations, gene knock-outs, and modeling of the gene regulatory map. A number of mechanisms governing size-regulation and self repair are uncovered, such as: subtle timing of gene activations, stable regulation based on attractor points, opportunistic use of the environment and information content replication.

Categories and Subject Descriptors

I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search—*Developmental Systems*

General Terms

Experimentation, Measurement

Keywords

Open-box analysis, gene knockout, gene regulatory map, ablation, self repair, genetic regulatory network, development

1. INTRODUCTION

A lofty engineering goal is to build complex systems (structures or behaviors) while at the same time freeing the human creator from the task of specifying explicitly such complexity. Developmental systems hold the potential of achieving just that: they use simple building blocks and may lead to results of great complexity [3, 2, 6, 8, 4]. The resulting artifact is oftentimes analyzed in a black-box manner, showing that it *does* solve the problem, but not *how* it *internally*

achieves its task. In more rare instances, an open-box analysis is tackled such as in the work of Andersen and al. [1]. The authors present a gene regulatory map and use it to explain the mechanisms for growth and self-repair. But this avenue, while clearly of benefit, may not be enough when the number of genes increases (in their case there were only three). In other research, gene knock-outs were performed with not much bearing towards understanding their function. For example in [7] genes are knocked out one at a time and the effects are measured in terms of phenotypic variability, with no attempt at identifying precisely the effect each gene induces.

As it stands, in the field of developmental systems, there is no established methodology for performing a thorough open-box analysis that identifies the roles each gene carries and that subsequently identifies the inner workings of the evolved mechanisms that solve the given task. This research tries to reduce this gap, by proposing the use of the following investigative methods as they transfer from the field of biology: ablation of environmental features, chemical concentration monitors, *in silico* simulations of sub-systems, modeling of the gene regulatory map, and gene knock-outs. The latter are done individually, as well as in combinations, while keeping track qualitatively of the types of pathologies they induce. Finally, if biologists can benefit from *in silico* simulations, the method can be applied here too, by separating aspects of the whole and simulating them in isolation. This method sheds light on the dynamics of the system as far as existence of attractor points, stationary regimes, and sensitivity analysis around it.

The paper presents such an open-box analysis of an evolved solution to the problem of self-repair. The following two sections summarize portions of [5] in order to provide the background knowledge and create the context. Namely, section 2 briefly describes the developmental system used and then section 3 summarizes the details of the self-repair experiment. Finally, in section 4 the obtained results are presented and the detailed analysis is performed.

2. ADS: THE ARTIFICIAL DEVELOPMENTAL SYSTEM

Evolutionary algorithms draw their inspiration from nature: they encode potential solutions in a genotype and harness the power of selective breeding in order to obtain ever-improved solutions. Yet from a biological point of view they are so far removed from the natural model that indeed, they don't have much to say to the biologist. Effectively, evolutionary algorithms can be seen as mere engineering tools

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO '07, July 7–11, 2007, London, England, United Kingdom.
Copyright 2007 ACM 978-1-59593-697-4/07/0007 ...\$5.00.

for solving practical problems. In much the same way, when attempting to harness the power of development, one ponders what should be the appropriate level of abstraction? ADS was introduced ([5]) with the goal of providing a good paradigm that abstracts the natural development model and makes it practical for engineering tasks. The rest of this section summarizes what ADS *is* and – more importantly – *is not*, in effect highlighting the overall simplicity of the system.

ADS manages cells occupying locations in a discrete substrate, and protein concentrations located in the cells (internal protein concentrations), as well as in the substrate locations (external protein concentrations). Under the control of an evolved genome, cells can divide into an unoccupied location, die, or modify the internal and external protein concentrations. The substrate features protein diffusion among locations, with diffusion coefficients that are also evolved for each protein individually. ADS achieves its simplicity by abstracting many aspects of biological development: there are no chemical reactions, no cell membrane to model, no elaborate division process, no explicit cell differentiation, no transcription, no signaling factors, no complex gene activation process.

The evolved genome is a variable-length list of genes. Each gene has a condition and an action. The condition portion is either always true, or tests one or two protein concentrations against encoded thresholds. If the condition holds true, than the action fires. The action encoded in a gene enables the cell to divide, die, or modify protein concentrations internally or exchange protein content between its internal and external environment. At its very core, the genome governing the cells activity is in fact a simple rule system augmented with internal state. It can also be viewed as a cellular automaton, due to the fact that it involves a grid of locations and its functionality is governed by local interactions. Finally it is an instance of gene regulatory network in the sense that it models a many to many interaction pattern between genes: one gene regulates multiple genes, while itself being regulated by many.

3. THE SELF-REPAIR EXPERIMENT

This experiment is structured on a two-dimensional 27×27 grid, seeded with an egg-cell in the center, and cranking development for 150 time steps. Periodically, the environment acts in a destructive way: it randomly chooses one location and kills any existing cells located within 3 units from it. The system is configured to handle four proteins (a number chosen arbitrarily) and the evolutionary algorithm attempts to identify a set of diffusion coefficients and the genome to be contained in the egg-cell. The purpose of the genome is to govern the developmental phase such that it yields a stable colony of cells matching the desired pattern (fig 1). Also, it must detect and self-repair damage randomly inflicted by the environment.

The evolutionary algorithm is based on ES(2 + 16), but working on a binary representation and bit-flip mutation operator, augmented with primitives for varying the size of the genome. The evaluation phase compares the current shape with the desired one, accounting for differences at each developmental step and averaging them out. Finally the EA attempts to minimize this 'error'. In this manner there is incentive for the colony to quickly grow into the desired shape and also to quickly restore damage inflicted.




Figure 1: Target Shape To Be Maintained

There is noise in this evaluation, stemming from two sources. Firstly, ADS itself has a measure of stochasticity built in: the placement of the daughter cell after division is governed by a roulette drawing process. Secondly, the adversarial aspect of the environment is inherently random. In order to deal with this noise, the described evaluation is repeated 4 times and the averaged result is used by the EA as fitness.

4. RESULTS AND ANALYSIS

Fig 2 shows, for each of the 5 runs executed, the fitness of the best individual in each generation of the EA. All these runs converge at solutions whose fitness is around 0.2, which means that the resulting shape differs from the intended one by about twenty percent, as averaged across all steps of the developmental process. This fitness is not perfect (zero), because of two reasons. The first one is that the colony needs to grow into its shape from an just an egg cell, and also repair damage inflicted throughout its lifetime. Secondly, the resulting shape (even in the absence of damage) does not match perfectly the target shape. If the two shapes were only to differ in the points along the perimeter, then the fitness would measure about 0.36. As such a, 0.2 fitness signifies a high degree of overlap between the desired and the obtained shapes.

However a good fitness value in itself is not revealing of *how* the self-repair is actually achieved. For this reason, the best individual in run 4 is the subject of further scrutiny. The choice for this particular individual is two-fold. On one hand it has the best overall fitness across all the runs. On the other hand, it has the smallest number of genes among the individuals best in their run (24 to 145 genes). Fig. 3 shows a few consecutive steps that serve as a visual assessment of the quality of this solution. The color red (intermediate shade) marks cell existence, violet (lightest shade) depicts a cell that would like to divide and black signifies newly born cells. Notice the two types of damage: one completely internal and one affecting the edge. This individual is capable of repairing both types in the quickest possible time.

Unfortunately, even though it only contains 24 genes, a simple reading through the genome of this evolved individual does not reveal the mechanisms it employs to satisfy the self-repair task. It is at this point that more powerful investigative methods are needed, methods that are targeted and more revealing in the amount of detail. The inspiration comes again from the study of the natural organisms and processes. Precisely, the following avenues are pursued: ablation of environmental features, gene knock-outs (in isolation and combinations) and qualitative identification of




Figure 2: Results in Self-Repair Experiment

pathologies they induce, chemical concentration monitors, and, finally, inference and modeling of regulatory pathways.

As the underlying mechanisms gradually emerge, a number of hypotheses can be formulated regarding the roles that individual genes may serve. Such hypotheses are further validated or contradicted, providing an even clearer understanding of the inner working of the organism analyzed.

4.1 Internal and External Protein Concentration (Globally)

A first investigation focuses on the time evolution of the 4 protein concentrations both externally and internally at a global level. Subsequently, the proteins are named with the prefix 'p' followed by their zero based index. Proteins p3 and p1 have no significant presence either internally or externally throughout development. That is not to say that they have *no* effect, but only that their concentrations sharply decrease toward zero and stay there.

In contrast, proteins p0 and p2 have a much more interesting trajectory. Fig 4 and 5 show the internal and external concentrations of protein p0 and p2 early in the run (development step 13) and late (development step 113). Stronger shades denote higher concentrations. Early in the run protein p0 is localized on the footprint of the colony, with the internal concentrations being higher around the border. This higher level does not form a continuous border though. And in time, as can be seen late in the run, is maintained in the same fragmented manner. However, the external concentration of p0 spreads around the substrate eventually overcoming it.

In contrast p2 has a much more local footprint. Both internally and externally, it establishes itself in a gradient maximal at the center of the colony, and decreasing towards the edge. This pattern is maintained throughout the run, with the only change being that the gradient is shrinking both in maximal peak concentration and spatial extent. Incidentally, a quick look at the genome shows that the only gene responsible for triggering division, gene 20, is conditioned by protein p2 being above a certain value. The pattern illustrated in fig 5 thus elucidates why cells at the periphery of the colony lack the ability to divide, and effectively establish a perimeter that stops growth.

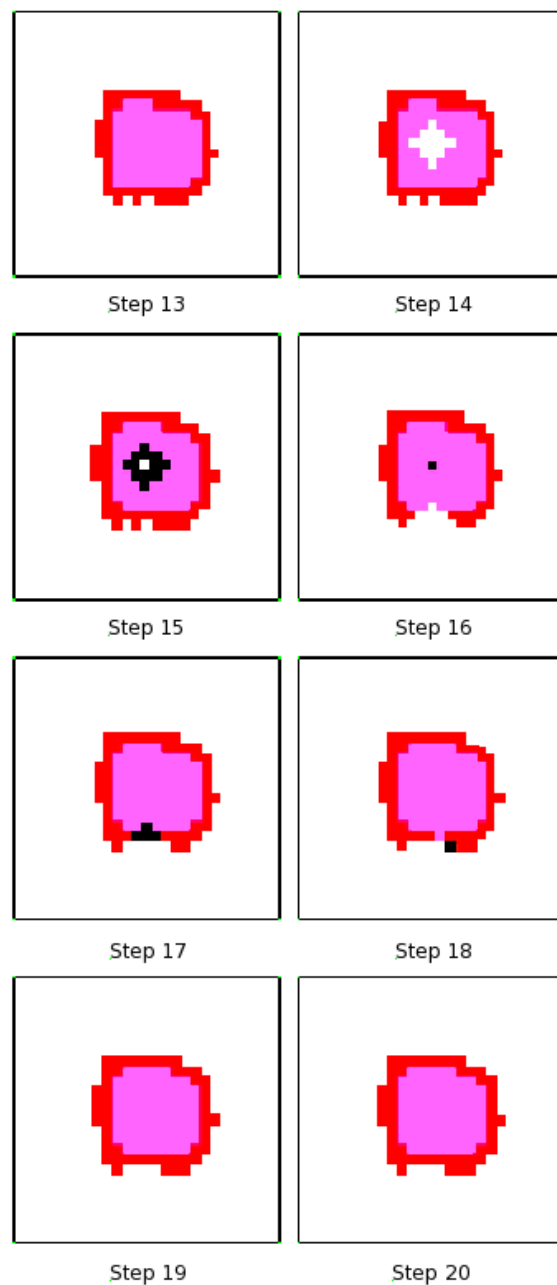


Figure 3: Self-Repair At Work

4.2 Internal and External Protein Concentration (Locally)

Having seen the overall pattern of protein concentrations, a next question is to look at the same data, but localized to a single cell. Thus a number of cells were chosen and the internal and external concentrations were plotted in time. For just one such cell, that is typical for all the others, fig 6 shows internal protein concentrations and fig 7 shows external protein concentrations.

The interesting points are about proteins p0 and p2. Precisely, p2 seems to have its internal *and* external concentrations moving in sync either higher or lower as if they are




Figure 4: Protein p0 Concentration: Internal (left) and External (right); Early In Development (up) and Late in Development (down)




Figure 5: Protein p2 Concentration: Internal (left) and External (right); Early In Development (up) and Late in Development (down)

scaled versions of each other. In contrast protein p0 has a more surprising and interesting behavior. While externally it goes to maximal concentration, internally it is extremely well regulated. Its concentration is very well controlled to be around 0.17. While the number is not relevant in an absolute sense, it is worth pointing out that it is higher in some other cells. As seen in section 4.1, some of the border cells have a higher p0 internal concentration (around 0.28). In either case, the concentration is very stable. The occasional spikes mark birth and death events, either of the cell itself, or one of its neighbors – one that this particular cell divides into




Figure 6: Typical Internal Protein Concentrations




Figure 7: Typical External Protein Concentrations

in order to repair damage. Notice that in such instances, the internal protein concentrations are sharply halved. Protein p0 then gets a kick upwards, and then it quickly settles down in its normally controlled steady state.

4.3 The Effects of Diffusion

As observed in sections 4.1, the protein p0 spreads throughout the environment rather quickly. A quick look at its diffusion coefficient shows a very large value (316.228), explaining such spread. The question arises whether this plays a role in the self-repair mechanism. Switching off diffusion (an environmental feature), proves that it plays *no role whatsoever*. The individual behaves identically whether the diffusion mechanism is on or off.

4.4 Gene Knockout Studies

By means of visual inspection and simple one-gene knockout studies, the 24 genes are categorized as in table 1. Each gene is named with the prefix 'g' followed by its 2 digit, 0-based index. Three categories need further explanation. The *syntactic no-op* means that the action associated with a certain gene is the *no operation*. Naturally knocking out such genes will have absolutely no effect. The *quasi stillborn* means that the colony grows to a size of 3 cells and then stops growing. Incidentally, it does not have the ability to

Table 1: Individual Genes Effects

Effect	Gene (1) knocked out
Syntactic no-op	g01, g13, g17
No visible effect	g00, g02, g08, g11, g15, g18, g19
Minor shape differences	g03, g04, g05, g07, g09, g10, g12
Stillborn	g20
Quasi stillborn	g23
Cancerous grown	g06, g14
Premature aging	g16, g21, g22

Table 2: Gene Pairs Effects (Select Results)

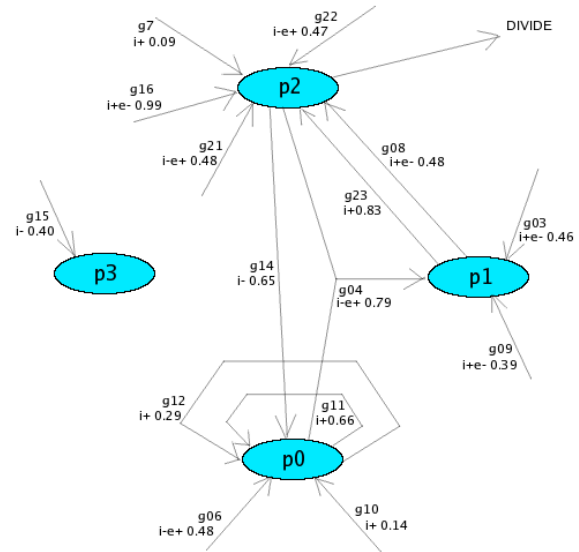
Effect	Genes (2) knocked out
Premature aging	(g03, g09)
Delayed Cancerous Growth	(g06, g23)
Cancerous growth inhibited	(g02, g06) (g06, g10) (g02, g14) (g10, g14)

regenerate either. Finally, *premature aging* means that while there is growth to the right shape and some regeneration, there is visible decrease in the regeneration ability throughout the 150 developmental steps. Incidentally, running the unaltered individual past 150 steps does show *aging*, as in it gradually loses its ability to regenerate, ultimately being overtaken by the environment and completely destroyed. Losing the regenerating ability *throughout* the 150 developmental steps is called *premature aging* and obviously incurs a semantical penalty on the fitness judgment scale.

The next step in the investigation is to knock out *pairs* of genes. Genes g16, g21 and g22 are responsible for staving the aging process. When knocked out in combination with themselves or with other genes, they accelerate the aging process. Similarly, genes g06 and g14 are responsible for inducing cancerous growth. When knocked out in combination with themselves or other genes, they continue to do so. However there are some surprising exceptions, summarized in table 2. It is not immediately obvious why g03 and g09, when knocked-out in tandem, lead to premature aging. Section 4.7 explains this manifestation as a deviation from the normal regulation of the growth process. Nor is it straightforward how does the absence of g23 delay the effect of knocking-out g06 (whose absence induces uncontrolled growth). This and the unexpected inhibition of uncontrolled growth caused by knocking out either of g02 or g10 in combination with either g06 or g14. are revealed in section 4.8.

4.5 Protein Regulation Map For Normal Individual

In this section, the causes and effects that proteins have over genes are investigated. Fig 8 displays the protein regulation pattern. It centers around the 4 proteins present in the system. Arrows are labeled with the gene number and a short indication of the effect: internal/external increase/decrease and by what factor. The value of the factor is not immediately obvious, but the intuition is as follows: the larger the factor (towards 1) the larger the effect; the smaller the factor (towards 0) less is the effect; the least effect is a no-op. Some genes always fire because their condition is the semantic equivalent of 'always true'. Other genes always fire because their condition holds true for the range

**Figure 8: Protein Regulation Map**

of protein concentrations that occur during the lifetime of the cells. Such genes label arrows that have a 'free' starting point. However, if the gene sometimes fires and other times doesn't, depending on some protein concentration, then the arrow starts from that particular protein. And in the case of 2 proteins cooperating to conditionally fire a certain gene (as is the case in gene g04) then multiple corresponding arrows are joined into one.

This graphical representation helps in further understanding the genes that participate in the mechanisms governing the behavior of the colony. For instance, protein p3 does not affect any other proteins, nor does it affect division. Hence, gene g15, the only one that has anything to do with p3 has no overall effect – a fact already known from previous analysis. But in addition, it becomes clear that if protein p0 has any bearing on the division, it is via gene g04, protein p1, genes g08 and g23 and finally p2. However, previous studies showed knocking out g04 and g08, separately and together has no effect. Thus the conclusion that p0 must itself have no effect. In turn, the effect of protein p1 on division if not via g08, may only be via g23. And there must be some effect, as knocking out g23 leads to quasi stillborn manifestation. A closer look at g23 reveals that it fires when p1 has an internal concentration higher than a very low value, as is the case early in the development phase and subsequently only sporadically. The effect of g23 is to increase (rather dramatically by a factor 0.83 – close to maximal value of 1) the internal concentration of p2. The egg cell starts with a maximal concentration of p2, but through division this concentration quickly decreases. The role of g23 during initial growth is to replenish the supply of p2. Thus, one of the effects of g23 can be summarized as being a *very well timed enabling of the growth process*. A second role of g23 manifests itself during the aforementioned sporadic activation and will be detailed in section 4.7 dealing with the combined effects of g07, g16, g21 and g22 of regulating p2. How genes g06, g10, g11, g12 and g14 work together to regulate p0, is the subject of the next section.

Table 3: Gene Effects on Protein p0

Gene	Effect	Factor
g06	eliminate	0.46875
g10	produce	0.148438
g11	produce	0.664062
g12	produce	0.296875
g14	consume	0.65625

4.6 Stabilizing p0

In this section, the focus is on genes g06, g10, g11, g12 and g14. They work together to regulate the concentration of protein p0. Table 3 presents the chained effects of these genes. If e and i denote the external and internal protein p0 concentration at a certain time step and e' and i' the same at the subsequent time step, the following simultaneous equations describe the transition:

$$\begin{cases} i' = 0.274607 + 0.0367324 * i + 0.032411 * e * i \\ e' = e + 0.46875 * i - 0.46875 * e * i \end{cases} \quad (1)$$

In order to understand the dynamical effects of equations 1, sub-system *in silico* simulations are performed. Start with random (i, e) values in $[0, 1] \times [0, 1]$ and iterate the above equations, plotting the resulting trajectory. Fig 9 shows the result of 100 such random trajectories, iterated for just 5 steps. Wherever the starting point may be, the internal




Figure 9: Regulating p0 (g06, g10, g11, g12, g14)

protein concentration gets close to about 0.28 in just one time step. Subsequently, the trajectory climbs upward towards $(i, e) = (0.295005, 1)$ point that is an attractor for the equations.

However, from gene activation maps it turns out that g11 only fires (and then it does so every time) only in border cells that also maintain a higher protein p0 internal concentration. In other cells, g11 does not fire at all. Here are the resulting equations governing the system, without the effects of g11:

$$\begin{cases} i' = 0.137928 + 0.109343 * i + 0.096479 * e * i \\ e' = e + 0.46875 * i - 0.46875 * e * i \end{cases} \quad (2)$$

The effect of iterating equations 2 is almost identical. The only difference is that the attractor point is changed to

Table 4: Gene Effects on Protein p2

Gene	Effect	Factor
g07	absorb	0.09375
g08	absorb (sporadic)	0.09375
g16	absorb	0.992188
g21	eliminate	0.484375
g22	eliminate	0.476562
g23	produce (sporadic)	0.835938

$(i, e) = (0.173674, 1)$, thus explaining the lower concentration of p0 in non-border cells.

Thus the mechanisms of regulating protein p0, both externally and internally, rely on gene activations attractor points that result from the parameters encoded in the gene actions. However, as effective as this control mechanism is, protein p0 plays no effective role in neither shape formation, nor self-repair. It can only be speculated that protein p0 may have played some role in one of the ancestors that led to the evolution of this particular individual.

4.7 Controlled growth and self-repair: the key is p2

This section focuses on genes g07, g16, g21 and g22, regulating protein p2. Gene g23 and g08 also affect protein p2, but they fire during the initial growth and after that only sporadically. The effects of g08 and g23 are thus corrective in nature, and therefore do not influence the fundamental dynamical behavior of the concentrations in protein p2 in the 'adult' colony. The actions associated with these genes are presented in table 4. Notice that if the sporadic effect of g23 is ignored, all these remaining genes affect p2 via either 'absorb' or 'eliminate' operations. Both these operations are characterized by a conservation of protein concentrations. In other words, using the formalism previously introduced, $i + e$ is an invariant for both 'absorb' and 'eliminate' operations. More importantly, because neither 'produce' nor 'consume' operations affect the total amount of p2 (g23 being ignored at this time) it would stand that p2 must be conserved overall throughout the lifetime of the colony. The egg cell starts out with an internal concentration of p2 set to 1, and that is all there is. However, when a cell is killed, its internal proteins are lost. In this way some p2 is removed from the system, and there is no way to replenish it. It stands to reason that the cause of eventual aging observed is as follows: the overall concentration of p2 decreases, more and more cells lose the potential to divide and eventually the colony loses the ability to self-repair. However, as each cell division effectively halves the concentration of p2, the immediate intuition would be that the colony would spend its ability to self-repair rather quickly. Obviously this is not the case, as it is able to withstand the adversarial environment quite well for the required 150 developmental steps.

There are two mechanisms at play and an understanding of the first one comes from an experiment similar to that performed in section 4.6. Considering only g07, g16, g21 and g22 (ignoring the sporadic effects of g08 and g23), the equations governing the system are quite complicated: polynomial whose highest degree is $i^8 * e^8$. Yet, fig 10 sheds light on this matter. This time there is no single attractor point, but there are an infinity of equilibrium points along a concave curve linking the (0,0) point with the (1,1)




Figure 10: Regulating p2 (g07, g16, g21, g22)

point. Again, the control mechanism is powerful in the sense that a point away from this curve, takes only a few steps to reach the equilibrium. But what is the use of this curve of equilibrium points? The answer reveals itself by intuitively following the trajectory of the protein p2 concentrations of a cell through the process of being killed and then regenerated. When a cell is killed its internal protein p2 is lost. However the external protein p2 survives unchanged. A new cell is spawned in the same location, with an internal concentration of p2 of about half of what it has been. Then the internal/external concentration balance is restored to a new equilibrium point, somewhere on the curve, closer towards the (0,0) point. But the concave shape of the equilibrium curve serves to increase the number of such steps (kill-respawn) needed to effectively lose all of the protein p2 (reach the (0,0) point). This in fact is the key to the primary mechanism conferring the resilience of the colony to externally inflicted damage. As the colony grows, it lays outside in the environment a template gradient of p2. This serves double purpose. On one hand serves to restore an internal p2 concentration after respawning close to what it originally was (thus regulating which cells maintain division potential and which not). On the other hand the external template of p2 gradient serves to somewhat 'isolate' p2 from the irreversible loss induced by cells being killed, thereby extending in time its regenerative ability.

From gene knockout experiments it was observed that removing either g21 or g22 brings upon premature aging, the effect being even more pronounced if both are knocked out simultaneously. Why is that? Fig 11 shows what happens when g21 is knocked out: the equilibrium curve loses its concave property. An almost identical effect is observed when g22 is knocked out, and when both g21 and g22 are suppressed in combination, the curve becomes sharply convex. These findings strengthen the role attributed to the concavity of the equilibrium curve and validate the understanding of the mechanism.

And finally, before concluding this section, a few words about the strange result noticed when g03 and g09 are both knocked out: premature aging is observed. Further investigations reveal a second mechanism crucial to achieving and maintaining the self-repair capability. One of the effects of g03 and g09 is to increase the concentration of p1 internally. In other words, they both slow down the decrease of




Figure 11: Regulating p2 (g21 Suppressed)

internal p1 concentration. Interestingly, this slowing down is enough to allow g23 to keep firing sporadically after the initial growth, and thus to replenish some of the reserve of p2 existent in the system. Knocking out either g03 or g09, reduces the rate of sporadic g23 activation, but in tandem they completely eliminate this aspect. Consequently by having both g03 and g09 knocked out a second mechanism essential for self-repair is identified: the *sporadic* firing of g23. It is also essential that this firing be sporadic. If g23 were to fire more often, this would be one cause leading to uncontrolled growth.

4.8 'Cancers': Causes and Cures

From gene knock-out studies it was revealed that suppressing gene g14 leads to uncontrolled growth. Similarly, knocking g06 produces the same undesirable result. What is puzzling is that the gene regulation map shows both g14 and g06 only have an effect on protein p0: they both actively contribute to lower it. But analysis so far revealed that, while highly regulated, normally p0 has no effect neither on growth, nor on self-repair. The secret lies on gene g02. In an unmodified individual (no genes knocked out) gene g02 never fires, because it requires high levels of p0. However, if the regulating mechanism of p0 is damaged, and p0 levels increase, g02 kicks in. Its effects are to increase unconditionally the level of p2, thus leading to uncontrolled growth. It may very well be that this mechanism is an atavism kept from an ancestor, but in this incarnation its triggering is detrimental. In order for this to be avoided it is essential to either control the internal concentration of p0 below a certain threshold, or disrupt the mechanism triggered by g02. As it turns out, knocking out g02 does indeed stop the cancerous growth, the reason being obvious. However, knocking out g10 also stops cancerous growth, by stopping p0 from reaching levels that would activate g02. Interestingly, one would think that similarly to knocking out g10, knocking out either g11 or g12 (or in tandem) would have the same effect. But it doesn't. The difference between g10 on one side and g11 and g12 on the other side is that g10 fires all the time unconditionally, while g11 and g12 only fire if p0 itself is below certain levels. When either g14 or g06 are out, the concentration of p0 'escapes' beyond those levels required for g11 and g12 to fire, thus thwarting any controlling effect they might otherwise induce.

Another interesting aspect is the combination of both g23 and g06 being absent. While a g23 knock-out leads to quasi stillborn pathology by means of not replenishing the concentration of p2 early during the growth process, lacking p06 leads to the aforementioned escalation of p0 concentration and further to cancerous growth. When both g23 and g06 are absent, the observed result is a delayed cancerous growth. In other words, the growth effect induced by lack of g06 overcomes the growth inhibition caused by lack of g23. The cause of this is that g14 only fires when the concentration of p2 is above some threshold, which doesn't happen quite as often without g23. Consequently, the rate of increase in p0 induced by missing g06 is much slower without g23. Slower, but not enough to avert the onset of cancerous growth (damage of the p0 regulation and activation of g02). Thus it can be said that g23 has yet another role: in the absence of p06 it indirectly guards against cancerous growth.

5. CONCLUSIONS

This work illustrates how investigative methods that are traditionally associated with the field of biology can successfully transfer to the field of abstract developmental systems and be used to perform open-box analysis of evolved artifacts.

In the instance analyzed, the evolved mechanisms that solve the problem of self-repair turn out to be of significant complexity: finely tuned timing of gene activations, stable regulation based on attractor points, opportunistic use of the environment and information content replication. Some of these mechanisms clearly address the task of self-repair. Others are more subtle, as they seem to have no effect. But if disturbed by genes knock-outs they lead to pathological manifestations such as aging and cancerous growth. Fixing these pathologies is possible by even more gene knock-outs, but care must be exercised as not all hypotheses turn out to hold as expected. To get a deep understanding of how different mechanisms are interlocked in a regulatory network, an open-box analysis is needed. The methods appropriate for such investigation are those typically encountered in wet labs: ablation of environmental features, chemical concentration monitors, subsystem simulations, modeling of the gene regulatory network and gene knock-outs. Gene suppression can be done in isolation and in combination, and is associated with a qualitative assessment of pathologies induced.

Although this study focused the analysis on just one instance of complexity, namely an evolved solution to the self-repair problem, the methods used are general and can be applied to analyze other evolved dynamics too. Evolution is particularly well suited to address the creative part of discovering new ways to solve problems, but then the practical goal is to take the evolved design and refine it by means of manual optimization. The value of open-box analysis and the investigative methods illustrated in this paper is significant, as they enable the researcher/designer to identify primary mechanisms, redundancies and potential vulnerabilities.

6. REFERENCES

- [1] T. Andersen, R. Newman, and T. Otter. Development of virtual embryos with emergent self-repair. Technical Report FS-06-03, Developmental Systems; AAAI Fall Symposium, 2006.
- [2] P. J. Bentley and S. Kumar. Three ways to grow designs: A comparison of embryogenies for an evolutionary design problem. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, editors, *Genetic and Evolutionary Computation Conference – GECCO-1999*, pages 25–43. Morgan Kaufmann, 1999.
- [3] P. Eggenberger. Evolving morphologies of simulated 3d organisms based on differential gene expression. In *Fourth European Conference on Artificial Life*, 1997.
- [4] D. Federici. Evolving a neurocontroller through a process of embryogeny. In S. Schaal, A. Ijspeert, A. Billard, S. Vijayakumar, J. Hallam, and J. Meyer, editors, *From Animals To Animals 8: SAB 2004*, pages 373–384, 2004.
- [5] A. Grajdeanu and S. Kumar. A novel developmental system for the study of evolutionary design. Technical Report FS-06-03, Developmental Systems; AAAI Fall Symposium, 2006.
- [6] G. S. Hornby. *Generative Representations for Evolutionary Design Automation*. PhD thesis, Brandeis University, 2003.
- [7] S. T. Psujek and R. D. Beer. Phenotypic variability in canalized development systems. In *Proceedings of ALIFE9*, pages 415–420, 2004.
- [8] D. Roggen and D. Federici. Multi-cellular development: is there scalability and robustness to gain? In X. Yao and al. ed., editors, *Proceedings of PPSN VIII*, pages 391–400, 2004.