

Information Integration in Evolutionary Processes

Informatie Integratie in Evolutionaire Processen
(met een samenvatting in het Nederlands)

Proefschrift

ter verkrijging van de graad van doctor aan de Universiteit Utrecht
op gezag van de Rector Magnificus, Prof. Dr. H.O. Voorma,
ingevolge het besluit van het College voor Promoties
in het openbaar te verdedigen op
woensdag 1 december 1999 des middags te 12.45 uur

door

Ludo Wilhelmus Petrus Pagie

geboren op 23 april 1967
te Delft.

Promotor:

Prof. Dr. P. Hogeweg
Faculteit Biologie
Universiteit Utrecht.

The studies described in this thesis were performed at the department of theoretical biology and bioinformatics at Utrecht University. The investigations were financially supported by the Research Council for Earth and Lifesciences (ALW), which is subsidized by the Netherlands Organization for Scientific Research (NWO).

Aan tante Jeanne en Berend Marius

ISBN: 90-393-2277-5

Omslag: John Couwenberg en Ludo Pagie

Druk: Optima, Grafische communicatie

Contents

1	General Introduction	1
1.1	Coevolution	2
1.2	Structured genotype-phenotype mappings	3
1.3	This thesis	5
2	Evolutionary consequences of coevolving targets	9
2.1	Introduction	10
2.2	The Model	11
2.3	Results	13
2.3.1	Success Rates	14
2.3.2	Generalizability	14
2.3.3	Mutational stability	18
2.3.4	Higher dimensional target functions	19
2.4	Discussion	20
2.4.1	Coevolution and optimization	20
2.4.2	Side-effects of variable problem sampling	25
2.4.3	Conclusion	25
3	Colicin diversity: a result of eco-evolutionary dynamics	27
3.1	Introduction	28
3.2	model	29
3.3	results	31
3.3.1	Typical temporal and spatial dynamics	31
3.3.2	Hyperimmunity and multitoxicity: a phase transition	33
3.3.3	Plasmid invasions	37
3.4	Discussion	38
3.4.1	Plasmid diversity	38
3.4.2	The cost of defence	39
3.4.3	Experimental data	39
3.5	Conclusion	40
4	Individual- and population-based diversity in restriction-modification systems	43
4.1	Introduction	44
4.2	The model	45
4.3	Results	46
4.3.1	Individual- and population-based diversity	47
4.3.2	Transition between modes	50
4.4	Discussion	53
4.4.1	Model predictions and comparison to experimental data	53
4.4.2	Frequency-dependent selection and pairwise infectability	54

4.4.3	Cooperation across levels	55
4.4.4	Conclusion	56
5	Coevolutionary dynamics: information integration, speciation, and red queen dynamics	67
5.1	Introduction	68
5.2	The model	69
5.2.1	Spatial embedding and local dynamics	70
5.2.2	Some (evolutionary) properties of the density classification task	71
5.3	Results	73
5.3.1	Two typical simulations	74
5.3.2	Information integration and red queen dynamics	76
5.3.3	Eco-evolutionary side-effects	81
5.4	Discussion	89
5.4.1	Specialists and generalists; neither ‘outfits’ a queen	89
5.4.2	Red queen dynamics; quick, quick, slow	92
5.4.3	Conclusion	93
6	Discussion	95
6.1	A short review	95
6.2	Sparse fitness evaluation	97
6.3	Speciation, optimisation, and red queen dynamics	99
6.4	Information integration	99
6.5	Conclusion	101

1

General Introduction

In studies of evolution different points of view are expressed which impose distinctive characteristics on the evolving system and lead to various dénouements of the evolutionary process. Evolution “*by means of natural selection*” was originally proposed as the origin of species (Darwin, 1859). Since then, as a result of the phrase “*survival of the fittest*”, evolution has been studied as an optimisation process (Holland, 1975; Maynard Smith, 1978). On the basis of the notion of reciprocal adaptations of multiple species, evolution has also been viewed as a source of continuous change, i.e. “*Red Queen dynamics*” (Van Valen, 1973) or “*Arms-Races*” (Dawkins & Krebs, 1979).

In its basic form, i.e. a process of reproduction and selection, evolution is the only theory that provides a potential basis for the overwhelming diversity and complexity that we can observe in nature everywhere. Despite the diversity of phrases describing evolutionary processes, traditional evolutionary theory, e.g. population genetics and optimisation theory (Roughgarden, 1979; Maynard Smith, 1978), has yet not been able to explain the origin, maintenance, and renewal of biological diversity and complexity in a satisfactory way (Maynard Smith & Szathmáry, 1995*a*). This may be largely a result of the fact that evolution is still treated as a single-level process that takes place in the context of static, low-dimensional genotype-fitness mappings which are generally taken to be of a linear nature. The selection of individuals is often based on fitness values that express the individual’s lifetime fecundity in all possible contexts and in all possible situations. Although biological organisms may undergo many ‘evaluation events’ in their lifetime, this number will in general fall well short of ‘full’ fitness evaluation. In studies of coevolution individual fitness is no longer defined a priori but is dependent on the state of a ‘coevolving’ population. However, in coevolutionary studies too the genotype-fitness mappings are unstructured, linear, or absent. Thus, studies of evolutionary systems generally rule out the possibility that structuring will occur in the evolving system and regard structuring merely as a direct consequence of reproduction and selection. Structuring may result from the takeover of a population by individuals with higher fitness, or from the maintenance of polymorphic populations due to coevolutionary selection pressure.

More recent studies of evolution take into account some aspects of biological evolution that are typically lacking in traditional studies, e.g. the multi-level character of biological systems (Hogeweg, 1994; Maynard Smith & Szathmáry, 1995*a,b*; Hogeweg, 1998), self-structuring of biological systems and its consequences for evolving systems (Hogeweg & Hesper, 1991; Boerlijst & Hogeweg, 1991; Savill *et al.*, 1997), non-linear

genotype-phenotype mappings (Schuster, 1989; Huynen *et al.*, 1993; Kauffman, 1993; Huynen & Hogeweg, 1994), or the occurrence of *neutrality* in genotype spaces (Huynen *et al.*, 1996; Huynen, 1996; Van Nimwegen *et al.*, 1999). These studies show that evolution is better viewed as a multi-level informatic process, i.e. a process in which “*information and novelty are created and processed across and between the self-structured multiple levels*” (Savill, 1997).

In this thesis we study an aspect of evolution, i.e. *information integration*, that is concerned with ways in which information present in a system becomes integrated into evolving entities and how this integrated information reflects on the evolving system. Thus, we focus particularly on the storage of information in evolving systems. The studies described in this thesis revolve around three themes. First we study how the process of information integration relates to the degree of completeness of information presentation. Thus, whereas in traditional evolutionary modelling fitness is often (assumed to be) based on a full presentation of environmental conditions per generation, we study here the evolutionary case of very sparse presentation of information, i.e. sparse fitness evaluation. Second, we allow information to be stored in various ways in the system. Thus, a population of evolving individuals that are confronted with a set of environmental challenges can ‘solve’ the problems individually, or as a population. Alternatively, they can restructure the environment and thereby reshape the problem such that they can solve it in ways that were originally not feasible. The third theme that lingers in the projects is the side-effects that the process of information integration can have on the evolving system. Side-effects can occur at the level of individual genomes, e.g. in terms of mutational sensitivity, or at the level of populations, e.g. in terms of the degree of invadability.

1.1 Coevolution

Biological organisms live in an environment that is defined by (the action of) other organisms that evolve themselves. When this aspect of the environment is taken into account the selection pressure that is experienced by organisms can change continuously (Roughgarden, 1979; Futuyma & Slatkin, 1983*a*; Maynard Smith, 1989), making current adaptations possibly futile in future situations. The term “*coevolution*” was first used by Ehrlich & Raven (1964) to describe the mutual evolutionary influence of phytophagous insects and plants. Darwin (1859), however, already recognised the idea of mutual evolutionary adaptations in his discussion of pollination by insects.

In general, the term coevolution is used for systems in which a small number of species impose direct selection pressure on each other (Futuyma & Slatkin, 1983*b*). Maynard Smith (1989) distinguishes three types of coevolution defined on the basis of the interaction between the coevolving species: competitive coevolution, mutualistic coevolution, and exploitative coevolution. The latter type includes predator-prey systems, host-parasite systems, and plant-herbivore systems.

Competitive coevolution occurs between species that are limited by the same resource. Competing species will evolve such that the resource overlap will decrease so as to lower the interspecific competition until the stabilising selection pressure balances the coevolutionary selection pressure (Roughgarden, 1983). This process can occur between

species that have been different since their origin but it can also lead to speciation, the classic example of which are Darwin's finches on the Galapagos Islands (Abbott *et al.*, 1977; Dieckmann & Doebeli, 1999; Sato *et al.*, 1999). Interestingly, Van Valen (1965) already notes that full utilisation of a resource seems to occur either due to utilisation by different species, with each species utilising only a limited variety of resource, or due to one species showing differentiation with respect to the resource such that the complete resource is utilised by the same species. In terms of information integration, the full information content of the resource-variation is reflected either in the variety of species, or in the variety within a species.

Some mutualistic interactions have originated from antagonistic interactions (Barrett, 1983; Feinsinger, 1983; Pellmyr & Leebens-Mack, 1999), e.g. plant-pollinator systems. In the case of such interactions, mutualism between partners can be based on an underlying intrinsic conflict which makes the interaction unstable and seemingly vulnerable to 'cheaters' (Bogler *et al.* (1995), but see also e.g. Pellmyr *et al.* (1996), and Couwenberg (1997)). On the other hand, mutualists that benefit mainly from the relation may evolve such that they maximise the opportunity to form new associations. Thus, once the interaction is established it does not evolve any more (Maynard Smith, 1989). Exploitative interactions may seem most likely to result in continued evolutionary change, i.e. "*Red Queen dynamics*" (Van Valen, 1973) or "*Arms-Races*" (Dawkins & Krebs, 1979) but they may also lead to stable coexistence. In fact, parasite-host relations, for instance, may show a decrease in the virulence of the parasites if within-host populations are monomorphic and the association between the parasite and the host is long lasting (May & Anderson, 1990; Herre, 1993; Lenski & May, 1994) (however, see also (Ebert, 1998; Bergstrom *et al.*, 1999)).

A strict definition of coevolution differentiates between "*diffuse*" and "*pairwise*" coevolution. It describes the degree to which a coevolutionary response in one species depends on one other species or a number of other species. Coevolution is said to be pairwise if traits of two species evolve simultaneously, each in response to the (evolving) state of the trait of the reciprocally evolving species (Janzen, 1980). Otherwise coevolution is diffuse. Sometimes, pairwise coevolution is considered to be the only true type of coevolution (Rothstein, 1990). Biological systems that show strict pairwise coevolution are, however, rare (Farrell & Mitter, 1992). On the other hand, diffuse coevolutionary systems, such as predator-prey systems, or plants and their associated microbial pathogens, are often considered as a source of stable polymorphisms of prey or host populations (Levin & Segel, 1982; Fry, 1990; Iwao & Rausher, 1997; Kraaijeveld *et al.*, 1998; Vogel *et al.*, 1999). The studies in this thesis, however, show that the mode of coevolution, i.e. pairwise or diffuse, can evolve in the same system, starting from a pairwise interaction between two species (chapters 3, 4, 5).

1.2 Structured genotype-phenotype mappings

In traditional theoretical evolutionary studies, e.g. population genetics, the structure of genotypes are generally discarded. Thus, different genotypes are equated with differences in fitness values. However, the relation between a genome and the characteristics that are eventually selected for, or selected against, is often not at all linear. Small

changes at the level of the genotype can result in large differences in fitness and vice versa. Also, some changes in the genotype do not have any effect on the fitness at all.

In work on the evolutionary properties of RNA secondary structures Huynen *et al.* (1993) showed that the mapping from the RNA primary structure to the RNA secondary structure is very heterogeneous (see also Schuster (1989), Huynen & Hogeweg (1994), Fontana & Schuster (1998)). The primary and secondary structures determine the information content and the enzymatic functionality of the RNA molecule respectively (the secondary structure of RNA molecules is highly indicative for the tertiary structure of RNA molecules). RNA molecules show a many-to-many mapping; a particular secondary structure can result from very different RNA primary structures, and closely related primary structures can fold into very different secondary structures. The effect is that almost any RNA primary structure can evolve from almost any other RNA primary structure without changing the secondary structure of the molecule significantly, and vice versa. These properties seem to make RNA primary-secondary mapping very attractive as a genetic coding scheme.

Selfstructuring of genotypes can occur due to the redundancy which is prevalent in the RNA primary-secondary mapping. For instance, RNA molecules which experience a selection for a particular secondary structure may also evolve their primary structure. The latter may result from as a side-effect of the evolutionary dynamics rather than from direct selection. In Huynen & Hogeweg (1994), for instance, it was shown that as a result of selfstructuring RNA molecules evolve toward flatter parts of the genotype landscape, i.e. parts of the landscape where the RNA molecule becomes less sensitive to changes in its primary structure. Other examples of the self-structuring of genotypes are, for instance, Takumi & Hogeweg (1998), Hogeweg (1994), and the studies described in this thesis in chapters 2 & 5.

With the advent of genetic algorithms (Holland, 1975; Goldberg, 1989; Mitchell, 1996) and evolutionary algorithms (Rechenberg, 1973; Schwefel, 1977) evolutionary principles are used as optimisation technique. In such explicit genotypic simulation models structured genotypes are used. That is, the genotype is a structured set of genes, generally of fixed size. The genes can assume different 'trait-values', which can change through genetic operators such as point-mutation or cross-over. A particular realization of a genotype is then mapped to a fitness value by a 'fitness-function'. By performing subsequent rounds of selection, growth, and mutation to such structured genotypes, the genotype space is searched for genotypes with high fitness values.

From the point of view of evolution as an optimisation process the fitness-function encompasses all knowledge concerning the problem for which a solution is searched. However, the particular choice and structuring of the genes that make up the genotype of the individuals may strongly influence the 'success' of the evolutionary process as an optimisation process. This choice determines the fitness landscape that is to be searched, and thereby the ease of searching. Also, the simulation parameters such as population size (which is generally fixed in evolutionary optimisation models), mutation rates, or selection schemes can determine how successful the evolutionary process will be in terms of finding individuals with optimal genotypes.

1.3 This thesis

In this thesis a number of eco-evolutionary models are studied that focus on different aspects of the integration of information in evolving systems. All models employ a form of coevolution, ranging from diffuse coevolution to pairwise coevolution. Also, all studies are based on individual-oriented, spatially explicit models. The spatial embedding allows for spatial pattern formation and, thus potential for the emergence of extra levels of selection. We study two biological systems (chapters 3 & 4) and two artificial systems (chapters 2 & 5). The former allow for variable population sizes, resulting in more realistic population dynamics. The latter two models are based on evolutionary optimisation models; they are based on fixed population sizes and individuals interact solely on the basis of their behaviour with respect to externally defined, artificial ‘evolutionary goals’. In the artificial systems we employ structured representations of the genomes, i.e. functional lisp expressions in chapter 2, and bit strings in chapter 5. The corresponding genotype-phenotype mappings are non-linear, redundant, and allow for selfstructuring at the level of the genome. In chapters 3 & 4 we study bacterium systems in which individual bacteria can acquire genetic elements, e.g. plasmids, which code for various types of colicins (chapter 3) or restriction modification systems (chapter 4). A collection of genetic elements can be interpreted as (an extension of) the bacterium genome. Here, we can directly study the integration of information in terms of the accumulation of ‘genes’.

In chapter 2 we study information integration in evolution under coevolving, sparse fitness evaluation and under static, complete fitness evaluation. The ‘search space’ in which evolution takes place, i.e. the ‘evolutionary goal’ or the ‘set of fitness cases’, is identical in both cases. Under sparse fitness evaluation individuals, or ‘solutions’, are evaluated on the basis of small subsets of the complete set of fitness cases. Complete fitness evaluation is characterised by a ‘full’ evaluation on the basis of all fitness cases during each evaluation event. The evolutionary goal, which ‘drives’ the evolutionary process, is an optimisation problem which is chosen arbitrarily; it has no effect on the behaviour of the individuals; it only serves as a basis for the fitness evaluation. In the case of sparse fitness evaluation the coevolution of the subsets, or ‘problems’, is based on previous work by Hillis (1992); the solutions and the problems have an antagonistic coevolutionary relation. First of all, we show that sparse evaluation of fitness leads to the integration of the complete evolutionary goal in individuals, even if only 9 out of 65,536 fitness cases are presented per evaluation event. Secondly, we show that individuals that evolve under static or under sparse fitness evaluation differ in a number of properties; sparsely evaluated individuals are more robust at the phenotypical level, but less robust at the genotypic level. These results show that different instantiations of the process of information integration can approximately integrate the same information in individuals but it can have side-effects that substantially influence the evolutionary process. Finally, we show that when the magnitude of the total amount of information contained in the environment becomes larger, coevolution rather than random selection of fitness cases is necessary if sparse evaluation is to lead to complete information integration. Thus, evolution under sparse fitness evaluation can integrate an amount of information that is at least of the same order of magnitude as the amount of information that can be integrated in evolution under complete fitness evaluation; sparse fitness evaluation does

not hinder evolution in this respect a priori. Moreover, sparse fitness evaluation has side-effects that potentially help the evolutionary process, e.g. decreased mutational stability or increased phenotypical robustness.

In the next two chapters we focus on the second theme of the thesis, i.e. the possibility of storing information in various ways in an evolving system. In chapter 3 we study the evolutionary dynamics of a bacterium population in which the bacteria can carry different types of colicin complexes, i.e. gene-complexes that code for a toxin and a corresponding antidote. Colicinogenic bacteria can kill competing bacteria when the latter do not carry the corresponding antidote. On the other hand, bacteria suffer a reduction in their growth rate for each colicin type they carry. The two processes that define the benefits and costs of the colicin complexes for the bacteria, i.e. killing competitors or suffering a decrease in growth rate, respectively, occur on different time scales; the ability to kill neighbours gives rise to much more pronounced differences in competitive propensity than differences in growth rate. When we study the evolutionary dynamics of the model for different values of the growth penalty per colicin type we find two qualitative modes of information integration, an individual-based mode and a population-based mode. The individual-based mode, which occurs for low values of the growth penalty, is characterised by a bacterium population in which all bacteria carry all antidote types that are present and only a small number of toxins or no toxins at all. The population-based mode is characterised by a heterogeneous population in which all bacteria carry only complete colicin complexes, i.e. the toxin and the antidote parts. In the population-based mode, per bacterium the number of different colicin types is much less than the total number present in the population. In the two modes, however, the total number of colicin types present in the population is equal. Thus, although at the level of the population the total amount of ‘information’, i.e. the number of colicin types, is equal in both modes, at the individual level it differs greatly. In addition, the different modes of information storage influence several aspects of the ecological dynamics, and therewith the evolutionary disposition of the bacteria.

The model that we study in chapter 4 is similar to the colicin model: the evolutionary dynamics of restriction-modification (RM) systems in a bacterium population that is infected by phage. In this model bacteria can accumulate different RM-types to protect themselves against phage-infected neighbouring bacteria. Similarly, phages can accumulate specific modifications that render them insensitive to corresponding RM-types. Contrary to the colicin model, however, in this model the phage population and the bacterium population can integrate independently the modifications and RM systems, respectively. The individual-based mode and the population-based mode occur also in this model, but, whereas the colicin model shows a phase transition between the two modes dependent on the value of the growth penalty per colicin type, the RM model shows bi-stability of the two modes. For a large parameter range the bacteria exhibit both modes as a possible ‘answer’ to the threat of phage infections. The individual-based mode can be interpreted as a form of optimisation of individual protection, whereas the population-based mode is a form of speciation of bacteria as a response to environmental circumstances. Surprisingly, in the individual-based mode, the phages are not hindered at all by the RM systems carried by the bacteria; the phages have also accumulated all modifications. In the population-based mode, on the other hand, the bacterium population is polymorph with respect to RM-types, which prevents the phage from ac-

cumulating modifications; phages are fully sensitive. Although individual bacteria seem to be much more vulnerable they are in fact much less infected by the phage. In the colicin model and in the RM model we find an individual-based and a population-based mode of information integration, although in the first model only one of the two modes, dependent on the growth penalty, is stable whereas in the RM model the two modes are bistable, apparently because in the presence of the phage population the latter model has one more dynamical level.

In chapter 5 we study the different ways in which a system which consists of two co-evolving species can self-structure as a response to each other's evolution. We find that optimisation, speciation, and red queen dynamics all occur in the same model for small parameter changes, or for particular structural changes of the model. Again we place the evolutionary process in the context of artificial coevolutionary optimisation, similar to chapter 2. The external goal is the density classification task of cellular automata, which make up the first population, with respect to their initial conditions, the second population. Traditional evolutionary studies tend to classify the outcome of evolutionary dynamics in terms of either speciation (i.e. differentiation), optimisation (i.e. maximisation), or continued evolutionary change (i.e. red queen dynamics). In this model we find that these outcomes are often connected. Also, the distinction between *pairwise* and *diffuse* coevolution fades when we examine the interactions of individuals as they occur in the model. The primary distinction that we study in the model is evolution under spatial pattern formation and evolution under global mixing of the populations. In the first case generalised individuals evolve that implement reasonable solutions to the external evolutionary goal, whereas in the second case we find red queen dynamics in which both populations evolve from one 'dumb' phenotype to the other 'dumb' phenotype. Although the 'smartness' of the cellular automata, as defined by the externally imposed criteria, differs very much in the two models, the 'local fitness', as it 'emerges' in the evolutionary process, is the same in both models when averaged over time. In the first model general solutions evolve in the presence of a speciated population of initial conditions, but also the population of cellular automata shows speciation. In fact, in the population of initial conditions shows speciation to occur at more than one level. It seems that it is precisely the speciation, or localised specialisation, that steers the evolution of the cellular automata towards general solutions. In addition, the local specialisations appear to evolve continuously, as in red queen dynamics. On the other hand, in the mixed model evolution produces 'optimal' red queens, i.e. queens that evolve fast. Thus, speciation, optimisation, and continued evolutionary change play a role in both cases, although in different degrees. Nevertheless, the outcomes of the two evolutionary processes are quite different.

2

Evolutionary consequences of coevolving targets

Ludo Pagie and Paulien Hogeweg
*Theoretical Biology and Bioinformatics group,
Utrecht University, Padualaan 8,
3584 CH Utrecht, The Netherlands.*

Evolutionary Computation 5(4) (1997), 401-418

abstract

Most evolutionary optimization models incorporate a fitness evaluation which is based on a predefined static set of test cases or problems. In the natural evolutionary process selection is of course not based on a static fitness evaluation. Organisms do not have to combat every existing disease during their lifespan; organisms of one species may live in different or changing environments; different species coevolve. This leads to the question how information is integrated over many generations.

This study focuses on the effects that different fitness evaluation schemes have on the types of genotypes and phenotypes that evolve. The evolutionary target is a simple numerical function. The genetic representation is in the form of a program (i.e. a functional representation, as in genetic programming). Many different programs can code for the same numerical function. In other words, there is a many-to-one mapping between 'genotypes' (the programs) and 'phenotypes'. We compare fitness evaluation based on a large static set of problems and fitness evaluation based on small coevolving sets of problems. In the latter model very little information is presented to the evolving programs regarding the evolutionary target per evolutionary time step. In other words the fitness evaluation is very sparse. Nevertheless the model produces correct solutions to the complete evolutionary target in about half of the simulations. The complete evaluation model on the other hand does not find correct solutions to the target in any of the simulations. More important, we find that sparse evaluated programs are better generalizable compared to the complete evaluated programs when they are evaluated on a much more dense set of problems. In addition, the two evaluation schemes lead to programs that differ with respect to mutational stability; sparse evaluated programs are less stable than complete evaluated programs.

2.1 Introduction

Evolutionary optimization processes are based on the biological evolutionary process (Goldberg, 1989; Holland, 1992). Most artificial evolutionary models however include a static fitness evaluation function, which clearly does not exist in the natural evolutionary process. In nature the fitness of an individual depends in many ways on non-static features. Organisms live in different environments and interact and coevolve with other organisms.

In 1991 Hillis presented a spatially embedded coevolutionary optimization model in which sorting algorithms coevolved with sorting problems (Hillis, 1992). The sorting algorithms were evaluated on the basis of local problems instead of on the basis of a globally defined static set of problems. Since the 'evolutionary goals' of the sorting problems were opposed to those of the sorting algorithms but were nevertheless dependent on each other, Hillis called the sorting algorithms *hosts* and the sorting problems *parasites*. He found that the incorporation of the coevolving problems resulted in better (i.e. faster) sorting algorithms than those that evolved with a static set of sorting problems. Hillis attributed the difference in the success of the two evaluation schemes to two properties of the coevolving scheme. First, the coevolution of the parasites prevented the population of hosts from 'getting stuck' on local optima. Second, the search process was more efficient because the coevolving parasites focused on those problems that had not yet been solved correctly.

Several authors have studied optimization models in which the fitness evaluation of individuals depends on other individuals in the same population or in other populations. They have reported that such models yield higher fitness values and involve lower computational costs than traditional evolutionary optimization models or other optimization techniques. Three main forms of coevolutionary models can be distinguished, although many variants are used.

- Host-parasitoid models in which candidate solutions are evaluated on the basis of small subsets of a data set which defines the evolutionary target. The subsets coevolve with the candidate solutions (Hillis, 1992; Paredis, 1994, 1995).
- In competitive evolutionary models candidate solutions compete with each other in game-like tournaments. The fitness of the solutions depends on the ratio of wins and losses in these tournaments¹ (Angeline & Pollack, 1993; Rosin & Belew, 1997; Sims, 1994; Juillé & Pollack, 1996).
- In cooperative evolutionary models individuals of several different (coevolving) species are combined before they are evaluated with respect to an evolutionary target¹ (Husbands, 1994; Potter & De Jong, 1994; Vafaie & De Jong, 1996; Potter *et al.*, 1995).

Other studies, however, have shown that in some cases coevolution does not lead to better results (Thompson, 1996).

¹Both competitive and cooperative evaluation schemes are used in models in which the evolutionary target is predefined and thus static, and in models in which the evolutionary target is defined solely with respect to the behavior of the opponent or the cooperator.

Here we present results of a study in which we compare static fitness evaluation and sparse, coevolving fitness evaluation of candidate solutions, the latter being similar to the model studied by Hillis. Our model is based on a simple evolutionary optimization process. We specify an external optimization problem or *evolutionary target* which is defined with respect to a so-called 'complete' set of test cases or *problems*. The fitness evaluation of the statically evaluated solutions is based on this 'complete' set of problems, whereas the partial fitness evaluation of the coevolving solutions, or *hosts*, is based on coevolving subsets of the 'complete' set. The evolutionary process is placed in a 2-D space which leads automatically to a tournament-like selection process as in the competitive evolutionary models mentioned above. Since we use a static evolutionary target we can easily compare static fitness evaluation and coevolving sparse fitness evaluation in terms of, for instance, optimization time and correctness of solutions. Another important advantage of using a static evolutionary target instead of a more open-ended target is that we can easily study how information regarding the target is integrated over evolutionary time.

The 'complete' set of problems does not change during evolution. The parasites on the other hand can mutate, changing some of the problems. It is important to note that the problems contained in the parasites are elements of the 'complete' set of problems. Thus in both fitness evaluation regimes the solutions can attain maximum fitness by solving all problems of the complete set. However, the (partially evaluated) hosts can also attain maximal fitness by solving only those problems of the complete set on which they are actually evaluated. Of course the parasites that specified those problems will have minimum fitness and will thus quickly be outcompeted or mutated.

We code the solutions which are to approximate the evolutionary target in functional form (i.e. as a program) as in genetic programming. Such coding leads to a multitude of implementations of any one function. The resulting many-to-one mapping between genotypes, i.e. programs, and phenotypes, i.e. programs evaluated on a certain set of problems, influences the evolutionary process considerably (c.f. Schuster (1989); Huynen *et al.* (1993); Altenberg (1994); Hightower *et al.* (1995)). The reverse is also the case: different evolutionary processes may consistently lead to different types of genotypes (Huynen & Hogeweg, 1994). In our model the two fitness evaluation schemes lead to programs which markedly differ at the genotypic level in terms of mutational stability, as well as at the phenotypic level in terms of generalizability.

2.2 The Model

We studied the two different fitness evaluation schemes in the context of several evolutionary targets. In most of these studies the same trend is visible. First we will describe a model that incorporates one particular evolutionary target. In section 2.3 we will discuss results obtained from this one model. After that we will briefly describe results of models with higher dimensional evolutionary targets.

The evolutionary target that we used in the following model is a simple 2-D numerical function: $\frac{1}{(1+X^4)} + \frac{1}{(1+Y^4)}$, see fig. 2.1. The problems on which the fitness evaluation of the solutions is based are simply X, Y values. The problems of the complete set are regularly distributed over the problem domain; 26×26 problems in the




Figure 2.1: A plot of the target function, with a view from a) the side, b) the bottom and c) from the top. The X,Y domain is $[-5,5][[-5,5]$, the maximum of the function approaches 2.0, the minimum ($X=0, Y=0$) is 0.

domain $X = -5.0, 5.0$ and $Y = -5.0, 5.0$ with an interval of 0.4. The problems of the coevolving sets of problems are elements of the complete set. Obviously the target function does not represent any biological function. However, using this artificial numerical function we can easily study the two different fitness evaluation schemes. In addition, the generalizability of evaluated solutions can easily be studied by changing the set of problems for which they are evaluated. It is also clear that the 'complete' and the coevolving sets of problems (the latter contains only nine problems, see below) 'cover' the domain of the target function in different ways. Whereas the 'complete' set of problems covers the domain adequately, the coevolving sets clearly do not.

We embedded the populations of solutions and parasites in space. We used a 2-D toroidal square lattice, with one solution and one parasite per grid cell. The size of the lattice is 50×50 cells, giving population sizes of 2500. Competition for growth is local in space. In each 3×3 neighborhood the solutions are ranked based on their fitness value in ascending order. The i^{th} ranked solution is selected with a probability of $(\frac{1}{2})^i$. The selected solution will grow into the central cell of the nine cells under consideration. The same growth procedure affects the parasites in the coevolutionary case, except that there the ordering is reversed; parasites are ranked in descending order.

The fitness of a solution is defined as the mean of the absolute differences between the target function and the solution over all problems on the basis of which it is evaluated. A solution is considered completely 'correct' if, for all 676 problems in the 'complete' problem set used in the static evaluation scheme, the absolute difference between solution and target function is less than 0.01 (this is a so-called hit).

In the static evaluation scheme the fitness of solutions is based on all 676 problems of the complete set of problems, whereas in the coevolving evaluation scheme the fitness of solutions is based on the 9 problems of the parasites in the surrounding 3×3 neighborhood. The fitness of a coevolving parasite is defined as the absolute difference between the target function and the solution (in the same grid cell as the parasite) evaluated on the basis of that parasite. Since the parasites are confronted with only one host they are more likely to be affected by random fluctuations, such as mutations, in that host. The fitness evaluation of hosts, on the other hand, is based on nine parasites, so changes in one parasite have a less drastic effect on the fitness evaluation of the host. We found that

this asymmetric fitness evaluation gives better results with respect to optimization time, than the symmetric fitness evaluation.

The parallelization of the evolutionary process resulting from the spatial embedding leads to an increase in the genetic diversity of each population and thereby possibly to enhanced performance of the optimization process (Collins & Jefferson, 1991). Combined with the localized interactions within and between populations the spatial embedding can lead to specialization of the coevolving populations with respect to each other (Husbands, 1994) and thus to the natural incorporation of features such as niching and sharing (Rosin & Belew, 1997; Mahfoud, 1995). Clearly, sharing also occurs automatically in the statically evaluated model as a result of the spatial embedding.

The genetic representation of the solutions is based on genetic programming. That is, the function that we use set is composed of: $\{+, -, *, \%\}$. The division operator $\%$ is said to be protected in the sense that division by zero gives 1.0. In genetic programming the division operator is normally implemented in this way to ensure that the programs maintain syntactic closure under the genetic operators (Koza, 1992). The terminal set is composed of: $\{X, Y, \mathfrak{R}\}$, where \mathfrak{R} is the ephemeral random constant (Koza, 1992). Note that one does not have to use constants in order to create a correct program. The constant 1.0 in the target function can easily be obtained by dividing one variable by itself. In fact, not all functions in the function set are needed to create a correct program; the target function can be implemented with only division, plus either addition or subtraction. The use of superfluous function and terminal sets increases the number of possible implementations of a correct solution, and thus the freedom of the evolutionary process to 'choose' a program. We used crossover and point mutations as genetic operators, with probabilities of 40 and 20 per cent respectively.

The genotypes of the parasites, which in this model specify only one X,Y-problem, are simply the values of the variables. Mutation of a parasite means that one of the variable values is changed into a neighboring value, i.e. plus or minus 0.4. Ten per cent of all parasites are mutated every time step. The genotype space of the parasites is not toroidal: parasites with extreme variable values (-5.0 or 5.0) can only mutate in one direction.

Simulations are started with small, randomly created programs of maximum depth 3 and, in the case of the coevolutionary regime, parasites with the values $X=0.2, Y=0.2$. Simulations were stopped either when no correct solution was found within 500 time steps, or when a correct solution was found and retained in the population for 50 time steps. The demand for retention of the correct solution for 50 time steps was based on the idea that the coevolving evaluation scheme might produce a correct solution but could not keep it in the population due to lack of selection for complete correctness. As it turned out, this did not happen in any of the simulations; in fact solutions defined as correct were often (slightly) improved upon during these final 50 time steps.

2.3 Results

We report on 20 simulations for each of the two types of fitness evaluation schemes with the 2-D evolutionary target: complete static and coevolving sparse evaluation. During the simulations we recorded the fitness of the best solution based on the 'complete'

evaluation scheme	size of problem set	success ratio	mean number of nodes in final program
static	676	0%	68
coevolving	9 out of 676	45%	44

Table 2.1: The success ratios and mean size of the final solutions for the different evaluation schemes.

problem set (the so-called total fitness) at that time step (best in the sense of total fitness) in order to compare the different evaluation schemes. The success rates of the two evaluation schemes are described in the following section. Thereafter, we will investigate results concerning the generalizability of the evaluated programs of the different evaluation schemes and describe the differences in their mutational stability.

2.3.1 Success Rates

Figure 2.2 shows the fitness curves of the best-of-generation solution for all simulations of the two evaluation schemes. Table 2.1 shows the percentage of simulations that produce correct solutions for the two evaluation schemes. The correct solutions found by the coevolving evaluation scheme come in two varieties. A number of solutions have fitness values ranging from 10^{-15} to 10^{-17} , while the others have fitness values ranging from 10^{-2} to 10^{-4} . These fitness values reflect whether the solution is an exact or only an approximate implementation of the target function. Solutions that implement the target function exactly can be easily rewritten in the same form as the target function. Although such 'perfect' solutions may still contain constants, these constants are not functional in the evaluation of the program. They are either multiplied by zero or two equal constants are subtracted from each other. The constant 1.0 in the target function is formed by the term X/X or Y/Y . The fact that the fitness values are larger than zero is a consequence of the finite numerical precision of the fitness calculation. Solutions which approximate the target function (i.e. correct solutions with fitness values in the range 10^{-2} , 10^{-4}) still contain constants that affect the evaluation of the solution.

The two evaluation schemes differ considerably in the number of problems evaluated per fitness evaluation; completely evaluated solutions compute $676/9 \approx 75$ times more problems than the coevolved solutions. Furthermore, the completely evaluated solutions are larger than the coevolved solutions (table 2.1); thus every single evaluation takes longer. On the other hand, the coevolution of problems takes some time too, as does the periodic computation of the total fitness of the coevolving solutions. Nevertheless, we found that simulations with complete fitness evaluation required 5 to 10 times more computer time than the simulations with coevolution.

2.3.2 Generalizability

We studied the generalization capabilities of the evolved solutions by increasing the sampling density of the program evaluation. The standard 'complete' static evaluation is based on 26×26 problems; the dense evaluation is based on 100×100 problems.

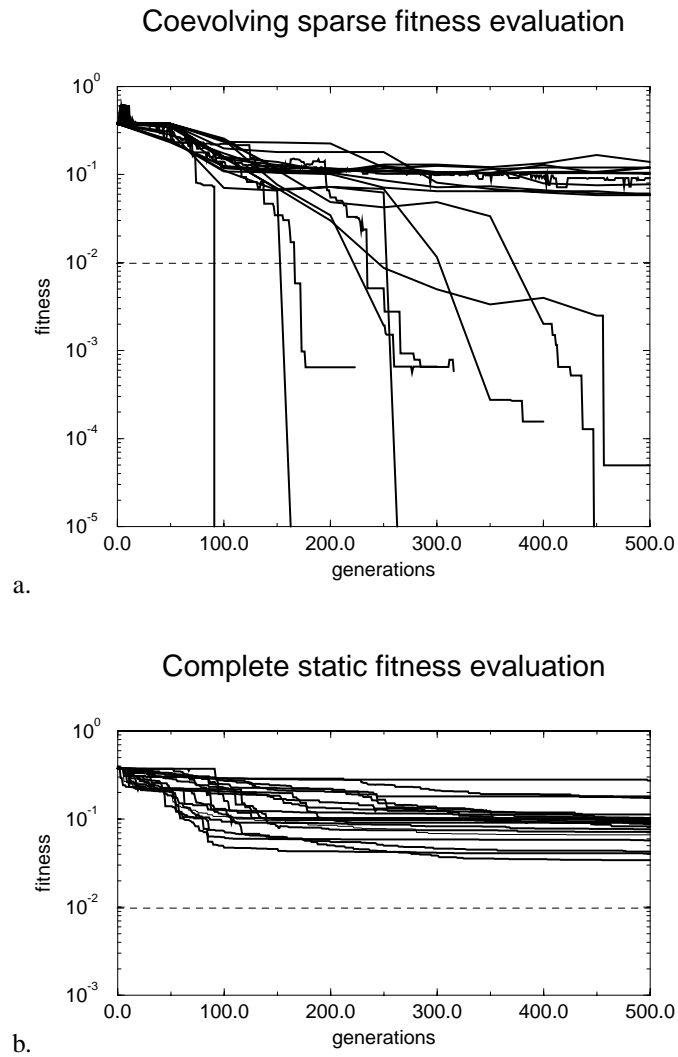


Figure 2.2: Fitness curves of the best-of-generation solution for coevolving a) and complete static problem evaluation b). Fitness is based on the 'complete' problem set that consists of 26×26 problems. The fitness curves that drop below 10^{-5} go to values between 10^{-15} and 10^{-17} . The horizontal dotted lines give the value of the hit criterium (see text).




Figure 2.3: Three typical final solutions produced by coevolving fitness evaluation. The left plots are based on 26×26 evaluated problems, the right plots on 100×100 evaluated problems. a) and b) are two correct solutions that approximate the target function c) is an incorrect solution. All solutions generalize well on the 100×100 problems.

Perfect solutions implement the target function exactly. Thus, plots of a perfect solution evaluated based on the dense set of problems are identical to plots of the target function (see fig. 2.1). Figures 2.3a and b show the standard (left) and dense (right) calculated evaluations of correct solutions from two simulations with the coevolving evaluation scheme. Neither solution is an exact implementation of the target function but is an approximation thereof. The right plot of fig. 2.3a shows that although the solution is not perfect it generalizes very well over the points that are not included in the set of problems used in the simulation. The dense plot of fig. 2.3b shows that the solution does not generalize very well for a small subset of the dense set of problems (i.e. $X = 0$). Figure 2.3c shows the standard and dense calculated evaluations of an incorrect solution. Although the solution is not correct it is surprising to find that it nevertheless generalizes so well.

If we study the generalization capabilities of evolved solutions from simulations with the static evaluation scheme we get quite different results. Figure 2.4a and b show plots of standard (left) and dense (right) calculated evaluations of the best solutions of two simulations. Although the left plots show that the programs yield approximately correct values for the set of X, Y values on which they evolved, the right plots show that they generate absurd values for intermediate X, Y values.




Figure 2.4: Two typical final solutions produced by static fitness evaluation. The left plots are based on 26×26 evaluated problems, the right plots on 100×100 evaluated problems. Neither solution is correct.

The solutions produced by the static evaluation scheme seem to get trapped in the complexity of the functions they implement. In the process of further adaptation to the target function, due to the 'complete' sampling of the problem domain, the solutions are forced to conserve any adaptation already achieved. The 'complete' sampling results in a severe selection against the occurrence of errors. Although the increase in the numerical complexity of the solutions gives the programs the opportunity to adapt to individual problems and thus increase their fitness the increased complexity reduces the generalizability of the solutions.

The solutions produced by the coevolving evaluation scheme have more freedom to make errors. As long as an error in the solution is not exposed by the local problems the solution can remain in the population. However, errors in solutions are simple evolutionary targets for the coevolving problems. Thus, sooner or later, errors in solutions will be selected against and these solutions will be expelled from the population.

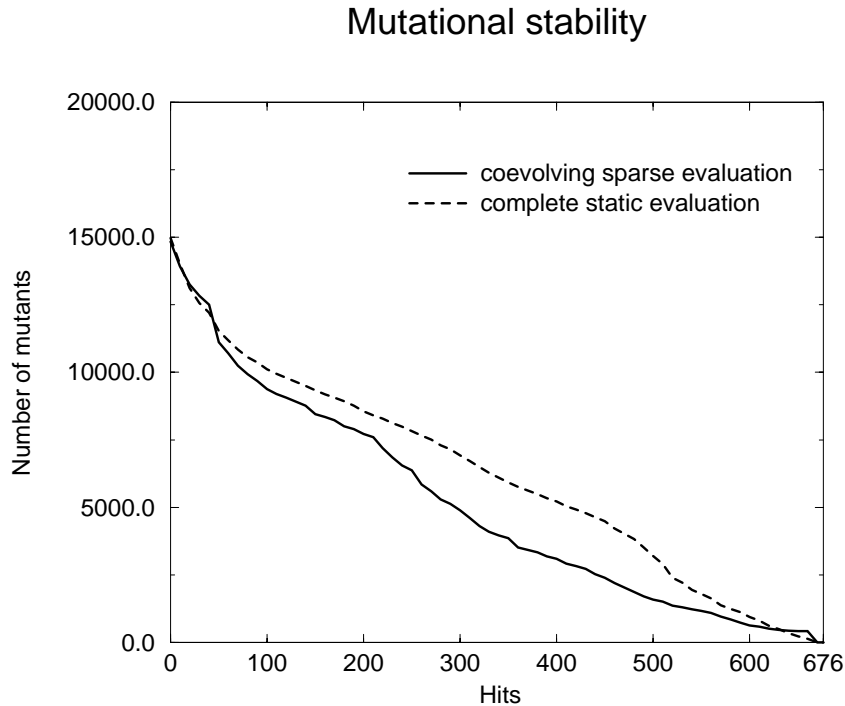


Figure 2.5: Histogram of the number of one-point mutants having at least x number of hits.

2.3.3 Mutational stability

In order to study the effects that the two evaluation schemes have on the structures of the programs, we compare the final programs with 1000 of their one point mutants, i.e. programs that differ from the final program by one point mutation. The original and mutant programs are compared for all 676 problems in the standard problem set. If the absolute difference between the original program and the evaluated one-point mutant is less than the hit criterion for a certain problem (an absolute difference of less than 0.01) the mutant scores a hit on that problem. Thus, if many mutants score many hits the original program is phenotypically stable under mutations. For each evaluation scheme we looked at 1000 one-point mutants of the final program for all simulations. In figure 2.5 we plot the similarity between all 20,000 mutants and their respective final programs for the two evaluation schemes. On the horizontal axis we plot the number of problems for which a mutant scores a hit, or in other words: the number of problems for which the mutant is near identical to its original final program. On the vertical axis we plot the number of mutants that score at least x number of hits. It is clear that statically evaluated programs are more stable than the programs of the coevolving evaluation scheme (significance < 0.01 on the Kolmogorov-Smirnov test).

It is not true that the difference in stability depends on the fact that the coevolving

evaluation scheme produces correct programs whereas the static evaluation scheme does not. If we look only at the coevolved programs we find no difference in the stability of correct and incorrect programs. We also compared the mutational stability of ten incorrect programs for both evaluation schemes but found no significant difference between those and the results shown here. The difference in the mutational stability of the programs that have evolved under the different evaluation schemes does not depend on the difference in the size of the programs either. If we compare the mutational stability of the final programs for one evaluation scheme we find no correlation between size and stability.

2.3.4 Higher dimensional target functions

We also studied the 3-, and 4-D extensions of the 2-D evolutionary target which are simply extensions on the numerical function: $\frac{1}{(1+X^{-4})} + \frac{1}{(1+Y^{-4})} + \frac{1}{(1+Z^{-4})} + \dots$, etc. In order to economize on computational time and resources we had to limit the problem domain to $X = -3.0, 3.0$ and $Y = -3.0, 3.0$. The interval was 0.4 again, thus limiting the number of values along one dimension to 16. In the 3-D case the total number of problems is 4096, in the 4-D case the total number is 65536. In the 3-D case we increased the lattice size to 100×100 cells, in the 4-D case to 150×150 cells. We found that the coevolutionary evaluation scheme did not find correct solutions in the smaller field of the 2-D target but it did in the larger lattices. Except with respect to these changes the model was identical to the 2-D model. With the 3-D target we ran 5 simulations of the static evaluation scheme and 10 runs of the coevolutionary evaluation scheme. The 4-D case was run only for the coevolutionary evaluation scheme since complete evaluation is extremely time consuming. These coevolutionary simulations were run for 500 time steps without complete evaluation. After that we performed a sparse evaluation on the basis of random sets of problems in order to get an indication of the total fitness of the solutions. Only if this indication was positive did we search for a correct individual.

In general the results of these simulations are very similar to those with the 2-D target. The static evaluation scheme does not produce any correct solutions in the 3-D case, whereas the coevolutionary evaluation scheme produces correct solutions in 5 out of 10 runs. The generalizability and mutational stability of the solutions in the 3-D case give similar results to the 2-D case. Statically evaluated solutions are much less generalizable and are mutationally more stable than coevolved solutions. However we found in rare cases that coevolutionary evaluation may also evolve to a mutationally stable program. This happens, for instance, in very unsuccessful runs if the program yields a constant value independent of the problem values on the basis of which it is evaluated. Contrary to the mutationally stable programs that evolved under static evaluation the coevolved stable program is very generalizable since it simply yields a constant value.

We performed only three runs of the coevolutionary evaluation scheme with the 4-D evolutionary target, due to long simulation times. In these three runs we found two correct solutions. We have not attempted any further extensions of this evolutionary target since this would require an extensive increase in computational resources. It is clear, however, that coevolutionary fitness evaluation can still produce correct solutions under an increase in the dimensionality, and thereby in complexity, of the evolutionary

target. Furthermore, the evolution of the different types of solutions is independent of the dimensionality of the evolutionary target. In all models we see that the coevolved solutions are more generalizable and mutationally less stable.

2.4 Discussion

In the previous section we showed that the evaluation based on small coevolving sets of problems and evaluation based on a large static set of problems differed in their success rate; the coevolving evaluation gives correct solutions in roughly half of the simulations, whereas 'complete' static evaluation does not produce correct solutions in any of the simulations. Even more importantly we showed that the types of solutions that evolved under the different evaluation schemes differ markedly. The coevolved solutions are more generalizable, less complex and mutationally less stable than the statically evaluated solutions. These differences between the solutions reflect the multiplicity of the coding of phenotypes in genotypes; genotypes which differ in generalizability, mutational stability and complexity can map to similar phenotypes, i.e. phenotypes that approximate the same evolutionary target. Although the degree to which the solutions differ in these properties may depend on model-specific properties, the qualitative difference is consistent in all simulations.

In the following sections we will discuss the results of the model. First we will discuss some aspects of the coevolution of solutions and problems in terms of an evolutionary optimization process. Thereafter we will discuss the effect of the two evaluation schemes on the type of solutions that evolved.

2.4.1 Coevolution and optimization

Several authors have suggested that coevolutionary fitness evaluation is more successful than 'complete' static fitness evaluation because the coevolving problems sample the problem domain more efficiently. (Hillis, 1992; Paredis, 1994, 1995). The idea is that the parasites sample particularly those problems in the problem domain that have not yet been solved by the solutions; as a result the fitness evaluation process becomes focused on 'hard' problems. In order to study the effect of this focusing on the coevolving problems in our model we also studied a variant of this model. In this variant the solutions are evaluated on nine problems (as in the coevolving model) which are randomly chosen from the 'complete' set at every fitness evaluation. Thus here the evaluation is sparse but does not coevolve. First, however, we will discuss briefly a few results from a similar but much simpler model in order to establish a baseline with respect to the efficiency of coevolutionary optimization. In this model the genotype-phenotype mapping is a one-to-one mapping; contrary to the previous model, the evolutionary target here can be represented in only one way in the genotype.

Coevolution towards a simple linear evolutionary target

In this model we use bitstrings as the representation for the solutions, an arbitrarily chosen bitstring as evolutionary target, and a simple additive fitness function for the




Figure 2.6: Coevolving (solid line), static (dashed line) and random (dotted line) fitness evaluation of a simple linear problem. Optimization time in terms of bit evaluations. Data are averaged over 5 runs.

fitness evaluation. In other words this is a simple imitation function. Except with respect to these aspects this model is the same as the original model.

The complete set of problems here consists simply of all bits in the string, each specified by its position. In the complete static evaluation scheme² solutions are evaluated on the basis of all positions in the bitstring. The corresponding fitness landscape is smooth with one global maximum. In the coevolving evaluation scheme each parasite specifies three positions in the bitstring. Hosts are evaluated based on 9 parasites again, thus they are evaluated on the basis of 27 positions. We also studied random sparse evaluation. In this model each solution is evaluated based on 27 randomly chosen positions at every fitness evaluation. The latter two models have fitness landscapes that are more rugged compared to the first fitness landscape. The ruggedness reflects the sparse and dynamic nature of the fitness evaluation in these models.

We performed simulations of this evolutionary model for target strings of different lengths; 256, 512, 1024 and 2048 bits, each with several population sizes. The data are averaged over five runs. For all fitness evaluation regimes we record the evolutionary time required until a bitstring is found that is completely correct. Obviously in terms of

²Note that complete evaluation in this model is truly complete.

evolutionary time the static evaluation scheme needs fewer steps than the coevolutionary sparse evaluation scheme. However, the number of bit evaluations, and thus the computational cost per evolutionary step, differ greatly for the different evaluation regimes. For the coevolutionary and the random evaluation scheme the number of bit evaluations per evolutionary step is fixed. For the static evaluation scheme the number of bit evaluations is equal to the length of the target bitstring.

In fig. 2.6 we plot for each target string length the number of bit evaluations required until a correct bitstring is found. Now we see that for each target string length there is a population size for which the coevolutionary evaluation scheme is more efficient than the most efficient population size of the static evaluation scheme. In fact the range of population sizes for which this is true is quite large. The static evaluation scheme on the other hand is characterized by the very small range of relatively good population sizes. The random evaluation scheme shows the same trend as the coevolutionary runs for shorter string lengths. However, for longer strings random evaluation performs less efficiently. For the longest target string, i.e. 2048 bits, the random evaluation scheme does not find a correct solution in any of the runs.

For the shortest string length, i.e. 256 bits, a single random evaluation evaluates 10% of the complete problem set. Note that solutions are evaluated independently but selected with respect to each other. Also they have a certain lifetime during which they are evaluated on the basis of different random sets of problems. Thus, for shorter strings the random evaluation scheme resembles more and more the complete evaluation scheme. The parasites in the coevolving evaluation scheme on the other hand are much more structured, in time as well as in space. The temporal change in parasites is gradual via mutations. The spatial structuring is a result of local growth. Both properties lead to a variation in the evaluated problems, which is expected to be much smaller in the coevolutionary case than in the random evaluation case.

On the other hand, for longer target strings the random evaluation performs less efficiently. The failure of the random evaluation scheme in the 2048 bit target string runs is due to the fact that the population of solutions cannot retain already acquired information about the target function. The average fitness of the population of solutions levels off fairly quickly without ever producing a completely correct solution. In fact for longer target string lengths the solutions rise above the error threshold (Eigen *et al.*, 1989). It is not clear to what extent the coevolving parasites introduce extra information on which the evolutionary selection pressure can act. The results presented here show that the efficiency increase due to 'focused' evaluation (which happens in the coevolving but not in the random evaluation scheme) is visible only for larger evolutionary targets.

A surprising result is that for both the linear and the nonlinear models (see 2.3.4) the coevolutionary evaluation scheme seems to work best if the population size is of the same order as the size of the 'complete' set of problems. We have not pursued this finding further but think that future studies should investigate this relation.

From this simple model we can draw several conclusions. First, evaluation on only a small part of the evolutionary target can nevertheless lead to integration of the complete evolutionary target in the solutions, given a sufficiently large population and local competition. Second, although the coevolving evaluation scheme needs much larger populations than the population in the static evaluation regime the difference in the computational cost nevertheless favors the coevolutionary evaluation scheme due to the sparse

evaluation per individual per time step. Third, random sparse evaluation performs similarly to coevolutionary sparse evaluation with respect to smaller evolutionary targets. For large target strings random partial fitness evaluation *falls* over the error threshold. Coevolving evaluation on the other hand can still provide a high enough selection coefficient in these cases.

Sparse and dense random fitness evaluation; efficiency and focusing

In the original model we also studied a sparse random evaluation scheme. In the sparse random evaluation model the fitness of the solutions was evaluated on the basis of nine random problems. The random problems were chosen from the standard 'complete' problem set, and were chosen anew for every fitness evaluation. In 20 simulations 7 correct solutions were found; thus the difference between the success rate of this random evaluation scheme (35%) and that of the coevolving evaluation scheme (45%) is small. The final solutions of this random evaluation scheme are comparable to the final solutions of the coevolutionary evaluation scheme in the sense that both sets have good generalizing capabilities and are mutationally unstable. In the 2-D model the coevolving evaluation scheme is not much more efficient than the random evaluation scheme. In the previous section we showed that for relatively small evolutionary targets random fitness evaluation performs similarly to coevolutionary fitness evaluation with respect to optimization time. For larger targets coevolutionary evaluation outperforms random evaluation due to the fact that the coevolving parasites focus on hard problems. In fact, the random evaluation scheme in the original model with the 3-D evolutionary target cannot find correct solutions for this larger evolutionary target.

Although in the coevolutionary 2-D model the focusing of the parasites does not lead to a large increase in success rate, the selection of fit parasites does lead to focused sampling of the problem domain. Figure 2.7 shows that the coevolving problems focus on difficult regions during evolution. After allowing for an initial transient of 100 time steps we counted the number of times that a particular problem was present in the problem population and plotted these values for all problems in the domain. The plot shows that particularly the centre of the domain and the regions around $X=0$ and $Y=0$ are sampled by the coevolving problems. The effect on the success rate caused by the focusing of the parasites is probably negligible in the 2-D model because it is counteracted by the poorer variety of problems that are present in the parasite population, as mentioned previously.

The solutions that have evolved under the coevolutionary and random evaluation scheme are more generalizable and mutationally less stable than the solutions evolved under the complete static evaluation scheme. However, the evolution of generalizable and mutationally unstable solutions does not necessarily result in correct solutions.

A second variant that we studied is an evaluation scheme in which the solutions are evaluated on the basis of 676 randomly chosen problems which are not elements of the 'complete' set of problems. These randomly chosen problems are uniformly distributed over the domain. In this case evaluation is not sparse but it does vary over time. The mutational stability of the programs produced in these simulations is even lower than the stability of the coevolved programs and the generalizability is similar to that of these programs. Nevertheless, only two correct solutions were found in 20 simulations. Thus,




Figure 2.7: The coevolving problems focus on difficult regions in the problem domain. The height of each point depicts the total number of occurrences of that problem in the problem population during the evolution. The problems focus on the regions near $X=0$ and $Y=0$.

generalizability and mutational stability can evolve without increasing the success rate appreciably.

An important factor of the coevolving and the random evaluation scheme is that the solutions are evaluated on the basis of a very small subset of the 'complete' set of problems. This increases the freedom of the evolutionary process to traverse the space of possible solutions. The dense random evaluation scheme samples the evolutionary target to such an extent that solutions cannot make large mistakes easily. A clear example of the freedom in the evolutionary process is the following simulation taken from the coevolving evaluation scheme. During this simulation we traced the ancestry of the correct final solution. The final solution was perfect in the sense that it was a direct implementation of the target function $\frac{1}{(1+X^{-4})} + \frac{1}{(1+Y^{-4})}$. The parent of this final solution had as second term: $\frac{1}{(1+Y^{-3})}$, the grandparent had as second term: $\frac{1}{(1+Y^{-2})}$. The parent of the final solution has a total fitness of the order of 10^{15} . In fact, at the moment when the final solution was produced, the parent had a local fitness³ about 10^{15} . Before that, however, the parent had a low local fitness of the order of 10^{-3} . This is the result of the sparse sampling of the problem domain. The grandparent and the parent solutions were produced at a moment when the local problems were at the edge of the domain, i.e. in the flat regions. In this part of the domain the transition of the second term from $\frac{1}{(1+Y^{-2})}$ to $\frac{1}{(1+Y^{-3})}$ does indeed result in lower fitness values. From that point onwards parasites mutated towards the centre of the domain, resulting in the exposure of the error in the second term. By that time the final solution had

³Local fitness is based on only the (nine) problems on the basis of which a solution is evaluated in the coevolutionary evaluation scheme.

been produced with the correct second term. Thus the sparseness of the evaluation helps (rather than hinders) the search process.

2.4.2 Side-effects of variable problem sampling

Clearly, generalizability is a very important property of evolving entities, be they solutions to optimization tasks or biological organisms. In our model we see a clear relation between the generalizability of the solution and the type of fitness evaluation we use. Generalizability appears to be a side-effect of the evolutionary dynamics resulting from fitness evaluation on the basis of a varying set of problems. Generalizability is of course a good strategy for coping with changing environmental conditions, here in the form of changing sets of problems. Thus, although there is no direct selection for generalizability, variable problem sampling may indirectly select for solutions which are generalizable.

Just as generalizability is not directly selected for in the coevolving and the random evaluation schemes, mutational stability is not directly selected for either. Nevertheless, solutions that are evaluated on the basis of coevolving or randomly selected problems are consistently mutationally less stable than statically evaluated solutions.

Mutational instability is a second possible strategy for counteracting changing environmental conditions. A mutationally unstable program is better able to adapt to new conditions than a mutationally stable program (this is also discussed by Thompson (1996) in the context of error correction by evolving entities). A decrease in genetic stability as a response to changing environmental conditions in RNA landscapes has been reported previously by Huynen et al. (1993). It is interesting to see that evaluation under variable “environmental” conditions (i.e. problem sets) produces solutions which implement the two strategies for coping with variable evaluation, namely mutational instability and generalizability.

Note that mutational stability and generalizability are two properties that render a solution robust. Mutational stability reflects genetic robustness; small changes in the genotype have little effect on the phenotype. Generalizability is a phenotypic measure of robustness; small changes in environmental input produce similar phenotypic responses. In our model we see that solutions are either robust in the sense of being generalizable, or they are robust in the sense of being mutationally stable, but not both. It is not clear whether mutational instability and generalizability are properties that are necessarily linked. Even if they are not it might be that neither of these properties is attainable without the other.

2.4.3 Conclusion

With respect to evolutionary optimization processes we see that in our model the coevolution of problems and solutions does indeed yield better results than complete static fitness evaluation. This implies of course that sparse dynamic fitness evaluation can result in the complete integration of an evolutionary target. This is not at all a trivial finding. An important aspect of the success of coevolutionary evaluation is in fact the sparse sampling of the problem domain. We have shown that sparse sampling gives the solutions the opportunity to make large errors with respect to some problems as long as

the solutions are not evaluated on the basis of these problems. This gives the evolutionary process greater freedom to explore the space of possible solutions.

More interesting than the difference in the success rate is our finding that the co-evolutionary and the static evaluation schemes evolve different types of solutions. The coevolutionary evaluation scheme leads to solutions that are more generalizable, mutationally less stable and less complex than the solutions produced under the static evaluation scheme. The differences in the properties of the solutions that evolve under the different fitness evaluation schemes are side-effects of evolutionary dynamics; none of these properties has a direct effect on the fitness of the solutions and thus is not directly selected for. An evolutionary process which acts on genotypic representations that incorporate a many-to-one genotype-phenotype mapping can mould the genotypes in different ways. Nevertheless, the different genotypes can implement phenotypes that approximate the evolutionary goal to a similar extent. Such side-effects can occur only in genetic coding schemes that incorporate a multiple mapping from genotypes to phenotypes, as is the case in ,for instance, genetic programming or the natural genetic code, but is generally not the case in classical genetic algorithms. With respect to evolutionary optimization models which do incorporate such a multiple genotype-phenotype mapping this result suggests that a simple change in fitness evaluation can produce more generalizable solutions.

With respect to the biological evolutionary process, it is clear that neither the 'complete' static, nor the sparse variable set of problems serves as a good approximation of the fitness evaluation in nature. However, we consider the static completeness and the variability of fitness evaluation as a continuous transition from static total sampling to variable sparse sampling. Our results suggest that this transition has a large impact on the genotypic structures that evolve. Static sampling leads to complex and mutationally stable solutions with low generalizing capabilities. Variable sparse sampling on the other hand leads to much less stable and simpler solutions with high generalizing capabilities. With regard to the natural evolutionary process, our results show that not all properties of evolving entities are the result of direct selection. Many properties, such as being robust on one level or the other, can be side-effects of the genotypic structuring that results from several aspects of the evolutionary process.

Acknowledgements:

The authors thank Ms. S.M. McNab for linguistic advice. They also thank the referees for constructive criticism and helpful suggestions. The investigations were supported by the Life Sciences Foundation (SLW), which is subsidized by the Netherlands Organization for Scientific Research (NWO).

3

Colicin diversity: a result of eco-evolutionary dynamics

Ludo Pagie and Paulien Hogeweg
*Theoretical Biology and Bioinformatics group,
Utrecht University, Padualaan 8,
3584 CH Utrecht, The Netherlands.*

J. Theor. Biol. **196** (1999), 251-261

abstract

Colicins are plasmids that are carried in *Escherichia coli*. They code for a toxic protein and for proteins that confer on the host immunity against this toxin. When bacteria carry plasmids their growth rate is reduced. At the same time, the production of toxins makes it possible for colicinogenic bacteria to invade bacterium strains that are not immune. In natural bacterium populations there is a high diversity of colicin types. The reason for the maintenance of this diversity has been the subject of much recent debate.

We have studied a simple eco-evolutionary model of the interaction of bacteria with colicins and show that high diversity of colicins is to be expected. We find two different dynamical modes each with a high diversity: a *hyperimmunity mode* and a *multitoxicity mode*. Bacteria are immune to most toxins in the first mode but in fact produce very few toxins. In the second mode bacteria are immune only to those toxins that they actually produce. In the second mode toxin levels per bacterium are much higher, whereas immunity levels per bacterium are lower.

3.1 Introduction

Colicin plasmids are carried in *Escherichia coli* and related bacteria (Pugsley, 1984; Konisky, 1982; Luria & Suit, 1987; Hardy, 1975). The plasmids code for a toxin that kills bacteria. The plasmids also code for proteins that neutralise the toxic protein, thereby conferring immunity on bacteria that carry the plasmid. Many types of colicin plasmids exist; they are generally cross-sensitive (but see (Cooper & James, 1984; Lau *et al.*, 1984) for exceptions).

Natural *E. coli* populations are often found to carry colicin plasmids (Hartl & Dykhuizen, 1984; Riley & Gordon, 1992; Achtman *et al.*, 1983; Singh *et al.*, 1989). The degree of polymorphism of colicins found in natural bacterium populations is much higher than one might expect (James *et al.*, 1996; Riley, 1993; Dykes & Hastings, 1997). On the one hand, plasmids reduce the lowered maximum growth rates of their host. As a result plasmid-carrying bacteria should be outcompeted by wild-type bacteria. On the other hand, colicinogenic bacterium strains produce toxins that kill closely related bacteria. This gives colicin producing bacterium strains an advantage in that they produce ‘anti-competitor’ molecules that kill bacteria competing in the same environment for the same resource. The question remains why the observed polymorphism of colicin plasmids is so high. Many attempts to answer this question are based on the additional fitness-enhancing properties that are supposedly carried by the plasmids (Bouma & Lenski, 1988; Dykes & Hastings, 1997; Feldgarden *et al.*, 1995). Coexistence of several types of colicinogenic bacteria would in this case be possible in habitats with different ecological niches (Frank, 1994*b*).

Previous studies on the competition between bacteria that carry different plasmids have shown that the outcome of competition in an aqueous culture, i.e. a well-mixed environment, depends on the initial frequency of the competing strains (Chao & Levin, 1981; Levin, 1988). Competition, however, always yields one winner. In spatially explicit environments, on the other hand, colicinogenic bacteria have an overall competitive advantage over wild-type bacteria, even if the colicinogenic bacteria are initially rare (Chao & Levin, 1981; Iwasa *et al.*, in press). Of course, the natural habitat of *E. coli* (the lower intestine of warm-blooded animals) is certainly a spatially structured environment (Hartl & Dykhuizen, 1984).

Durrett & Levin (1997) have shown that in a spatially explicit but otherwise single niche model it is possible for wild-type bacteria and colicinogenic bacteria to co-exist. The condition for coexistence is that different types of bacteria exist such that each bacterium strain is capable of outcompeting another strain. In the model developed by Durrett and Levin two colicinogenic bacteria co-exist with one wild-type bacterium. The two plasmid types are cross-immune but one plasmid type produces less or no toxin and thereby imposes a lower growth penalty on its host bacterium. The latter, *cheating* plasmid outcompetes the more toxic plasmid but is in turn outcompeted by the wild-type bacterium. Although this transitive cycle of bacterium states does not result in the stable coexistence of the three bacterium types in a well-mixed environment, stable coexistence is easily achieved in a spatial explicit environment.

Experimental data suggest that many bacteria are immune to colicins that they do not produce (Riley & Gordon, 1992). This immunity without toxin production may be the result of carrying plasmids that code for a second immunity protein but do not code

for the corresponding toxin (Cooper & James, 1984; Lau *et al.*, 1984), or bacteria can become insensitive to colicins through an alteration in the membrane receptors (Luria & Suit, 1987; Hardy, 1975). Bacteria that are immune to toxins they do not produce themselves can play a role that resembles that of the ‘cheating’ plasmids in the model of Durrett and Levin.

We study a spatially explicit multi-plasmid eco-evolutionary model of bacteria and colicin plasmids in order to investigate further whether space is a sufficient condition for the maintenance of a high diversity of colicin plasmids and how the interaction between plasmids and bacteria affects their ecological dynamics. We will show that the dynamics is such that a high diversity of plasmids is easily maintained in a single niche.

3.2 model

The model is based on the interactions between bacteria and colicin plasmids, implemented in a cellular automaton. The model-bacteria are very simple entities that grow and die, depending on their own growth dynamics and on the presence of plasmids within the bacteria and in the 8 neighbouring bacteria. The model-plasmids grow and die along with their host bacterium. In addition they can be generated and deleted by means of mutations.

The model contains 9 different types of plasmids, i.e. 9 different toxin-immunity pair. We allow no cross-immunity between plasmid types. If a bacterium carries a plasmid type we do not specify how many plasmids of that type are present per bacterium. Plasmids are generated through mutation with a rate u_g and are deleted with a rate u_d (see Table 3.1).

A single plasmid specifies two genes: a toxin gene and an immunity gene. The two genes can mutate independently, i.e. the state of a gene can be active or inactive. Thus, plasmids can be in four states, only two of which will be considered. In a plasmid both toxin and immunity genes can be active or a plasmid can have an active immunity gene only. In the first case the host bacterium produces the toxic protein and the proteins that confer immunity. In the latter case the bacterium produces only the immunity proteins. The case in which only the toxin gene is active is not taken into consideration since it is not a viable situation; the bacterium host would be killed by the plasmid instantaneously. Since the case in which both genes are inactive is equivalent to the case in which the plasmid itself is absent the cases will be treated identically. Thus, a bacterium can carry between 0 and 9 toxin genes and between 0 and 9 immunity genes. Of course, a viable bacterium will carry at least as many immunity genes as it carries toxin genes and possibly more.

Although in the remainder of the paper we will generally talk about (plasmids carried by) single bacteria the model could just as well be interpreted at the level of monomorphic bacterium strains. In fact, *natural* bacterium hosts that secrete toxin die in the process as result of cell lysis. Thus, being toxic to conspecifics could only be attained by ‘suicide’. However, in a *natural*, monomorphic strain of colicinogenic bacteria only a small proportion of the bacteria produce toxin and are thereby killed (Pugsley, 1984; Luria & Suit, 1987). This renders the strain as a whole toxic to (non-immune) neighbours but imposes only a relatively small penalty on the growth rate per bacterium since

Parameter	Value
grid-size	300×300
bacterium death rate	0.1
add. death per toxin	0.3
immunity growth penalty (P_i)	0.02 - 0.20
toxin growth penalty (P_t)	0.02 - 0.07
bacterium growth rate	$1 - G_t \times P_t - G_i \times P_i$
degenerative mutation rate (μ_d)	10^{-4}
generative mutation rate (μ_g)	10^{-7}
colicin types	9

Table 3.1: Default values for the parameters used in the simulations. G_t is the number of toxin genes in a bacterium, G_i is the number of immune genes per bacterium. See text for explanations of parameters.

only a few bacteria in the strain actually produce toxin and die. This situation is effectively modelled here; by assuming that the growth dynamics and mutation events affect the strains as a whole we simplify the model structure by interpreting the residents in single sites as single bacteria. Any group selective effects that single bacteria may experience as members of a strain are incorporated in the model through the effect of local growth which leads to patches of identical bacteria¹.

The cellular automaton is a 2-dimensional rectangular grid of 300×300 sites with periodic boundary conditions. The neighbourhood in which interactions between bacteria take place consists of the 8 nearest neighbours of a site. A site is either empty or occupied by a bacterium. An empty site can be colonised by the 'offspring' of a bacterium in a neighbouring site, resulting in two identical bacteria. The growth rate R_b of (a strain of) wild-type bacteria, i.e. bacteria that do not carry any plasmids, is equal to one. If bacteria carry plasmids the growth rate is lowered by the number of toxin genes and immunity genes multiplied by the toxin penalty and the immunity penalty respectively (see also Table 3.1). Actually, the effect of plasmids on the growth rate of their host may take many forms but in order to make the model as general as possible and to avoid any a priori non-linearities we used simply a linear relation.

The probability that a bacterium will grow in an empty site is equal to its growth rate divided by 8. We have used different rules to model the colonisation of empty sites; a cumulative colonisation rule and an probabilistic colonisation rule but found no differences in the qualitative results of the model. The results reported here are based on the cumulative colonisation rule (probability of colonisation is equal to the sum of the growth rates of all 8 neighbours).

There is a 10% probability that bacteria will die regardless of the plasmids that they carry. If in the neighbourhood of a bacterium toxins are produced for which the bacterium does not carry immunity genes its probability of death is increased by 30% per

¹Colicinogenic bacteria are sometimes considered a model for altruistic behaviour; toxin producers commit suicide for the benefit of their conspecifics. Here, as in other models of altruistic behaviour, the spatial embedding results in the possibility of the evolution and maintenance of altruistic behaviour.

toxin gene. The results of the model do not depend much on the actual parameter values. We also studied the model with for instance bacterium death rates as low as 1% and as high as 30% and mutation rates that differed several orders of magnitude from the one used here. However, in all cases the qualitative result of the model was the same as we report here. Table 3.1 gives the default parameter values, or range of parameter values that we used in the simulations reported here.

3.3 results

We have studied the model for a large variety of parameter values. We found that the behaviour of the model can best be characterised in terms of two dynamical modes which we call *hyperimmunity* mode and *multitoxicity* mode. The two modes are determined primarily by the value of the growth penalty parameter (P_i). In section 3.3.2 we will show that the transition from one mode to the other with P_i is very abrupt. First, however, we will describe two typical simulations of the model.

3.3.1 Typical temporal and spatial dynamics

Figure 3.1 shows time plots of simulations with growth penalties a) $P_i = 0.06$ and b) $P_i = 0.10$ ($P_t = 0.02$ for both simulations). The results are representative for the two dynamical modes.

The two upper panels in each figure show the bacterium population size. In panel 4 the population is partitioned into bacteria that carry 0,1,...,9 number of immunity genes, in panel 3 the population is partitioned into bacteria that carry 0,1,...,9 number of toxin genes. The two lower panels show cumulatively the filled plots of the number of toxin genes (panel 1) and the number of immunity genes (panel 2) of all plasmid types. Simulations are started with 3 out of 9 plasmid types present in the population. The plasmids have either only an immunity gene or both immunity and toxin genes.

After the transients (in Fig.3.1a after $t \approx 60000$, in Fig.3.1b after $t \approx 5000$) several distinctive features are visible. In Fig.3.1a ($P_i = 0.06$) a bacterium typically carries between 6 and 9 immunity genes and between 0 and 3 toxin genes. Thus, bacteria carry many more immunity genes than toxin genes. In the lower panels we can see that at the population level the amount of toxin is very low and the immunity level is relatively high (note the difference in scale). This situation is very stable; all immunity gene types are always present, toxin gene types may temporarily disappear from the population but this happens very rarely.

In Fig.3.1b ($P_i = 0.10$) the situation is quite different; a bacterium carries on average between 4 and 6 immunity genes and between 3 and 6 toxin genes. Here, bacteria carry just about as many immunity genes as toxin genes. In the lower two panels we see that compared to Fig.3.1a the total level of toxin genes is much higher and the total level of immunity genes is much lower in the population. The temporal dynamics of plasmid types is much less stable; toxin genes, and even immunity genes, frequently disappear from the population.

Thus, in the first case (Fig.3.1a) all bacteria carry (almost) all immunity types but carry only a few or even no toxin genes. In the second case bacteria carry far fewer

Figure 3.1a: Transient to hyperimmunity mode

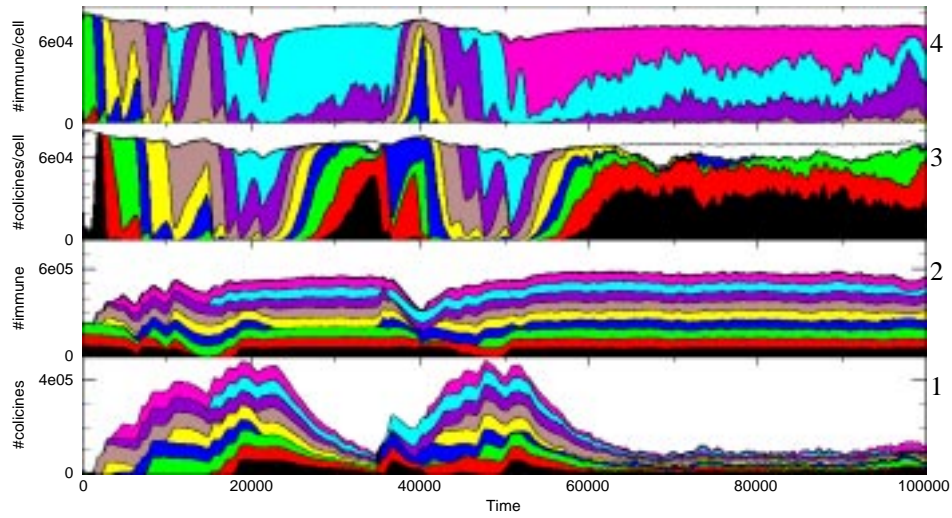


Figure 3.1b: Transient to multitoxicity mode

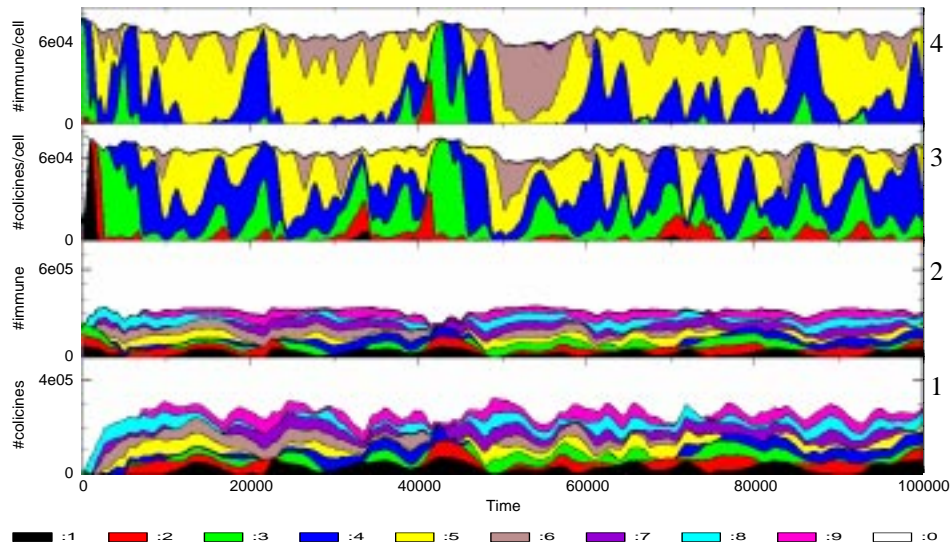


Figure 3.2: Spacetime plots of hyperimmunity and multitoxicity




Figure 3.1: See also colour plate 1 (page 59). Temporal dynamics in a) hyperimmunity mode ($P_i = 0.06, P_t = 0.10$) and b) multitoxicity mode ($P_i = 0.10, P_t = 0.10$). The two lower panels show cumulatively the filled plots of the number of toxin genes (panel 1) and the number of immunity genes (panel 2) per plasmid type (1 through 9). The two top panels show the bacterium population size, partitioned into bacteria that carry the specified number of toxin genes (panel 3) and immunity genes (panel 4), i.e. the distribution of toxin load and immunity load per bacterium (0 through 9). In panel 1 & 2 the colours denote different plasmid types, in panel 3 & 4 the colours denote the number of (toxin/immunity) genes carried per bacterium. In the hyperimmunity mode the transient is relatively long compared to the transient in the multitoxicity mode.

immunity genes but many more toxin genes than in the first case. The two dynamical modes also differ in how they are affected by invasions of new plasmids, in the transient from the initial situation, and in the temporal and spatial dynamics of different bacterium types. We return to these issues in Sect.3.3.3. First we show that the dynamical modes as shown here reflect instances of two different phases in the model behaviour which are separated by a sharp transition.

3.3.2 Hyperimmunity and multitoxicity: a phase transition

In Fig.3.3 we plot the size of the bacterium population and the mean number of toxin and immunity genes per bacterium for a range of growth penalties P_i . The figure shows three panels for different values of P_t . For all values of P_t a phase transition occurs for $P_i \approx 0.07$. The dynamics in the model to the left and right of the phase transition is characterised by:

- the bacterium population size; on the left it decreases slightly with P_i while on the right it is independent of P_i .
- the mean number of immunity genes per bacterium; this number is independent of P_i on the left but decreases with P_i on the right.
- the mean number of toxin genes per bacterium; this number increases somewhat with P_i on the left but decreases with P_i on the right.
- the difference between immunity level and toxin level per bacterium; bacteria on the left have a much higher level of immunity than level of toxin; on the right the two are about the same.

Figure 3.2: See also colour plate 1 (page 59). Space-time plot showing invasions of colicins. Different colours denote bacteria that carry different combinations of plasmid genes. a): Invasion of plasmid results in complete turn over of the population (parameters as in Fig.3.1a). b): Invasion remains localized (parameters as in Fig.3.1b). c): Same state as in b) except all bacteria are white; boundaries between different patches are often large gaps as a result of mutual killing.

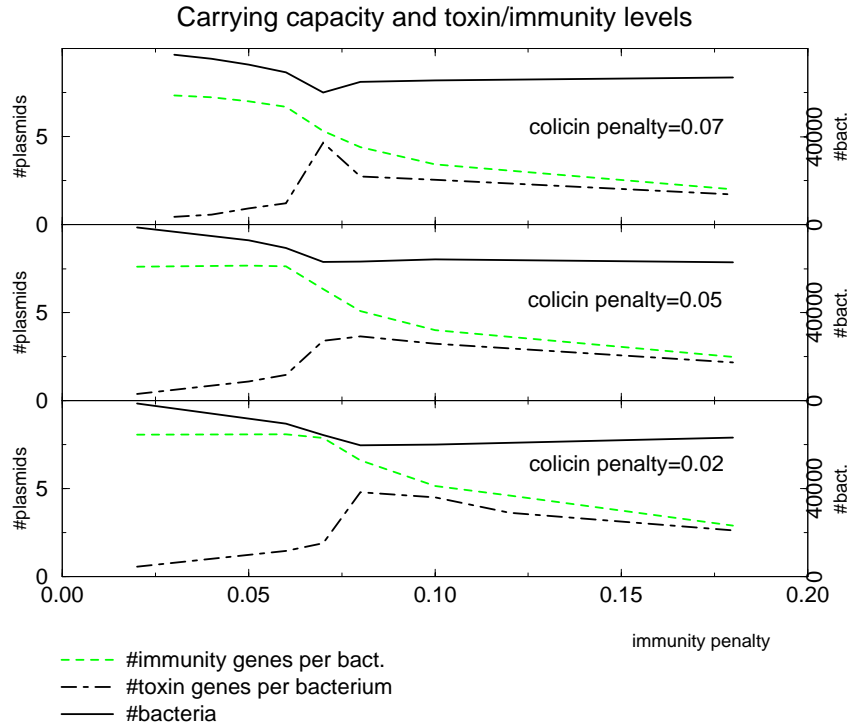


Figure 3.3: Population size, immunity levels and toxin levels for a range of immunity growth penalties P_i . From top panel to bottom panel: $P_T = 0.07$, $P_t = 0.05$, $P_t = 0.02$. For low P_i the bacteria are *hyperimmune*; bacteria have high immunity levels and low toxicity levels. For high P_i the bacteria are *multitoxic*; bacteria have lower immunity levels but higher toxicity levels.

- at the transition the toxin level rises abruptly.

Thus, for low P_i an increase in the growth penalty of the plasmids results in a decrease in carrying capacity of the bacteria while the number of toxin genes per bacterium rises and the number of immunity genes per bacterium remains equal. Here, it are the bacteria that 'pay' the price of an increasing cost of the plasmids. For high P_i , however, the increase in penalty leads to a decrease in immunity while the carrying capacity of the bacterium population remains equal. In this case it are the plasmids that pay the price.

In Fig.3.4 we plot the diversity of bacterium types. A bacterium type is defined by the ensemble of plasmid genes that it carries. Given that a plasmid can be in three states (absent, toxic-immune, or immune-only) the upper bound for total number of possible bacterium types is $3^9 = 19,683$. However, this upper bound includes bacteria with a very small number of immunity genes and bacteria with a very large number of toxin genes, neither of which occur in the model in either of the two modes. We plot the average number of bacterium types per time-step (solid line) and the total number of bacterium types found in 10 samples taken 10,000 time-steps apart (dashed line). Both

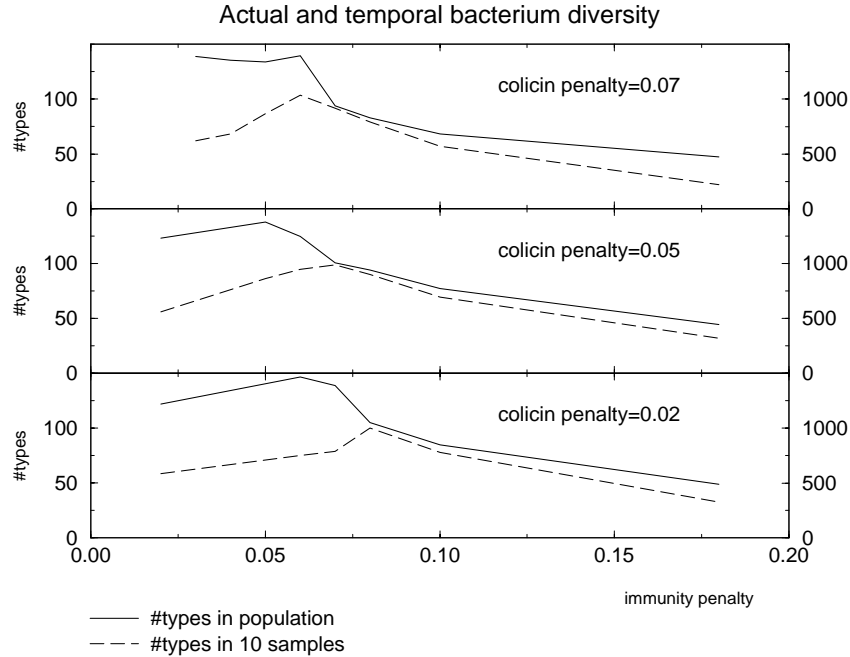


Figure 3.4: Bacterium diversity for a range of immunity growth penalties. Shown are the average number of bacterium types per generation and the total number of bacterium types over 10 samples. From top panel to bottom panel: $P_T = 0.07$, $P_t = 0.05$, $P_t = 0.02$. For all P_i average and total diversity is high. For low P_i the difference between the two lines indicate that a large proportion of the bacteria are residents. For high values of P_i all bacteria are transients.

lines indicate that a high diversity of the bacterium population is easily maintained at any one time and over longer time periods. Moreover, we see a phase transition occurring for $P_i \approx 0.07$ again. Interestingly, the difference between the average number of types and the actual number of types changes from a factor two to almost zero. Thus, for lower values of P_i a large proportion of the bacteria in the population are residents. For larger values of P_i on the other hand almost all bacterium types are transient; in each sample the population is completely replaced by new types of bacteria.

Carrying capacity, growth-rate and the phase transition

The carrying capacity of a population that grows in space depends non-linearly on the growth-rate per individual. The phase transition in the behaviour of the model can be understood if this dependence is taken into account.

If we measure the carrying capacity of a population in a simple spatially explicit growth model² we get the relation as it is plotted in Fig.3.5. For high growth rates

²In fact, in a mean field approximation (i.e. $\frac{dX}{dt} = bX(S - X) - dX$, with b =birth rate, d =death rate,




Figure 3.5: Carrying capacity of a simple locally growing population versus growth rate. The x-axis corresponds to the x-axis of Figures 3.3 and 3.4 for bacteria that carry 8 immunity genes (penalty values between brackets) and 1 toxin gene ($P_t = 0.02$).

the carrying capacity is almost independent of the growth-rate. In a relatively small parameter region, however, the carrying capacity of the population collapses rapidly with decreasing growth-rate.

The phase transition in the colicin model corresponds to the collapse in the carrying capacity in Fig.3.5. The x-axis in this figure is scaled such that it corresponds to the growth rate of a bacterium that carries 8 immunity genes and 1 toxin gene ($P_t = 0.02$, see the growth rate equation in table 3.1). The dependence of the carrying capacity on the growth rate leads to the following observation. If the growth rate has little effect on the carrying capacity, i.e. $P_i < 0.07$, bacteria maintain high numbers of immunity genes per bacterium, independent of the growth penalty per gene. For $P_i > 0.07$ this option is no longer viable. However, instead of gradually lowering the number of immunity genes in order to maintain a sufficient growth rate the strategy of bacteria in coping with the plasmids changes dramatically. For high values of P_i the bacteria carry simply as many complete plasmids (i.e. both toxin and immunity genes) as possible. In this regime the number of plasmids carried per bacterium drops with P_i .

The dynamics for low values of P_i and for high values of P_i are very typical. We will refer to these two modes as the *hyperimmunity* and the *multitoxicity* mode. Of course many variations on the parameter settings that we used for these simulations are possible, but we found that in all simulations that we performed the dynamics could best be characterized with respect to these two modes. In short, hyperimmunity is characterized by high immunity levels, low toxicity levels and relatively stable dynamics. Multitoxicity is characterized by levels of immunity and toxicity which are about equal per bacterium and the dynamics are less stable; both toxin genes and immunity genes may disappear temporarily from the population.

S =size of space) the carrying capacity shows the same, non-linear relation to the growth rate; $\tilde{X} = 1 - \frac{d}{bS}$.

3.3.3 Plasmid invasions

The transients in the two simulations in Figures 3.3 and 3.4 help to characterize the two modes further. If a certain plasmid type is completely absent from the population, or if a plasmid is of a novel type, invasions of bacteria that carry the plasmid can have drastic effects on the dynamics since none of the bacteria present in the population is immune to the invading plasmid. However, the *hyperimmunity* and the *multitoxicity* modes react differently to such invasions.

In Fig.3.1a up to $t \approx 60,000$ invasions of new plasmid types lead to large disruptions in the population. An invasion at $t \approx 35,000$ is particularly representative. At $t \approx 20,000$ the type 4 (i.e. blue) plasmid disappears from the population. After that the dynamics tend to settle down in the characteristic behaviour; all remaining immunity and toxin gene types are present, the immunity level per bacterium is high and toxin level per bacterium goes down. But at $t \approx 35,000$ plasmid type 4 comes back by generative mutation. The bacterium strain that carries this plasmid takes over the whole population. The result of the invasion is that all other toxin types, except type 1 (i.e. black) and 2 (i.e. red), are expelled from the population because the invading strain does not carry these toxin genes but it does carry the immunity genes. The immunity gene types of the expelled toxin types now no longer come into contact with the corresponding toxins, so they too are slowly expelled from the population. More invasions follow until at $t \approx 60,000$ the dynamics is stabilized.

In Fig.3.1b we see that the transient is much shorter. Contrary to the previous case, in the multitoxicity mode bacteria do not carry all immunity genes. In general, the bacteria carry only those immunity gene types for which they also carry the toxin genes. Thus, patches of bacteria that carry different ensembles of plasmids are often sensitive to some of each other's toxins. At the boundary between patches this leads to mutual killing of bacteria of the different patches.

In Fig.3.2 we show spacetime plots of invasions by plasmids of the bacterium population in the hyperimmunity mode (a) and multitoxicity mode (b). In the spacetime plot we draw horizontal cross-sections of the grid for every tenth time-step. The different colours represent different bacterium types (see Sect.3.3.2).

In Fig.3.2a a red patch appears; it denotes a bacterium strain that carries a plasmid that is new in the population. The patch takes over the population completely on a short time-scale. The subsequent competition between this bacterium type and derived bacterium types that lack toxin genes (but not immunity genes) is on a much longer time-scale.

In Fig.3.2b a red patch also appears, again denoting a bacterium strain that carries a new plasmid³. In this case the invading bacterium strain does not take over the population. In fact, in one direction it grows at roughly the same speed as the invading patch in the hyperimmunity mode, but in the other direction it is itself taken over by a bacterium strain that was already present in the population.

Figure 3.2c is the same spacetime plot as Fig.3.2b except that all bacteria are coloured white. Many boundaries between patches in Fig.3.2b are actually large gaps that have resulted from the mutual killing. Since an invading bacterium strain does not carry

³The red, invading bacteria in Figures 3.2a and b are of different types. Both are drawn with the same colour for clarity.

all immunity genes it will at some point encounter bacteria that produce toxins against which it is not immune. As a result invasions will not spread out over the whole population. Although in the hyperimmunity mode the immunity level per bacterium is very high the multitoxicity mode seems more robust to invasions.

Once the hyperimmunity mode is stabilised it is independent of the creation of plasmids. In the multitoxicity mode plasmid types often disappear from the population. Only by continuous generation of new plasmids are the dynamics maintained in the long run. If absent plasmids are not reintroduced the bacteria keep losing plasmids up to the point where they can attain hyperimmunity mode behaviour with the available plasmid types. That is, because fewer plasmid types 'exist' bacteria can carry all immunity genes that are present while maintaining a viable growth rate. The rate at which plasmids are generated, i.e. the generative mutation rate μ_g , and the rate at which plasmids genes are mutated, i.e. the degenerative mutation rate μ_d , only weakly influence the results we have reported. In fact, the mutation rates scale with the size of the spatial grid. A minimum grid size, or μ_d , is necessary for all plasmid types to fit in stably. If either is too small the hyperimmunity mode will not remain stable because of continued invasions and expulsion of plasmids in the population, as mentioned before.

3.4 Discussion

3.4.1 Plasmid diversity

In our model we see that at the level of the bacterium population a high diversity of plasmids is easily maintained under a wide variety of parameter settings. An important feature of our model in this respect is the spatial embedding of the population of bacteria. In fact, if we perform global mixing in the model the population becomes completely homogeneous due to global competitive exclusion.

If we look at the competition between *complete*⁴ plasmids in the spatial model we see that plasmids cooperate rather than compete. In the parameter region in which we might expect to find the hyperimmunity mode we find that a decrease in the growth rate of the bacteria leads to an increase in the number of plasmids. For this phenomenon it makes no difference whether the decrease in the growth rate of the bacteria is imposed directly or as a result of increasing the number of plasmids the bacteria carry or increasing the growth penalty per immunity gene (e.g., see Fig.3.3). The bacteria decrease their carrying capacity with increasing cost per immunity gene rather than reduce the number of immunity genes. Thus, the bacteria can be said to integrate information with respect to colicin types over many generations. Although individual bacteria encounter only a small number of colicin types in their lifetime they nevertheless maintain a high immunity level as a result of spatial and temporal information integration (see Pagie & Hogeweg (1997)).

In the parameter region in which bacteria cannot carry the immunity genes of all types, i.e. the multitoxicity mode, plasmids are better off if they are in the presence of other plasmids as well. If we remove plasmid types in this parameter region the state

⁴Complete plasmids consist of both the toxin gene and the immunity gene. Immunity-only plasmids are in themselves non-viable in the long run.

of the bacterium population changes from the multitoxicity mode to the hyperimmunity mode, where there are high numbers of immunity genes but low numbers of toxin genes, or complete plasmids. Thus, from the point of view of the complete plasmids it is always advantageous to be with many other types of plasmids rather than only a few.

In the hyperimmunity mode plasmids only interact with each other via the growth rate of the host bacterium. In the multitoxicity mode, on the other hand, plasmids cooperate in a more direct manner. Here, plasmids are members of an ensemble of plasmids that are carried in one bacterium which therefore produces a combination of toxins. In Sect.3.3.3 we showed that competition between different ensembles of plasmids can lead to mutual killing of bacteria. As a result plasmids, by being members of an ensemble rather than being alone, have additional protection against other competing plasmid ensembles. Of course competition, here, at the level of plasmid ensembles is harsh. In the hyperimmunity mode plasmid ensembles are small and hardly ever come into contact.

3.4.2 The cost of defence

We found that plasmids in general fare well if their corresponding immunity genes are readily 'lost' by the bacterium host. Clearly, complete plasmids will always be in the neighbourhood of bacteria that carry immunity-only plasmids since these are one of the viable mutants of the complete plasmids. Thus, bacteria that carry complete plasmids will always find themselves in the neighbourhood of bacteria that are insensitive to their toxin and are better competitors. Unless toxin-producing bacteria can invade patches of sensitive bacteria the toxin-producing bacteria will be outcompeted. Toxin-producing bacteria invade patches of sensitive bacteria very fast. The time-scale of the competition between the remaining combinations of bacterium variants (e.g. sensitive versus immunity-only bacteria) is much longer because it is based only on differences in growth rate. The speed at which the immunity-only bacteria are overtaken by the sensitive bacteria is the primary limiting factor for the number of toxin-producing bacteria. Thus, a plasmid increases its chance of long-term survival by making it very advantageous for immunity-only bacteria to lose their immunity gene also. The plasmids can do this by imposing on the host a high growth penalty as result from carrying the immunity gene. Under what circumstances a high immunity penalty can evolve remains an open question since both bacteria and plasmids experience a short-term benefit of low costs. From other systems we know, however, that as a result of spatial dynamics long-term benefits may prevail over short-term benefits (Savill *et al.*, 1997; Boerlijst & Hogeweg, 1991).

3.4.3 Experimental data

The previous sections have suggested that it is advantageous for the plasmids to impose high growth penalties on the host. From the point of view of the bacterium it seems advantageous to have a large growth rate. However, the reduction in the carrying capacity due to a decrease in growthrate is small (see Fig.3.3 and 3.5).

Experimental data suggest that in natural isolates the proportion of bacterium strains that produce one or more colicins is high (51% in Achtman *et al.* (1983) and 35% in Riley & Gordon (1992)), and resistance to colicins is also high (22% of the bacteria in Riley & Gordon (1992) are resistant to all colicin types tested). In our model we find that

in the hyperimmunity mode the percentage of bacteria that produce one or more colicins varies between 38% ($P_i = 0.02, P_t = 0.02$) and 99% ($P_i = 0.07, P_t = 0.02$). The immunity levels in the hyperimmunity mode are high; bacteria are on average immune to 90% of the plasmid types and 32% of the bacteria are immune to all plasmid types ($P_i = 0.02, P_t = 0.02$). In the multitoxicity mode, on the other hand, we find that each bacterium produces several colicins. Also, the immunity level per bacterium is much lower than in the hyperimmunity mode (e.g. see Fig.3.3). Finally, in the multitoxicity mode the bacteria are generally not immune to any more colicin types than they actually produce.

These data suggests that in natural circumstances the growth penalty imposed by plasmids falls well within the hyperimmunity mode. In fact, it is shown that a significant number of bacteria do not carry any plasmids (14% in (Riley & Gordon, 1992), 8% in (Caugant *et al.*, 1981), and 4% in (Achtman *et al.*, 1983)). Assuming that the immunity against colicins is coded on plasmids, the absence of plasmid makes sense in the (local) absence of colicin producing plasmids only. The results of our model differs most from the experimental data with respect to the number of wild-type bacteria, i.e. bacteria that do not carry any plasmid at all. One possible mechanism that may influence this discrepancy between our results and the experimental data is the active removal by bacteria of all plasmids at one time through 'spitting'. This would 'convert' the bacteria directly into a wild-type state. However, if we include in our model the possibility 'spitting' we see that the numbers of wild-type bacteria are in fact very low ($< 3\%$; $P_i = 0.02, P_t = 0.02, \text{spit-rate}=0.01$). The number of colicin-producing bacteria, however, increases from 38% without spitting to 73% for high spit-rates. In the multitoxicity mode, on the other hand, the inclusion of spitting in the model does not have any effect other than a slight increase in the death rate of bacteria due to unwarranted loss of immunity.

Other than the difference in number of wild-type bacteria mentioned above the agreement between the experimental data and our results is striking, although experimental data are still somewhat scarce. Most experimental studies report only on colicin production and not on potential additional immunity. Moreover, in general they do not differentiate between immunity and other means of defence. Finally, it is not clear whether natural bacterium populations in different nutritional conditions adopt different strategies with respect to colicin plasmids; our results suggest that this might be the case.

3.5 Conclusion

We have studied the interaction between multiple colicin plasmid types and bacteria in a spatially explicit model. We have shown that a high diversity of colicin plasmids is easily maintained under many parameter settings. The most important finding in our study is the occurrence of a phase transition in the colicinogenic dynamics. This transition is dependent on the growth penalty caused by immunity production. We have shown that if the cost of immunity is such that a bacterium can carry most types of immunity without a large reduction of its carrying capacity, it will actually do this. Simultaneously this will minimise the number of toxin genes per bacterium. We call this hyperimmunity. When the cost of carrying all immunity gene types becomes too high for the bacteria to remain

viable they will then switch to the multitoxicity behaviour. In this mode bacteria carry as many complete plasmids as possible. At the phase transition the immunity level per bacterium drops whereas the toxin level per bacterium rises (Fig.3.3).

We have shown that in our model it is advantageous for plasmids to impose a high growth penalty on their host that results from the defence mechanism that defends the host to the rather unpleasant action of the plasmid. The high cost of the defence mechanism ensures that a plasmid-variant that only defends its host against toxins without coding for the toxin will quickly be outcompeted by wildtype bacteria. The presence of wildtype bacteria gives the plasmid an opportunity to remain in the population. Other systems that also comprise a threat plus defence should also benefit from the high cost of the defence mechanism.

Acknowledgements:

The authors thank Ms. S.M. McNab for linguistic advice. They also thank the referees for constructive criticism and helpful suggestions. The investigations were supported by the Life Sciences Foundation (SLW), which is subsidized by the Netherlands Organization for Scientific Research (NWO).

4

Individual- and population-based diversity in restriction-modification systems

Ludo Pagie and Paulien Hogeweg
*Theoretical Biology and Bioinformatics group,
Utrecht University, Padualaan 8,
3584 CH Utrecht, The Netherlands.*

Submitted, Bull. Math. Biol.

abstract

Restriction-modification (RM) systems are cognate gene complexes that code for an endonuclease and a methylase. They are often thought to have developed in bacteria as protection against invading genetic material, e.g. phage DNA. The high diversity of RM systems, as observed in nature, is often ascribed to the coevolution of RM systems (which 'invent' novel types) and phages. However, the extent to which phages are insensitive to RM systems casts doubts on the effectiveness of RM systems as protection against infection and thereby on the reason for the diversity of RM systems.

We present an eco-evolutionary model in order to study the evolution of the diversity of RM systems. The model predicts that in general diversity of RM systems is high. More importantly, the diversity of the RM systems is expressed either at individual level or population level. In the first case all individuals carry RM systems of all sequence specificities, whereas in the second case they carry only one RM system or no RM systems at all. Nevertheless, in the second case the same number of sequence specificities are present in the population.

4.1 Introduction

Restriction-modification (RM) systems are cognate gene complexes that code for an endonuclease and a methylase. They were discovered in the early 50s when it was found that phages grown on a particular bacterium strain were unable to infect a different strain (Luria & Human, 1952). The endonuclease of RM systems cuts DNA molecules at or near specific short nucleotide sequences unless it is methylated at these sequences. The DNA of the host bacterium is protected against the endonuclease activity by the methylase which recognises (and methylates) the same sequence as the endonuclease. In natural bacterium populations a large number of different 'RM types' are found, i.e. RM systems that recognise different nucleotide sequences (for reviews see: Bickle & Kruger (1993); Barcus & Murray (1995); Redaschi & Bickle (1996)).

It is generally believed that the main function of RM systems is to protect bacterium cells against foreign DNA, i.e. phage infections. If a bacterium is infected by a phage the DNA of the phage is cut by the endonuclease and the infection is aborted. However, the probability that the endonuclease of a typical RM system will fail to cut an invading DNA molecule ranges from 10^{-2} to 10^{-6} (Wilson & Murray, 1991; Korona *et al.*, 1993). As a result the invading phage DNA becomes 'inadvertently' methylated, or *modified*, and thereby protected against the endonuclease activity just like the DNA of the host bacterium. Only when the modified phage infects a bacterium that lacks the RM system does the phage lose the methylation patterns and become sensitive again.

Previous theoretical models have shown that under well-mixed conditions novel RM systems can invade existing bacterium-phage communities as a result of frequency-dependent selection (Levin, 1988). The common RM types are assumed always to be accompanied (and limited) by correspondingly modified phages, whereas novel RM types initially provide their host with total protection. Once RM systems are established they will remain in the population due to frequency-dependent infection. The (high) diversity of the RM-systems is then a consequence of subsequent invasions by RM-systems and theoretically is limited only by the bacterium population size. Frank (1994a), however, showed that by increasing the number of RM systems in the bacterium population the diversity of bacteria that carry different RM systems can reach a level at which a phage population can no longer be sustained. This is caused by a decrease in the number of *effective* phage infections, i.e. infections by modified phage of bacteria that carry the corresponding RM type. With the disappearance of the phage population the selection for novel RM systems is lost as well.

Data from natural and laboratory bacterium strains show that many bacteria are resistant to many phages in ways other than through restriction. Also, many phages appear to be insensitive to many RM types (Korona *et al.*, 1993). Moreover, in experiments with sensitive laboratory bacterium strains RM systems seem to provide the bacteria with only transient protection from the phage. Quickly, bacteria arise that are immune to the phage, due to evolution of receptor-based resistance (Korona & Levin, 1993). Experiments on surface cultures suggest that RM systems only favour bacteria in colonisation events in spatially explicit environments in which phages are already present (Korona & Levin, 1993).

Most work on the evolution of the diversity of RM systems has been concerned with the invasion of bacterium populations by bacteria that carried only single RM systems

(Levin, 1988; Korona & Levin, 1993; Frank, 1994a). We study the eco-evolutionary dynamics of RM systems in bacterium-phage communities on longer time-scales in which we explicitly allow for the competition or cooperation between RM systems in individual bacteria, as well as at the level of the bacterium population. We find that, as a result of the frequency-dependent selection of novel RM systems, the diversity of RM systems increases if the bacterium population is infected by phages. More importantly, we find that the diversity is expressed in two modes. Both modes occur as stable attractors if the system is viable, i.e. if phages, bacteria and RM systems are present. Although each mode may contain large numbers of RM types the ecological dynamics of the two modes differ greatly. The selection pressure of novel RM systems also differs as a result of the different ecological dynamics.

4.2 The model

We modelled the interactions of the bacteria in a spatially explicit, discrete event and discrete variable model with probabilistic updating. The bacterium population is a partially open system (e.g. an intestine) which can be infected by bacteria, phage and RM systems from external sources (see Sect. 4.4.1). The bacteria in the model can be interpreted as individual cells or as monomorphic bacterium strains (see also Pagie & Hogeweg (1999a)). Phages are not modelled explicitly but as infected bacteria. RM systems are modelled as independent genetic elements that can be carried by plasmids or on the host genome. We assume that the RM systems do not code for any functionality influencing the behaviour of the bacterium host other than the endonuclease and methylase activity. They impose a penalty on the growth rate of their host due to the forced maintenance of endonuclease and methylase levels in the cytoplasm plus the maintenance of the RM system itself. RM systems can be of different, discrete types, i.e. recognise different nucleotide sequences. In all other respects, such as growth penalty or endonuclease efficiency, RM systems of different types are identical.

RM systems are acquired and lost independently of the rest of the genome of the host bacterium. The processes by which RM systems are transferred within the bacterium population and by which RM systems are acquired from external sources are not modelled at a detailed level. In natural bacterium populations genetic material is readily exchanged within and between bacteria populations through a variety of processes, e.g. conjugation, transduction, and transformation. We generalised with respect to these processes by assuming a single process to account for the exchange of existing RM systems between bacteria within the population, with rate μ_t , and a single process to account for the acquisition of ‘novel’ RM systems that are not yet present in the population, with rate μ_i . New RM systems are acquired in a bacterium in addition to any RM systems that the bacterium may already carry. A scenario can be that phages from external sources carry a novel RM type (e.g. by transduction) and are cut up during the infection process by the endonuclease of the RM systems which are present in the bacterium cell. The novel RM system is incorporated in the bacterium genome by homologous recombination. Of course, RM systems can also disappear, e.g. by mutation or by segregation, with rate μ_d . We do not take into account the number of RM systems per type that are carried by a bacterium. Also, we assume that all RM systems are compatible, i.e. a single bacterium

can carry all RM types (provided the total growth penalty does not get too large).

The space in the model is a 2-dimensional, square, regular grid of 300*300 sites with a neighbourhood consisting of the 8 nearest sites (i.e. the ‘Moore-neighbourhood’). Each site is either empty or occupied by a bacterium. Empty sites can be colonised by bacteria from the 8 neighbouring sites, each with a probability based on their growth rate. A wild-type bacterium, i.e. an uninfected bacterium that does not carry any RM systems, has a growth rate $r = 1.0$. Infected bacteria do not grow; they have a growth rate $r = 0.0$. The growth rate of bacteria that carry RM systems is lowered linearly according to the number of RM systems they carry; $r = 1.0 - C * P$, where C is the number of RM systems and P is the growth penalty per RM system. The linearity between growth rate and number of RM systems is chosen in order to avoid the occurrence of a priori non-linearities in the interaction between bacteria and the number of RM systems that they carry. The death rate d_w of bacteria is 0.1 and the death rate d_i of infected bacteria is 0.45.

Phages are modelled as infected bacteria. A susceptible bacterium is infected if it has an infected bacterium in its neighbourhood. A bacterium is susceptible if it is not protected by an ‘effective’ RM system (see below). Phages can escape the endonuclease activity of an effective RM system during infection with probability e and subsequently acquire the same methylation patterns as the host bacterium. Such phages are no longer sensitive to the corresponding RM types. The latter are no longer ‘effective’ against these modified phages. The probability that a bacterium that carries N effective RM types will be infected by a phage is e^N . If a phage that carries methylation patterns infects a bacterium that lacks (some of) the corresponding RM types the phage loses these methylation patterns and once more becomes sensitive to these RM systems. We used an influx of phages at rate i per bacterium in order maintain a bacterium-phage-RM system interaction under all circumstances (see Sect.4.3.1).

4.3 Results

In order to study the evolutionary consequences of the interaction of multiple RM systems in a bacterium population we performed a large number of simulations, for a wide range of parameter values. In general, the simulations were started with a wild-type bacterium population, i.e. bacteria without RM systems, infected by phages. The acquisition of RM types by the bacteria occurred on a long time-scale.

We find that the first few novel RM types are integrated into all bacteria in the population. This means that during this period the bacterium population remains homogeneous with respect to the combination of RM types that each bacterium carries. In this period novel RM types that are introduced into the bacterium population at first confer an advantage on their host because they provide protection against the phages. Thus, although they impose an (additional) growth penalty on their host, bacteria that carry the novel RM type out-compete the established bacteria. When they become more abundant phages will arise with methylation patterns that correspond to this RM type. The ‘new’ bacteria are still better competitors than the previously established, or ‘old’, bacteria because the latter can be infected by phages that are produced in the ‘new’ bacteria. Conversely, the ‘new’ bacteria cannot be infected by phages produced in the ‘old’

bacteria unless inadvertent methylation occurs. This *unidirectional infectability* results in a complete take-over by the new bacteria.

All RM types that are introduced into the bacterium population are integrated into all bacteria. However, the phages accumulate all corresponding modifications. Thus, except during the transients following the introduction of a novel RM type, the phage population is not hindered by the RM systems carried by the bacteria.

Of course, RM systems cannot continue to accumulate indefinitely in bacteria because the growth rate of bacteria is finite. At a certain point, following the introduction of additional RM types into the population, the accumulated diversity of RM systems in each individual bacterium breaks down; this leads to a situation in which bacteria carry only one RM system or even no RM system. The total number of RM types in the bacterium population, however, is maintained at the same level during and after this breakdown. The different RM types are now distributed over the bacterium population. During the breakdown event the phage population is greatly reduced and loses all modifications, or it even dies out completely.

Simulations like the one described above showed the transition at arbitrary numbers of RM systems present in the population suggesting bi-stability between the two modes. We found that the model, in fact, shows bi-stable behaviour for a large range of RM types. In Fig.4.1 we plot the average bacterium population-size against the number of RM types in the population of simulations that were started with a homogeneous or a heterogeneous bacterium population. Both modes are stable except at the extremes. If only two RM systems are present in an initially heterogeneous bacterium population, exchange of RM systems within the bacterium population generates bacteria that carry both RM systems. The latter bacteria out-compete bacteria that carry only one of the two RM systems because of unidirectional infectability (see above). On the other hand, nine RM types can occur only in a heterogeneous bacterium population due to growth limitations of the bacteria. Note also that the size of the bacterium population increases with increasing number of RM systems per bacterium. Although the growth rate per bacterium decreases, the phage population size decreases as well (see also Fig.4.3). The latter leads to the increase of the bacterium population.

Thus, the model can best be described in terms of the two dynamical modes which we denote as *individual-based diversity* and *population-based diversity*. In the first mode the bacterium population is almost homogeneous with respect to the composition of RM types, i.e. all bacteria carry all RM types that are present in the population. In the second mode the bacterium population is heterogeneous, i.e. bacteria carry none or only one of the RM types present in the population.

In order to understand the transition between the bi-stable modes we discuss simulations in which novel RM types are introduced explicitly at regular time intervals. First, however, we will show the characteristics of the individual- and population-based modes in ecological simulations, i.e. simulations which have already been initialised in either of the two modes and which do not include the introduction of novel RM types.

4.3.1 Individual- and population-based diversity

In Fig.4.2 we show snapshots and space-time plots of two simulations which have equal parameter values but which were initialised differently. In both simulations five different

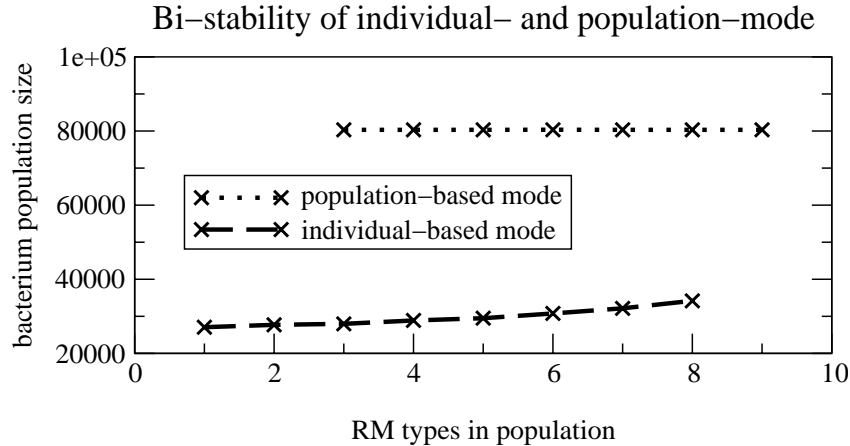


Figure 4.1: Bacterium population sizes of individual-based mode and population-based mode versus the number of RM systems in the bacterium population; $d_w = 0.1$, $d_i = 0.45$, $P = 0.05$, $\mu_d = 10^{-4}$, $\mu_g = 10^{-8}$, $\mu_t = 10^{-6}$, $e = 10^{-4}$, $i = 10^{-6}$.

RM types are present in the bacterium population. The first simulation (Fig.4.2a) is initialised with wild-type bacteria and with bacteria that carry all five RM systems, i.e. the individual-based mode. The second simulation (Fig.4.2b) is initialised with wild-type bacteria and bacteria that carry only one of the five RM types, i.e. the population-based mode. Colours other than white (empty space), red (wild-type bacteria) and black (infected bacteria) denote uninfected bacteria that carry different combinations of RM systems.

In Fig.4.2a we see that almost all bacteria are either infected (i.e. black), or carry all five RM systems that were present in the initial population (i.e. pink). The wild-type bacteria have been expunged from the population due to the unidirectional infectability (see above). A few bacteria have lost one RM system; they have different colours (blue, orange, purple). The space-time plot shows that these bacteria, despite their higher growth rate, do not take over the population again due to the unidirectional infectability. The ecological dynamics and the spatio-temporal patterns are typical for spatial embedded host-parasite systems with a high infection rate: here, a large phage population size, a relatively small host population size, and turbulent wave patterns (Savill *et al.*, 1997; Johansen, 1996).

In Fig.4.2b we see that the bacteria are highly polymorphic with respect to the RM systems that they carry. A small number of bacteria are wild-type (red) whereas the other bacteria carry only one RM system (blue, purple, yellow, pink, green). In addition to the polymorphism at the population level the bacteria with different RM systems are distributed over the whole field in small patches. These patches are fairly stationary in time; competition between bacteria that carry the same number of RM systems is neutral. As a result of the local bacterial heterogeneity the phage population is unable to maintain the modifications that render them insensitive to an RM type. Although phages still acquire methylation patterns through inadvertent methylation by RM systems during infection,




Figure 4.2: See also colour plate 2 (page 62). Snapshots and space-time plots (i.e. horizontal cross-sections of the grid of every fifth time-step) of a. individual-based and b. population-based diversity. Different colours denote different combinations of RM systems. Parameters as in Fig.4.1.

they also readily lose the modification. Due to the small patches with monomorphic bacteria, modified phages will quickly infect bacteria that lack the corresponding RM type. By doing this they lose the methylation pattern. The overall result is that the phage population lives primarily on the wild-type bacterium population and only occasionally succeeds in infecting bacteria that carry RM systems. At this point, if the number of RM types in the bacterium population is small the phage population is still viable. If the number of RM types is large the phage population is no longer viable and depends on the influx for its preservation. Note, however, that the field size is important; larger




Figure 4.3: See also colour plate 2 (page 62). Temporal dynamics of bacteria (bottom panel), RM types (middle panel) and phage population (top panel). The bacterium population is split into bacteria that carry the same number of RM systems. The diversity of RM systems increases until $t = 9000$ and the individual-based diversity is maintained. After the breakdown at $t \approx 10.000$ the bacterium population settles into the population mode with no RM systems or only one RM system per bacterium. The total number of RM types remains high. Parameters as in Fig.4.1.

fields can more easily maintain viable phage populations for larger number of RM types in the bacterium population.

4.3.2 Transition between modes

Here, we present the results of a simulation in which we start with a homogeneous wild-type bacterium population that is infected by phages and then we introduce novel RM systems at regular time intervals.

In Fig.4.3 we plot different characteristics of the simulation. In the bottom panel we plot the size of the bacterium population divided into bacteria that carry $0, 1, \dots, N$ RM systems. At $t = 0$ we start with the wild-type bacterium population, after which we introduce a novel RM system every 1000th time step. Per RM type, in the middle panel, we plot the number of RM systems that are present in the bacterium population. In the upper panel we plot the total phage population.

The first six RM systems that come into the population become integrated into the entire bacterium population. In this period all bacteria carry the same number of RM types (bottom panel), the number being equal to the total number of RM types in the population. Thus, individual-based diversity remains whereas the diversity of RM sys-

tems in the bacterium population increases, as described above. The transient to full integration of a new RM type into the population, however, takes longer and longer. The seventh RM type never completely succeeds in taking over the population. Although the introduction of the eighth RM system still leads to a substantial number of bacteria that carry all eight RM systems no phage ever acquires all eight modification types; the breakdown of the individual mode has definitely set in.

The space-time plots in Fig.4.4 show the breakdown of the individual mode in more detail. Into a population that was in a stable individual-based mode with eight RM types we introduced a novel RM type in the middle of the field. Figure 4.4a shows that the patch containing the novel RM type grows to cover the whole field. In the growing patch new combinations of RM types arise continuously.

In Fig.4.4a we have used different colours to denote different combinations of RM types, with white, red and black denoting empty space, wild-type bacteria and phages respectively. In Fig.4.4b we have used different colours to denote bacteria with different numbers of RM systems in order to show the transient from individual-based to population-based diversity. Bacteria that carry the ninth RM type in addition to other RM types have shades of a blue-green-yellow colour-ramp; bacteria that do not carry the ninth RM type are depicted in different shades of grey.

Initially, the patch consists of bacteria that carry the ninth RM type; they cannot readily be infected by bacteria from outside the patch because no correspondingly modified phage yet exist. In the patch bacteria start to lose RM types due to mutation and segregation. Because locally phage are still absent the ‘mutant’ bacteria are not yet experiencing the unidirectional infectability (sect. 4.3), and are therefore not yet outcompeted by bacteria that carry the ‘complete’ set of nine RM types. The breakdown process becomes irreversible as soon as, what we call, *mutually uninfected* groups of bacteria arise, i.e. groups of bacteria that carry one, or more, RM systems that the bacteria of another group lack. As a result, one such group cannot be infected by phages resident in another such group, and vice versa. Even if, at this point, phages arise that are modified in accordance with the ninth RM type, the presence of the mutually uninfected groups of bacteria has the same result as the polymorphism described in the population-based mode; phages can only infect bacteria that carry the same combination of RM types (or a subset thereof) as the combination of methylation patterns of the phage. Because bacteria that carry only subsets arise continuously by loss of RM systems the phage population loses more and more modifications. The mutual uninfectedability quickly restricts phages to small patches of monomorphic bacteria, which eventually leads to the extinction of the phage population.

The phage population dies out when the patch has grown over the whole field. At that moment the bacteria carry 1-9 RM types. The maximum number of possible combinations with 1-9 RM types is $\sum_{i=1}^9 \binom{9}{i} = 511$. At the time of phage extinction ($t = 900$) 247 different RM types combinations exist, the maximum number of different types being 451 and occurring at $t = 2000$. At this point phages cannot re-infect the bacterium population. Thus, bacteria continue to lose RM systems until the system settles into the population-based mode where bacteria carry at most one RM system. When the number of wild-type bacteria is sufficient for the maintenance of a viable phage population (given the influx; see Sect.4.3.1) the system stabilises. If phages do not infect the population after the transition the bacteria eventually lose all RM systems which leads the




Figure 4.4: See also colour plate 3 (page 63). Space-time plots of the breakdown of individual-based mode to population-based mode. In a) different colours denote different combinations of RM systems. In b) shades of grey denote different numbers of RM systems per bacterium. Parameters as in Fig.4.1.

bacteria back to the beginning of the individual-based mode.

Thus, the irregular occurrence of the breakdown of the individual-based mode in the population-based mode, which we found in section 4.3, is caused by a change event, i.e. the arising of mutually uninfected bacteria. The subsequent explosion of polymorphic bacteria only stops when the population-based mode is reached.

4.4 Discussion

In section 4.3 we showed that a population of bacteria that carry RM systems can be in one of two stable modes; *individual-based diversity* and *population-based diversity*. The individual-based diversity mode is characterised by a monomorphic bacterium population; all bacteria carry all RM systems that are present in the population and have thus developed maximal protection against infection. However, also the phages are fully protected; all phages are modified, and thus insensitive, with respect to all RM types. As a result the population dynamics in the individual-based mode are very similar to those of wild-type bacteria that are infected by phages with high a infection rate; the bacteria are strongly limited by the phages. Only if *mutually uninfected* bacteria arise will the individual-based diversity break down and settle in the population-based mode.

The population-based diversity mode on the other hand is characterised by a heterogeneous bacterium population in which bacteria carry only one RM system or no RM system at all. Also the spatial heterogeneity is very high; bacteria with different RM systems reside in small patches that are in neutral competition. In this mode the phage population, which is now fully sensitive to the RM systems that are present in the population, often is only viable as a result of influx from external sources and thus do not limit the bacterium population to any great extent.

The finding that two stable modes exist in which the diversity of RM systems is expressed is very robust. We can modify the model for instance by performing global mixing of the bacterium population, incorporating independent mutation of the two genes that code for the endonuclease and for the methylase (see 4.4.3, or increasing the longevity of the methylation patterns on the phage. All these (structural) changes of the model do not substantially change the results that we report here, i.e. the bi-stability of the two modes, although the two modes are most pronounced in the original model. In the modified models we see that bacteria in the population-based mode can carry more RM systems per bacterium than in the unmodified model but retain the other properties of the population-based mode, i.e. the very small phage population. In the individual-based mode more bacteria are present that carry not quite all RM systems but rather one or two RM systems less but still bacteria are present that carry all RM systems and the phage population is very large. The spatial scale of our model also does not have an appreciable effect on the results either. However, the dependence of the phage population in the population-based mode on the influx is influenced by the spatial scale of the model (Sect.4.3.1).

4.4.1 Model predictions and comparison to experimental data

The bacterium population in our model is an open system in which influx of phage and RM systems occur with low rates. Thus, our population can be compared to (part of) a single bacterium population in for instance an intestinal tract but also to (biofilm-) communities in the soil or sewage. Both environments show high diversity of bacteria (Caugant *et al.*, 1981; Falk *et al.*, 1998) and phages (Hantula *et al.*, 1991; Korona *et al.*, 1993) and in both environments infections and invasions of local populations occurs (Caugant *et al.*, 1981; Savage, 1977).

The model predicts that populations of bacteria that carry RM systems can be in two

modes. Both modes show a high diversity of RM systems in the population. Studies of natural bacterium populations show that RM systems in bacterium populations are highly diverse and that many bacteria carry RM systems, and often more than one (Wilson & Murray, 1991; Barcus & Murray, 1995). It is, however, not clear if single bacterium populations are mostly homogeneous as they are in the individual-based mode (Barcus & Murray, 1995) or heterogeneous as in the population-based mode. However, given the results of the modified models mentioned above, the fact that bacteria in natural circumstances often carry more than one RM system does not mean a priori that the population cannot be in the population-based mode.

The model also predicts that the bacterium population is either strongly limited by phages (i.e. the individual-based mode) or not (i.e. the population-based mode). In the first case all phages are modified with respect to all RM types that are present in the population while in the second they are modified to only a very few RM types or not modified at all. Also in this respect the available experimental data is not conclusive. Although phages occur in many microbial communities the extend to which bacteria are limited by the phages is often unclear (Havelaar *et al.*, 1986; Proctor & Furhman, 1990), although it seems to be related to the physical condition of the host of a microbial community (Cornax *et al.*, 1994). It would be interesting to see if the presence of the individual-based or population-based mode, since both can occur, is related to the physical condition of the host as well. Although many phages are insensitive to many RM systems by means other than having acquired methylation patterns (Kruger & Bickle, 1983; Sharp, 1986), still, many phages are also sensitive to several RM systems (Hantula *et al.*, 1991; Korona *et al.*, 1993). Most importantly, however, experimental data do not specify to what extent phages in natural communities are insensitive (by being resistant or modified) to RM systems that are present in those communities. Given the available experimental data it is at this moment not possible to designate natural bacterium populations to be in the individual-based mode or in the population-based mode. A survey of the extend to which phages are insensitive to the RM systems that are actually present in the bacterium population can give more insight to whether the populations are in the individual-based mode or in the population-based mode. A subsequent comparison between the extend to which phages have developed resistance to RM systems and the extend to which they carry additional modifications which correspond to other RM types can give additional insights in the (evolutionary) interaction pressures between the phages and bacteria.

4.4.2 Frequency-dependent selection and pairwise infectability

Previous theoretical studies have focussed on the importance of frequency-dependent selection as possible explanation for the diversity of RM systems (Levin, 1988; Korona & Levin, 1993; Frank, 1994a). Novel, and thus rare, RM types are favoured because no modified phage yet exist. Common RM types, on the other hand, are quickly infected by correspondingly modified phages. A high diversity of RM systems is then achieved by successive invasions of novel RM types which are subsequently also maintained in the population by frequency-dependent selection. The maximum number of RM types in the bacterium population in that case only depends on the population size. Studies by Korona & Levin (1993), on the other hand, suggest that bacteria easily evolve

resistance to phage by altering or eliminating certain membrane receptors. Sensitive bacteria that carry RM systems are limited by correspondingly modified phages. Thus, bacterium populations are dominated by resistant bacteria if resource levels are high (but see Fuhrman (1999)). Frank (1994a) showed, however, that an increase in the diversity of RM systems in the bacterium population could result in a transition from a resistance dominated population to a population that is dominated by bacteria that carry RM systems. This transition is caused by the decrease in the probability that a phage is ‘correctly’ modified when it infects a bacterium that carries an RM system. A further increase in diversity of RM systems can even lead to the extinction of the phage population.

In the models discussed above bacteria carry at most one RM system. In these models, states of the bacterium population in which a high number of RM types is present mostly resembles the population-based diversity mode in our model. However, bacteria in our model can carry as many RM systems as their growth rate allows. As a result we find two stable modes with a high diversity of RM systems. In addition to the population-based mode we also find the individual-based mode. Although frequency-dependent selection plays a role in our model as well, it plays different roles in the two modes. In the individual-based mode RM types are maintained by uni-directional selection while frequency-dependent selection only affects the invasion by novel RM systems. In the population-based mode it plays a minor role in the invasion by novel RM types because of the highly reduced phage population with high numbers of RM types present in the population. Only when the numbers of RM types in the population is low does frequency-dependent selection result in long-term integration of novel RM types into the population. However, the maintenance of RM types does depend on frequency-dependent selection, despite the small phage population.

Frequency-dependent selection can lead to higher numbers of RM types in the individual-based mode compared to the population-based mode. Thus, in the population-based mode a high degree of diversity can be attained via the individual-based mode which is impossible to achieve in the population-based mode only. Also, a possibility for fluctuations between the two modes exists if the influx of phage is very small or only periodically present. If the population-based mode breaks down due to (temporal) absence of phage the bacteria will revert to a wild-type bacterium population; a highly attractive state for a spurious phage infection and the starting point of the individual-based mode.

The bi-stability of the system is caused by the stabilising processes in the two modes, i.e. unidirectional infection in the individual-based mode and frequency-dependent selection in the population-based mode. During the transition of the individual-based mode to the population-based mode the number of different combinations of RM types in bacteria grows very large. Due to the mutual uninfectedness of these polymorphic bacteria they cannot be infected by phage and the system is then unstable.

4.4.3 Cooperation across levels

In multi-level evolutionary models such as the ones described above it is not easy to determine who gains by cooperating with whom and what the evolutionary outcome will be. Recently we have studied the evolution of the diversity of colicin plasmids

in bacterium populations (Pagie & Hogeweg, 1999a). We reported that the behaviour of the model could also be described in terms of either an individual-based mode or a population-based mode. In the colicin model the colicins appeared to cooperate with each other at the expense of the bacteria, rather than acting as a ‘mere’ sum of gains and losses at the level of bacterial competition. The colicin model, however, was not bi-stable like the RM model discussed here; the system settled in one of the two modes, depending on the growth penalty parameter of the colicins and the number of colicins in the bacterium population. In the RM model the phage population acts as a third level in the system by means of which the functionality encoded by the RM systems comes about. The fact that the RM model has one more level than the colicin model, apparently brings about the bi-stability of the system.

Recently, two groups proposed an alternative mechanism for the maintenance of RM systems in terms of selfishness ((Naito *et al.*, 1995; Kulakauskas *et al.*, 1995) but see also (O’Neill *et al.*, 1997)). The selfish character of RM systems results from post-segregational host killing. Competition between selfish RM systems enhances the differentiation of RM types. However, RM systems that incorporate post-segregational host killing will at first raise the death rate of the bacteria that carry these RM systems. Thus, the bacteria that carry these selfish RM systems will be outcompeted by wild-type bacteria; post-segregational host killing does not increase the stability of RM systems in this situation (Monod, 1992).

Nevertheless, in our model post-segregation host killing favours the RM systems in the individual-based mode. In the model bacteria can lose RM systems by decay and segregation but we did not include post-segregational host killing. The main effect of losing RM systems, however, is the generation of faster growing, phage-sensitive bacteria. In the population-based mode these competitors already exist, whereas in the individual-based mode they are outcompeted by bacteria that carry all RM systems because of unidirectional infection (Sect.4.3). However, in the individual-based mode a higher ratio of bacteria that carry fewer than all RM systems increases the probability that mutually uninfected bacteria will be generated (Sect.4.3.2) and the population-based mode will break down. Because post-segregational host killing will lower the rate at which bacteria arise that have lost one or more RM systems (they are killed) it may prevent the breakdown of the individual mode and thereby maintain high levels of RM systems per bacterium. Indeed, lowering the rate at which bacteria lose RM systems (μ_d) increases the stability of the individual-based mode somewhat. In the individual-based mode, although bacteria seem maximally protected against phage infection, it is only the RM systems and phages that prosper.

4.4.4 Conclusion

The system that we studied here, i.e. RM systems in a bacterium population that is infected by phages, has characteristics of other systems in which individuals can build up a diverse set of responses or actions to guard against parasites or predators, or to compete with conspecifics. A general question in such systems is if a diverse ensemble of interactions arises between the antagonists, and to what extent individuals should strive for individual ability at high cost, or choose to live cheap and dangerous but differently.

We have shown that in our model the two extreme cases turn out to be two stable

states. In fact, that the two modes are stable does not depend on the cost of carrying RM systems. Only for unrealistically high growth penalties is the individual-based mode not stable. Thus, the 'cost-benefit optimisation' of an individual bacterium is only a minor factor in determining the fate of the bacterium. We hypothesise that systems like the RM systems in bacterium-phage communities, such as MHC molecules in vertebrate immune systems or the toxic defence repertoire of plants, may also exhibit the bi-stability of an individual-based mode and a population-based mode. Although individuals may prefer maximal individual ability this is certainly not the only, or most agreeable evolutionary avenue.

Acknowledgements:

The authors thank Ms. S.M. McNab for linguistic advice. The investigations were supported by the Life Sciences Foundation (SLW), which is subsidized by the Netherlands Organization for Scientific Research (NWO).

5

Coevolutionary dynamics: information integration, speciation, and red queen dynamics

Ludo Pagie and Paulien Hogeweg
*Theoretical Biology and Bioinformatics group,
Utrecht University, Padualaan 8,
3584 CH Utrecht, The Netherlands.*

Preliminary version

The outcome of evolutionary processes is studied from different points of view. First of all, evolution was proposed as the origin of species. Later, it was also studied as an optimisation process, and as a source of red queen dynamics, or ‘arms-races’. Researchers with different points of view impose different structural properties on the evolutionary process and study results with different search-images. Here, we present one model in which all such evolutionary outcomes can be seen, depending on minor parameter changes. The model encompasses a coevolutionary system of two species that have an antagonistic interaction. The interaction between individuals of the two species depends on an explicit, non-linear genotype-phenotype mapping. The populations are embedded in space and the individuals of both species interact and compete locally in this space.

The outcome of the evolutionary process in simulations in which individuals remain localised through time is compared to simulations in which individuals are globally mixed every time step. In the first case we see information integration, i.e. evolution of a general solution which covers circumstances which are encountered over many generations. In the second case we see red queen dynamics, i.e. a continued evolutionary change in both species. If we use a somewhat different fitness function, in the first case we see speciation into separate specialised species. In the second case we also see red queen dynamics, although now we see optimisation of the red queen (she can run faster), apparently again as result of information integration.

5.1 Introduction

Biological evolution is generally characterised by sparse fitness evaluation; during their lifetime individual organisms do not experience all (types of) environmental circumstances which may influence their fitness. For instance, individuals do not encounter all possible diseases or predators or types of resources. The question that then arises is how they integrate evolutionary adaptations to these separate challenges, especially when they experience only a small number of selection events. This question becomes even more stringent if the environment consists of organisms that are evolving themselves. The sampling of the set of possible environmental circumstances is not only sparse but the set changes over time as well.

If the selection experienced by two organisms of different species depends on the other species and the resulting evolutionary process occurs simultaneously in both species they *coevolve*. Coevolution is often classified as either *diffuse* coevolution or *pairwise* coevolution (Janzen, 1980). Coevolution is pairwise if the coevolving traits in the two species evolve independently of the presence of other species and if the coevolving traits can change independently of other traits that the individuals express. In all other cases coevolution is said to be diffuse. Although some authors would claim that only pairwise coevolution is ‘true’ coevolution (Rothstein, 1990) coevolution is only seldom strictly pairwise (Farrell & Mitter, 1992). In fact, there is a gradual, rather than sharp, transition from adaptation to a constant environment to adaptation to an antagonistically coevolving population.

Previous studies have shown that pairwise coevolution can lead to three evolutionary outcomes (Dawkins & Krebs, 1979; Dieckmann *et al.*, 1995). One, or both of the species can die out; the coevolution of the two species can result in a stable coexistence where the species do not evolve further; or the two species can show a continued evolutionary change which can be of an oscillatory nature or which is best described as a runaway process. The latter evolutionary outcome, i.e. the continued evolutionary change of the two species, is often referred to as “red queen dynamics” (Van Valen, 1973) or an “arms race” (Dawkins & Krebs, 1979). An evolutionary runaway process is often thought to be unrealistic because it results in evolution toward unbounded character trait values, although it may lead to mutualistic interactions with extreme forms of obligatory reciprocal dependency (Pellmyr *et al.* (1996), see also Blaney & Miller (1995)).

Theoretical studies suggest that the occurrence of continued evolutionary change may be enhanced by an increase in the mutation rate (Dieckmann *et al.*, 1995) or by asymmetry in the “incentive-to-win” between the two species (Gavrilets, 1997) (e.g. the “life-dinner” principle; Dawkins & Krebs (1979)), or the existence of stabilising selection acting more strongly on the ‘victim’ of the two species than on the ‘exploiter’ (Gavrilets, 1997).

In the context of evolutionary optimisation techniques some studies show that coevolution leads to an increase in the performance or efficiency of the optimisation process (Paredis, 1995; Husbands, 1994; Rosin & Belew, 1997). In these models coevolution is often compared to predator-prey or host-parasite interactions, i.e. a reciprocal antagonistic interaction (Bullock, 1995). However, coevolution does not always lead to general solutions of the optimisation problem; red queen dynamics may hinder the optimisation process (Paredis, 1997), the coevolving species may speciate (Hillis, 1990), or settle into

“mediocre stable states”(Ficici & Pollack, 1998).

Hillis (1990) studied a coevolutionary optimisation model in which sorter algorithms coevolved with sorter problems. He found that coevolution of algorithms and problems resulted in a much more efficient process that led to faster sorter algorithms than algorithms found in traditional evolutionary optimisation processes. In addition to the coevolutionary, antagonistic relation between algorithms and problems Hillis embedded the evolutionary process in a spatial setting; algorithms and problems were situated on a 2-dimensional grid and interacted only locally. Similar, spatially embedded models were studied by Husbands (1994) and Pagie & Hogeweg (1997). In all cases an improvement of the optimisation process was reported.

Paredis (1997) studied a coevolutionary optimisation model which was not embedded in space. He found that the system showed continued evolutionary cycling of the species rather than evolution of a generalised solution (see also Juillé & Pollack (1998b)). In other non-spatial coevolutionary optimisation models additional techniques are used to ensure diversity of both antagonistic species and longevity of ‘good’ individuals (Collins & Jefferson, 1991; Paredis, 1995; Rosin & Belew, 1997; Juillé & Pollack, 1998b). The increased longevity of solutions and the ensuring of diversity of both species help to prevent evolutionary cycling. The effects of such techniques, however, are automatic side-effects of local dynamics such as occur in spatial evolutionary systems (Husbands, 1994; Mahfoud, 1995; Pagie & Hogeweg, 1997; Rosin & Belew, 1997).

We present results of a study of a spatially explicit coevolutionary model in which two species have an antagonistic interaction. We compare two cases. The first case depicts coevolution in a spatial environment in which individuals interact and compete locally with each other so that spatial pattern formation occurs and influences the local environment of the individuals and therewith the evolutionary process. The second case depicts coevolution in the same model except that the individuals of both populations are globally mixed every time step. In this case spatial pattern formation does not occur. In the first model the evolutionary process leads to individuals that have integrated adaptations to separate selection events into a general solution. We call this *information integration*. In the second model, in which the individuals are mixed, we see typical cyclic *red queen dynamics*. In both models, however, the individuals have approximately the same time-average fitness. Thus, from a biological point of view neither outcome is a priori good or bad; in both situations the individuals are well adapted to the environmental conditions which they help to shape.

5.2 The model

We study the coevolutionary process in the context of the optimisation of a computational task. Although the task is chosen rather arbitrarily it lends itself easily for embedding in a two-species system with antagonistic interactions. The genetic encoding of the task is characterised by a non-linear genotype-phenotype mapping with strong epistatic interactions. We use a individual-based, discrete space, discrete time model with synchronous updating. The general structure of the model is very similar to the structure of the models that were studied by Hillis (1990) and Pagie & Hogeweg (1997). The two

species present in the model are called CAs and ICs.

The CAs are 1-dimensional, binary state cellular automata next-state rule-tables with a neighbourhood size 3 (Wolfram, 1984; Toffoli & Margolus, 1987), the ICs are initial conditions of the cellular automata and are of length 149. Both CAs and ICs are represented as bitstrings. The interaction between a CA and an IC, and therewith the basis on which their fitness is calculated, is based on the density-classification task of cellular automata (Mitchell *et al.*, 1994). In the density classification task the CAs must classify ICs on the basis of the number of 0s and 1s in the bitstring of the IC. If the IC has a majority of zeros in its bitstring it belongs to class 0, otherwise it is class 1¹. The CA is allowed to iterate for maximally 320 time steps, starting with the IC as initial condition. If the CA settles into a homogeneous state of all zeros it classifies the IC as being of class 0. If the CA settles into a homogeneous state of all ones it classifies the IC as being of class 1. If the CA does not settle into a homogeneous state it answers "don't-know", and does not receive a fitness reward. Only if the CA classifies an IC correctly does it receive a fitness reward of 1. In all other cases the IC receives a fitness reward of 0 (see below).

This particular task for cellular automata and its evolutionary optimisation is studied extensively by the EvCA-group in the Santa Fe Institute (see Mitchell *et al.* (1996) for a review). Coevolutionary models using this task were previously studied by Paredis (1997) and Juillé & Pollack (1998b). The latter, however, used an intricate coevolutionary scheme incorporating global feedback strategies to prevent the occurrence of red queen dynamics (see also (Werfel *et al.*, 1999) for additional studies in that context). Here, we use the task of density classification primarily to study the process of coevolution between two antagonistic species. The (evolution of the) task itself is of little importance for this study although we are interested in its properties as evolutionary 'goal'. Below we will discuss some of these properties.

5.2.1 Spatial embedding and local dynamics

Individuals of both species are distributed in space which is a 2-dimensional regular grid of 30 by 30 cells with periodic boundary conditions. Each cell contains one CA and one IC, giving population sizes of 900 individuals. The CAs and ICs are evaluated with respect to each other locally in this space. The fitness of a CA is based on the ICs in its Moore adjoining, i.e. the eight cells directly neighbouring the middle cell plus the middle cell itself. The fitness of an IC is based only on the CA in the same cell. This asymmetric fitness evaluation procedure was found to improve the evolutionary optimisation process Pagie & Hogeweg (1997). The fitness evaluation scheme is characterised by a very sparse evaluation of the objective function, i.e. a general IC classification algorithm. Sparse evaluation is in fact unavoidable because the total number of ICs is 2^{149} and the total number of CAs is 2^{128} . Moreover, in (Pagie & Hogeweg, 1997) we showed that sparse fitness evaluation can help the evolutionary process rather than hinder it (see also (Hillis, 1990)). We call the fitness of CAs and ICs that they receive during fitness evaluation *local fitness*. In order to compare CAs from different populations we calculate a general fitness measure (see below) which we call *performance fitness* (Mitchell *et al.*,

¹The bitstring of the ICs have an odd length, so the majority is always defined

1994).

After fitness evaluation in each cell of the grid a selection procedure is performed between locally present CAs and between locally present ICs, and growth of the selected CA and IC in the cell. Selection is based, probabilistically, on the rank order of the nine individuals in the Moore neighbourhood. The probability for an individual to be selected is 0.5^{rank} , where $rank = 1..8$. The last ranked individual (i.e. $rank = 9$) also has a probability 0.5^8 for being selected. Note that we have constant population sizes. Although this is usual in evolutionary optimisation models it is of course less realistic from a biological point of view.

After selection and growth we apply mutations to the CAs and the ICs. We only use bit-flip mutations with rate 0.2 per CA and rate 0.5 per IC. The use of the bit-flip operator introduces a strong mutational bias, in terms of the density of bitstrings, towards density values of 0.5. The presence of this bias appears to have a large influence on the evolutionary dynamics in the context of the task that we study here (see also (Mitchell *et al.*, 1994; Paredis, 1997)). For the initial conditions this bias pushes them directly towards the phenotype phase-transition in genotype space where it is easy to be difficult (see below).

The two models that we study in this paper are as described above except that in the second model, i.e. the *mixed model*, we globally mix the individuals of both populations every time step. In the first model, i.e. the *base model*, spatial patterns can form and influence the evolutionary process (e.g. see (Boerlijst & Hogeweg, 1991; Savill & Hogeweg, 1997)).

5.2.2 Some (evolutionary) properties of the density classification task

The majority classification task has been studied extensively in the context of evolutionary optimisation models in the EvCA group at the Santa Fe Institute (Mitchell *et al.*, 1994; Crutchfield & Mitchell, 1995; Mitchell *et al.*, 1996) as a paradigm of a local computational algorithm for a global task and as a paradigm for evolutionary processes. Cellular automaton rule-tables have a very non-linear genotype-phenotype mapping; small changes in the rule-table can have small or large influences on the phenotype of the cellular automaton. In addition, for the task that we study here, many neutral paths exist in the genotype space, i.e. many rule-tables result in the same fitness value. The presence of neutral paths in a genotype-phenotype mapping influences the evolutionary process considerably by increasing the freedom of individuals to search the space of genotypes (Huynen *et al.*, 1996; Huynen, 1996; Fontana & Schuster, 1998; Van Nimwegen *et al.*, 1999). Although the task of classifying initial conditions (which are essentially bitstrings) is in itself trivial the implementation of the task in cellular automata is interesting from the point of view of embedding computations in parallel algorithms. Handwritten cellular automata rules that show reasonable performance on the density classification task have been known for some time, particularly the GKL rule. It has been proven, however, that no cellular automaton next-state rule-table exists that can correctly classify all possible initial conditions (Land & Belew, 1995).

The *performance fitness* of a cellular automaton is defined as the number of correct classifications out of 10,000 randomly created initial conditions that have an unbiased density distribution (i.e. a binomial distribution around 0.5). We use this fitness mea-

sure, or *performance fitness*, when we compare CAs of different populations. Initial conditions with a density of approximately 0.5 are the most difficult to classify because bitstring that are almost equal (e.g. differ on only one bit position) can belong to different density classes. In fact, the performance of a good cellular automaton, like for instance the GKL rule, decreases rapidly if it is evaluated on the basis of initial conditions whose density approaches 0.5 (Mitchell *et al.*, 1994; Juillé & Pollack, 1998b). A ‘good’ cellular automaton has a fitness value of about 0.8 (e.g. the GKL rule; 0.81), although cellular automata have been found recently with fitness values of up to 0.86 (Juillé & Pollack, 1998b).

As an evolutionary optimisation task evolving good cellular automata appears to be difficult; in only a small number of evolutionary runs are cellular automata found with fitness values in the same range as the fitness of the handwritten cellular automata (Mitchell *et al.*, 1996). In the evolutionary optimisation models studied by the EvCA group cellular automata evolved with respect to their performance on the basis of initial conditions which have a flat density distribution. Evolution in the context of random initial conditions only (i.e. initial conditions with a unbiased binomial density distribution) appeared to be too difficult for the first populations of cellular automata (but see also (Andre *et al.*, 1996)). An important impediment in finding good cellular automata appeared to lie in the breaking of symmetries in the strategies that cellular automata employ early in the evolutionary process; all individuals in the population handled the task in the same, asymmetric way (Mitchell *et al.*, 1994). The evolution of the density classification task generally showed the same sequence of strategies as that used by the cellular automata; default strategies (i.e. classification always the same, i.e. class 0 or class 1), with fitness typically around 0.5; block-expanding strategies, with fitness values between 0.50 and 0.65, and embedded-particle strategies with fitness values between 0.65 and 0.80. The first two strategies are asymmetric. All cellular automata that are known to perform well on the density classification task show embedded-particle strategies.

In a coevolutionary setting this same task was studied by Paredis (1997). His model is based on globally interacting and competing populations of cellular automata and initial conditions, whereas we embed the populations in space and thus have local fitness evaluation and local competition. Paredis found that the two populations showed cyclic evolutionary dynamics; the population of initial conditions was mostly homogeneous with respect to the density class to which they belonged. As a result, the cellular automata evolved such that they always classified initial conditions, irrespective of the actual state of the latter, into one density class. Once the cellular automata had converged to this behaviour the initial conditions switched to the other density class, en masse, and the cellular automata eventually followed. The cellular automata evolved in this way have a performance fitness of around 0.5. With respect to the coevolving population of initial conditions, however, the cellular automata can have very high fitness values.

An important property of the genetic coding of the initial conditions is that they can easily evolve to that part of their genotype-space where they are maximally difficult to classify (i.e. where they have a density of 0.5), and most easily evolve from one density class to the other one (i.e. by flipping as little as a single bit). The ease of evolution of the initial condition towards that part of the genotype space is enhanced by the mutational bias introduced by the point-mutation operator. The effect of the phase-transition in the phenotype of the initial conditions in the genotype-space is inherent in the coding of the




Figure 5.1: IC fitness function Φ . The fitness f an IC gets if it is not correctly classified depends on its density. As result we get stabilising selection toward minimal or maximal density values, which are the ‘easy’ ICs.

initial conditions; at the boundary a single, bitflip can change the phenotype of the initial condition into the only other possible phenotype.

Initially, we studied the model with an IC fitness function similar to the function used by Paredis, i.e. the fitness reward f that an IC receives when it is not correctly classified is equal to 1 (see also sect5.4). In this case we found red queen dynamics in the mixed model, similar to the results of Paredis (1997). We did not find evolution of general classifiers, because, it seemed, ICs could become too difficult too easily. Therefore we introduced a cost function for the ICs. The idea is that being simple is easy and therefore cheap. Being difficult, on the other hand, should be costly. We simply embedded this idea in the fitness function Φ of the ICs. If IC gets a fitness reward f the reward is dependent on its density (fig.5.1). This fitness function implements selection towards minimum (i.e. 0.0) and maximum (i.e. 1.0) density values. The actual values of f do not matter, only the symmetry of Φ around density = 0.5 and the fact that f increases monotonically when it approaches the minimum and maximum density values. In fact we simply used $\Phi(IC_i) = |(density(IC_i) - \frac{\max\ density}{2})|$.

5.3 Results

In this section we will describe the results we obtained by running the model described in the previous section. We will present our results by describing two typical simulations of the model, one simulation of the base model and one in which we apply, in addition, global mixing of the CA and IC populations. We found that the two simulations are typical for the possible outcomes of the evolutionary process in the model. The precise parameter settings do not influence the general results to a great extent. The values that we used in the simulations that we describe here were actually chosen rather arbitrarily, e.g. we did not optimise our results in any particular way. However, the two simulations discussed below are run with the same parameter values.

5.3.1 Two typical simulations

Simulations are started with randomly created CAs, i.e. CAs with a density around 0.5; the ICs have an initial density of 0.0. The first variables that we observed were the local fitness values of the CAs and of the ICs and their densities. In Fig.5.2 we show time-plots of the base model (A) and of the mixed model (B). We plot the average density of the CAs and the ICs, and the average of their local fitness values. All averages are normalised between 0.0 and 1.0, but the true ranges are given in the legend. The time-plots clearly show different dynamics in the long term. Figure 5.2A shows stabilisation of the dynamics. Figure 5.2B, on the other hand, shows continued large amplitude fluctuations of the average density values. The average local fitness value of the CAs in the mixed model is generally close to maximum but shows frequent spikes of very low fitness values.

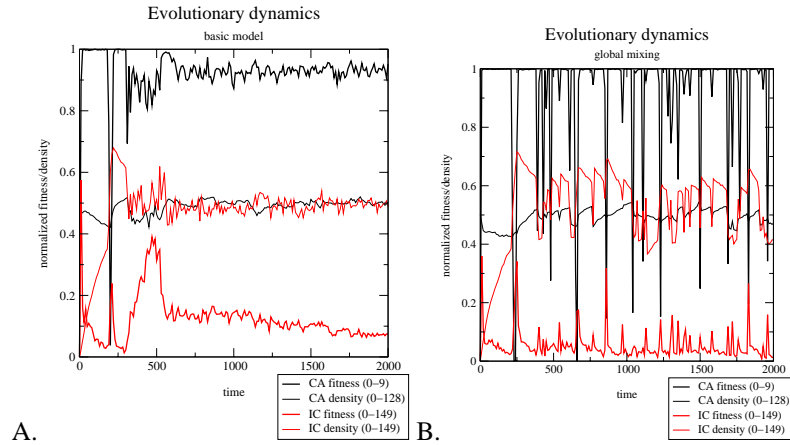


Figure 5.2: Evolutionary dynamics of basic model (A) and the model with global mixing (B). The simulations start with the same parameter values and with the same initial state.

For both simulations the initial transient shows roughly the same picture; large fluctuations of the average fitness values of the CAs and the ICs together with large fluctuations of the average densities of the CAs and the ICs. The simulations start with ICs that have a density of 0.0 which are very easy to classify correctly. Indeed, the average local fitness of the CA population quickly increases toward maximum values. As a result of mutations, ICs will arise with density values higher than 0.0. But initially these ICs are still very easy to classify correctly and the CAs maintain the high local fitness values.

The subsequent evolution of the IC population towards ICs with still higher densities increases the difficulty of the ICs. However, even when the density of the ICs approaches 0.5 the CAs in the population still classify them correctly and maintain high local fitness values. This is because the CAs simply settle into a homogeneous state of zeros independent of the state of the IC. Up to this time this strategy of the CAs in fact performs perfectly and this behaviour is easily evolved and easily maintained.

At $t \approx 200$ ICs arise that have a density larger than 0.5. Now the CAs have a problem; settling into a homogeneous state of zeros is no longer the correct behaviour. In-

deed, the average local fitness of the CAs drops to very low values. During this stage the IC population experiences strong stabilising selection as a result of which they evolve towards ICs with very large density values as a result of the IC fitness function Φ (fig.5.1). Soon after the switch in the average density of the ICs, however, we see that the average local fitness of the CA population rises again to very high values in both simulations. At this point the same general behaviour can be seen as at the beginning of the runs, except that the density of the ICs is now larger than 0.5.

From this point, the dynamics of the two simulations diverge. The mixed model continues to show fluctuations in the average density of the ICs and sharp drops in the average local fitness of the CAs for short periods of time. In the base model a different evolutionary phase unrolls. The fluctuations in the average IC density value become smaller, as do the fluctuations in the average local fitness of the CA population. The CAs, however, no longer attain maximum local fitness although they did initially, and continue to do so in the mixed model. In both models, however, the CA populations have approximately the same *local* fitness when we average over time (≈ 0.9). The IC populations do better in the base model; they have a time-averaged local fitness of 0.08 in the base model whereas in the mixed model they have a time-averaged local fitness of 0.04. The IC fitness value, however, also depends on the density values of the ICs in the population. Seen as a biological system the CAs do equally well in both models. Of course, in these models we do not take into account the population dynamics which may alter the results in this particular respect.

Although the CAs in the base model seem to classify correctly a large number of locally available ICs, because the average local fitness is very high, this does not mean that the CAs are general classifiers, i.e. that they can classify a set of randomly generated initial conditions correctly. It is possible that CAs and ICs are distributed locally such that CAs perform well only with respect to the locally present ICs. Below we will compare the two simulations in terms of performance fitness and will see that the CAs in the base model evolve such that they become good classifiers in a general sense rather than only in a local sense.

Of course, from the point of view of optimisation of density classification the most important variable is the performance fitness. In fig.5.3 we plot the evolution of the performance fitness of the best CA in the population in the base model (solid line) and of the best CA in the population of the mixed model (dashed line). The performance fitness of the best individual in the mixed model fluctuates between 0.50 and 0.55. Even the best CAs in this model do not classify random initial conditions much more accurately than random classification into class 0 or class 1.

The performance fitness of the best CA in the base model initially increases and then fluctuates between 0.70 and 0.75. These values for performance fitness of the CAs are in the same range as the performance fitness values for the best cellular automata found in the evolutionary optimisation models studied by Mitchell *et al.* (1994), Crutchfield & Mitchell (1995), and Paredis (1997). Clearly, they are much more general than the CAs from the mixed model. Following the concepts of Crutchfield & Mitchell (1995), and Hordijk *et al.* (1998), the CAs use particle-based strategies in order to compute the density of ICs, as does, for instance, the rule GKL.

In the base model we see that the CAs evolve a generalised classification algorithm whereas in the mixed model the performance fitness ≈ 0.5 . Next, we will further describe

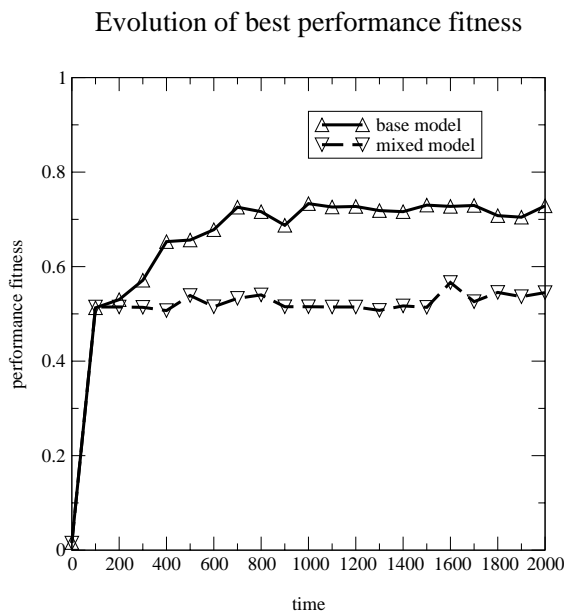


Figure 5.3: Evolution of absolute fitness of the best individuals in the population in the base model (solid line) and the model with global mixing (dashed line). Whereas the CAs in the base model show an increase in the fitness of the best individual the best CA in the globally mixed model remains around 0.55.

the dynamics in the two models, and show that red queen dynamics dominate in the mixed model whereas in the base model information integration occurs which leads to general density classifiers.

5.3.2 Information integration and red queen dynamics

In order to understand how CAs in the mixed model can attain near maximum local fitness values although their performance fitness is only 0.5, we look at the distribution of ICs in the population.

Figure 5.4 shows the distribution of the densities of all ICs in the mixed model (A) and the base model (B) between $t=2100$ and $t=2200$ and the average local fitness of the CAs. In fig.5.4A the population of ICs switches back and forth between high and low density values. At $t=2100$ the IC population has just switched from an average density value larger than 0.5 to one smaller than 0.5. The density distribution at this point is very narrow and rapidly decreases to lower values. All CAs still classify the ICs incorrectly, as can be seen from the average local fitness. The ICs, therefore, experience only the selection pressure imposed by the fitness function Φ (fig.5.1). As soon as CAs arise that classify the ICs correctly, here $t \approx 2110$, the density distribution of the IC population starts to broaden considerably. This is due to the combined effects of a large reduction in selection pressure towards low density values, plus the effect of the mutational bias

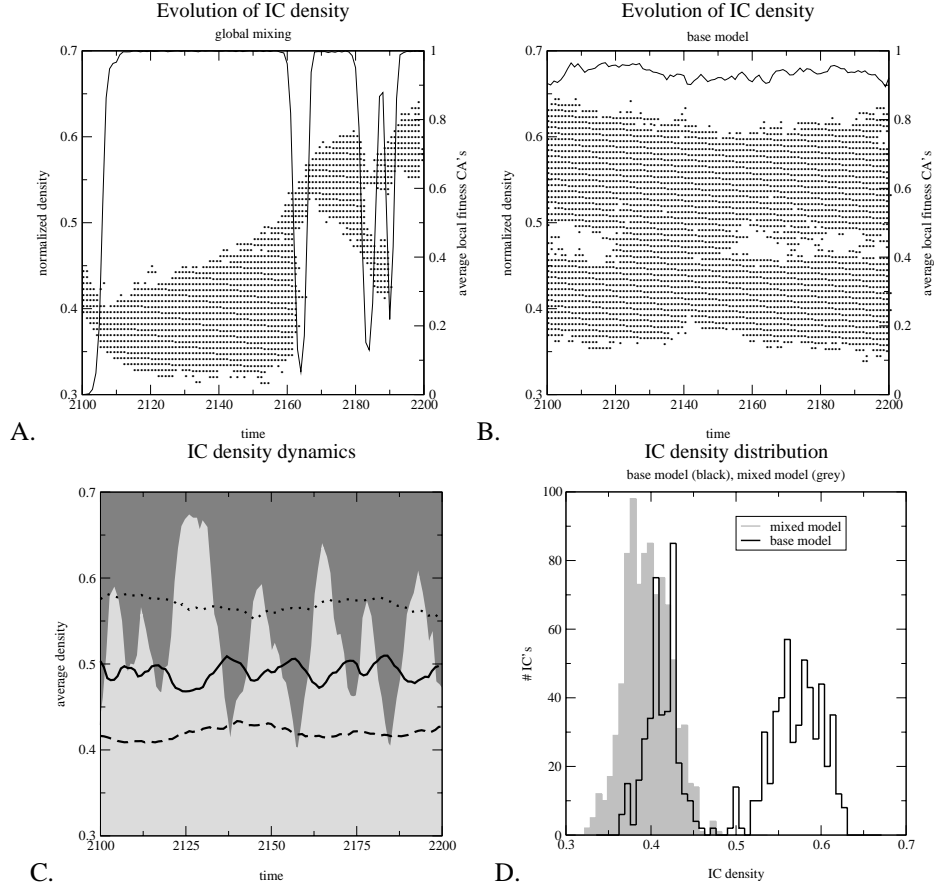


Figure 5.4: Evolution of IC density over 100 time steps in mixed model (A), and base model (B). The densities of all individuals are plotted as dots, the average local fitness values are plotted as thin lines. The population of ICs in the mixed model switches en masse from their density class, also characterised by short drops in average CA fitness. In the base model two subpopulations of ICs exist. In (C) the IC population of the base model is split into ICs with a density lower than 0.5 (density class 0; light-grey) and ICs with a density higher than 0.5 (density class 1; dark-grey). The black solid line denotes the IC density averaged over all ICs, the dashed line denotes the average IC density over all ICs of density class 0, the dotted line denotes the average IC density over all ICs of density class 1. Whereas the ICs show a relatively stable distribution per subpopulation, the size fluctuations of the subpopulations bring about the fluctuations in the average IC density of the whole population. In (D) the density distributions of ICs of the base model (black solid line) and of the mixed model (grey) at $t=2150$ are plotted.

Figure 5.5: CA phenotype distribution: histograms of classifications made by CAs on a set of initial conditions with an unbiased density distribution (black) and initial conditions with a flat density distribution (grey). (A): CA population from mixed model, (B): CA population from base model. From lower to upper panel: number of correct classifications (panel 1), number of classifications into density class 0 (panel 2), number of classifications into density class 1 (panel 3), number of classifications into class ‘undefined’ (panel 4). CAs from the mixed model classify all initial conditions ‘single-mindedly’ into density class 0, whereas CAs from base model classify initial conditions to one of the two density classes, and often correctly.

toward medium density values resulting from the bitflip mutation operator. As soon as ICs arise with a density larger than 0.5 the IC population jumps from class 0 to class 1 en masse and the same picture is seen again. Thus, what we might have expected from the global dynamics depicted in fig.5.2 is in fact what happens; the population of ICs switches back and forth between density values below 0.5 and density values above 0.5. At $t \approx 2080$ we see that only a small subpopulation of ICs switches from its density-class, resulting in a temporary coexistence of the two density-classes in the IC population. This coexistence does not last long however.

In order to get an indication of the behaviour of the CAs in the mixed model we plot in fig.5.5 the behaviour of the better half of a single generation of CAs based on two large sets of random initial conditions. The first set (black line) is a set of initial conditions with a binomial density distribution around 0.5. The second set is a set of initial conditions with a flat density distribution ranging from [0.0 .. 1.0] (grey). This set mainly consists of relatively ‘easy’ initial conditions. The different panels are histograms of CAs that in panel 1: correctly classify ‘x’ initial conditions; in panel 2: classify ‘x’ initial conditions as class 0; in panel 3: classify ‘x’ initial conditions as class 1; in panel 4: classify ‘x’ initial conditions as “don’t-know”.


Figure 5.5A shows the histograms for CAs from the mixed model at $t=2000$. Panel 1 shows that the CAs in the mixed model classify only about half of the initial conditions correctly whether they are easy or difficult. When we look at the behaviour of the CAs we see that they classify nearly all initial conditions as belonging to class 0 (panel 2), i.e. they are absolutely *single-minded*. Almost none of the initial conditions are classified as belonging to class 1 (panel 3), or as undefined (panel 4). Given the current state of the population of ICs ($t = 2000$) this behaviour of the CAs is very sensible. We found that in the mixed model if the CA population has near maximum average fitness the CAs are *single-minded* in their classification of initial conditions.

Now it is clear how CAs in the mixed model can have near maximum local fitness as well as having a performance fitness which is similar to the performance fitness of a randomly classifying cellular automaton; in the context of a homogeneous population of ICs (in terms of their density) a single-minded strategy is very successful. In the context of a diverse set of initial conditions, however, this strategy does not perform better than a random one.

In fig.5.4B we see that in the base model the IC population has speciated into two distinct subpopulations of ICs, with densities around 0.4 and 0.6, which stably coexist.

CA phenotype distribution


mixed model



A.

CA phenotype distribution

base model



B.

In fig.5.4C we have split the IC population into these two subpopulations and have plotted the average density of the total IC population (black line), the average density of the two IC subpopulations (dotted and dashed lines) and the sizes of the two subpopulations (light- and dark-grey surfaces). The fluctuations of the average density of the total IC population is caused mainly by the population fluctuations of the two subpopulations, rather than by fluctuations in the density distribution within each subpopulation. There is also a fluctuation in the average density per subpopulation but on a much larger time scale than that of the density fluctuations of the total population and with much smaller amplitude. Although fig.5.4B suggests that the two subpopulations merge at $t \approx 2040$ fig.5.4D shows that the average densities of the two subpopulations remain very far apart (see also fig. 5.8B).

Figure 5.5B shows the CA phenotype distributions in the base model, at $t=2000$. Panel 1 shows that initial conditions with a density close to 0.5 are classified correctly in about 70% of the cases. With respect to the initial conditions with the flat density distribution, however, the CAs classify them as being nearly perfect. If we look at the behaviour of the CAs in terms of the number of times that they classify an initial condition as belonging to class 0 (panel 2) or to class 1 (panel3) we see that the behaviour of the CAs is centred around 50% for both classes. Thus, the CAs classify initial conditions into one of the classes (panels 2 & 3) and often the correct one (panel 1), they hardly ever “don’t-know” (panel 4).

In the base model we see that the CAs evolve toward general classifiers; they get high performance fitness values. Below we will go into some of the causes and effects of the process of information integration which leads to this outcome (sect.5.3.3). In the mixed model we see that the evolutionary dynamics do not stabilise; the IC population continues to oscillate between ICs with a density lower than 0.5 and ICs with a density higher than 0.5. At every such switch the local fitness of the CAs drops to very low values but the CAs quickly recover by changing their “single-minded” behaviour in accordance with the state of the IC population. This is typical for red queen behaviour; both species evolve towards a state that is beneficial given the state of the other population, which puts the other population in a bad state again.

The switching of the CA population in the mixed model from one density class to the other can be seen in fig.5.6. There we plot how long CA ancestries remain in the population in the base model (lower panel) and the mixed model (upper panel). At $t = 2000$ we assign a unique number to all individual CAs. During the subsequent evolution all offspring get the same number as their parent. Now we can track the descendants of the ancestors at $t = 2000$ over a period of 100 time steps. In addition we plot the average local fitness values of the CA populations. In the base model (lower panel) we see that ancestries constantly disappear at a relatively high rate; the selection pressure is strong. In the mixed model (upper panel), on the other hand, we see that initially the rate at which ancestries are lost is much lower than in the base model; initially the selection pressure is relatively low. But at $t \approx 2050$ the IC population switches from one density class to the other density class. This event acts as an evolutionary bottleneck; suddenly all ancestries except one die out. This loss of ancestors occurs at the moment when the average local fitness of the CAs starts to rise again, i.e. when a single CA found the correct strategy and outcompetes all other CAs.

The relatively low rate of loss of ancestries initially in the mixed model is caused

Persistence of CA ancestries




Figure 5.6: Persistence of ancestries in base model (lower panel) and mixed model (upper panel). The average local fitness of the CA population is drawn as well. In the base model ancestries die out exponentially, in the mixed model ancestries die very quickly after birth or when the population goes through a bottleneck, e.g. at $t \approx 2070$.

by the homogeneity of the IC population (in terms of density class) and the particular classification strategy that all CAs use. While the ICs remain in their homogeneous state there is no selection pressure on the CAs to do better simply because they cannot do any better. As we will show below, in the base model the ICs succeed in exerting a constant selection pressure on the CAs to evolve with respect to the current state of the IC population. In fig.5.6 we already see the consequence of this: a high loss of ancestries.

5.3.3 Eco-evolutionary side-effects

In the previous section we showed that in the base model general density classification algorithms evolve. In the mixed model the local fitness of the CAs is generally near maximum but frequently drops to very low values for short periods of time. The only difference between the two models is the possibility of spatial pattern formation in the base model and the absence of this possibility in the globally mixed model. In this section we study some side-effects of the eco-evolutionary dynamics in the two models. We will show that ICs in the base model challenge the CAs in a variety of ways, thereby exerting a differentiated selection pressure on the CAs. In the mixed model, however, both populations take the easy way out. Both exploit the weaknesses of the other but thereby make themselves easily exploitable.

Coevolving with easy and difficult IC populations

In both models the CAs are evaluated on the basis of ICs that show a broad density distribution around values well above or below 0.5 (fig.5.4). Thus, the CAs are rarely evaluated on the basis of the initial conditions that are most difficult, i.e. that have a density ≈ 0.5 , whereas performance of ‘good’ cellular automata in the density classification task deteriorates rapidly if the density of the initial condition approaches 0.5 (Mitchell *et al.*, 1994; Juillé & Pollack, 1998b).




Figure 5.7: Fitness of 3 CAs with respect to different IC populations taken from the simulation at $t=2000 \dots 3000$, every 100th time step. The CAs evolved at $t=2500$. The fitness is calculated with respect to the ICs taken from the simulation (solid lines) and with respect to the same set of ICs after the ICs were randomised (dashed lines). The CAs perform worse with respect to coevolved ICs than with respect to random ICs with the same density. The same statistics are plotted for the rule GKL. Thus, the ICs are specialised with respect to the CAs.

The ICs in the base model appear to be less difficult to classify than ICs with density values around 0.5, because they appear in two subpopulations that have average density values around 0.4 and around 0.6 (fig.5.4C). However, the coevolved ICs turn out to be more difficult than was expected on the basis of their densities. In fig.5.7 we plot the fitness of three CAs with respect to eleven IC populations of different generations. The CAs are taken from the population at $t=2500$, the IC populations are from $t=2000$ to $t=3000$, every 100th time step. The three CAs are chosen because they will give rise to all individuals after $t=3000$. However, they are not the best individuals at this point in terms of performance fitness. The best individual at this point has a performance fitness of 0.77 but it dies out before $t = 2600$. As a comparison we also plot the same statistics of the GKL rule. The solid lines denote the fitness of the CAs with respect to the evolved ICs. The dashed lines denotes the fitness of the CAs with respect to the same set of ICs

but after the bitstring of each IC has been shuffled, which corresponds to random ICs with the same density.

We see that the CAs perform significantly better on the basis of random initial conditions than on coevolved ICs of the same density. This is not the result of a particular evolution of the CAs, as can be seen by considering the statistics of the rule GKL. The IC population exploits not only the density-dimension in its coevolution with the CAs, it also exploits a ‘difficulty-dimension’ which is independent of the density of the ICs.

By coevolving with the CAs ICs can explore on a small scale different bitstring configurations that make life difficult for the CAs that have to cope with a global property of the ICs. For instance, ICs can evolve long stretches of zeros in the bitstring while keeping the total density higher than 0.5. This will typically mislead CAs that use block-expanding strategies which are functional precisely because they expand large homogeneous blocks of ones or zeros in the initial bitstring.

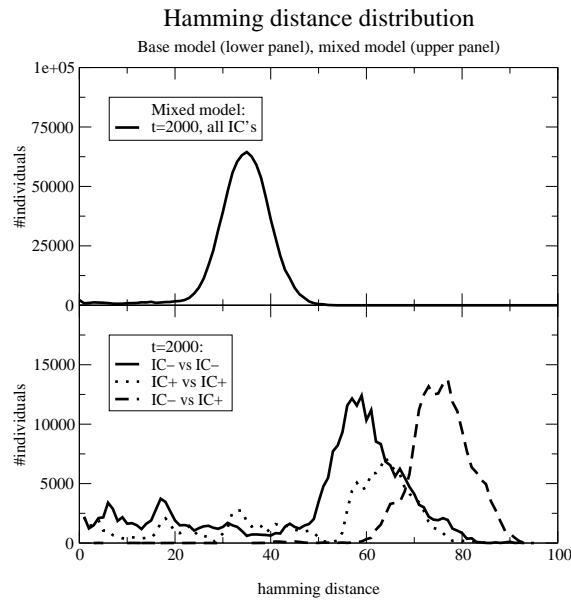


Figure 5.8: Hamming distance between ICs of base model (lower panel) and between ICs of mixed model (upper panel), at generation $t=2000$. Lower panel: hamming distances are calculated between all ICs with a density lower than 0.5 (solid line), all ICs with a density higher than 0.5 (dotted line), and between ICs of different subpopulations (dashed line). Upper panel: hamming distances are calculated between all ICs of the population. ICs of the base model within a subpopulation have a hamming distance that is almost as large as ICs of different subpopulations and much larger than that of ICs of the mixed model.

In addition to the evolved difficulty of the ICs in the base model the ICs are also very diverse at the population level. In fig.5.8 we have plotted the hamming distance of ICs within and between subpopulations of different density-classes. Whereas the hamming

distance between ICs of different generations and of different subpopulations within one generation is expected to be relatively large, the hamming distance of ICs of the same subpopulation of one generation is also very large. In fact, the distance between the latter ICs peaks near the distance which is expected between two random ICs, i.e. 75 bits. Thus, in the base model the IC population maintains a very large diversity. The CAs are consequently evaluated with respect to very different ICs. In the top panel we have plotted the hamming distance of all individuals in a population of ICs in the mixed model at $t = 2150$ (see also fig.5.4B). Clearly, the ICs are less diverse in the mixed model than in the base model. This is all the more surprising if we look at fig.5.4D in which we plot a histogram of the density distribution of the CA populations used in fig. 5.8. The CAs of the base model are clearly split into two subpopulations, whereas the CAs of the mixed model constitute a single population. The latter, however, form a density distribution which is relatively broad, compared with the two distributions of the CAs of the base model.

ICs in the mixed model do not show the effect of ‘extra’ difficulty. CAs evaluated on the basis of evolved ICs have approximately the same fitness as when they are evaluated on the basis of shuffled ICs. Also, ICs evolved in the mixed model do not show the high diversity of hamming distances between individual ICs (fig.5.8). In the mixed model CAs and ICs seem to coevolve only on the basis of the density of the ICs. If all ICs in a population are classified correctly by all CAs because all ICs are of the same density class and all CAs ‘single-mindedly’ classify ICs always as belonging to that class, the only evolutionary way out for the ICs is to switch its density class. After the switch, whereas all CAs classify all ICs incorrectly, the only thing the ICs can do to increase their fitness is to evolve further away from the density = 0.5 region, until CAs arise that immediately classify all ICs correctly. In the mixed model a take-over of the population by a newly arisen individual that has high fitness will occur on a short time scale due to the global mixing. The time scale on which populations can evolve from one density class, or single-mindedness type, to the other is much longer. Populations do not get the opportunity to retain information from previous adaptations and individuals are not selected with respect to a diverse environment; taking the easy way out is always a good strategy.

Spatial and temporal distribution of ICs

The question that arises is how the base model can maintain the diversity of the ICs and the large hamming distances between individuals of the same density classes. Figure 5.9A shows five snapshots of consecutive time steps of the spatial distribution of the ICs of different density classes. We used different shades of red to colour the ICs of class 0, and different shades of green to colour the ICs of class 1. The ICs are distributed in many small patches rather than in only a few large patches. In fig.5.9B we show a space-time plot of the ICs over a period of 180 time steps in which we plot a vertical cross-section of the grid at consecutive time steps. The space-time plot shows that complex wave patterns are present; patches of red ICs grow into patches of green ICs, and vice versa. As a consequence, at any one point in space ICs of the two density classes alternate frequently. This alternation of the two density classes is not primarily a result of mutation, which causes the global oscillations of the average IC density in the mixed




Figure 5.9: See also colour plate 4 (page 66). A: Snapshots of spatial distribution of the IC population at 5 consecutive time steps in the base model. ICs of density class 0 are different shades of red, ICs of density class 1 are different shades of green. B: Space-time plot of IC population over 180 time steps, time going from left to right. ICs of different density classes are distributed in complex wave patterns which overtake each other continuously.

model, but it is a result of spatial dynamics. In the base model, ICs ‘chase’ CAs not only in genotypes-space, as in the mixed model, in addition they ‘chase’ them in space-space. As a result of these spatial dynamics in the base model individual CAs, or CA-lineages, “see”, i.e. are evaluated on the basis of, the whole spectrum of IC density classes. In fig.5.10 we plot the distribution of the number of ICs of density class 1 on which a CA is evaluated per time step over a period of 51 time steps. The three CAs for which we have

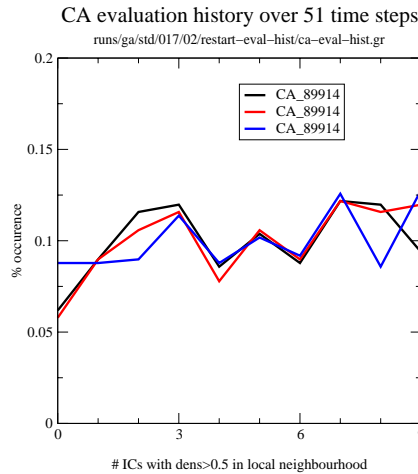


Figure 5.10: See also colour plate 4 (page 66). The evaluation history of three CA-lineages over 51 time steps. CAs are evaluated 0..9 IC of class 1 per time step; over 51 time steps the evaluation-history of CAs is approximately a flat distribution of occurrences of ICs of class 1. The final CAs of the lineages are the same as those used in fig.5.7.

plotted the evaluation history are actually ancestry-lineages that give rise to the three CAs of fig.5.7. Per time step the number of class 1 ICs a CA "sees" can vary between 0 and 9. In the mixed model the CAs are, per time step, almost always exclusively evaluated on the basis of ICs of only one density class. Thus this distribution will peak at low values and high values in the mixed model. Figure 5.10 shows that over time in the base model CAs are evaluated on the basis of an approximately flat distribution of ICs density classes rather than on the basis of a peaked distribution. CAs are evaluated, repeatedly, on the basis of ICs of different density classes. Moreover, they "see" ICs with a mix of density classes as often as they "see" one density class exclusively. Of course, this diversity in evaluation helps the CAs in evolving general classification algorithms.

Mutational stability

If CAs in the base model can evolve general density classification algorithms, why don't the CAs in the mixed model do the same? Because it is too difficult to do it right and because it is too easy to do it differently. The difficulty of evolving good classification algorithms in cellular automata is clear from the results given above. The ease of evolving different algorithms can be understood if we consider the mutational stability of CAs in the mixed model.

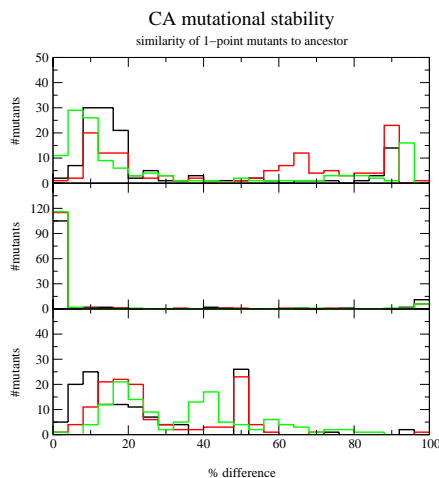


Figure 5.11: See also colour plate 4 (page 66). Mutational stability of three CAs in the base model (panel 1), the mixed model just before a switch in the average IC density (panel 2), the mixed model just after a switch in the average IC density (panel 3).

Previous studies have already shown that in coevolutionary systems the mutational sensitivity of individuals may evolve such that they can quickly evolve from or towards a new evolutionary 'goal' (Huynen & Hogeweg, 1994; Pagie & Hogeweg, 1997). Here we can study the mutational stability of a cellular automaton by comparing its behaviour to the behaviour of all, i.e. 128, its one-point mutants. We compare the behaviour of two cellular automata (the mutant and its 'ancestor') by evaluating them both on the basis

of the same set of 100 random initial conditions with an unbiased density distribution. ‘Similarity’ is defined as the ratio of this set of initial conditions that the two cellular automata classify the same, i.e. as being of class 0, class 1, or as “don’t-know”. In fig.5.11 we plot a histogram based on the difference (1.0 - similarity) between the CA and its mutants. Note that we do not consider the correctness of the classification in this case.

We compare the mutational stability of CAs from the base model (panel 1), CAs from the mixed model just before a switch of the IC population (panel 2), and CAs from the mixed model just after a switch of the IC population (panel 3). The most striking feature in the plot is the difference between the two sets of ICs from the mixed model. Before the switch almost all CAs are absolutely stable; all 1-point mutants show the same behaviour as the original CA. After the switch the stability is much reduced. The effect that we see here corresponds the persistence of ancestries that we already found in fig.5.6. Under neutral selection, i.e. just before a switch, evolution is expected to push individuals toward flatter parts of the genotype landscape. This corresponds to what we see here. Immediately after a switch the selection pressure is much higher and individuals are expected to be much less stable.

However, the most important feature of the ICs just before the switch is that a small number of 1-point mutants are 100% different. Most of these mutants have become CAs that do not settle in a homogeneous state anymore, i.e. they classify ICs always as ‘undefined’. However, some of the one-point mutants that are 100% different are, again, ‘single-minded’, but now with respect to the other density class. They are mutant CAs that switch from one ‘single-mindedness’ to the other. The two types of ‘single-mindedness’ appear to be easily accessible from one another in genotype-space. Thus, following the red queen scenario, in the mixed model the CAs evolve ‘single-minded’ behaviour, but also evolve such that they can easily change their mind.

CAs from the base model are about as stable as CAs in the mixed model just after the switch. Here, also, the CAs are subjected to a relatively high selection pressure. In contrast to the mixed model many 1-point mutants exist that are 50% different rather than 100%. Many of these mutants that show different behaviour on 50% of the evaluations in fact are also “single-minded” CAs. Apparently, “single-minded” CAs are also easily accessible from good CAs. We do not know to what extent “single-minded” CAs are distributed over the entire space of possible cellular automata, but we do know that randomly created cellular automata rarely show sensible behaviour; they classify most initial conditions as ‘undefined’.

In figure 5.12 we have plotted the behaviour of a CA from the mixed model and a 1-point mutant CA with respect to three initial conditions of different density. The difference in the CAs leads to a switch of the direction of the particle that expands a block of ones in the ancestor CA such that the block of ones shrinks in the mutant CA. As a consequence, the ancestor CA will classify initial conditions as belonging to class 1 while the mutant CA will classify the same initial conditions as belonging to class 0.

The mutational stability results of the CAs in the mixed and in the base model show that it is easy for CAs in the mixed model to switch between single-minded behaviours. Also for CAs in the base model it is easy to switch to single-minded behaviour. Thus, if the population of ICs oscillates between states that are homogeneous in terms of density class it is not only beneficial for the CAs to employ a single-minded strategy, it is very




Figure 5.12: Space-time plots of 1-point mutants evaluated on the basis of three initial conditions. In the top row the ancestor CA is evaluated on the basis of initial conditions with densities of 0.44, 0.54, and 0.62 resp. The bottom row is the 1-point mutant CA evaluated on the basis of the same initial conditions.

easy to do so.

5.4 Discussion

In general, individuals can respond in several ways to a evolutionary selection pressures imposed by their environment. In a diverse environment a population may speciate, producing several species that are specialised with respect to only a part of the environmental ‘potential’. In contrast, individuals may evolve a generalised behaviour that is well adapted to all environmental circumstances. Finally, individuals and the environment can specialise with respect to each other, potentially resulting in a continued adaptive change of the individuals and the environment.

Whereas evolution is generally considered to lead to one of the aforementioned outcomes only, in the previous section we showed that they can all occur in the same evolutionary system. We have seen that evolution in the context of spatial pattern formation leads to the evolution of generalised individuals whereas in the context of global mixing the evolutionary process enters a continuous oscillatory regime. Both evolutionary processes occur with an equal basic ‘interaction structure’, i.e. a population of cellular automata and a population of initial conditions that have an antagonistic interaction. Also the basis of the interaction, i.e. the density classification of initial conditions by the cellular automata is equal, as is the genetic encoding, the genetic operators, etc. But before we conclude we briefly compare our results with those obtained from earlier work.

5.4.1 Specialists and generalists; neither ‘outfits’ a queen

The optimisation task that we use in this model to drive the evolutionary process is studied before in a non-spatial coevolutionary model (Paredis, 1997; Juillé & Pollack, 1998a). From those studies it appeared that the initial conditions easily evolve toward a part of the genotype space where they are most difficult to classify correctly by the cellular automata. This evolution is brought about by selection pressure as well as directed mutational drift introduced by the genetic operators, most notably the point mutation operator. The resulting evolutionary process shows typical red queen behaviour; the initial conditions switch from density class as soon as the cellular automata evolve such that they start to classify the initial conditions correctly (Paredis, 1997). Paredis suggested that the long periods of neutral drift that cellular automata experience when they classify every initial condition incorrect results in the loss of previous adaptations, hence, they evolve only ‘single-minded’ behaviour.

We use a fitness function for the initial conditions that gives high fitness values to initial conditions that have extreme density values (i.e. close to 0.0 or 1.0) and low fitness values to initial conditions with medium fitness values (fig. 5.1). The rationale behind this fitness function is that ‘being difficult’ comes with a cost, e.g. due to the development of more intricate mechanisms to attack the host, or to circumvent its immune system. This fitness function introduces a selection pressure toward initial conditions with extreme density values in addition to the selection pressure due to the coevolution with the cellular automata. As we have seen in the previous section the interplay between the two selection pressures brings about a number of patterns that are characteristic for the evolutionary dynamics as they occur in the two models, e.g. the dynamics of the density distribution of the initial conditions (fig. 5.4).

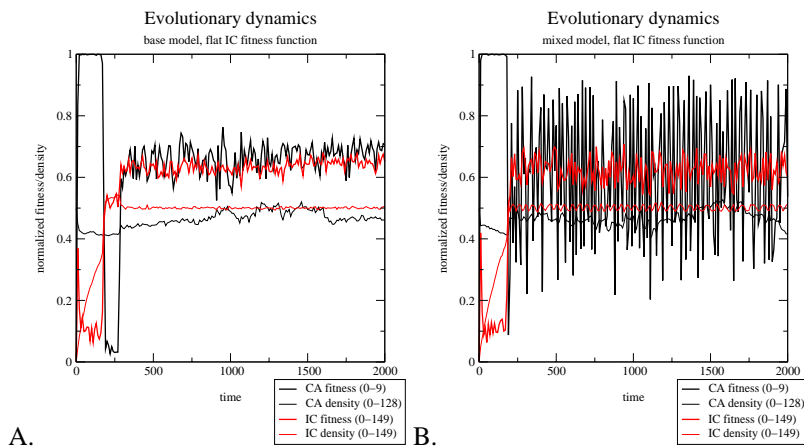


Figure 5.13: Evolutionary dynamics with a flat IC fitness function of basic model (A) and the model with global mixing (B). Compare to fig. 5.2. Here, the base model shows speciation of CAs. The mixed model shows red queen dynamics again, but on a very short time scale. Again, the time-average local fitness of the CAs is equal in both models (≈ 0.65).

Initially, we studied our model with a flat fitness function instead of the peaked function of fig.5.1. The flat fitness function corresponds to the fitness function used by Paredis (1997) and Juillé & Pollack (1998a). Using the flat fitness function we do not find evolution of general density classifying cellular automata, either in the base model or in the mixed model. Instead, in our model also the initial conditions evolve directly towards that part of the genotype space where they have medium density values, i.e. are most difficult to classify correctly and most easily mutate back and forth between the two phenotypes. In the well-mixed model we find red queen dynamics again, as in the original model, although the evolutionary dynamics are on much shorter time scales (fig. 5.13.B).

Using the flat fitness function in the base model we find comparable stabilising evolutionary dynamics as in the original model (fig. 5.13.A), but now stabilisation is caused not by the evolution of a general ‘solution’ but rather by the evolution of many ‘specialists’. Some of the specialised CAs are single-minded, as in the well-mixed model. Other specialised CAs classify initial conditions seemingly randomly into density classes 0 or 1. In fig. 5.14 we plot the distribution of the behaviour of such specialised cellular automata with respect to two sets of random initial conditions with an unbiased density distribution (black solid line) and a flat density distribution (grey), as in fig. 5.5. We took a single generation of cellular automata and filtered all single-minded ones out, i.e. the cellular automata that classified over 95% of the initial conditions into one density class. The remaining CAs constituted approximately half of the population.

The specialised nature of the cellular automata can best be seen by considering the behaviour of the cellular automata with respect to the first set of initial conditions (black). The cellular automata classify approximately 50% of the initial conditions cor-

rectly, thus, similar to ‘single-minded’ CAs they perform no better than random classifiers. But if we look at the distribution of their behaviour in terms of classification into density class 0 (panel 2), into density class 1 (panel 3), or into the class ‘undefined’ (which, as in the models with the peaked fitness function, almost never happens; panel 4) they do not behave in a ‘single-minded’ manner. Many CAs classify initial conditions into both density classes according to a particular ratio, some even close to a fifty-fifty ratio, although the total number of ‘correct’ classifications does not exceed the 50%; they err with respect to initial conditions of both density classes. It is not clear yet how this form of specialised behaviour comes about. It is clear that it is a viable strategy; approximately half of the population of CAs follows this strategy.

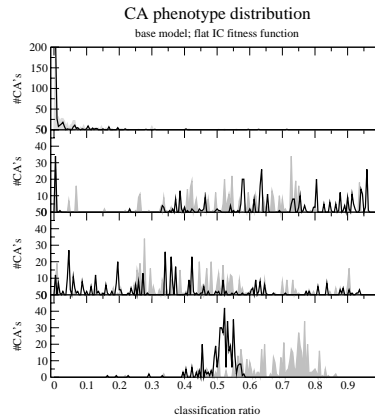


Figure 5.14: Phenotype distribution of ‘specialised’ CAs from the base model with a flat IC fitness function after all ‘single-minded’ CAs have been filtered out. Black line: CA phenotype distribution with respect to initial conditions with unbiased density distribution; Grey area: CA phenotype distribution with respect to initial conditions with a flat density distribution. The CAs have performance of approximately 0.5 but classify initial conditions. Compare with fig. 5.5B.

In the base model using the peaked IC fitness function the time-averaged local fitness of the CAs is 0.9, but this value is approximately 0.65 when we use the flat fitness function in the base model. Surprisingly, also in the mixed model with the flat fitness function the time-averaged local fitness of the CAs is approximately 0.65. For both fitness functions, the peaked function and the flat function, we find very different evolutionary dynamics in the base model, i.e. evolution of generalists and specialists respectively, and in the mixed model, i.e. red queen dynamics in both cases. Nevertheless, the lifetime fecundity of individual CAs, averaged over time, is equal in both models. That the lifetime fecundity of CAs is much lower in the models with a flat fitness function is the result of the much higher difficulty of the ICs that are present in the population; they are much more centred around density values of 0.5. It is unclear why the local fitness of the CAs is almost independent of the spatial context in which they evolve and whether this result holds in other systems. Clearly, this deserves further study as well.

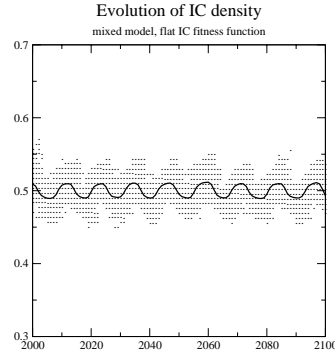


Figure 5.15: Evolution of IC density over 100 time steps in mixed model (A) with a flat IC fitness function. ICs fluctuate around density = 0.5. Compare to fig. 5.4A. The average of the local fitness is plotted in the solid line.

5.4.2 Red queen dynamics; quick, quick, slow

Evolutionary red queen dynamics are characterised by two populations that show continuous change in terms of their behaviour with respect to each other. Normally, it is taken that the populations mutate such as to change their behaviour. However, when this change is of oscillatory nature, rather than being similar to a ‘run away’ process, the continuous ‘change’ can also result from population dynamics. When a ‘coevolving’ population is composed of two subpopulations, each with one of the two behavioural types, fluctuations in the subpopulations as response to changes in the other population will exhibit dynamics similar to red queen dynamics. Figure 5.15 shows the evolution of the density distribution of the ICs in the mixed model with the flat IC fitness function. In contrast to fig. 5.4A, here we see that the density distribution oscillates very close around the value 0.5. In fact, at all times ICs of both density classes are present. However, mutational change also occurs in this model; ICs with density values further from the value 0.5 come and go with the oscillations of the density distributions. For the ICs in this model both processes occur; behavioural oscillatory dynamics due to ecological dynamics and due to evolutionary dynamics. However, if we continue a simulation with the IC mutation rate set at 0.0 the evolutionary dynamics quickly collapse due to homogenisation of the IC population; mutation is required to maintain the dynamics. Also the CAs in this model are mainly single-minded in their behaviour with respect to the ICs present in the population. And, again, these CAs can easily mutate between the two types of ‘single-mindedness’. Contrary to the mixed model with the peaked IC fitness function, when we use the flat fitness function the CA population does not always converge to one type of ‘single-mindedness’. Also with respect to the CAs in this model a mixture of ecological and mutational factors seems to occur, although it is not clear to what extent each plays a role in the oscillations. This shows that although the seemingly easy strategy for adapting to the oscillating IC population, i.e. the strategy based on ecological dynamics, is possible, and even present, the seemingly more difficult strategy, i.e. the mutationally based strategy, evolves as well. Moreover, neither strategy takes

over completely.

In the studies by Paredis (1997) similar combinations of mutational and ecological origins of red queen behaviour seemed to be present in the IC population. The CA population, however, was fully dependent on the mutational strategy. In one of the studies by Juillé & Pollack (1998*a*), in which they also found red queen like behaviour, both CA and IC populations seem to follow the ecological strategy. In our mixed model with the peaked IC fitness function both populations rely completely on the mutational strategy. In that case clearly the stabilising selection pressure pushes the ICs away from the density=0.5 value, results in a homogeneous IC populations. In our model the CAs not only mutate back and forth in order to adjust to the current state of the population of ICs, they have evolved so that they can do this by means of a very small number of point mutations. The latter feature enables the CAs to switch from phenotype very quickly, whereas in the model of Paredis it took a very long time for CAs to switch. It is not clear why we find easily switching CAs while Paredis does not.

5.4.3 Conclusion

We have studied a coevolutionary model of two antagonistically interacting species. We compared the evolutionary dynamics that occur if individuals remain localised in space, i.e. when spatial pattern generation occurs, and the evolutionary dynamics that occur if individuals are globally mixed every time step. In the first case we find that individuals evolve a generalised response to environmental circumstances, whereas in the second case the systems exhibit evolutionary oscillatory dynamics. In that case we see the evolution of much simpler behaviour, which is optimised with respect to one of the possible states of the other species. This strategy makes them easily exploitable, however. As a result we see red queen dynamics where both coevolving species oscillate between two states.

If we remove the cost for one species of being difficult the main effect is that the environment of the other species becomes more difficult; general strategies used by individuals of the latter species fail more often in that case. As a consequence these individuals change their strategy from a general one to a specialised one in which they can cope with only a few opponents. As a consequence speciation occurs in the population; different individuals specialise on different niches which are defined by the other species. We found a similar speciation process occurring due to changes in cost functions in (Pagie & Hogeweg, 1999*a,b*). In the mixed model we also find red queen dynamics under high costs. Now, individuals optimise their queenyness.

6

Discussion

This thesis describes a number of model-studies of evolutionary systems. All models are characterised by the concurrent evolution of a number of different types of replicators. Thus, in one form or another, the systems comprise coevolutionary processes. The genotypes of the replicators show different degrees of structure, ranging from simple but variably sized genotypes to genotypes composed of open-ended, functional structures with a non-linear genotype-phenotype mapping. Also, all model-studies are embedded in a spatial setting and incorporate different modes of local or global mixing.

During the studies described in the previous chapters we developed some concepts we considered to be important for understanding the evolutionary processes that occurred in the model-systems. In this final chapter of the thesis these concepts are discussed again and reviewed in the broader context of the thesis. First, however, we will give a short overview of the studies in which the different concepts were developed and thereafter we will discuss the concepts more thoroughly.

6.1 A short review

Chapter two is a model-study of coevolution in the context of evolutionary optimisation and introduces the concept of sparse fitness evaluation. In the model, the evolutionary process is driven by an externally defined, artificial fitness function. We compared the evolutionary dynamics resulting from a fitness evaluation procedure based on a full set of ‘fitness cases’ to fitness evaluation based on a small, coevolving set of ‘fitness cases’. The latter evaluation procedure is sparse in the sense that individual solutions “see” only a small proportion of all possible fitness cases during their lifetime. In the first evaluation procedure all individuals “see” all cases at every fitness evaluation; evaluation is said to be ‘full’. With finding complete solutions as ‘goal’ sparse fitness evaluation may not seem recommendable, this study, however, shows it to be effective in evolving complete solutions.

We use the term information integration to designate the accumulation of adaptations to the environment. This study shows that information integration can occur over many generations, and even lead to a complete solution, if the environment to which individuals adapt is presented only in a sparse manner. In fact, the evolved solutions comply better to the externally defined evolutionary goal if evolution occurs under sparse fitness evaluation than when it occurs under full fitness evaluation. In addition, a number of side-effects result from different forms of fitness evaluation, e.g. high phenotypic gen-

eralisability and low mutational stability in sparse evaluated individuals compare to full evaluated individuals. When the size, or complexity, of the evolutionary goal becomes larger sparse evaluation does not lead to a complete solution if evaluation is based on random samples of the ‘fitness cases’. Rather, a coevolutionary coupling between individuals and their environment, i.e. the solutions and the fitness cases, is necessary for successful information integration in these cases.

In chapter three we describe a model of the interaction between bacteria and colicins and study how the diversity of colicins evolves. Colicins are gene complexes that code for a toxin and a unique, corresponding antidote. Bacteria are killed by the toxins if they do not carry the proper antidotes. As a consequence, colicinogenic bacteria can invade sensitive bacterium populations. Here, it are the bacteria that ‘coevolve’ with the (ensembles of) colicins. In contrast to the previous model-system, here we employ a more realistic population dynamical view; population sizes are not fixed and thus populations can die out. On the other hand, we use a much simpler genetic structure to model the colicins; each colicin encodes two genes that are either active or inactive. On the level of the bacteria the colicins act as genes that can be acquired independently, i.e. a variable-sized bacterium genome.

We find that under all circumstances a high diversity of colicins evolves easily but this diversity can be expressed on one of two levels; the individual bacterium level or on the bacterium population level. In the first case individuals bacteria carry all antidote gene types that are present in the system but only a very limited number of active toxin genes. In the second case bacteria carry only complete colicin complexes, i.e. both toxin gene and corresponding antidote gene are active. Now, although the global diversity of colicin types is equal in both modes, the number of colicins per bacterium is much less than the total number present in the population. On which level the diversity of colicins is expressed depends on the growth-cost per antidote that is imposed on the bacteria, and the number of colicin types in the system. Here, we find that information integration can occur on different levels, the choice of which has a large impact on the population dynamics and further evolutionary dynamics.

In the fourth chapter a somewhat similar model is studied; the eco-evolutionary dynamics of restriction-modification (RM) systems in a bacterium-phage community. In this model we also use realistic population dynamics and simple genetics, as in chapter three. The bacteria and phages constitute a host-parasite system. The RM systems can serve as a defence mechanism for the bacteria that protects them to phage infections. Phage, however, can become insensitive to specific RM systems through acquisition of modification patterns. The system shows similarity to the colicin system; phage can be interpreted as ‘toxins’ and RM systems act as ‘antidote’, although the former can become ‘immune’ here as well. Also, contrary to the two genes of a colicin complex in the colicin system the phage and the bacteria act on independent levels in this system. Now, an environment on the bacterium level is defined in terms of the presence of phages, plus their modification patterns. But also an environment on the level of the phage is defined, based on the presence of bacteria and the RM systems that they carry. As a consequence, information integration now occurs on two levels rather than one as in the colicin-bacterium model. Again, the model is characterised by an individual-based mode and a population-based mode, but now as two attractors of a bi-stable system.

The relation between the individual-, and population-based modes and the sparse or

full evaluation of the environment is reversed in the two modes in this model compared to the model of the bacterium-colicin system. In the individual-based mode both bacterium and phage populations are homogeneous; both are fully equipped with RM systems, respectively modification patterns. As a result, bacteria experience full evaluation in terms of phage and their modifications and the phages experience full evaluation in terms of bacteria and the RM systems. In the population-based mode, on the other hand, the bacteria carry only a minimal number of RM systems and the phages are fully sensitive or even completely absent; in this mode evaluation is sparse on both levels.

In chapter five a coevolutionary process is studied between cellular automata and initial conditions. In this model we use again an artificial fitness function, i.e. the density classification task, to drive the evolutionary process, as in chapter two as well as structured genotypes and fixed population sizes. In a strict sense fitness evaluation in this model is necessarily sparse; the total number of initial conditions, i.e. the complete set of fitness cases, is 2^{149} . In a local sense, however, evolutionary dynamics can enforce sparse fitness evaluation as well as full fitness evaluation on the individuals. We compare evolutionary dynamics that occurs in the context of spatial pattern formation to evolutionary dynamics occurring under global mixing of the populations. In this system we find optimisation, speciation, and red queen dynamics all in the same system for small changes of the parameters.

Although a complete solution is not possible in the context of this artificial fitness function, we find that under spatial pattern formation the cellular automata evolve toward generalised solutions, i.e. they optimise individual solutions. On the population level, the cellular automata and initial conditions show speciation, each to different extents. For instance, speciation of initial conditions occurs on at least two levels. Two main subpopulations of initial conditions evolve with densities below or above 0.5. Within these main ‘phylogonies’ many disparate taxons exist, with similar density values but far apart in genotype space. When we apply global mixing the system shows typical red queen dynamics; both populations show continued oscillations between two phenotypical states. As a result, the populations are mostly homogeneous and speciation, and also optimisation of the density classification task does not occur. Nevertheless, we find information integration to occur such that individuals can more easily chase the opponent population in the evolutionary red queen race.

6.2 Sparse fitness evaluation

Fitness evaluation in natural systems, i.e. selection events over a lifetime, often conveys only a very limited amount of information concerning all possible environmental contingencies. Thus, evolution seems to have little information available for devising competent responses. Despite the sparseness of fitness evaluation of the environment in evolution, we find beauty, complexity and diversity in nature, everywhere. Apparently, the natural evolutionary process is not hindered by the sparseness of the fitness evaluation. In fact, studies in this thesis suggest that it is *thanks to* sparse fitness evaluation that evolution can integrate information.

In chapter two, the coevolutionary optimisation model, the extent to which fitness evaluation is sparse or full is imposed on the individuals and the resulting evolutionary

dynamics of the two cases are compared. In the other models the extent to which evaluation is sparse or full is a result of the population and evolutionary dynamics. Both chapters two and five show that sparse evaluation, i.e. sparse presentation of the environment during fitness evaluation, does not stand in the way of the evolution of solutions that can cope with the *complete* environmental repertoire in a generalised manner. Also in chapter three we see that in the individual-based mode bacteria integrate and maintain all antidote genes although they hardly ever ‘see’ the various toxins. Below we will consider other possible outcomes of such an evolutionary ‘quest’, e.g. speciation or red queen dynamics, in more detail. Also, in chapter two, solutions that evolve under sparse fitness evaluation are shown to be more capable of generalising over novel environmental circumstances. Solutions that evolve under full fitness evaluation become specialised on the fitness cases presented, but show highly aberrant behaviour on novel cases. In chapters three, four, and five sparse fitness evaluation is found to occur both in a sparse manner as well as in a full manner, depending on the eco-evolutionary dynamics into which the system settles. Also, in the bacterium-colicin model sparse evaluation occurs in the individual-based mode, whereas in the RM model sparse evaluation occurs in the population-based mode; similar eco-evolutionary dynamics can result in different manners of evaluation. The density-classification model, chapter five, shows sparse and full evaluation at different levels simultaneously; in terms of density classes evaluation is fuller than in terms of evaluation of single ‘taxonomies’.

In the models that we studied in this thesis fitness evaluation can be sparse or full in a single system, depending on the parameter values or on behavioural regimes. In chapter five both modes of evaluation occur simultaneously. The extent to which fitness evaluation in natural evolutionary processes is sparse or full varies as well. The dichotomy between diffuse and pairwise coevolution shows the range of fitness evaluation in natural systems. Pairwise coevolution occurs for instance in some host-parasite interactions and some mutualisms. Such interactions can evolve such that participants always experience only one partner. In that case fitness evaluation may become relatively full. Diffuse coevolution on the other hand is exemplified by the co-occurrence of a variety of different interactions. If we compare diffuse coevolution to pairwise coevolution, we see the evolution of specialised organisms mainly in pairwise coevolution (e.g. (Rothstein, 1990)), whereas diffuse coevolution seems more often to lead to generalised organisms (e.g. (Boyes *et al.*, 1996)). In addition, sparse evaluation may trigger the evolution of complex systems. For instance, Huynen *et al.* (1999) have studied the tricarboxylic acid cycle (TCA) in organisms for which the complete genomes are published. The study shows that in different species the cycle is present in highly variable forms but is often incomplete. The observed variation shows that parts of the TCA have an independent functionality on their own, thereby suggesting how the TCA might have evolved. The complete metabolic pathway may be the result of the integration of different ecological and evolutionary circumstances rather than a solution to one, single ‘problem’.

6.3 Speciation, optimisation, and red queen dynamics

In chapter two we study evolution in the context of evolutionary optimisation. There, we consider speciation as the ‘mere’ occurrence of a variety of non-general, partial solutions and therefore we do not study speciation as such in chapter two. However, the results concerning the phenotypic generalisability suggest that speciation occurs more easily under full fitness evaluation than under sparse evaluation. If we were to make two sets of fitness cases that partially overlap and subsequently let two subpopulations evolve in the context of these two sets of cases they might speciate in terms of their respective behaviour in each others ‘habitat’. Individuals that undergo full evaluation evolve specialised behaviour in their own habitat and aberrant behaviour in unknown cases, i.e. the habitat of the second species. Sparsely evaluated individuals, on the other hand, may show phenotypically very similar behaviour because they generalise well over the cases they have not “seen” but that are part of the other’s set of fitness cases; speciation would not then occur.

Chapters three and four show how information integration can result in two dynamical modes: the individual-based mode and the population-based mode. The individual-based mode is characterised mainly by the homogeneity of the population; no speciation occurs in this case. The population-based mode, on the other hand, is characterised by the occurrence of many subpopulations, each with the same ‘genetic’ information content. The latter occurs easily, despite genetic information exchange between individuals through horizontal plasmid transfer, which plays a similar role to sex. The different species that occur in the system are, however, only a small subset of all possible species. Thus, speciation as such is not imperative, but population-based diversity is imperative. Speciation and the presence of a large number of species are ‘merely’ side-effects.

Chapter five, which describes the coevolution between cellular automata and initial conditions in the context of the density classification task, shows the full range of optimisation, speciation, and red queen behaviour. The latter only occurs under global mixing of the system. In that case, the populations are homogeneous and switch back and forth between two phenotypes. In all other cases the density classification task, as defined by the initial condition present at that moment, is solved either by the evolution of cellular automata that implement generalised solutions or by speciation (or niche differentiation) in the population of cellular automata such that different specialised ‘species’ *solve* different subsets of the initial conditions that are present in the population. In fact, the evolution of optimal solutions concurs with the presence of a large number of different species in the population of initial conditions and in the population of cellular automata. Speciation at this level seems to help the evolutionary quest for optimal solutions by giving rise to continued competition between alternative solutions.

6.4 Information integration

Chapter two concerns the process of information integration under sparse and full fitness evaluation. In the latter case, information integration can be interpreted as a type of tuning process; at every selection event individuals that conform most to the ‘evolutionary goal’ are selected. Any mutation that is detrimental with respect to any one of

the ‘fitness cases’ is eliminated. Under full evaluation, if we assume that the ‘complete’ environment does not change over time, evolution can be interpreted as occurring on a fitness landscape (Wright, 1967; Kauffman & Levin, 1987). A fitness landscape is a mapping from genotypes to fitness values and is generally depicted in a 2-dimensional graph, in which mutationally close genotypes lie close together on the genotype axes. In this landscape individuals evolve uphill, towards genotypes with higher fitness values (Calvin, 1986). When individuals “see” only a small number of all possible fitness cases the fitness landscape becomes dynamic. Genotypes that were previously located on hills or in valleys in the fitness landscape may now change their character and thereby change from being ‘attractive’ to being ‘repulsive’, and vice versa. As a result, the evolutionary dynamics may enter a regime of continuous change, i.e. red queen dynamics, in which previously acquired adaptations in individuals can get lost while these individuals adapt to new circumstances. Now, ‘information integration’ seems less likely to occur.

In chapter two, however, we show that sparse fitness evaluation can lead to complete information integration at the individual level and may even be beneficial rather than detrimental to the integration process. In addition, information integration has side-effects that depend on the particular form of the integration process. In chapter five, which concerns the coevolution between cellular automata and initial conditions, we actually show that complete information integration at the individual level can occur, i.e. optimal *individual* solutions are evolved in an evolutionary system in which other outcomes of a coevolutionary process occur as well, i.e. red queen dynamics and speciation. In fact, the system shows evaluation to be most sparse in the case where information integration leads to general solutions at the individual level.

Successful information integration, i.e. the accumulation of adaptations to the environment, is generally considered to occur at the individual level, as described in the previous paragraph. This thesis shows that information integration can occur at different levels in the same system, e.g. the individual level and the population level in chapters three and four. Chapter three, i.e. the bacterium-colicin system, shows that information integration can occur in an individual-based and a population-based mode. Although at the individual level bacteria have much less information in the population-based mode than in the individual-based mode, at the population level the number of colicin types in the two modes is the same. Note that whereas above the population-based mode was interpreted as speciation, here it is interpreted as information integration at the population level as an alternative to integration at the individual level.

Chapter four again shows the occurrence of an individual-based mode and a population-based mode, but now as a bi-stable attractor. Individual-based information integration occurs simultaneously in the bacterium and the phage population or in neither population, in which case both populations showed either population-based integration or the phages became extinct. Here, the systems can settle into the different modes of information integration not as a result of a change in parameter values but as a result of the population dynamics. Whereas ‘optimisation’ is often considered at the individual level only, chapters three and four show that information integration at the population level may actually be more beneficial for individuals. In chapter three bacteria in the population-based mode were less vulnerable to invasions by novel colicin types. In the RM model the phage population was much less viable when the bacteria were also in the population-based mode. Thus, individuals may do better by giving up individual ability.

However, information integration can occur at other levels as well. In chapter five we saw that cellular automata do not integrate a complete solution directly in the phenotype but they make it easily *accessible* at the level of the genotype via a very small number of mutations. In the same chapter initial conditions show an evolution towards ‘being difficult’, not only at the level of the density value but also at the level of gene-structures. Similarly, by evolving to different parts of a genotype landscape (Huynen & Hogeweg, 1994), or to different parts of an interaction graph (Hogeweg, 1994) individuals and populations or communities can change their local evolutionary surroundings so as to cope with environmental challenges in alternative ways.

6.5 Conclusion

Information integration in evolutionary processes may at first seem to be centred around the question: “Did we end up with the optimal individual?”. This question has been posed before and has attracted much research. This thesis shows that the question posed as such is a meagre one, covering only a small part of the concept. We have shown that information integration does not occur only at the individual level, it can also occur at the level of the genotype or at the level of the population. We have also shown that information integration can occur under seemingly informationally poor conditions. In fact, the latter regime may favour generalised integration of information. Finally, we showed that different forms of information integration, resulting from sparse fitness evaluation or from full fitness evaluation, have different side-effects. Given that natural evolution tends to build and rebuild on the basis on what evolved before, different side-effects at this stage will influence the whereabouts of eco-evolutionary dynamics later.

Bibliography

- ABBOTT, I.; ABBOTT, L. K. & GRANT, P. R. (1977); Comparative Ecology of Galapagos Ground Finches (*Geospiza* Gould): Evaluation of the Importance of Floristic Diversity and Interspecific Competition. *Ecol. Monogr.* **47**(2):151–184.
- ACHTMAN, M.; MERCER, A.; KUSECEK, B.; POHL, A.; HEUZENROEDER, M.; AARONSON, W.; SUTTON, A. & SILVER, R. P. (1983); Six widespread bacterial clones among *Escherichia coli* K1 isolates. *Infect. Immun.* **39**:315–335.
- ALTENBERG, L. (1994); The evolution of evolvability in genetic programming. In: Kinnear, Jr., K. E. (ed.), *Advances in Genetic Programming*, MIT Press, chap. 3. pp. 47–74.
- ANDRE, D.; BENNETT III, F. H. & KOZA, J. R. (1996); Evolution of intricate long-distance communication signals in cellular automata using genetic programming. In: *Artificial Life V: Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*. MIT Press, Nara, Japan, vol. 1.
- ANGELINE, P. J. & POLLACK, J. B. (1993); Competitive environments evolve better solutions for complex tasks. In: Forrest, S. (ed.), *Proceedings of the 5th International Conference on Genetic Algorithms*. Morgan Kaufmann, San Mateo, CA, pp. 264–270, pp. 264–270.
- BARCUS, V. A. & MURRAY, N. E. (1995); Barriers to recombination: restriction. In: Baumberg, S.; W., Y. J. P.; R., S. S. & H., W. E. M. (eds.), *Population genetics of bacteria*, Cambridge U.P, Cambridge, U.K. pp. 31–58.
- BARRETT, J. A. (1983); Plant-fungus symbioses. In: Futuyma, D. J. & Slatkin, M. (eds.), *Coevolution*, Sinauer Associates, Inc., Sunderland, MA. pp. 137–160.
- BERGSTROM, C. T.; MCELHANY, P. & REAL, L. A. (1999); Transmission bottlenecks as determinants of virulence in rapidly evolving pathogens. *Proc. Natl. Acad. Sci. U.S.A.* **96**:5095–5100.
- BICKLE, T. A. & KRUGER, D. H. (1993); Biology of DNA restriction. *Microbiol. Rev.* **57**:434–450.
- BLANEY, C. & MILLER, J. A. (1995); Moths escape from evolutionary dead end. *Bio-science* **45**(11):741–742.
- BOERLIJST, M. C. & HOGEWEG, P. (1991); Selfstructuring and selection: spiral waves as a substrate for evolution. In: Langton, C. G. (ed.), *Artificial Life*. Addison-Wesley, Redwood City, CA, vol. 2, pp. 255–276, pp. 255–276.
- BOGLER, D. J.; NEFF, J. L. & SIMPSON, B. B. (1995); Multiple origins of the yucca-yucca moth association. *Proc. Natl. Acad. Sci. U.S.A.* **92**:6864–6867.

-
- BOUMA, J. E. & LENSKE, R. E. (1988); Evolution of a bacteria/plasmid association. *Nature* **335**:351–352.
- BOYES, D. C.; MCDOWELL, J. M. & DANGL, J. L. (1996); Plant pathology: many roads lead to resistance. *Curr. Biol.* **6**:634–637.
- BULLOCK, S. G. (1995); Co-evolutionary design: Implications for evolutionary robotics. Cognitive Science Research Paper 384, School of Cognitive and Computing Sciences, University of Sussex, Brighton.
- CALVIN, W. H. (1986); *The river that flows uphill*. Sierra Club Books, San Francisco.
- CAUGANT, D. A.; LEVIN, B. R. & SELANDER, R. K. (1981); Genetic diversity and temporal variation in the E. coli population of a human host. *Genetics*. **98**:467–490.
- CHAO, L. & LEVIN, B. R. (1981); Structured habitats and the evolution of anticompetitor toxins in bacteria. *Proc. Natl. Acad. Sci. U.S.A.* **78**:6324–6328.
- COLLINS, R. J. & JEFFERSON, D. R. (1991); Selection in massively parallel genetic algorithms. In: Belew, Richard K.; Booker, L. B. (ed.), *Proceedings of the 4th International Conference on Genetic Algorithms*. Morgan Kaufmann, San Diego, CA, pp. 249–256, pp. 249–256.
- COOPER, P. C. & JAMES, R. (1984); Two new E colicins, E8 and E9, produced by a strain of Escherichia coli. *J. Gen. Microbiol.* **130**:209–215.
- CORNAX, R.; MORINIGO, M. A.; GONZALEZ-JAEN, F.; ALONSO, M. C. & BORREGO, J. J. (1994); Bacteriophages presence in human faeces of healthy subjects and patients with gastrointestinal disturbances. *Zentralbl. Bakteriologie*. **281**:214–224.
- COUWENBERG, J. (1997); *The fate of cheaters on ecological and evolutionary timescales*. Master's thesis, Dep. Theor. Biol./Bioinformatics, Utrecht.
- CRUTCHFIELD, J. P. & MITCHELL, M. (1995); The evolution of emergent computation. *Proceedings of the National Academy of Sciences* **92**:10742–10746.
- DARWIN, C. (1859); *On the origin of species*. John Murray, London.
- DAWKINS, R. & KREBS, J. R. (1979); Arms races between and within species. *Proc. R. Soc. Lond. B. Biol. Sci.* **205**:489–511.
- DIECKMANN, U. & DOEBELI, M. (1999); On the origin of species by sympatric speciation [see comments]. *Nature* **400**:354–357.
- DIECKMANN, U.; MARROW, P. & LAW, R. (1995); Evolutionary cycling in predator-prey interactions: population dynamics and the red queen. *J. Theor. Biol.* **176**:91–102.
- DURRETT, R. & LEVIN, S. (1997); Allelopathy in Spatially Distributed Populations. *J. Theor. Biol.* **185**:165–171.
- DYKES, G. A. & HASTINGS, J. W. (1997); Selection and fitness in bacteriocin-producing bacteria. *Proc. R. Soc. Lond. B. Biol. Sci.* **264**:683–687.

-
- EBERT, D. (1998); Experimental evolution of parasites. *Science* **282**:1432–1435.
- EHRlich, P. R. & RAVEN, P. H. (1964); Butterflies and plants: a study in coevolution. *Evolution* **18**:586–608.
- EIGEN, M.; MCCASKILL, J. & SCHUSTER, P. (1989); The molecular quasi-species. *Adv Chem Phys* **75**:149–263.
- FALK, P. G.; HOOPER, L. V.; MIDTVEDT, T. & GORDON, J. I. (1998); Creating and maintaining the gastrointestinal ecosystem: what we know and need to know from gnotobiology. *Microbiol. Mol. Biol. Rev.* **62**:1157–1170.
- FARRELL, B. & MITTER, C. (1992); Diversification at the insect-plant interface. *Bio-science* **42**(1):34–42.
- FEINSINGER, P. (1983); Coevolution and pollination. In: Futuyma, D. J. & Slatkin, M. (eds.), *Coevolution*, Sinauer Associates, Inc., Sunderland, MA. pp. 282–310.
- FELDGARDEN, M.; GOLDEN, S.; WILSON, H. & RILEY, M. A. (1995); Can phage defence maintain colicin plasmids in *Escherichia coli*? *Microbiology*. **141**:2977–2984.
- FICICI, S. G. & POLLACK, J. B. (1998); Challenges in coevolutionary learning: Arms-race dynamics, open-endedness, and mediocre stable states. In: Adami, C.; Belew, R. K.; Kitano, H. & Taylor, C. (eds.), *Proceedings of the 6th International Conference on Artificial Life (ALIFE-98)*. MIT Press, Cambridge, MA, USA, pp. 238–247, pp. 238–247.
- FONTANA, W. & SCHUSTER, P. (1998); Continuity in evolution: on the nature of transitions. *Science* **280**:1451–1455.
- FRANK, S. A. (1994a); Polymorphism of bacterial restriction-modification systems: the advantage of diversity. *Evolution* **48**(5):1470–1477.
- FRANK, S. A. (1994b); Spatial polymorphism of bacteriocins and other allelopathic traits. *Evol. Ecol.* **8**:369–386.
- FRY, J. D. (1990); Trade-offs in fitness on different hosts: evidence from a selection experiment with a phytophagous mite. *Am. Nat.* **136**(5):569–580.
- FUHRMAN, J. A. (1999); Marine viruses and their biogeochemical and ecological effects. *Nature* **399**:541–548.
- FUTUYMA, D. J. & SLATKIN, M. (eds.) (1983a); *Coevolution*. Sinauer Associates, Inc., Sunderland, MA.
- FUTUYMA, D. J. & SLATKIN, M. (eds.) (1983b); Introduction. In: Futuyma, D. J. & Slatkin, M. (eds.), *Coevolution*, Sinauer Associates, Inc., Sunderland, MA. pp. 1–13.
- GAVRILETS, S. (1997); Coevolutionary chase on exploiter-victim systems with polygenic characters. *J. Theor. Biol.* **186**:527–534.

-
- GOLDBERG, D. E. (1989); *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading. Addison-Wesley, Massachusetts.
- HANTULA, J.; KURKI, A.; VUORIRANTA, P. & BAMFORD, D. H. (1991); Ecology of bacteriophages infecting activated sludge bacteria. *Appl. Environ. Microbiol.* **57**:2147–2151.
- HARDY, K. G. (1975); Colicinogeny and related phenomena. *Bacteriol. Rev.* **39**:464–515.
- HARTL, D. L. & DYKHUIZEN, D. E. (1984); The population genetics of *Escherichia coli*. *Annu. Rev. Genet.* **18**:31–68.
- HAVELAAR, A. H.; FURUSE, K. & HOGEBOM, W. M. (1986); Bacteriophages and indicator bacteria in human and animal faeces. *J. Appl. Bacteriol.* **60**:255–262.
- HERRE, E. A. (1993); Population structure and the evolution of virulence in nematode parasites of fig wasps. *Science* **259**:1442–1445.
- HIGHTOWER, R. R.; FORREST, S. & PERELSON, A. S. (1995); The evolution of emergent organization in immune system gene libraries. In: Eshelman, L. J. (ed.), *Proceedings of the 6th International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers, San Francisco, pp. 344–350, pp. 344–350.
- HILLIS, D. W. (1990); Co-evolving parasites improve simulated evolution in an optimization procedure. *Physica D* **42**:228–234.
- HILLIS, W. D. (1992); Co-evolving parasites improve simulated evolution as an optimization procedure. In: Langton, C. G.; Taylor, C.; Farmer, J. D. & Rasmussen, S. (eds.), *Artificial Life II*, Addison-Wesley, Santa Fe Institute, New Mexico, USA, vol. X of *Santa Fe Institute Studies in the Sciences of Complexity*. pp. 313–324.
- HOGEWEG, P. (1994); Multilevel evolution: replicators and the evolution of diversity. *Physica D* **75**:275–291.
- HOGEWEG, P. (1998); On searching generic properties in non-generic phenomena: an approach to bioinformatic theory formation. In: Adami, C.; Belew, R.; Kitano, H. & Taylor, C. E. (eds.), *Artificial Life VI*. MIT press, pp. 285–294, pp. 285–294.
- HOGEWEG, P. & HESPER, B. (1991); Evolution as pattern processing: Todo as substrate for evolution. In: Meyer, J. & Wilson, S. W. (eds.), *From Animals to Animats*. MIT Press, pp. 492–497, pp. 492–497.
- HOLLAND, J. H. (1975); *Adaptation in natural artificial systems*. University of Michigan Press, Ann Arbor.
- HOLLAND, J. H. (1992); *Adaption in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT Press.
- HORDIJK, W.; CRUTCHFIELD, J. P. & MITCHELL, M. (1998); Mechanisms of emergent computation in cellular automata. In: *PPSN-V*.

-
- HUSBANDS, P. (1994); Distributed coevolutionary genetic algorithms for multi-criteria and multi-constraint optimisation. In: Fogarty, T. (ed.), *Evolutionary Computing, AISB Workshop Selected Papers*. Springer-Verlag, vol. 865 of *Lecture Notes in Computer Science*, pp. 150–165, pp. 150–165.
- HUYNEN, M. A. (1996); Exploring phenotype space through neutral evolution. *J. Mol. Evol.* **43**:165–169.
- HUYNEN, M. A.; DANDEKAR, T. & BORK, P. (1999); Variation and evolution of the citric-acid cycle: a genomic perspective. *Trends Microbiol.* **7**:281–291.
- HUYNEN, M. A. & HOGEWEG, P. (1994); Pattern generation in molecular evolution: exploitation of the variation in RNA landscapes. *J. Mol. Evol.* **39**:71–79.
- HUYNEN, M. A.; KONINGS, D. A. & HOGEWEG, P. (1993); Multiple coding and the evolutionary properties of RNA secondary structure. *J. Theor. Biol.* **165**:251–267.
- HUYNEN, M. A.; STADLER, P. F. & FONTANA, W. (1996); Smoothness within ruggedness: the role of neutrality in adaptation. *Proc. Natl. Acad. Sci. U.S.A.* **93**:397–401.
- IWAO, K. & RAUSHER, M. D. (1997); Evolution of plant resistance to multiple herbivores: quantifying diffuse coevolution. *Am. Nat.* **149**(2):316–335.
- IWASA, Y.; NAKAMARU, M. & LEVIN, S. A. (in press); Allelopathy of bacteria in a lattice population: Competition between colicin-sensitive and colicin-producing strains. *Evol. Ecol.* .
- JAMES, R.; KLEANTHOUS, C. & MOORE, G. R. (1996); The biology of E colicins: paradigms and paradoxes. *Microbiology.* **142**:1569–1580.
- JANZEN, D. H. (1980); When is it coevolution? *Evolution* **34**:611–612.
- JOHANSEN, A. (1996); A simple model of recurrent epidemics. *J. Theor. Biol.* **178**:45–51.
- JUILLÉ, H. & POLLACK, J. B. (1996); Co-evolving intertwined spirals. In: Fogel, L. J.; Angeline, P. J. & Baeck, T. (eds.), *Evolutionary Programming V: Proceedings of the Fifth Annual Conference on Evolutionary Programming*. MIT Press, pp. 461–467, pp. 461–467.
- JUILLÉ, H. & POLLACK, J. B. (1998a); Coevolutionary learning: a case study. In: *Proc. 15th International Conf. on Machine Learning*. Morgan Kaufmann, San Francisco, CA, pp. 251–259, pp. 251–259.
- JUILLÉ, H. & POLLACK, J. B. (1998b); Coevolving the ideal trainer: Application to the discovery of cellular automata rules. In: Koza, J. R.; Banzhaf, W.; Chellapilla, K.; Deb, K.; Dorigo, M.; Fogel, D. B.; Garzon, M. H.; Goldberg, D. E.; Iba, H. & Riolo, R. (eds.), *Genetic Programming 1998: Proceedings of the Third Annual Conference*. Morgan Kaufmann, University of Wisconsin, Madison, Wisconsin, USA, pp. 519–527, pp. 519–527.

-
- KAUFFMAN, S. & LEVIN, S. (1987); Towards a general theory of adaptive walks on rugged landscapes. *J. Theor. Biol.* **128**:11–45.
- KAUFFMAN, S. A. (1993); *The origins of order : self-organization and selection in evolution*. Oxford U.P., New York.
- KONISKY, J. (1982); Colicins and other bacteriocins with established modes of action. *Annu. Rev. Microbiol.* **36**:125–144.
- KORONA, R.; KORONA, B. & LEVIN, B. R. (1993); Sensitivity of naturally occurring coliphages to type I and type II restriction and modification. *J. Gen. Microbiol.* **139**:1283–1290.
- KORONA, R. & LEVIN, B. R. (1993); Phage-mediated selection and the evolution and maintenance of restriction-modification. *Evolution* **47**(2):556–575.
- KOZA, J. R. (1992); *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA.
- KRAAIJEVELD, A. R.; VAN ALPHEN, J. J. & GODFRAY, H. C. (1998); The coevolution of host resistance and parasitoid virulence. *Parasitology*. **116**:0.
- KRUGER, D. H. & BICKLE, T. A. (1983); Bacteriophage survival: multiple mechanisms for avoiding the deoxyribonucleic acid restriction systems of their hosts. *Microbiol. Rev.* **47**:345–360.
- KULAKAUSKAS, S.; LUBYS, A. & EHRLICH, S. D. (1995); DNA restriction-modification systems mediate plasmid maintenance. *J. Bacteriol.* **177**:3451–3454.
- LAND, M. & BELEW, R. K. (1995); No perfect two-state cellular automata for density classification exists. *Physical Review Letters* **74**(25):5148–5150.
- LAU, P. C.; ROWSOME, R. W.; WATSON, R. J. & VISENTIN, L. P. (1984); The immunity genes of colicins E2 and E8 are closely related. *Biosci. Rep.* **4**:565–572.
- LENSKI, R. E. & MAY, R. M. (1994); The evolution of virulence in parasites and pathogens: reconciliation between two competing hypotheses. *J. Theor. Biol.* **169**(3):253–265.
- LEVIN, B. R. (1988); Frequency-dependent selection in bacterial populations. *Philos. Trans. R. Soc. Lond. B. Biol. Sci.* **319**:459–472.
- LEVIN, S. A. & SEGEL, L. A. (1982); Models of the influence of predation on aspect diversity in prey populations. *J. Math. Biol.* **14**:253–284.
- LURIA, S. & HUMAN, M. (1952); A nonhereditary, host-induced variation of bacterial viruses. *J. Bacteriol.* **64**:557–569.
- LURIA, S. K. & SUIT, J. L. (1987); Colicins and col plasmids. In: Neidhardt, F. C. (ed.), *Escherichia Coli and Salmonella Typhimurium: cellular and molecular biology*, American Society for Microbiology, Washington, D.C., vol. 2, chap. 102. pp. 1615–1624.

-
- MAHFOUD, S. W. (1995); A comparison of parallel and sequential niching methods. In: Eshelman, L. J. (ed.), *Proceedings of the 6th International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers, San Francisco, pp. 136–143, pp. 136–143.
- MAY, R. M. & ANDERSON, R. M. (1990); Parasite-host coevolution. *Parasitology* **100**:0.
- MAYNARD SMITH, J. (1978); Optimization theory in evolution. *Annual Review of Ecology and Systematics* **9**:31–56.
- MAYNARD SMITH, J. (1989); *Evolutionary genetics*. Oxford U.P., New York.
- MAYNARD SMITH, J. & SZATHMÁRY, E. (1995a); *The major transitions in evolution*. W.H. Freeman Spektrum, Oxford.
- MAYNARD SMITH, J. & SZATHMÁRY, E. (1995b); The major transitions in evolution. *Nature* **374**:227–232.
- MITCHELL, M. (1996); *An Introduction to Genetic Algorithms*. MIT Press.
- MITCHELL, M.; CRUTCHFIELD, J. P. & DAS, R. (1996); Evolving cellular automata with genetic algorithms: A review of recent work. In: *Proceedings of the First International Conference on Evolutionary Computation and its Applications (EvCA'96)*.
- MITCHELL, M.; CRUTCHFIELD, J. P. & HRABER, P. T. (1994); Evolving cellular automata to perform computations: Mechanisms and impediments. *Physica D* **75**:361–391. SFI Working Paper 93-11-071.
- MONOD, J. A. (1992); Theoretical implications for the evolution of postsegregational killing by bacterial plasmids. *American Naturalist* **139**(4):677–689.
- NAITO, T.; KUSANO, K. & KOBAYASHI, I. (1995); Selfish behavior of restriction-modification systems [see comments]. *Science* **267**:897–899.
- O'NEILL, M.; CHEN, A. & MURRAY, N. E. (1997); The restriction-modification genes of *Escherichia coli* K-12 may not be selfish: they do not resist loss and are readily replaced by alleles conferring different specificities. *Proc. Natl. Acad. Sci. U.S.A.* **94**:14596–14601.
- PAGIE, L. & HOGEWEG, P. (1997); Evolutionary consequences of coevolving targets. *Evolutionary Computation* **5**(4):401–418.
- PAGIE, L. & HOGEWEG, P. (1999a); Colicin Diversity: a Result of Eco-evolutionary Dynamics. *J. Theor. Biol.* **196**:251–261.
- PAGIE, L. & HOGEWEG, P. (1999b); Individual- and population-based diversity in restriction-modification systems. Submitted.
- PAREDIS, J. (1994); Steps towards co-evolutionary classification neural networks. In: Brooks, R. A. & Maes, P. (eds.), *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems Artificial Life IV*. MIT Press, Cambridge, MA, USA, pp. 102–108, pp. 102–108.

-
- PAREDIS, J. (1995); Coevolutionary computation. *Artif. Life* **2**:355–375.
- PAREDIS, J. (1997); Coevolving cellular automata: be aware of the red queen! In: Baeck, T. (ed.), *Proceedings of the 7th Int. Conference on Genetic Algorithms (ICGA 97)*. pp. 393–400, pp. 393–400.
- PELLMYR, O. & LEEBENS-MACK, J. (1999); Forty million years of mutualism: evidence for eocene origin of the yucca-yucca moth association. *Proc. Natl. Acad. Sci. U.S.A.* **96**:9178–9183.
- PELLMYR, O.; LEEBENS-MACK, J. & HUTH, C. J. (1996); Non-mutualistic yucca moths and their evolutionary consequences. *Nature* **380**:155–156.
- POTTER, M. A. & DE JONG, K. A. (1994); A cooperative coevolutionary approach to function optimization. In: Davidor, Y.; Schwefel, H.-P. & Manner, R. (eds.), *Proceedings of the Third Conference on Parallel Problem Solving from Nature*. Springer-Verlag, Jerusalem, Israel, vol. 866 of *Lecture notes in computer science*, pp. 249–257, pp. 249–257.
- POTTER, M. A.; DE JONG, K. A. & GREFENSTETTE, J. J. (1995); A coevolutionary approach to learning sequential decision rules. In: Eshelman, L. J. (ed.), *Proceedings of the 6th International Conference on Genetic Algorithms*. Morgan Kaufmann Publishers, San Francisco, pp. 366–372, pp. 366–372.
- PROCTOR, L. M. & FURHMAN, J. A. (1990); Viral mortality of marine bacteria and cyanobacteria. *Nature* **343**(6253):60–62.
- PUGSLEY, A. P. (1984); The ins and outs of colicins. Part II. Lethal action, immunity and ecological implications. *Microbiol. Sci.* **1**:203–205.
- RECHENBERG, I. (1973); *Evolution strategy: Optimization of technical systems by means of biological evolution*. Fromman-Holzboog, Stuttgart.
- REDASCHI, N. & BICKLE, T. A. (1996); Dna restriction and modification systems. In: Neidhardt, F. C. (ed.), *Escherichia Coli and Salmonella: cellular and molecular biology*, American Society for Microbiology, Washington, D.C., vol. 1, chap. 52. 2nd edn., pp. 773–781.
- RILEY, M. A. (1993); Molecular mechanisms of colicin evolution. *Mol. Biol. Evol.* **10**:1380–1395.
- RILEY, M. A. & GORDON, D. M. (1992); A survey of Col plasmids in natural isolates of *Escherichia coli* and an investigation into the stability of Col-plasmid lineages. *J. Gen. Microbiol.* **138**:1345–1352.
- ROSIN, C. D. & BELEW, R. K. (1997); New methods for competitive coevolution. *Evolutionary Computation* **5**(1):1–29.
- ROTHSTEIN, S. I. (1990); A model system for coevolution: Avian brood parasitism. *Annual Review of Ecology and Systematics* **21**:481–508.

-
- ROUGHGARDEN, J. (1979); *Theory of population genetics and evolutionary ecology, an introduction*. Macmillan, New York.
- ROUGHGARDEN, J. (1983); Coevolution between competitors. In: Futuyma, D. J. & Slatkin, M. (eds.), *Coevolution*, Sinauer Associates, Inc., Sunderland, MA. pp. 383–403.
- SATO, A.; O’HUGIN, C.; FIGUEROA, F.; GRANT, P. R.; GRANT, B. R.; TICHY, H. & KLEIN, J. (1999); Phylogeny of Darwin’s finches as revealed by mtDNA sequences. *Proc. Natl. Acad. Sci. U.S.A.* **96**:5101–5106.
- SAVAGE, D. C. (1977); Microbial ecology of the gastrointestinal tract. *Annu. Rev. Microbiol.* **31**:107–133.
- SAVILL, N. J. (1997); *Eco-evolutionary consequences of spatial pattern formation*. Ph.D. thesis, Dep. Theor. Biol./Bioinformatics, Utrecht.
- SAVILL, N. J. & HOGEWEG, P. (1997); Evolutionary stagnation due to pattern-pattern interactions in a co-evolutionary predator-prey model. *Artificial Life* **3**:81–100.
- SAVILL, N. J.; ROHANI, P. & HOGEWEG, P. (1997); Self-reinforcing spatial patterns enslave evolution in a host-parasitoid system [published erratum appears in *J Theor Biol* 1997 Oct 21;188(4):525-6]. *J. Theor. Biol.* **188**:11–20.
- SCHUSTER, P. (1989); Optimization of RNA structure and properties. In: Perelson, A. S. & Kauffman, S. A. (eds.), *Molecular evolution on rugged landscapes: proteins, RNA and the immune system*, Addison Wesley. pp. 47–71.
- SCHWEFEL, H. P. (1977); *Numerische Optimierung von Computer Modellen mittels der Evolutionsstrategie*, vol. 26 of *Interdisciplinary systems research*. Birkhäuser Verlag, Basel.
- SHARP, P. M. (1986); Molecular evolution of bacteriophages: evidence of selection against the recognition sites of host restriction enzymes. *Mol. Biol. Evol.* **3**:75–83.
- SIMS, K. (1994); Evolving 3D morphology and behavior by competition. In: Brooks, R. A. & Maes, P. (eds.), *Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems Artificial Life IV*. MIT Press, Cambridge, MA, USA, pp. 28–39, pp. 28–39.
- SINGH, M.; AHMAD, C. M. & YADAVA, J. N. (1989); Drug resistance and colicin production among various serotypes of *Escherichia coli* of animal origin. *Indian. J. Pathol. Microbiol.* **32**:161–166.
- TAKUMI, K. & HOGEWEG, P. (1998); Evolution of the immune repertoire with and without somatic DNA recombination. *J. Theor. Biol.* **192**:308–317.
- THOMPSON, A. (1996); Evolutionary techniques for fault tolerance. In: *Proceedings of UKACC International Conference on control 1996 (CONTROL’96)*. The Institution of Electrical Engineers, IEE conference publications, London, UK.

-
- TOFFOLI, T. & MARGOLUS, N. (1987); *Cellular Automata Machines: A new environment for modeling*. Scientific Computation. MIT Press, Cambridge, Massachusetts.
- VAFSAIE, H. & DE JONG, K. A. (1996); Genetic algorithms as a tool for restructuring feature space representations. In: *Proceedings of the International Conference on Tools with AI*. Institute of Electrical and Electronics Engineers, IEEE Computer Society Press, Los Alamitos, CA.
- VAN NIMWEGEN, E.; CRUTCHFIELD, J. P. & HUYNEN, M. (1999); Neutral evolution of mutational robustness. *Proc. Natl. Acad. Sci. U.S.A.* **96**:9716–9720.
- VAN VALEN, L. (1965); Morphological variation and the width of the ecological niche. *Amer. Natur.* **99**:377–390.
- VAN VALEN, L. (1973); A new evolutionary law. *Evolutionary theory* **1**:1–30.
- VOGEL, T. U.; EVANS, D. T.; URVATER, J. A.; O’CONNOR, D. H.; HUGHES, A. L. & WATKINS, D. I. (1999); Major histocompatibility complex class I genes in primates: co-evolution with pathogens. *Immunol. Rev.* **167**:327–337.
- WERFEL, J.; MITCHELL, M. & CRUTCHFIELD, J. P. (1999); Resource sharing and coevolution in evolving cellular automata. Submitted to IEEE Trans. Evol. Comp.
- WILSON, G. G. & MURRAY, N. E. (1991); Restriction and modification systems. *Annu. Rev. Genet.* **25**:585–627.
- WOLFRAM, S. (1984); Cellular automata as models of complexity. *Nature* **311**:419–424.
- WRIGHT, S. (1967); “Surfaces” of selective value. *Proc. Natl. Acad. Sci. U.S.A.* **58**:165–172.

Samenvatting

Biologische evolutie heeft geleid tot een grote verscheidenheid aan organismen welke een hoge mate van complexiteit vertonen. Ofschoon evolutie al lange tijd bestudeerd wordt hebben we nog geen verklaring voor het ontstaan van deze verscheidenheid en complexiteit of hoe ze in stand worden gehouden. Een mogelijke oorzaak hiervoor is dat in veel studies evolutie nog steeds wordt beschouwd als een selectie proces dat zich afspeelt op een enkel niveau. Daarbij wordt de relatie tussen het genotype en het resulterende gedrag van het individu meestal op zeer eenvoudige wijze gemodelleerd. En als laatste, 'fitness evaluatie', d.w.z. de evaluatie van het succes van individueel gedrag, wordt meestal als volledig veronderstelt: individuele fitness wordt uniek bepaald door het genotype van het individu.

Meer recente studies van evolutie nemen sommige van dergelijke aspecten van biologische evolutie wel in acht, zoals de meerlagigheid van biotische systemen, het optreden van zelf-structurering in dergelijke systemen, de niet-lineaire relaties tussen genotypen en fenotypen, en het voorkomen van neutraliteit in deze relatie. In dit proefschrift bestuderen we een aspect van evolutie, *informatie integratie*: op welke manieren wordt informatie, die aanwezig is in de (a-)biotische omgeving, geïntegreerd in evoluerende systemen en welke gevolgen heeft dit voor het gehele systeem. We onderzoeken dit aan de hand van drie verschillende thema's.

Ten eerste kijken we naar het effect van (zeer) beperkte, i.p.v. complete, fitness evaluatie per individu. Ofschoon meestal verondersteld wordt dat het (reproductieve) succes van een individu, ofwel de *fitness*, uniek bepaald wordt door het genotype van het desbetreffende individu zal het in werkelijkheid vaak gebeuren dat een individu slechts een (klein) deel van alle mogelijke omstandigheden, waarin het succes moet blijken, 'ziet'. Als een individu maar ten opzichte van een klein deel van alle mogelijke omstandigheden wordt geëvalueerd is het mogelijk dat het individu slechte eigenschappen behoudt en goede eigenschappen verliest, in beide gevallen door een 'incomplete' fitness evaluatie.

Het onderzoek dat is beschreven in dit proefschrift toont aan dat dit laatste niet altijd optreedt. Het blijkt dat informatie over verschillende *fitness evaluaties*, en over meerdere generaties geïntegreerd kan worden op een zodanige wijze dat individuen evolueren die succesvol zijn in veel meer omstandigheden dan alleen maar diegene waarop ze kort geleden (in evolutionaire tijd) geëvalueerd zijn.

Het tweede thema waar we in dit proefschrift naar kijken is de vrijheid van een evoluerend systeem om informatie op verschillende manieren op te slaan. Traditioneel wordt verondersteld dat informatie wordt opgeslagen in de genotypen van individuen, maar informatie kan ook worden 'opgeslagen' op, bijvoorbeeld, het niveau van de populatie.

Het laatste thema dat we behandelen zijn zogenaamde neveneffecten van informatie integratie in evoluerende systemen. Hierbij moeten we denken aan verschillen in gevoeligheid van individuen voor mutaties van het genotype of aan verschillen in gevolgen van invasies van een populatie door individuen met nieuwe eigenschappen.

In hoofdstuk twee bekijken we specifiek het effect van zeer beperkte fitness evaluatie op de evolutionaire dynamica. We modelleren het evolutionaire proces zodanig dat de

fitness van individuen bepaald wordt aan de hand van een expliciete fitness functie die vooraf is gedefinieerd. Deze fitness functie is bepaald met betrekking tot een vast aantal 'fitness-cases'. We vergelijken de evolutionaire dynamica die optreedt als individuen altijd geëvalueerd worden met betrekking tot de volledige verzameling van fitness-cases, met de evolutionaire dynamica die optreedt als individuen per tijdstap op slechts een klein aantal van de fitness-cases wordt geëvalueerd. In het laatste geval onderscheiden we twee mogelijkheden; de fitness-cases waarop een individu wordt geëvalueerd worden willekeurig getrokken, of deze fitness-cases zijn co-evoluerende 'parasieten'.

In alle drie de gevallen, complete fitness evaluatie, beperkte fitness evaluatie op willekeurig bepaalde fitness-cases, en beperkte fitness evaluatie op co-evoluerende fitness-cases, vergelijken we het gedrag van de individuen ten opzichte van de complete fitness functie, d.w.z. de complete verzameling van fitness-cases. Ofschoon de individuen die evolueren onder beperkte evaluatie nooit alle fitness-cases 'zien' tijdens hun leven blijkt uit dit onderzoek dat desondanks individuen kunnen evolueren die het complete 'probleem' oplossen, hier zelfs beter dan individuen die evolueren onder complete evaluatie. Blijkbaar wordt er tijdens het evolutionaire proces informatie geïntegreerd over verschillende fitness evaluaties.

Als de fitness-functie complexer wordt zien we dat beperkte evaluatie alleen leidt tot individuen met een complete oplossing in het co-evoluerende geval, willekeurige evaluatie is dan onvoldoende. We vinden ook dat als gevolg van de verschillende evaluatie procedures, d.w.z. complete versus beperkte, genotypen evolueren die consistent verschillen in eigenschappen waarop ze niet direct geselecteerd worden. We vinden bijvoorbeeld dat individuen die evolueren onder complete fitness evaluatie relatief minder gevoelig zijn voor enkele mutaties van hun genotype en meer gevoelig zijn voor kleine veranderingen in de verzameling van fitness-cases waarop ze worden geëvalueerd.

In hoofdstuk drie en vier bestuderen we twee modellen van specifieke biologische systemen. Bacteriën kunnen makkelijk genetisch materiaal uitwisselen, bijvoorbeeld in de vorm van plasmiden. Uit experimenteel onderzoek blijkt dat een veel groter aantal verschillende colicin-complexen worden gecodeerd op plasmiden (hoofdstuk 3), en ook een veel groter aantal restrictie-modificatie systemen (hoofdstuk 4) dan op a priori gronden zou worden verwacht.

Een colicin-complex codeert zowel voor een eiwit dat giftig is voor bacteriën als voor een eiwit dat het giftige eiwit onschadelijk maakt zodat de bacterie zichzelf niet vergiftigd. Verschillende colicinen coderen voor verschillende gif-antigif paren. Bacteriën kunnen colicin-complexen 'gebruiken' om andere bacteriën te doden en zodanig meer ruimte of voedsel voor eigen groei te verkrijgen.

Restrictie-modificatie (RM) systemen zijn enigszins vergelijkbaar; zij coderen voor eiwitten die genetisch materiaal dat een bacterie-cel binnen dringt kapot knipt, en eiwitten die het genetisch materiaal dat de bacterie zelf draagt beschermt tegen het knip eiwit. Verschillende RM systemen knippen en beschermen genetisch materiaal op verschillende plaatsen. RM systemen worden geacht met name van belang te zijn voor bacteriën als bescherming tegen virus infecties. In dit model (hoofdstuk vier) hebben we dan ook twee populaties; de bacteriën en de virussen. Virussen kunnen overigens ook ongevoelig raken voor de RM systemen met als gevolg dat zij, naast de bacteriën, een tweede niveau van potentiële informatie integratie vormen.

In beide modellen vinden we dat een hoge diversiteit aan colicinen, respectievelijk RM systemen, makkelijk ontstaat door lokale competitie tussen individuele bacteriën. Veel verrassender is het dat in beide modellen deze diversiteit ‘opgeslagen’ wordt op twee verschillende manieren: op het individuele niveau, of op het niveau van de populatie. In beide gevallen zijn er evenveel verschillende typen colicinen of RM systemen aanwezig in de populatie, maar in het eerste geval draagt ieder individu alle verschillende typen terwijl in het tweede geval individuele bacteriën geen, of slechts één plasmide dragen.

In het eerste model blijkt dat de bacterie populatie altijd een van beide toestanden aanneemt, afhankelijk van de kosten per plasmide. In het tweede model zien we dat, onafhankelijk van deze kosten, beide toestanden kunnen voorkomen; het systeem vertoont bi-stabiliteit. Blijkbaar komt dit doordat in het tweede model een extra niveau van organisatie aanwezig is: de virussen.

De resultaten die we in deze modellen vinden tonen ook dat optimalisatie van individuele bescherming tegen infecties niet altijd gunstig is op het niveau van de populatie: een optimaal beschermd individu dwingt andere individuen tot eenzelfde optimalisatie of zij worden weggeconcentreerd. In beide gevallen leidt dit tot een homogene populatie wat een makkelijk doelwit kan zijn voor een evoluerend virus.

In hoofdstuk vijf, tenslotte, bestuderen we een model van co-evolutie van twee populaties. De populaties hebben een antagonistische relatie. Het gedrag van de individuen uit beide populaties wordt bepaald aan de hand van een extern gedefinieerde, artificiële functie: het lange termijn spatio-temporele gedrag van cellulaire automaten (de eerste populatie) beginnende vanuit een bepaalde initiële conditie (de tweede populatie).

In dit hoofdstuk ligt de focus van de studie op de mate waarin verschillende vormen van evolutionaire dynamica optreden in één en het zelfde proces. Specifiek kijken we naar: evolutie van algemene gedrag-strategieën, soortsvorming leidend tot meerdere soorten, en voortdurende evolutionaire verandering.

Alle drie de vormen zien we in het model, afhankelijk van kleine veranderingen in de model-structuur. Als we individuen in de ruimte laten bestaan, zodat er ruimtelijke patroon-vorming optreedt, zien we dat er algemene strategieën evolueren. Als we individuen van beide populaties elke tijdstap mixen in de ruimte zien we voortdurende evolutionaire verandering optreden: beide populaties tonen blijvende oscillaties tussen twee simpele strategieën. Als we het vormen van algemene oplossingen moeilijker maken dan zien we in het eerste geval soortsvorming optreden in beide populaties terwijl in het tweede geval eenzelfde soort voortdurende evolutionaire verandering optreedt als eerder. In dat laatste geval is de tijd-schaal van verandering echter veel sneller; individuen lijken hun vermogen tot verandering geoptimaliseerd te hebben.

Concluderend kunnen we zeggen dat informatie integratie in evolutie veel meer behelst dan het probleem van het vinden van optimale individuen. In dit proefschrift hebben we laten zien dat informatie niet alleen op het niveau van het individu wordt geïntegreerd maar ook, bijvoorbeeld, op het niveau van de populatie. Daarnaast hebben we laten zien dat informatie integratie ook kan optreden, en soms zelfs tot meer optimale individuen leidt, onder omstandigheden waarbij informatie betreffende ‘optimaal gedrag’ slechts beperkt wordt aangeboden. Tenslotte hebben we laten zien dat neveneffecten van verschillende vormen van informatie integratie kunnen leiden tot individuele eigenschappen die geen direct selectie voordeel hebben. Gegeven dat natuurlijke evolu-

tie voort bouwt op datgene wat er al eerder geëvolueerd is kunnen verschillen in neven-effecten die op dit moment optreden de latere eco-evolutionaire dynamica beïnvloeden.

Curriculum vitae

De auteur werd op 23 april 1967 geboren te Delft. In 1986 behaalde hij het ongedeeld VWO diploma aan het Erasmus College te Zoetermeer. Van 1986 tot 1988 studeerde hij medische biologie aan de Universiteit Utrecht. Van 1990 tot 1995 studeerde hij Cognitieve Kunstmatige Intelligentie aan dezelfde universiteit. Het doctoraalexamen Cognitieve Kunstmatige Intelligentie werd afgelegd in februari 1995, met bioinformatica als onderwerp. Van maart 1995 tot december 1999 was hij in dienst van de Nederlandse Organisatie voor Wetenschappelijk Onderzoek (NWO), werkzaam als OIO aan de vakgroep Theoretische Biologie/Bioinformatica van de Universiteit Utrecht, onder begeleiding van Prof. Dr. P. Hogeweg. Het onderzoek dat is beschreven in dit proefschrift is verricht in die hoedanigheid.

Dankwoord

Op deze plaats wil ik graag enkele mensen bedanken die mij geholpen hebben bij het tot stand brengen van dit proefschrift. Om te beginnen mijn promotor, Paulien Hogeweg. Paulien, zonder jouw wetenschappelijke inzicht, creativiteit en enthousiasme zou ik het nooit gered hebben en zou ik nooit gekomen zijn waar ik nu ben. Ben, jou wil ik graag bedanken voor de inspiratie, steun en visie die je me in de loop der jaren gegeven en getoond hebt.

Verder wil ik de gehele vakgroep bedanken voor de inspiratie en het plezier dat in overvloed aanwezig was. Met name wil ik Rob, Jose en Stan bedanken voor alle inzichten, geduld, warmte en antwoorden die ik zo vaak nodig had, en Andre en Tomoyuki voor alle plezier die zij daar nog bij voegden. Alle voormalige en huidige promovendi en studenten wil ik graag bedanken voor alle discussies, klets en pret, zowel wetenschappelijk als onwetenschappelijk, met name Nick, Katsuhisa, John en Roeland. Kirsten en Jelle wil ik graag bedanken voor het vijf-voor-twaalf werk.

I would like to thank Simon Levin, Stephanie Forrest, Rolf Hoekstra, and Phil Husbands for being members of my reading committee. Sheila McNab, I would like to thank you for working so hard on improving my texts, even at times when you were so busy with your own thesis.

Buiten de vakgroep zijn er veel mensen die mij op meer en minder cruciale momenten terzijde hebben gestaan om er plezier in te krijgen en te houden. Mijn familie heeft altijd achter mij gestaan, met name op momenten dat ik het zelf niet meer zag zitten. Frans, mijn eerste paranimf, jou ken ik bijna zo lang als ik leef en dat hoop ik nog lang te kunnen blijven zeggen. En natuurlijk bedankt dat je mijn paranimf bent, met alle daar (nog) bij komende zaken. Sascha, jou moet ik bedanken mijn vriendin te worden en blijven op één van de moeilijkste momenten in de afgelopen vier jaar en natuurlijk dat je ook mijn paranimf wilt zijn. Wilma, jou wil ik graag bedanken voor alle waardevolle momenten in de afgelopen jaren. Maarten en Barbara (en Gijs) jullie hebben mij zo vaak laten gaan en weer opgevangen; dank. Can, my old colleague, I would like to thank you for two perfect weeks and for all other times.