

Darwin's Continent Cycle Theory and Its Simulation by the Prisoner's Dilemma

Heinz Mühlenbein

GMD Schloss Birlinghoven,
D-5205 Sankt Augustin 1, Germany

Abstract. Theoretical biology has largely ignored Darwin's true evolution model. In his famous book about the origin of species, Darwin detailed the importance of a population structure on evolution. He conjectured that a large continent that exists for long periods in a broken condition will be the most favorable for the production of many new forms of life. We call this conjecture Darwin's *continent cycle theory*. In this paper we investigate some of Darwin's arguments in support of his theory by simulating an artificial ecology with the parallel genetic algorithm. The artificial ecology consists of a population playing the Iterated Prisoner's Dilemma. The major emphasis of this paper is on the methodological questions of the simulation. These are the genetic representation, the mapping of the genotypes to phenotypes, and the spatial population structure.

1. Introduction

The most complex systems we observe in nature are the results of evolutionary processes. This was conjectured by Charles Darwin in his famous book *On the Origin of Species by Means of Natural Selection* [5].

The idea of applying evolutionary strategies to computer algorithms dates back to the invention of the first stored program computers. John von Neumann already investigated the question of whether modelling evolution on a computer could solve the complexity problem of programming. He invented automata theory to research the question of whether the construction of automata by automata can progress from simpler types to increasingly complicated types [19].

Von Neumann stated the following important difference between natural and artificial systems: "Today's organisms are phylogenetically descended from others which were vastly simpler than they are, so much simpler, in fact, that it is inconceivable how any description of the later, complex organism could have existed in the earlier one." But in most artificial systems the situation is different: "Everyone knows that a machine tool is more complicated than the elements which can be made with it, and that, generally

speaking, an automaton A, which can make an automaton B, must have a complete description of B and also rules on how to behave while effecting the synthesis."

In the theory of cellular automata, von Neumann was able to construct an automaton that could reproduce itself. This result showed that reproduction was possible. Von Neumann was well aware of the other two important evolutionary processes—namely variation and selection. He decided that knowledge about these two processes was not yet sufficient to incorporate them into his theory of automata. "Conflicts between independent organisms lead to consequences which, according to the theory of *natural selection*, are believed to furnish an important mechanism of evolution. Our models lead to such conflict situations. The conditions under which this motive for evolution can be effective here may be quite complicated ones, but they deserve study."

Selection and variation are the basic components of the genetic algorithm (GA) developed by John Holland [9]. But the genetic representation consists just of a simple binary bit string, not a complex automaton. Nevertheless the GA has been successfully applied in many areas. For a number of technical reasons we have extended the GA to the parallel genetic algorithm (PGA). The PGA is a totally distributed algorithm running with maximal efficiency on parallel computers. The major changes are:

- the population is spatially structured;
- the individuals themselves select the partner for mating;
- the offspring replaces the parent;
- the individuals may actively improve their fitness during their life.

The PGA models natural evolution more realistically than the GA. The selection of individuals for mating is done within the population itself. The GA, in contrast, models evolution using a centralized selection scheme.

We hoped to find arguments in favor of population structures in biological textbooks. But we were very disappointed because most of the mathematical models abstract from the spatial population structure of the population. The importance of spatial population structures on evolution is still controversially discussed, as well in theoretical biology as in genetic algorithms. We will give a short survey of mathematical models of spatial population structures in section 3. But before doing so we will summarize the impressive discussion about the influence of population structures on evolution, which can be found in Darwin's *Origin*.

2. Darwin revisited

In the section "Circumstances favourable and unfavourable to Natural Selection," Darwin tried to describe the influence of intercrossing, isolation, and the number of individuals on the speed of evolution. Darwin mentioned that this is an extremely intricate subject. He argued very carefully, for example,

that "to ascertain whether a small isolated area or a large open area like a continent has been the most favourable for the production of new organic forms, we ought to make comparisons within equal times; and this we are incapable of doing."

Simulation makes such a comparison possible. It is of course impossible to simulate the real evolution of nature, so we have to find an artificial environment that is nevertheless complex enough to model the important aspects of evolution. As a first step we have decided to simulate an artificial population where each individual plays a two-person game against the other individuals. We have selected the Iterated Prisoner's Dilemma (IPD), because it is surprisingly complex. Furthermore it is possible to compare our simulation results with theoretical results and simulations done by other researchers.

In a series of papers we will try to follow step-by-step Darwin's reasoning by simulation. We can cite here only Darwin's major conclusion: "I conclude that a large continental area, which will probably undergo many oscillations of level, and which consequently will exist for long periods in a broken condition, will be the most favourable for the production of many new forms of life, likely to endure long and spread widely." Darwin argued as follows. In a large continent, there is severe competition. This leads to the extinction of over-specialized species. But it is highly improbable that something new arises on a large continent. This happens much easier on small islands. But if the islands are isolated for a long time, then over-specialized forms will develop. So Darwin postulated that the islands should reconvert to a large continent. There will again be severe competition that eliminates the specialized forms. This briefly sketches Darwin's true evolution model. We call Darwin's conclusion the *continent cycle theory*. The interested reader is referred to the above mentioned chapter in Darwin's *Origin*.

Mainstream science seems not to have noticed or to have deliberately neglected this section in Darwin's book. Darwin's true evolution model is a non-equilibrium model, whereas all the popular Darwinian, Neo-Darwinian, and Synthesis theories are equilibrium models. Darwin viewed evolution as a property of the organization of the organism, rather than as a property of the matter of an individual organism.

Darwin's arguments in favor of the continent cycle, with just some minor changes, can also be applied to other areas like the *invention of successful scientific ideas* or the *efficient organization of companies*. Take the organization of companies as an example. If the market (the environment) is stable, a large centralized company with severe internal competition is most effective. If the company has to adapt to a changing market, the large company should be subdivided into small companies that can adapt much faster.

It is this general aspect that gives Darwin's true evolution theory such a broad range of applications, ranging from artificial intelligence to sociology, economy, psychology, and philosophy.

The major emphasis of this paper will be on methodological questions for at least two reasons. First, we believe that methodological questions are of utmost importance in a scientific field where it is almost impossible

to compare simulation results with actual experiments. Second, it is not possible to explain in one short paper all the simulation experiments that are necessary to support or disprove Darwin's continent cycle theory.

3. Spatial population structures

Several researchers in biology have tried to investigate the importance of spatial population structures for evolution—without ever referring to Darwin. It is now accepted that a spatial population structure has more variety than a panmictic population. The importance of this fact on evolution, however, has been highly controversial.

Wright [22] has argued that the best way to avoid being hung up on a low fitness peak is to have the population broken up into many nearly isolated subpopulations. Wright's theory has three phases [23]. *Phase 1* consists of the *differentiation* of innumerable small local populations by more or less random processes that occasionally lead to higher peaks. *Phase 2* is the occupation of higher peaks by *local mass selection*. *Phase 3* is the *diffusion* of these successful subpopulations throughout the species, followed by the appearance of still more successful centers of diffusion at points of contact. Then the whole process starts again.

Fisher [8], in contrast, argued that no such theory is needed. In a highly multidimensional fitness surface, the peaks are not very high and are connected by fairly high ridges, always shifting because of environmental changes. According to Fisher, the analogy is closer to waves and troughs in an ocean than to a static landscape. Alleles are selected because of their average effects, and a population is unlikely ever to be in such a situation that it can never be improved by direct selection based on additive variance.

The difference between these two views is not as much mathematical as physiological. Does going from one favored combination of alleles to another often necessitate passing through genotypes that are of lower fitness? Fisher argued that evolution typically proceeded in a succession of small steps, leading eventually to large differences by the accumulation of small ones. According to this view, the most effective population is a large panmictic one in which statistical fluctuations are slight and each allele can be fairly tested in combination with many others alleles. According to Wright's view, a more favorable structure is a large population broken up into subgroups, with migration sufficiently restricted (less than one migrant per generation) and size sufficiently small to permit appreciable local differentiation.

Four different models for spatially structured populations have been investigated mathematically:

- the one-island model;
- the island model;
- the stepping-stone model;
- the isolating-by-distance model.

In the one-island model, an island and a large continent are considered. The large continent continuously sends migrants to the island. In the island model, the population is pictured as subdivided into a series of randomly distributed islands among which migration is random.

In the stepping-stone model migration takes place between neighboring islands only. One- and two-dimensional models have been investigated.

The isolation-by-distance model treats the case of continuous distribution where effective demes are isolated by virtue of finite home ranges (neighborhoods) of their members. For mathematical convenience it is assumed that the position of a parent at the time it gives birth relative to that of its offspring when the latter reproduces is normally distributed.

Felsenstein [6] has shown that the isolating-by-distance model leads to unrealistic clumping of individuals. He concluded that this model is biologically irrelevant. There have been many attempts to investigate spatial population structures using computer simulations, but they have not had a major influence. A good survey of the results of the different population models can be found in [7]. Population models with oscillations like Darwin's continent cycle have not been dealt with.

The issue raised by Wright and Fisher is still not settled. Phase 3 of Wright's theory has recently been investigated by Crow [4]. He concludes: "The importance of Wright's shifting-balance theory remains uncertain, but we believe whatever weaknesses it may have, they are not in the third phase."

The problem of spatial population structures is now reappearing in the theory of genetic algorithms. The plain GA is based on Fisher's model. It is a well-known fact that the GA suffers from the problem of premature convergence. In order to solve this problem, many genetic algorithms enforce diversification explicitly, violating the biological metaphor. A popular method is to accept an offspring only if it is genetically different from all the members of the population by more than a certain factor.

Our PGA tries to introduce diversification more naturally using a spatial population structure. Fitness and mating is restricted to neighborhoods. In the PGA we have implemented the isolation-by-distance model and the stepping-stone model. The three phases of Wright's theory can actually be observed in the PGA. But the relative importance of the three phases are different than Wright believed. The small populations do not find better peaks using random processes. The biggest changes of the population occur at the time after migration between the subpopulations. Recombinations between immigrants and native individuals occasionally lead to higher peaks, which were not found by any of the subpopulations during isolation. This behavior can easily be demonstrated in the application function optimization (see [16] for details). We can therefore state the following observation.

The creative forces of evolution take place at migration and for a few generations afterward. Wright's argument that better peaks are found just by chance in small subpopulations is wrong.

In our opinion the most important part of Wright's theory is what Wright postulated as "the appearance of still more succesful centres of diffusion at points of contact."

We have not yet used Darwin's continent cycle population structure for optimization problems because the results of the stepping-stone model with migration are already very good. We believe that the reason for this success lies in the static fitness function. The fitness of an individual is given in a closed form; it depends only on the genotype of the individual. The same is true for the mathematical models where almost trivial fitness functions are used.

We believe that static fitness functions cannot model real evolution. In a real environment the fitness of an individual depends on the outcome of its interactions with other organisms in the environment. The fitness cannot be specified in advance. In order to investigate Darwin's continent cycle we decided to simulate an artificial ecology. In our ecology the interactions of the individuals are modeled by a game. The fitness of the individual is the sum of the payoffs the individual gets during its lifetime. We have chosen the Iterated Prisoner's Dilemma because it has been investigated from a number of different viewpoints.

4. The artificial ecology

Over its 30-year lifespan, the Iterated Prisoner's Dilemma (IPD) has been one of the most frequently studied phenomena in economics, political science, sociology, and psychology (see Axelrod [1] for a survey). The basic Prisoner's Dilemma is a two-person game, with each player having a choice of either cooperating (C) or defecting (D). A typical set of payoffs is presented below.

| Move | C | D |
|------|-----|-----|
| C | 3/3 | 0/5 |
| D | 5/0 | 1/1 |

Given these payoffs, it is easily shown that mutual defection is the only Nash equilibrium. Of course, the intrigue of the Prisoner's Dilemma is that this unique equilibrium is Pareto inferior to the mutual cooperation outcome. If the basic Prisoner's Dilemma is iterated, the resulting supgame is an Iterated Prisoner's Dilemma. If the number of iterations is a known finite number, then a simple backward induction argument implies that the only equilibrium is mutual defection in every round. However, if the game is repeated a finite but unknown number of times, then cooperative behavior can theoretically emerge.

The *ecological* approach to experimental games has added another dimension to the study of conflict and cooperation in societies. John Maynard Smith [17] introduced *evolutionary game theory*, where the games are played by a population of individuals. The higher the payoff of an individual, the

more offspring he will get. In this manner the most effective strategies survive. A strategy is called *evolutionary stable* [17] if it cannot be invaded by a single mutant strategy. The theory assumes that the strategies are not changed during the course of evolution. In our simulations the strategies are coded by genes. The strategies are constantly changed by the parallel genetic algorithm, which uses mutation and crossing-over to generate offspring.

5. The method of investigation

There have been at least three attempts to investigate the IPD with genetic algorithms. The first simulation was performed by Axelrod [2]. Axelrod considered strategies where the moves are based on the game's past three-move history. The major focus of Axelrod's study was on strategies evolving against a fixed environment. Each individual played against eight representative strategies. Marks [10] extended the investigation to *bootstrap* evolution, where the individuals play against each other. Miller [12] used finite automata to represent strategies. Furthermore, he investigated the effect of informational accuracy on the outcome of the simulation. All three researchers used the plain genetic algorithm for evolving the population. They have been interested in equilibrium states and "optimal" strategies. We concentrate on the evolution of the behavior of the total population.

In our simulations we use the parallel genetic algorithm to simulate the ecology. This gives us the possibility of investigating the importance of spatial population structures on evolution. The PGA is defined as follows.

Parallel Genetic Algorithm

STEP0: Define a genetic representation of the problem.

STEP1: Create an initial population and its population structure.

STEP2: Each individual does local hill-climbing.

STEP3: Each individual selects a partner for mating in its neighborhood.

STEP4: An offspring is created with genetic crossover of the parents.

STEP5: The offspring does local hill-climbing. It replaces the parent, if it is better than some criterion (acceptance).

STEP6: If not finished, return to STEP3.

The PGA has been successfully applied to important optimization problems. The major enhancements to the plain genetic algorithm are the spatial population structure, the distributed selection, and the local hill-climbing. The individuals are active. They look for a partner for mating in their neighborhood. The partner is chosen according to the preference of the individuals. The best individual in a neighborhood has the chance to get as many offspring as the global best individual of the population. The PGA therefore

has a very “soft” selection scheme. Each individual has the chance that, on average, 50% of its genes are contained in the chromosome of an offspring. The offspring replaces the parent.

We have applied the PGA to many optimization problems. The PGA outperforms the GA on many large complex optimization problems by far; see [13, 15, 16] for combinatorial optimization and [14] for function optimization. In fact, the PGA has solved large optimization problems never solved before.

A similar observation has been reported by Tanese [18]. Tanese implemented a population structure that resembles the island model. The population is divided into subpopulations, and the subpopulations are connected by a hypercube network. At each generation a specified number of individuals migrate to one of the neighboring subpopulations.

These results already show the advantage of spatial population structures for solving optimization problems with a genetic algorithm. In order to investigate Darwin’s continent cycle theory we decided to simulate an artificial ecology where the fitness of the individuals is not statically specified by an optimization problem. In the simulations the individuals are not allowed to improve their fitness by learning.

We now turn to the problem of genetic representation of strategies.

6. The genetic representation

There are at least two obvious ways to represent strategies as genetic chromosomes: one is based on a simple table lookup, the other on a finite automaton. In this paper we will discuss deterministic table lookup strategies. A k -lookback strategy can be defined as a mapping of the outcome of the last k moves into a new move. In the simplest case of just looking one play back, a strategy can be defined by four entries in a table symbolizing the four possible moves of the last game—DD, DC, CD, CC. In addition two bits are necessary to specify the first move. The genetic representation of one-lookback thus consists of six bits. This gives 2^6 different genotypes. Three popular strategies are given below.

| | | | | | | Strategy |
|---|---|---|---|---|---|-------------|
| C | * | * | * | C | C | ALL-C |
| D | * | D | D | * | * | ALL-D |
| C | * | D | C | D | C | TIT-FOR-TAT |

The sign * denotes that the allele on this locus does not have any influence on the performance of the strategy. There are twelve different bitstrings that each define an ALL-C strategy. The problem of this straightforward genetic representation is that we have a distinction between the *representation* and the *interpretation*. The program that interprets the representation is not part of the genetic specification and therefore not subjected to the evolution process. But we have a clear distinction between genotype, phenotype, and behavior. The genotype is mapped into some phenotype, and the phenotype

together with the environment (in our case the other phenotypes) defines the strategy. Let us take the famous TIT-FOR-TAT as an example. In TIT-FOR-TAT the player makes the move the opponent made the game before. In an environment where only C is played, TIT-FOR-TAT cannot be distinguished from an ALL-C player. A different behavior can only be recognized if there exists an individual who occasionally plays D.

The mapping from genotype to phenotype is many-to-one. This makes a behavior-oriented interpretation of a given genetic representation very difficult. There exists no simple structure of the genotype space. The Hamming distance between two ALL-C genetic representations can be as large as four, whereas the Hamming distance between two very different strategies like ALL-C and ALL-D can be as small as one. An example is shown below.

| | | | | | | | |
|---|---|---|---|---|---|--|----------|
| | | | | | | | Strategy |
| C | C | D | D | C | C | | ALL-C |
| C | C | D | D | C | D | | ALL-D |

If we assume that the genetic operators *mutation* and *crossing-over* uniformly explore the genotype space, then strategies like ALL-C and ALL-D will have a much higher chance of being generated than other strategies that are less often represented. The genetic search is therefore biased by the genetic representation. We believe that this effect is not a shortcoming of the chosen representation, but that this feature models real-life evolution. The evolution must always work within the constraints it creates for itself.

The complex mapping between genotype and phenotype makes it difficult to estimate the outcome of a genetic operator. For example, a winning strategy may be crossed with a losing strategy, giving in most cases a new strategy. An ALL-D strategy that is crossed over with an ALL-C strategy gives a probability of 0.2 ALL-D and 0.2 ALL-C. With the remaining probability of 0.6 we get a strategy that is different from the strategies of the parents.

We believe that in our artificial ecology the crossover operator is too disruptive compared to real evolution. The same problem occurs if the genetic representation is based on a finite automaton. In order to solve this problem we have to find a genetic representation that is based on a more complex genetic machinery than simple bitstrings. It is outside the scope of this paper to discuss this genetic machinery. We only want to mention that we have to incorporate some ideas from models of self-reproduction proposed as early as the 1960s.

The influence of spatial population structures is independent of the genetic representation, therefore we will concentrate on this subject.

7. Structured populations

Before we discuss some of the simulation results in detail we want to show by a simple analysis how a spatial population structure influences the development of strategies.

For simplicity we assume that we have a population consisting of *inhabitants* playing strategy *I* and *invaders* playing strategy *J*. Let *s* be the proportion of invaders. We assume that *s* is very small; furthermore the invaders are clustered. We model this assumption by a clustering factor *k*. Let $P(I, J)$ denote the payoff of an individual playing strategy *I* against an individual playing strategy *J*. After invasion the fitness of the inhabitants can be approximately computed as

$$F(I) = \left(1 - s * \frac{1 - ks}{1 - s}\right) * P(I, I) + s * \frac{1 - ks}{1 - s} P(I, J) \quad (7.1)$$

The invaders have the fitness

$$F(J) = (1 - ks) * P(J, I) + ks * P(J, J) \quad (7.2)$$

We see that for $k = 0$ the invaders play only against the inhabitants; the case $k = 1$ gives the panmictic population normally considered in the theory of evolutionary games. In the case of $k > 1$ we have a clustering effect. For $k = s^{-1}$ the invaders and the inhabitants play only within their group.

A strategy is called *collective stable* if no strategy can invade it. A new strategy is said to *invade* if the newcomer gets a higher score than the native strategy. For small *s*, this means

$$P(I, I) < (1 - ks)P(J, I) + ks * P(J, J) \quad (7.3)$$

If we set $p = ks$ then the above inequality is identical to Axelrod's *p-cluster invasion* ([1], page 212). It is now easily seen that even ALL-C can withstand the invasion of ALL-D if there is a strong preference for each strategy to play only against each other. With our payoff values we find that ALL-C will not be invaded for $k > 0.5s^{-1}$.

The above model can be applied to a population structure consisting of groups of individuals playing the same strategy. Each individual in a group has the same number of contests with individuals of other groups. The number of contests between the inhabitants compared to the number of contests between inhabitants and invaders defines the clustering factor *k*. Then the fitness is given by formulas (7.1) and (7.2).

In a one-dimensional spatial population structure with fixed neighborhoods the situation is more difficult. The contest between the strategies happens at the boundary of the neighborhoods, whereas the individuals in the interior play only against members of their own group. In this spatial structure the success of the invasion is therefore totally determined by the outcomes at the boundary.

It is almost impossible to investigate realistic spatial population structures using analytical methods; one has to use simulations. This was first done by Axelrod ([1], pages 158–168). Axelrod investigated a simple two-dimensional structure where each player had four neighbors. The selection was very strong. If a player had one or more neighbors that had been more successful, the player converted to the strategy of the most successful of them.

Axelrod's major conclusion was that mutual cooperation can be sustained in a (not too highly connected) territorial system at least as easily as it can be in a freely mixing system. We will extend Axelrod's work. First, different population structures are compared and, second, the strategies evolve controlled by the genetic algorithm.

8. Simulation results

In our simulation we have investigated the following population structures

- a small panmictic population;
- a large panmictic population (500 individuals);
- a panmictic population with randomly selected neighbors (seven);
- a one-dimensional population (ring structure with four neighbors);
- a cycle between a ring population and a panmictic population;
- a cycle between ten-island populations and a panmictic population.

In a panmictic population each individual plays against each other. In a spatial population structure the individuals play only against their neighbors. In a panmictic population with randomly selected neighbors the individuals play only against these neighbors. Most of the experiments have been done with a small population of 50 individuals. Detailed simulation results can be found in [3]. We outline in this paper only the major facts supporting or disproving Darwin's argument. In our simulations we used two-lookback strategies, which can be coded by 20 bits. This gives 2^{20} different genotypes.

In the first set of simulations we started with a population where the initial strategies had been randomly generated. Figures 1 through 4 show the results. In the figures the minimum payoff, the maximum payoff, and the average payoff of the population are plotted and multiplied by a factor of 100. All four populations cooperate most of the time. The small panmictic population (50 individuals) arrives at generation 80 at cooperation. It occasionally may change to non-cooperative behavior. This happened in our simulations at least until generation 500. The large population (500 individuals) arrives at cooperation first. The equilibrium is stable, and the variety of strategies is smaller than for the small panmictic population. The ring population supports a large variety of strategies. The maximum payoff oscillates between 3 and 5, and the minimum payoff between 0.5 and 2. The average payoff is between 2.5 and 3. In figure 4 the continent cycle is simulated. Every 20 generations the continent breaks into 10 islands. After 20 generations the islands convert back into a continent. The simulation starts with 10 islands. The simulation shows that during the continent phase the variation of the strategies is reduced. The population arrives late at cooperation, but the average fitness does not drop below 200. Overall the population remains at cooperation.




Figure 1: Small panmictic population (50 individuals).




Figure 2: Large panmictic population (500 individuals).




Figure 3: Ring population.




Figure 4: Continent cycle.

Figures 1 through 4 show individual runs that are “representative.” Each individual experiment has been repeated 10 times. The results are qualitative as described. The figures are what we believe to be “average” runs. Because of the stochastic nature of the individual runs (especially the runs with small populations) it makes no sense to average the fitness over the 10 different runs. A small panmictic population, for instance, occasionally changes to a non-cooperative behavior for a certain time. The time when this happens cannot be predicted.

The above measure, the average fitness of the population, and its variance, are suitable for investigating Darwin’s arguments qualitatively. Batz [3] has also used a more difficult measure like the distribution of classes of strategies. We briefly discussed the problem of classification of genetic specified strategies in section 6; we postpone a thorough discussion of this problem to a later paper.

This first set of experiments shows the influence of the number of individuals and of the population structure on evolution. A large panmictic population has the smallest variety and nothing new happens. Evolution within a continent cycle seems to be very robust. The ring population has the largest variety.

The difference between a large panmictic population and a small one is also shown in figures 5 and 6. The simulation of the large population started with three predefined strategies—5 ALL-D, 490 ALL-C, and 5 TFT. The population arrives after 24 generations at non-cooperative behavior and stays there. Similarly the small population started with 5 ALL-D, 40 ALL-C, and 5 TFT. In figure 6 two different selection schemes are compared. With the PGA selection the average payoff of the small population decreases only a little and the small population arrives at cooperation at about generation 150. If a stronger selection scheme is used—in our case both parents are selected proportional to their fitness—then the small population first turns to non-cooperative behavior, but it arrives at cooperation earlier than with the PGA selection.

The result of this experiment can also be explained mathematically. The initial fitness of the strategies can easily be computed. In the small population the fitness of ALL-D is given by

$$F(\text{ALL-D}) = (4 + 200 + 5)/49 = 4.27$$

Similarly we obtain $F(\text{TFT}) = 2.84$ and $F(\text{ALL-C}) = 2.69$. For the large population we get $F(\text{ALL-D}) \approx 5$, $F(\text{TFT}) \approx 3$, and $F(\text{ALL-C}) \approx 3$. The difference of the fitness between ALL-D and TFT is larger in the large population. This explains the success of ALL-D in the large population. ALL-D has success only in the small population if a strong selection is used. The “soft” selection scheme of the PGA slows down the spreading of ALL-D individuals. In this case the population does not settle on non-cooperation, but takes longer to arrive at cooperation.

The ring population oscillates as shown in figure 3. The selection scheme of the PGA is too soft for this population structure. In order to




Figure 5: Large population with an initial population of 5 ALL-D, 490 ALL-C, and 5 TFT.




Figure 6: Two parents selected versus PGA selection: 5 ALL-D, 40 ALL-C, and 5 TFT.




Figure 7: Continent cycle with an initial population of 5 ALL-D and 45 ALL-C.

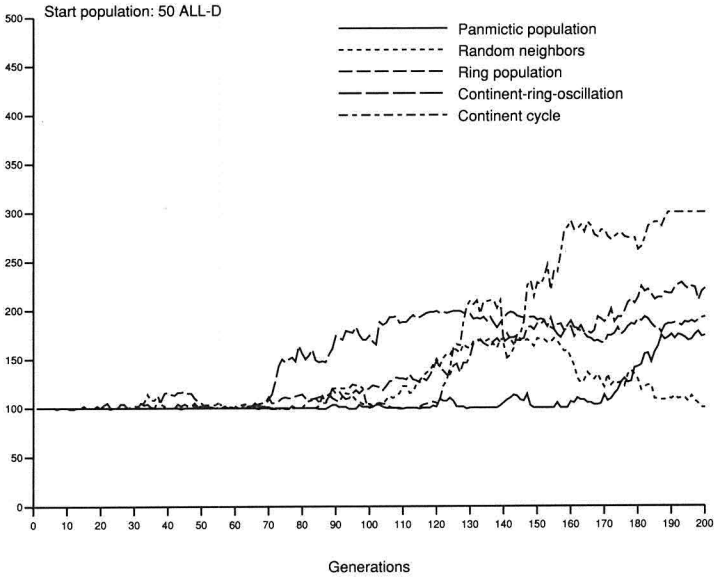


Figure 8: Start population of 50 ALL-D.

implement a higher selection pressure we introduced an acceptance test of the offspring. In the first scheme the offspring replaced the parent only if it won the IPD against the parent. The effect was dramatic. Now the population always settled on non-cooperative behavior. The situation changed with our second scheme. This extension we called the *family game*. Each mating produces two offspring. After the mating the family consisting of the two parents and the two offspring plays an IPD tournament. The winner replaces the parent. With this selection scheme the population settled on cooperative behavior. The explanation of the result is simple. In the IPD, non-cooperative strategies can be eliminated if the cooperative individuals stick together. In a single contest, ALL-D can never be beaten. It is outside the scope of this paper to compare the family game with *kin selection* proposed in sociobiology [21].

In figure 7 we investigate the invasion of an ALL-C population by a cluster of 5 ALL-D individuals. At the beginning the ALL-D population is living on a separate island, and nothing interesting happens. The average of the fitness of the population remains constant. After 20 generations the islands are converted to a continent. Now ALL-D invades the population, and the population rapidly changes to non-cooperation. But after only two continent cycles the population is again at cooperation. In a single panmictic population ALL-D successfully invades the population. Then the population stays at non-cooperative behavior if the population is large (e.g., 500 individuals).

We believe that our simulations confirm Darwin's arguments. This is also demonstrated in figure 8. Here the average fitness of the population is shown for five different population structures. The simulation started with a homogeneous ALL-D population, and we investigated whether the populations changed to cooperation. Figure 8 shows the result. We see that the population that is subjected to the continent cycle is the first to arrive at cooperation. This result was consistent over ten runs. A closer analysis of the strategies showed that the winning cooperative strategies are not naive like ALL-C, but rather resemble TIT-FOR-TAT.

In a further set of experiments we changed the game during the course of the simulation; for instance, we changed the IPD to the chicken game. The spatially structured populations adapted much faster to the new game than a large panmictic population. This is one of the extensions that have already been proposed by Axelrod for investigation ([1], page 221).

9. Conclusion

Theoretical biology has almost neglected Darwin's true evolution theory. We have investigated some of Darwin's arguments in support of his continent cycle theory by simulating an artificial ecology. The simulations are based on highly simplified assumptions. The populations are very small and the sexual reproduction has no sexual differentiation between male and female. The most important simplification is that each individual can mate with each other. Therefore we have only one species in our ecology. But the number of

different species is the original performance measure of Darwin's continent cycle theory.

We are now trying to extend our model. This model will be more ecologically oriented and contain food webs. It is part of the "conventional wisdom" of modern ecology that an increase in the number of links in the food web increases the stability of the ecosystem ([20], page 139). In order to derive this result, arguments similar to Darwin's arguments concerning the development of species on a large continent are used. The result is questioned by May, who showed that in mathematical models increasing complexity of the food web tends to beget diminished stability ([11], page 3). May suggests that "the interplay between migration and extinction in a number of local populations in a spatially heterogeneous environment can have a stabilizing effect." So we find the same problem of spatial population structures in these models.

But the main advantage of simulations can already be glimpsed from the experiments reported in this paper. Simulations provide a new intellectual perspective on evolution. Instead of having to rely only on observations of real biological systems (like Darwin) or standard mathematical models (like Wright), we are able to approach genetics and evolution as a theoretical design problem. The scientific value of the simulations is mainly given by the methodological questions, not by an individual simulation run.

Science in the past simplified evolution to obtain evolution models that could be dealt with mathematically. In using simulations there arises another problem: a single simulation run gives no scientific insight, so it has to be interpreted carefully. We have shown that a slight change in the selection scheme can dramatically change the result of the simulation.

Von Neumann pointed out that there are two ways in which a theory of automata (or any research in artificial life) might prove useful. The first way is that "some of the regularities which we observe in the organization of natural organisms may be instructive in our thinking and planning of artificial automata." The second is the converse of the first: "A good deal of our experiences and difficulties with our artificial automata can be to some extent projected on our interpretations of natural organisms." We try to use the parallel genetic algorithm in both ways.

Acknowledgment

The author thanks M. Batz, who performed the simulation experiments.

References

- [1] R. Axelrod, *The Evolution of Cooperation* (New York, Basic Books, 1984).
- [2] R. Axelrod, "The Evolution of Strategies in the Iterated Prisoner's Dilemma," pages 32-41 in *Genetic Algorithms and Simulated Annealing*, edited by L. Davis (Los Altos, Morgan Kaufmann, 1987).

- [3] M. Batz, "Evolution von Strategien des Iterierten Gefangenen Dilemma" (Master's thesis, Universität Bonn, 1991).
- [4] J. F. Crow, W. R. Engels, and C. Denniston, "Phase Three of Wright's Shifting Balance Theory," *Evolution*, **44** (1990) 233–247.
- [5] Charles Darwin, *On the Origin of Species by Means of Natural Selection* (London, John Murray, 1859).
- [6] J. Felsenstein, "A Pain in the Torus: Some Difficulties with Models of Isolation by Distance," *American Naturalist*, **109** (1975) 359–368.
- [7] J. Felsenstein, "The Theoretical Population Genetics of Variable Selection and Migration," *Annual Review of Genetics*, **10** (1976) 253–280.
- [8] R. A. Fisher, *The Genetical Theory of Natural Selection* (New York, Dover, 1958).
- [9] J. H. Holland, *Adaptation in Natural and Artificial Systems* (Ann Arbor, Univ. of Michigan Press, 1975).
- [10] R. E. Marks, "Breeding Hybrid Strategies: Optimal Behavior for Oligopolist," pages 198–207 in *3rd International Conference on Genetic Algorithms*, edited by H. Schaffer (San Mateo, Morgan Kaufmann, 1989).
- [11] R. M. May, *Model Ecosystems* (Princeton, Princeton University Press, 1974).
- [12] J. K. Miller, "The Coevolution of Automata in the Repeated Prisoner's Dilemma," Technical report, Santa Fe Institute (1989).
- [13] H. Mühlenbein, M. Gorges-Schleuter, and O. Krämer, "Evolution Algorithms in Combinatorial Optimization," *Parallel Computing*, **7** (1988) 65–88.
- [14] H. Mühlenbein, M. Schomisch, and J. Born, "The Parallel Genetic Algorithm as Function Optimizer," *Parallel Computing*, to appear.
- [15] H. Mühlenbein, "Parallel Genetic Algorithms, Population Dynamics and Combinatorial Optimization," pages 416–421 in *3rd International Conference on Genetic Algorithms*, edited by H. Schaffer (San Mateo, Morgan Kaufmann 1989).
- [16] H. Mühlenbein, "Evolution in Time and Space: The Parallel Genetic Algorithm," pages 316–337 in *Foundations of Genetic Algorithms*, edited by G. Rawlins (San Mateo, Morgan Kaufman, 1991).
- [17] J. Maynard Smith, *Evolution and the Theory of Games* (Cambridge, Cambridge University Press, 1982).
- [18] R. Tanese, "Distributed Genetic Algorithms," pages 434–440 in *3rd International Conference on Genetic Algorithms*, edited by H. Schaffer (San Mateo, Morgan Kaufmann, 1989).

- [19] J. von Neumann, *The Theory of Self-Reproducing Automata* (Urbana, University of Illinois Press, 1966).
- [20] E. O. Wilson and W. H. Bossert, *A Primer of Population Biology* (Sunderland, MA, Sinauer Associates, 1971).
- [21] D. S. Wilson and L. A. Dugatkin, "Nepotism vs Tit-for-Tat, or, Why Should You Be Nice to Your Rotten Brother," *Evolutionary Ecology*, **5** (1991) 291–299.
- [22] S. Wright, "The Roles of Mutation, Inbreeding, Crossbreeding and Selection in Evolution," pages 356–366 in *Proceedings of the 6th International Congress on Genetics* (1932).
- [23] S. Wright, "Factor Interaction and Linkage in Evolution," *Proceedings of the Royal Society of London B*, **162** (1965) 80–104.