

Evolving Fractal Gene Regulatory Networks for Graceful Degradation of Software

Peter J. Bentley

Department of Computer Science, University College London,
Malet Place London, WC1E 6BT, UK
P.Bentley@cs.ucl.ac.uk
<http://www.cs.ucl.ac.uk/staff/p.bentley/>

Abstract. Fractal proteins are an evolvable method of mapping genotype to phenotype through a developmental process, where genes are expressed into proteins comprised of subsets of the Mandelbrot Set. The resulting network of gene and protein interactions can be designed by evolution to produce specific patterns that in turn can be used to solve problems. In this paper, adaptive developmental programs, capable of developing different solutions in response to different signals from an environment, are investigated. The evolvability of solutions and the capability of these solutions to survive damage is assessed. Evolution is used to create a fractal gene regulatory network (GRN) that calculates the squareroot of the input (its environment). This is compared with a GP-evolved squareroot function and a human-designed squareroot function. The programs are damaged by corrupting their compiled executable code, and the ability for each of them to survive such damage is assessed. Experiments demonstrate that only the evolutionary developmental code shows graceful degradation after damage. This provides evidence that software based on gene, protein and cellular computation is far more robust than traditional methods. Like a multicellular organism, with its genes evolved and developed, it shows graceful degradation. Should it be damaged, it is designed to continue to work.

1 Introduction

Human designs are carefully-crafted, consciously-created fusions of experience and skill. Our designs usually work reliably and well under the conditions they were designed for. Unfortunately, they rarely work well under unforeseen conditions. A ship will sink in the wrong kinds of seas, a car will crash on the wrong kinds of roads. A program will fail in the wrong kind of software environment. And sadly for computer-users worldwide, the mess of different software on an average computer causes such complex environments that programs fail with tired regularity.

Natural systems are also carefully crafted, but there is no conscious mind or skill needed to produce her designs. Generations of past experience drives evolution to create robust, damage-tolerant solutions. Organisms don't fail if they sustain minor damage. A broken finger hardly slows us. The equivalent damage to a human-designed program would produce terminal failure.

The work described here continues an ongoing investigation into the use of developmental systems with evolutionary computation. Here, fractals are employed as a

computer representation of proteins. Earlier work has shown that fractal proteins are highly evolvable by a genetic algorithm (Bentley 2004, 2003c), that specific patterns of activation in a fractal gene regulatory network (GRN) can be evolved (Bentley, 2004, 2003b), that they can perform computational tasks such as function regression and robot control (Bentley 2003a), and that evolved fractal GRNs naturally show fault-tolerance (Bentley 2003c). This work now focuses on the evolution of developmental programs that display graceful degradation when damaged.

2 Background

Questions of reliability and graceful degradation occur frequently in fields focusing on embedded systems. To date, most solutions seem to depend on architectures that partition software into separate components, organised in such a way that the failure of non-critical components will not induce the failure of the whole system (Shelton & Koopman, 2001).

In Evolutionary Computation, scientists have been focussing on the ability of evolution, and more commonly developmental methods, to enable self-repairing behaviour and graceful degradation of solutions. The work of Andy Tyrrell and his group create fault-tolerant hardware inspired by ideas of embryology and immune systems (Jackson and Tyrrell, 2002). More recently, Julian Miller has described experiments evolving developmental programs to create “French Flag” patterns (Miller and Banzhaf, 2003). He shows that development is able to regenerate these patterns should some of their cells be removed. Current work by Mahdavi and Bentley (2003) demonstrates how adaptive evolutionary control can enable a “Smart Snake” to redevelop new movement strategies even after the loss of a crucial muscle (Nitinol wire).

In his research on fault-tolerant systems, Thompson (1997) describes how “graceful degradation for free” can be achieved in theory and in practice for robot controllers, “from the nature of the evolutionary process.” Thompson suggests that mutation-insensitive individuals will, in the long term, survive better, thus producing a pressure towards fault-tolerant solutions. More recently, the same results were demonstrated with fractal developmental processes (Bentley 2003c), where there are no direct mappings: pleiotropy and polygeny are prevalent, and genes are reused over many developmental iterations. It was shown that through the Baldwin Effect, solutions “naturally” became more efficient and fault-tolerant (Bentley 2003c). In more detail, the work demonstrated that if evolution was permitted to run for a further 1000 generations after a perfect solution had evolved, the fractal GRNs continued to evolve: the number of genes and proteins that made up the solution was reduced (so there is less to be damaged), and duplicate genes were added, which provide redundancy and protection against damage.

This paper extends this work, showing that damage directly to the executable code (and not just a gene in the system) can be survived by evolved developmental programs.

3 Fractal Proteins

Development is the set of processes that lead from egg to embryo to adult. Instead of using a gene for a parameter value as we do in standard EC (i.e., a gene for long legs),

natural development uses genes to define proteins. If expressed, every gene generates a specific protein. This protein might activate or suppress other genes, might be used for signalling amongst other cells, or might modify the function of the cell it lies within. The result is an emergent, asynchronous, parallel “computer program” made from dynamically forming gene regulatory networks (GRNs) that control all cell growth, position and behaviour in a developing creature (Wolpert et al, 2001).

Table 1. Types of objects in the model

<i>fractal proteins</i>	defined as subsets of the Mandelbrot set.
<i>Environment</i>	contains one or more fractal proteins (expressed from the environment gene(s)), and one or more <i>cells</i> .
<i>Cell</i>	contains a <i>genome</i> and <i>cytoplasm</i> , and has some <i>behaviours</i> .
<i>Cytoplasm</i>	contains one or more fractal proteins.
<i>Genome</i>	comprising <i>structural genes</i> and <i>regulatory genes</i> . In this work, the structural genes are divided into different types: <i>cell receptor genes</i> , <i>environment genes</i> and <i>behavioural genes</i> .
<i>regulatory gene</i>	comprising operator (or promoter) region and coding (or output) region.
<i>cell receptor gene</i>	a structural gene with a coding region which acts like a mask, permitting variable portions of the environmental proteins to enter the corresponding cell cytoplasm.
<i>environment gene</i>	a structural gene which determines which proteins (maternal factors) will be present in the environment of the cell(s).
<i>behavioural gene</i>	structural gene comprising operator and cellular behaviour region.

FRACTAL DEVELOPMENT

For every developmental time step:

For every cell in the embryo:

Express all environment genes and calculate shape of merged environment fractal proteins

Express cell receptor genes as receptor fractal proteins and use each one to mask the merged environment proteins into the cell cytoplasm.

If the merged contents of the cytoplasm match a promoter of a regulatory gene, express the coding region of the gene, adding the resultant fractal protein to the cytoplasm.

If the merged contents of the cytoplasm match a promoter of a behavioural gene, use coding region of the gene to specify a cellular function.

Update the concentration levels of all proteins in the cytoplasm. If the concentration level of a protein falls to zero, that protein does not exist.

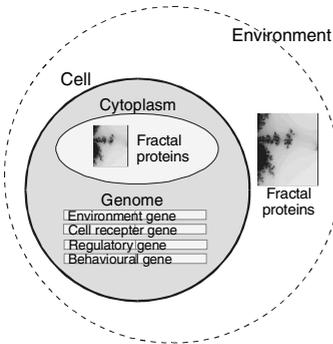


Fig. 1. Representation using fractal proteins

Fig. 2. The fractal development algorithm

In this work, a biologically plausible model of gene regulatory networks is constructed through the use of genes that are expressed into *fractal proteins* – subsets of the Mandelbrot set that can interact and react according to their own fractal chemistry. Further motivations and discussions on fractal proteins are provided in (Bentley, 2004 & 2003a,b,c). Table 1 describes the object types in the representation; Figure 1 illus-

trates the representation. Figure 2 provides an overview of the algorithm used to develop a phenotype from a genotype. Note how most of the dynamics rely on the interaction of fractal proteins. Evolution is used to design genes that are expressed into fractal proteins with specific shapes, which result in developmental processes with specific dynamics.

3.1 Defining a Fractal Protein

In more detail, a fractal protein is a finite square subset of the Mandelbrot set (Mandelbrot 1982), defined by three codons (x,y,z) that form the coding region of a gene in the genome of a cell. Each (x, y, z) triplet is expressed as a protein by calculating the square fractal subset with centre coordinates (x,y) and sides of length z , see fig. 3 for an example. In addition to shape, each fractal protein represents a certain *concentration* of protein (from 0 meaning “does not exist” to 200 meaning “saturated”), determined by protein production and diffusion rates.



Fig. 3. Example of a fractal protein defined by $(x=0.132541887, y=0.698126164, z=0.468306528)$

3.2 Fractal Chemistry

The model incorporates the notions of cell cytoplasm – a “container” which holds the proteins belonging to the corresponding cell) and (cellular) environment – the global “container” which holds proteins visible to all cells. In order to model complex protein-protein and protein-gene interactions, multiple fractal proteins are allowed to interact according to their fractal shapes. The interaction occurs by merging separate protein shapes to form new, complex compounds. The result is a product of their own “fractal chemistry” which naturally emerges through the fractal interactions.

Fractal proteins are merged (for each point sampled) by iterating through the fractal equation of all proteins in “parallel”, and stopping as soon as the length of any is unbounded (i.e. greater than 2). Intuitively, this results in black regions being treated as though they are transparent, and paler regions “winning” over darker regions. See fig 4 for an example.

3.3 Calculating Concentration Levels

The total concentration of two or more merged fractal proteins is the mean of the different concentrations seen in their merged product. For example, fig. 4 shows how fractal proteins are merged to form a new fractal shape. Figure 5 illustrates the



Fig. 4. Two fractal proteins (left and middle) and the resulting merged fractal protein combination (right)



Fig. 5. The different concentrations of the two fractal proteins (left and middle) and the concentration levels in their merged product (right)

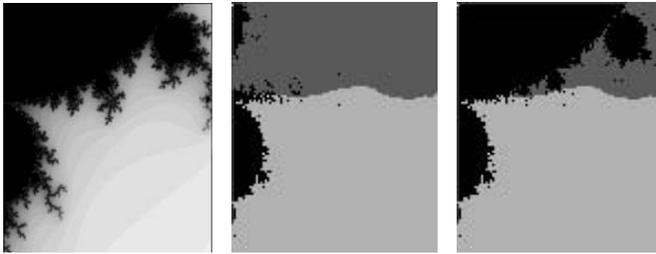


Fig. 6. The shape of the desired protein as defined by a promoter (left), the shape and concentration levels of merged proteins in the cytoplasm (middle) and the concentration levels seen on that promoter (right), where total concentration is taken as mean. Note that although a merged protein may decrease affinity (similarity) to the promoter, should the second protein have a higher concentration level to the first, it will boost overall concentration seen by the promoter, i.e., act like a catalyst to speed up (or slow down, if lower) the “reaction”

resultant areas of different concentration in the product. When being compared to the (xp, yp, zp) promoter region of a gene (the “conditional” part of the gene to be matched, see later section on genes), the concentration seen on that promoter is described by all those regions that “fall under” the promoter, see fig. 5. In other words, the merged product is masked by the promoter fractal, and the total concentration on the promoter is the mean of the resulting concentrations, see Fig. 6.

3.4 Updating Protein Concentration Levels

Every developmental time step, the new concentration of each protein is calculated (synchronously). This is formed by summing two separate terms: the previous concentration level after diffusion (*diffusedconc*) and the new concentration output by a gene (*geneoutputconc*). These two terms model the reduction in concentration of proteins over time, and the production of new proteins over time, respectively, where:

$$diffusedconc = prevconcentration \times (1 - 1/PROTEINDEC + 0.2)$$

(PROTEINDEC is a constant normally set to 5)

and:

$$geneoutputconc = totalconc \times \tanh((totalconc - ct) / CWIDTH) / CINC$$

where: *totalconc* is the mean concentration seen at the promoter,
ct is the concentration threshold from the gene promoter
 CWIDTH is a constant (normally set to 30)
 CINC is a constant (normally set to 2)

3.5 Genes

The environment gene, cell receptor gene, regulatory genes, and behavioural genes all contain 7 real-coded values:

<i>xp</i>	<i>yp</i>	<i>zp</i>	<i>Affinity threshold</i>	<i>Concentration threshold</i>	<i>x</i>	<i>y</i>	<i>z</i>	<i>type</i>
-----------	-----------	-----------	---------------------------	--------------------------------	----------	----------	----------	-------------

where (*xp, yp, zp, Affinity threshold, Concentration threshold*) defines the promoter (operator or precondition) for the gene and (*x,y,z*) defines the coding region of the gene. The *type* value defines which type of gene is being represented, and can be one or all of the following: *environment, receptor, behavioural, or regulatory*. This enables the type of genes to be set independently of their position in the genome, enabling variable-length genomes. It also enables genes to be multi-functional, i.e. a gene might be expressed both as an environmental protein and a behaviour.

When *Affinity threshold* is a positive value, one or more proteins must match the promoter shape defined by (*xp,yp,zp*) with a difference equal to or lower than *Affinity threshold* for the gene to be activated. When *Affinity threshold* is a negative value, one or more proteins must match the promoter shape defined by (*xp,yp,zp*) with a difference equal to or lower than $|Affinity\ threshold|$ for the gene to be repressed (not activated).

To calculate whether a gene should be activated, all fractal proteins in the cell cytoplasm are merged (including the masked environmental proteins, see later) and the combined fractal mixture is compared to the promoter region of the gene.

The similarity between two fractal proteins (or a fractal protein and a merged fractal protein combination) is calculated by sampling a series of points in each and summing the difference between all the resulting values. (Black regions of fractals are ignored.) Given the similarity matching score between cell cytoplasm fractals and gene promoter, the activation probability of a gene is given by:

$$activationprob = (1 + \tanh((matchnum - Affinity\ threshold - Ct) / Cs)) / 2$$

where: *matchnum* is the matching score,
Affinity threshold is the matching threshold from the gene promoter
Ct is a threshold constant (normally set to 50)
Cs is a sharpness constant (normally set to 50)

Regulatory Gene

Should a regulatory gene be activated by other protein(s) in the cytoplasm (which have concentrations above 0) matching its promoter region, its corresponding coding region (x,y,z) is expressed (by calculating the subset of the Mandelbrot set) and new concentration level calculated. To do this, the concentration of the resulting protein is modified by incrementing with *geneoutputconc*, the result of a function of the concentration threshold (*ct*) and the mean total concentration seen at the gene promoter (*totalconc*), as given in section 3.5. In this way, higher concentrations of protein on the promoter will cause an increased rate of output protein concentration growth, while lower concentrations (below the *ct* threshold) will increase the diffusion rate of the output protein (its concentration will decrease at a higher rate).

The cell cytoplasm, which holds all current proteins, is updated at the end of the developmental cycle.

Cell Receptor Gene

At present, the promoter region of the cell receptor gene is ignored, and this gene is always activated. As usual, the corresponding coding region (x,y,z) is expressed by calculating the subset of the Mandelbrot set. However, the resultant fractal protein is treated as a mask for the environmental proteins, where all black regions of the mask are treated as opaque, and all other regions treated as transparent. For an example, see fig. 7. If there is more than one receptor gene, only the first in the genome is used.

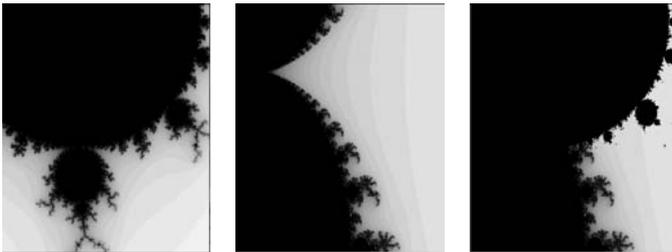


Fig. 7. Cell receptor protein (left), environment protein (middle), resulting masked protein to be combined with cytoplasm (right)

Environment Gene

Like the cell receptor gene, this gene is always activated. It produces environmental factors for all cells: fractal proteins of concentration 200. If there is more than one environmental gene, the expressed environmental proteins are merged before being masked by the receptor protein. If one or more values are being input to the system,

the concentration of the environmental fractal proteins are set to those values, i.e. an input to the system disturbs the environment during development.

Behavioural Gene

A behavioural gene is activated when other protein(s) in the cytoplasm match its promoter region (using the *affinity threshold*). For this application, a gradual activation between not activated and activated was required, using the x value of the coding region (x,y,z) triplet as a *fate* value to define a function, calculated as follows:

If the gene is being activated with a negative *Affinity threshold*,
 $output = output - (totalconcentration - concentrationthreshold) * fate$
 If the gene is being activated with a positive *Affinity threshold*,
 $output = output + (totalconcentration - concentrationthreshold) * fate$

Note how the total concentration of proteins seen on the promoter is offset against the *Concentration Threshold* gene and scaled by the *fate* gene (x value of the coding region), allowing evolution to adjust the range of values seen on the output, and used to specify behaviours. (If there is more than one behavioural gene, the change to *output* is averaged over all behavioural genes, each developmental step.)

3.6 Fractal Sampling

All fractal calculations (masking, merging, comparisons) are performed at the same time, by sampling the fractals at a resolution of 15×15 points. Note that the comparison is normally performed between the single fractal defined by (xp,yp,zp) of a gene and the merged combination of all other proteins currently in the cytoplasm. The fractal being compared is treated a little like the cell receptor mask – only those regions that are not black are actually compared with the contents of the cytoplasm.

3.7 Development

As was illustrated in figure 2, an individual begins life as a single cell in a given environment. To develop the individual from this zygote into the final phenotype, fractal proteins are iteratively calculated and matched against all genes of the genome. Should any genes be activated, the result of their activation (be it a new protein, receptor or cellular behaviour) is generated at the end of the current cycle. Development continues for d cycles, where d is dependent on the problem. Note that if one of the cellular behaviours includes the creation of new cells, then development will iterate through all genes of the genome in all cells.

3.8 Evolution

The genetic algorithm used in this work has been used extensively elsewhere for other applications (including GADES (Bentley 1999)). A dual population structure is employed, where child solutions are maintained and evaluated, and then inserted into a larger adult population, replacing the least fit. The fittest n are randomly picked as parents from the adult population. The degree of negative selection pressure can be controlled by modifying the relative sizes of the two populations. Likewise the degree of positive selection pressure is set by varying n . When child and adult population sizes are equal, the algorithm resembles a canonical or generational GA. When the

child population size is reduced, the algorithm resembles a steady-state GA. Typically the child population size is set to 80% of the adult size and $n = 40\%$. (For further details of this GA, refer to (Bentley 1999).)

Unless specified, alleles are initialised randomly, with (xp,yp,zp) and (x,y,z) values between -1.0 and 1.0 and *thresh* between -10000 and 10000. The ranges and precision of the alleles are limited only by the storage capacity of *double* and *long 'C'* data types – no range constraints were set in the code.

Genetic Operators

Genes are real-coded, but genomes may comprise variable numbers of genes. Given two parent genomes, the crossover operator examines each gene of parent1 in turn, finding the most similar gene of the same type in parent 2. Similarity is measured by calculating the differences between values of operator and coding regions of genes. One of the two genes is then randomly allocated to the child. If the genome of parent2 is shorter, the child inherits the remaining genes from parent 1. If the genomes are the same length, this crossover acts as uniform crossover.

Mutation is also interesting, particularly since these genes actually code for proteins in this system. There are four main types of mutation used here:

1. Creep mutation, where (xp,yp,zp) and (x,y,z) values are incremented or decremented by a random number between 0 and 0.5, *Affinity Threshold* is incremented or decremented by a random number between 0 and 16384 and *Concentration Threshold* is incremented or decremented by a random number between 0 and 200.
2. Duplication mutation, where a (xp,yp,zp) or (x,y,z) region of one gene randomly replaces a (xp,yp,zp) or (x,y,z) of another gene. (This permits evolution to create matching promoter regions and coding regions quickly.)
3. Gene mutation, where a random gene in the genome is either removed or a duplicate added.
4. Sign flip mutation, where the sign of *Affinity Threshold* is reversed.

Crossover is always applied; all mutations occur with probability 0.01 per gene.

4 Squareroot Function Regression

Previous work has demonstrated how evolution can generate specific fractal proteins that interact with each other in order to produce desired patterns of activation (Bentley 2004) or to produce a specific set of commands for a robot, to guide it past obstacles (Bentley 2003a). Here, the task is to produce the square root of a number. The input to the system is provided by setting the concentration of the first environment fractal protein (all others have a default value of 200). The output is produced by the behavioural gene(s) as described previously. Each genotype was developed ten times in succession with random input (concentration) values between 0 and 199. The fitness was the sum of the differences between the values obtained and true squareroot of the input.

To evolve the controllers, the fractal development system was initialised with a single cell, 2 environment genes, 2 receptor gene, 2 behavioural genes and 6 regulatory genes. (With variable length genomes, evolution was free to modify these gene

numbers). The operator and coding regions of the genes were randomly initialised with the alleles that defined 10 previously evolved protein fractals (Bentley, 2004). 8 developmental steps were employed (ten times, each with a different environmental protein concentration, corresponding to the ten random inputs), and the evolutionary algorithm used a population size of 100, running for 1000 generations.

To provide some assessment of how effective fractal proteins are in improving performance or evolvability, the same system was also run with all fractal proteins disabled. In this non-fractal version, the triplet of three real values were used directly (the affinity value now defined how small the sum of differences between the cis-region of the gene and the protein should be before the gene is activated). There were no protein-protein interactions; no fractal shapes were calculated, merged or compared. All other parameters were kept the same.

5 Squareroot Function Regression Results

Figure 8 shows the final fitness scores achieved by the fractal developmental system and the system using no fractal proteins, for thirty runs. Fitnesses below 40,000 achieved an acceptable accuracy. It should be clear that the system using fractal proteins achieved acceptable fitnesses in 20 out of 30 runs. The system without fractal proteins only achieved acceptable fitnesses in 7 out of 30 runs. In addition, solution quality often suffers without fractal proteins – no solutions achieved the same accuracy as those produced with fractal proteins.

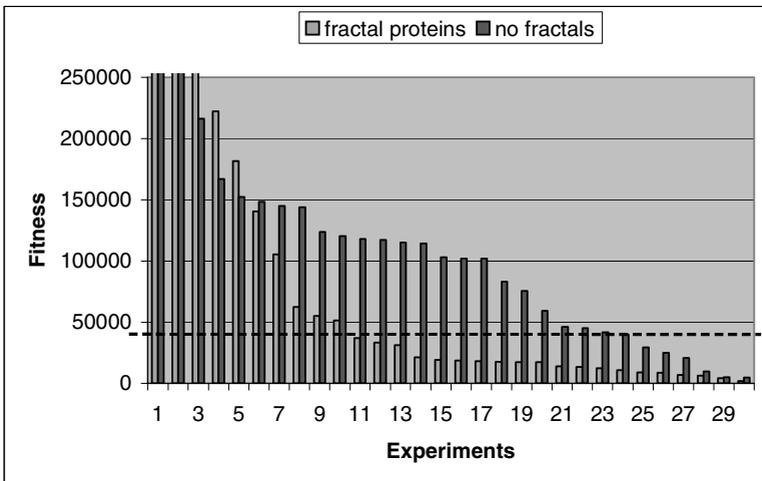


Fig. 8. Fitness of developmental squareroot program using fractal proteins and the developmental squareroot program without fractal proteins. Results are shown sorted into descending order of fitness for clarity and ease of comparison. The dotted line denotes a fitness of 40,000 – any solutions with fitnesses at or below this line are considered sufficiently accurate at calculating the squareroot of the input

The reason for the difference in performance seems to be *evolvability*. Without fractal proteins, solutions become trapped in local optima – GRNs that produce a linear output are common, instead of the required non-linear squareroot curve. Without the ability for new proteins and genes created by evolution to affect existing proteins and genes (through complex protein-protein and protein-gene interactions), there is no way for evolution to overcome the trap. With fractal proteins, evolution is free to add new genes which produce proteins that modify existing solutions subtly and in nonlinear ways. Evolvability is caused by the ability of this representation to enable gradual modifications to any solutions - not just by changing existing genes but also by adding new ones that act in combination with existing ones where necessary. This is evident during evolution as poor solutions gain large numbers of genes, and good solutions prune the genes down to more robust sizes. Without fractal proteins, each gene has a much more binary role - it is either critical to the GRN or has no effect at all - meaning evolution cannot make small changes quite as easily (despite still being able to duplicate genes).

Previous work (Bentley 2003c) has shown other aspects of evolvability: even after evolution has found a perfect solution, it continues to evolve, changing genes, proteins and entire GRNs constantly. This representation enables never-ceasing evolution, which also results in the solutions becoming compact and robust against damage.

6 Damage Tolerant Developmental Programs

Having shown that fractal proteins do convey a significant advantage for evolutionary development, a further experiment was performed in order to assess how fractal developmental programs show graceful degradation, when damaged. Using the results from the previous experiment, the fittest solution was picked (see figure 9a), the evolved fractal proteins were written into the code as a short list of real-valued constants (the x , y , z values described earlier), the genetic algorithm removed, and the resulting fractal developmental squareroot program compiled into a stand-alone executable. (The compiler was GCC 3.3, using xcode on a Mac Powerbook G4, Mac OS 10.3.2.)

6.1 Comparison Methods

Two other programs were used as comparisons in the experiments. The first was a fast squareroot function, written for speed of execution, provided by Hsieh¹. This is written in C and was simply compiled to produce a stand-alone executable. The second was evolved by Landon's simple GP (Langdon 1998)². This standard genetic programming engine used a function set comprising "+", "-", "*", and "/", population size of 100, max program size of 100 nodes, number of generations = 1000, probability of crossover 0.7, and mutation 0.01. Each individual was evaluated by presenting 10 random inputs and calculating the sum of the difference between the outputs and the true squareroot of the inputs. After twenty runs, only five solutions close to the

¹ <http://www.azillionmonkeys.com/qed/sqroot.html#fast>

² <http://www.cs.ucl.ac.uk/staff/w.langdon/ftp/gp-code/simple/simple-gp.c>

more calls to memory-handling routines and library functions, resulting in more code (which was corrupted more), and thus code more likely to fail. To overcome this, the three squareroot programs were combined into a single program. By changing a simple compiler directive, the program could be compiled in three different ways:

1. Output result of method 1 only if methods 2 and 3 executed correctly.
2. Output result of method 2 only if methods 1 and 3 executed correctly.
3. Output result of method 3 only if methods 1 and 2 executed correctly.

where method 1 was the fast squareroot function, method 2 was the GP-evolved function, and method 3 was the fractal developmental squareroot function. This way, all three programs contained the same code with the same susceptibility to damage, except that the code that generated the output was different in each program.

The three executables were corrupted 200 times using the method described previously. The corrupted programs were then executed and the results noted, see table 2.

Table 2. Results of running 200 corrupted executables for three squareroot programs. Graceful degradation is defined as solutions producing 10 non-zero values within 50 percent of the correct values

	Square root	GP square root	Ev. Dev. square root
Fail	197	198	177
Incorrect run	1	0	13
Graceful degradation	0	0	8
Perfect	2	2	2

8 Analysis

The results are fascinating. Despite all three programs suffering from the same proportion and type of errors (see fig. 11), there is marked difference in performance between the developmental squareroot program and the fast squareroot and GP-evolved programs. The latter both only manage to execute correctly 2 out of 200 corrupted executables. They display zero graceful degradation – they simply fail to execute (or in a single case, execute with incorrect results).

The developmental squareroot program manages to run 23 times out of 200. 13 of those produce incorrect results (usually all zeros). Only 2 produce perfect results (as good as the uncorrupted program). But in 8 cases, the developmental squareroot produces *approximately correct* solutions, fig 12. The damage to the executable has perhaps corrupted the genes or fractal proteins, and the developmental program recovers. As was described in section 2, evolution has not only evolved a good solution, it has created a solution that copes with damage. It seems that this protection even extends to damage done to the executable code, as well as simple mutation-driven damage to genes.

Note that the GP-evolved code does not display this property. It is conceivable that should the GP solution contain bloat (unused code), then it might survive damage more readily. However, this would not be the same phenomenon observed in the developmental system, which has no bloat (Bentley 2003c). In the developmental system, damage to the code *that is actually being used*, can be survivable (Bentley, 2003c). It seems probable that only highly evolvable developmental systems enable this kind of “natural” graceful degradation to emerge.

It should also be noted that all three squareroot programs were calculating the squareroot results according to their inputs. The developmental program did not, in any sense, have the answer “wired in” as constants – indeed it performed more calculation using the input to produce the results than the other two programs. The ability to survive damage arises because of the way the calculation was performed – the dynamic (gene) networks in the code are able to survive despite having “holes punched in them” by the corruption program.

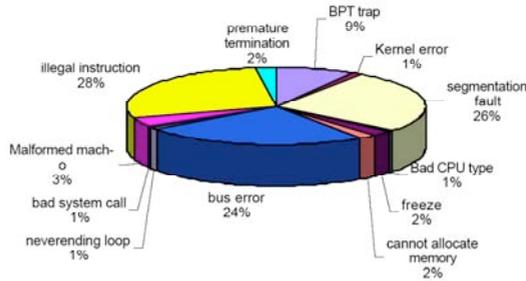


Fig. 11. Percentage & type of errors obtained in all runs of the corrupted programs

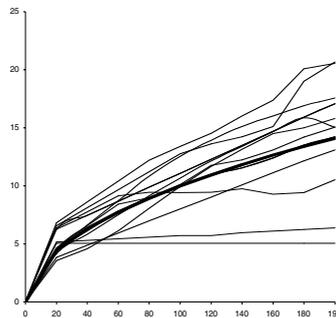


Fig. 12. Outputs produced by different runs of the damaged developmental squareroot function, true squareroot shown in bold. Note that most produce results that are approximately correct (within 2 of the true value), displaying remarkably graceful degradation

9 Conclusions

Development is the process used by evolution to construct complex, adaptive and robust forms. Computer algorithms based on development can share some of these properties. Here, experiments have shown that fractal proteins increase the evolvability of developmental programs by allowing new protein-protein and protein-gene interactions to incrementally modify solutions over several generations. Experiments have also shown that, unlike traditional software, evolved fractal developmental programs show graceful degradation after damage to their executable code. While surviving only around 14 bits (0.05%) of damage 10% of the time is not a great achievement compared to the robustness of natural systems, given the conventional (brittle)

nature of the programming language, compiler and hardware, it is still considered impressive. It seems likely that should computer science remove its brittleness and embrace evolutionary and developmental systems more fully, abilities such as graceful degradation will improve further.

Acknowledgments

This material is based upon work supported by the European Office of Aerospace Research and Development (EOARD), Airforce Office of Scientific Research, Airforce Research Laboratory, under Contract No. F61775-02-WE014. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author and do not necessarily reflect the views of EOARD. MOBIUS is an  project.

References

- [1] Bentley, P. J. Fractal Proteins. 2004. In Genetic Programming and Evolvable Machines Journal.
- [2] Bentley, P. J. Evolving Fractal Gene Regulatory Networks for Robot Control. 2003a. In Proceedings of ECAL 2003.
- [3] Bentley, P. J. Evolving Fractal Proteins. 2003b. In Proc. of ICES '03, the 5th International Conference on Evolvable Systems: From Biology to Hardware.
- [4] Bentley, P. J. Evolving Beyond Perfection: An Investigation of the Effects of Long-Term Evolution on Fractal Gene Regulatory Networks. 2003c. In Proc of *Information Processing in Cells and Tissues* (IPCAT 2003).
- [5] Bentley, P. J. From Coffee Tables to Hospitals: Generic Evolutionary Design. 1999. Chapter 18 in Bentley, P. J. (Ed) *Evolutionary Design by Computers*. Morgan Kaufmann Pub. San Francisco, pp. 405-423.
- [6] A.H. Jackson, A.M. Tyrrell Implementing Asynchronous Embryonic Circuits using AARDVArc. 2002. In *Proceedings of 2002 NASA/DoD Conference on Evolvable Hardware* (EH-2002), IEEE Computing Society, Alexandria, Virginia, pp. 231-240.
- [7] S. Kumar and P. J. Bentley. Computational Embryology: Past, Present and Future. 2003. Invited chapter in Ghosh and Tsutsui (Eds) *Theory and Application of Evolutionary Computation: Recent Trends*. Springer Verlag (UK).
- [8] Langdon, W. (1998) *Genetic Programming + Data Structures = Automatic Programming!* Kluwer Pub.
- [9] Mahdavi S. and Bentley P. J. Adaptive Evolutionary Motion of Smart Robots. 2003. In Proc. of EvoROB2003, 2nd European Workshop on Evolutionary Robotics.
- [10] Mandelbrot, B. *The Fractal Geometry of Nature*. 1982. W.H. Freeman & Company.
- [11] Miller, J. and Banzhaf, W. Evolving the Program for a Cell: From French Flags to Boolean Circuits. 2003. Invited chapter in Kumar, S. and Bentley, P. J. (Eds) *On Growth, Form and Computers*. Academic Press, 2003.
- [12] Shelton, C. & Koopman, P. Developing a Software Architecture for Graceful Degradation in an Elevator Control System. *Workshop on Reliability in Embedded Systems*.
- [13] Thompson, A. Evolving Inherently Fault-Tolerant Systems. 1997. In Proc. Instn. Mech. Engrs 1997.
- [14] Lewis Wolpert, Rosa Beddington, Thomas Jessell, Peter Lawrence, Elliot Meyerowitz, Jim Smith. *Principles of Development, 2nd Ed.* 2001. Oxford University Press.