Landscape Ruggedness in Evolutionary Algorithms

Krasimir Kolarov
Interval Research Corporation, 1801-C Page Mill Road, Palo Alto, CA 94304
E-mail: kolarov@interval.com

1. Introduction

This paper describes a model for the exploration of the dynamics of interaction in a population of individuals using an evolutionary approach. In particular we analyze the effect of the complexity of the fitness landscape and the population size on the performance of an evolutionary algorithm in terms of speed of fixation and fixation to suboptimal individuals. Our evolutionary model resembles those in the field of Population Genetics and approaches evolution as a process of adaptation rather than an optimization.

The simulation results from our experiments are justified with a theoretical probabilistic analysis of the dynamics of the population with and without recombination. The paper introduces several new insights into the interaction and roles of the different parameters of an evolutionary system.

There have been very few attempts to theoretically analyze the dynamics of interaction between the different operators and the performance of GA (e.g. Holland's *building blocks* and *schema* [8], [1]). The problem of finding the appropriate (and in some cases the optimal) population size have been discussed for certain applications in the GA literature [5]. There has also been research of the role of different selection schemes on GA performance [8]. Interesting work on analytical tools for GA include the approach in [13] based on the Perron-Frobenius theorem, and the Markov chain analysis described in [12].

We will give a detailed description of our model in Section 2. The experiment for which we vary the ruggedness of the landscape is described in Section 3. To better understand the results in this section, we performed a theoretical analysis of the model and derived recursive formulas describing the dynamics of interaction in the model. These results are explained in section 4. Section 5 draws some conclusions from our work.

2. The model

We performed experiments with a simple GA type system. Our model includes a finite number of diallelic haploid organisms. The chromosome consists of one string of bits and we chose each genotype or bit string to be of fixed length 20. Each allele can have two values: 0 or 1. The GA operators are crossover and mutation and there is random mating within non-overlapping generations. The initial population is selected at random with equal probability of 0's and 1's. Our model has fixed rates of recombination and mutation. We use single (one-point) recombination with the break point selected at random, uniformly across the chromosome. The offspring are subjected to mutation and selection. New offspring are accumulated until the fixed population size is reached. At that time we have formed the new generation and the current one becomes its parental generation. Thus we do not have explicit elitism - the new generation is entirely formed by applying the genetic operators to the previous one. For each case (i.e. random initial population) this process continues until all the individuals within the same generation have the same genotype - this is the condition for fixation. For statistical test we run each set of input parameter settings for 100 cases.

As explained in [9] this method of constructing the offspring generation where every member of the new population is added as a result of individual "selection" after application of the genetic operators to the previous generation, is a characteristic of population genetic analysis. In other words we implement fitness based proportional selection, independent of the fitness values for the rest of the population.

3. Fitness Landscape with Randomly Controlled Ruggedness

To control the ruggedness of the fitness landscape we use the same formalism as proposed by Bergman and Feldman 1992 [3], i.e.:

$$F(i) = \sum_{x=1}^{M} \frac{\Psi(x).\sin(x\pi \frac{i}{l})}{\sqrt{x}} \qquad (1)$$

where l is the number of loci in the genotype (20 in our case).

$\Psi(x)$, x = *1,2, ..., M* are random numbers that are uniformly distributed in the interval [0,1] . The *number of random elements M* is the landscape variable in this experiment. It determines the ruggedness of the fitness function landscape. When $M = 1$ there is only one wave of the sine function and its peak is the maximum we are looking for. The larger is M , the more rugged and complicated is the fitness function.

The fitness function is calculated for a particular genotype by counting the number of 1's in the genotype to give the index i used in formula (1) to find the fitness value for that individual. Thus it is clear that we can have maximum 21 different fitness values in our environment.

Note that the optimal genotype does not have all of its alleles equal to 1. In addition, as can be seen from the simulation results, for different cases of the same run, the genotypes at the fixed point can be different. For example if the optimal individual has ten 1's, i.e. i=10, the corresponding fixed point can be 000111110011100011100 in one case and 111110000111000011100 in another one. There are $\binom{l}{i}$ possible genotypes with the same optimal fitness.

In all experiments the mutation rate was set to zero, because our goal was to analyze fixation, and if there is any bi-directional mutation, fixation (as defined in our model) does not occur. However our framework can also be used for analyzing other processes that include mutation.

As we mentioned already for this first experiment we varied the PopSize, the number of random elements M and in few cases the recombination rate. The first two tables of results correspond to the case when the recombination rate is equal to 0.4 , i.e. there is 40% probability of one recombination event over the entire chromosome. This value for the recombination rate is often used in population genetics literature and allows us to compare our results to those in [2] and [3]. Table 1 displays the number of generations to fixation to any point in the search space for three population sizes: 100, 150 and 200 individuals, and six fitness landscapes: for *M=1, 6, 10, 13, 16* and *20*. The *mean*, *minimum* and *maximum* are taken over the total number of cases -> 100. The standard deviation for all cases in Table 1 and 2 follow the trend of the mean. We also performed χ^2 tests (see [4]) to analyze the results from all experiments. The tests confirmed the significance of the reported data with respect to the null hypothesis.

As we can see from Table 1 the more complicated the fitness function the faster fixation occurs. In addition a larger population generally takes longer to reach fixation.

For comparison we performed the same simulation with the Recombination Rate equal to zero. In this case we have neither mutation, nor recombination and the system performance is limited by the fitness of the most fit individual in the initial population.

Table 2 shows that for simple landscapes the model without recombination takes significantly less generations to fixate. For more rugged landscapes ($M > 6$) the recombination causes faster fixation. However if we look at the overall computing time the differences are not statistically significant for large M's.

Generations to fixation: mean (min, max) from 100 cases; Recombination - 0.4						
	Number of Random elements					
Population	1	6	10	13	16	20
100	335(159,648)	283(41,749)	193(28,643)	171(44,360)	160(37,508)	154(13,390)
150	493(174,104)	374(57,760)	244(33,628)	236(18,544)	226(71,676)	220(36,677)
200	611(288,1580)	441(48,1160)	294(35,885)	274(16,487)	257(21,630)	260(77,678)

Table 1. Generations to fixation for rugged landscape with recombination

Generations to fixation: mean (min, max) from 100 cases; No Recombination						
	Number of Random elements					
Population	1	6	10	13	16	20
100	173(46,745)	164(25,560)	153(31,617)	156(15,436)	180(48,471)	187(42,786)
150	259(71,813)	294(145,629)	255(49,1360)	277(71,998)	260(83,864)	280(68,820)
200	347(75,1357)	348(163,623)	379(58,1019)	333(82,976)	372(111,1493)	372(110,1206)

Table 2. Generations to fixation for rugged landscape without recombination

The next Table 3 analyzes the cases when the population fixates to a value that is different from the global optimum. The numbers in parenthesis indicate the number of 1's in the genotype on which different cases fix. For example for M=6 and Population Size of 100 most of the cases fixated to a genotype with eight 1's, while few cases (here the number is five) fix on a genotype with nine 1's and one case fixated to a genotype with two 1's. The total number of cases considered is again 100.

For all landscapes the model fixed on the local maximum that is closest to the mean of the distribution of the initial population. When there was only one random element in the fitness function (a bell-shaped or a U-shaped landscape), this local maximum happens to be the global maximum as well. Thus the numbers in Table 3 are smaller for this case compared to the rest of the table. However, the number of cases that do not fix on local maximum is smaller for more rugged landscapes ($M > 1$). One possible explanation of this phenomenon is that GA doesn't do well for functions with relatively flat tops when there is not enough selection pressure to drive it to the local optimum (it stops at $1-\varepsilon$ for some small ε rather than at 1).

From Table 3 we can also derive that, as the population size increases there is generally less fixation on suboptimum with all other conditions been the same.

A similar analysis for the case of no recombination is illustrated in Table 4. Clearly the presence of recombination leads to less fixation on suboptimum, *X*, especially in the case of a simple, non-rugged landscape ($M = 1$). In fact although as we pointed out above, it takes less generations for the model without recombination to fix in that case, it actually fixates more often on the "wrong" genotype. We can also compare the mean fitness of all 100 cases after fixation for the landscapes. Consistently this fitness is higher for the model with recombination than that without for all parameter sets. We would like to emphasize one feature of recombination in our model. Typically in GA it is assumed that recombination increases variability in the environment because it introduces new individuals. However in our experiments this is not the case and the reasoning goes as follows.

We start with a Gaussian distribution of the initial population, i.e. most of the individuals have the 10 ones (out of 20 maximal).

- If there is no recombination (and no mutation) the next generation will typically again have individuals with number of ones close to 10 and there will be a large variety of them. This trend will propagate throughout the simulation, thus maintaining large variability in the population. This is also the reason why it takes longer for this model to fixate and *X* is larger than when recombination is present.

Number of Cases that do not Fixate to the Global Optimum; Recombination - 0.4						
	Number of Random elements					
Population	1	6	10	13	16	20
100	28	97 (8,9,2)	98 (8)	100(13)	100(12,15,7)	99(12)
150	18	95 (8,9,2)	96 (8)	99(13)	100(12,15)	100(12)
200	7	93 (8,9,2)	94 (8)	99(13)	99(15,12,7)	100(12)

Table 3. Suboptimal Fixation for rugged landscape with recombination

Number of Cases that do not Fixate to the Global Optimum; No Recombination						
	Number of Random elements					
Population	1	6	10	13	16	20
100	48	100(8,9,2,3)	100(8,2)	100(13,1)	100(12,15,7)	100(12)
150	39	100(8,9,2,3)	100(8)	100(13)	100(12,15,7)	100(12)
200	25	100(8,9,3,2)	100(8)	100(13)	100(12,15)	100(12)

Table 4. Suboptimum fixation for rugged landscape without recombination

- With recombination, individuals with lower fitness are produced which because of the distribution of the initial population, are less statistically represented than the individuals with number of ones around 10. Thus in the second generation we will actually have less variability than in the initial one. In other words in future generations the rare individuals with higher and lower fitness than the median ones will be reinforced. As a result the population fixes faster and with higher overall fitness than the case without recombination. However recombination in this case reduces the variability.

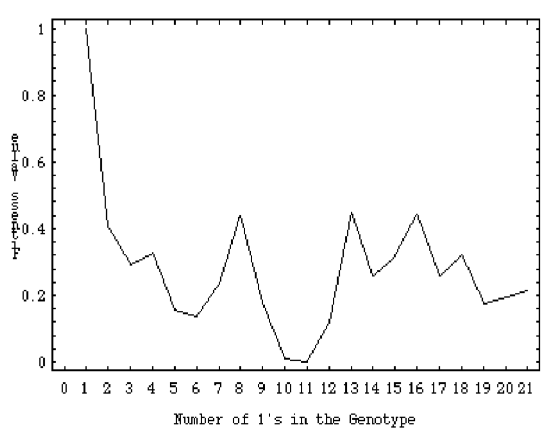

Figure 1. Fitness Landscape for the case $M = 16$

The advantage of recombination for faster fixation to an individual with higher fitness is further illustrated in the example in Figure 1. When the initial population has a 50% chance of 1's and 50% chance of 0's for the values of the alleles the population invariably fixes to an individual with the locally maximal fitness $F(13)$ (see eq.(1)) when $M = 16$. The predominant point of fixation moves as we vary the percentage of 1's in the initial population. For example if we have initially 90% 1's, the population fixes at $F(16)$ in most of the cases (with the model with recombination performing better than the one without it). As seen in Figure 1, this is a local maximum with a smaller fitness than $F(13)$. On the other side with 25% initial 1's the population with recombination fixes in 97% of the cases to the global optimum $F(1)$. Note that in the figure $F(1)$ (and resp. the rest of the indices) corresponds to the case of zero number of 1's in the genotype).

The corresponding model without recombination fixes to the global optimum only in 27% of the cases, with the rest fixing on a local maximum $F(8)$. Thus the mean fitness at fixation is lower than the case with recombination.

In addition to the analysis mentioned above, we did limited experiments with varying the recombination rate. It is not clear from this limited data what are the exact correlations. However, we can conjecture that with increase in recombination, the genetic drift decreases and the speed of fixation increases. The question of the role of the recombination is very interesting and we will present some results on that topic in a future publication. This question was extensively discussed for diploid individuals in [3].

4. Theoretical Analysis

In this section we will extend our discussion of the role of the shape of the fitness landscapes and derive iterative relationships that describe the dynamics of interaction in our system.

Let us denote with $F(s)$ the value of the fitness function for an individual with s number of 1's. With $a_k(s)$ we denote the expected number of individuals with fitness $F(s)$ in the population at generation k. As we have mentioned before N is the population size and l is the length of the genotype (in our examples $N = 100$ and $l = 20$). Because we always keep constant population size, we have:

$$\sum_{s=0}^{l} \varpi_k(s) = N \qquad \text{for every } k \qquad (2)$$

The average expected fitness of the population at generation k, is:

$$\Phi_k = \frac{1}{N}\sum_{s=0}^{l} F(s)\varpi_k(s) \qquad \text{for every } k \qquad (3)$$

We want to find a recursive relation describing the expected number of individuals with fitness $F(s)$ at generation $k+1$ in terms of the number of individuals with different fitness at generation k. For simplicity let us first consider the case with no recombination, i.e. the recombination rate $r = 0$.

4.1 Derivation without recombination.

The only way to add an individual with fitness $F(s)$ to the new generation without recombination is to choose it from the parent generation. The algorithm picks an individual with that fitness, compares it to a uniformly distributed random number and adds it to the new generation if it is larger than the random number. If it is not larger, the same cycle is repeated until an individual has been added to the new generation. We can derive the following formula:

$$\varpi_{k+1}(s) = \frac{F(s)}{\Phi_k}\,\varpi_k(s) \qquad (4)$$

This is the main recursive relationship that calculates the expected number of individuals with fitness $F(s)$ in a new generation as a function of the expected number of individuals with that fitness in the parent generation, the expected average fitness of the parent generation and the fitness $F(s)$. This relation is similar to the Fisher's Fundamental Theorem of Natural Selection (see [4] and [7]). In the GA literature an example of similar analysis is described in [6].

Using (4) we can express the expected number of individuals with fitness $F(s)$ at generation k , as a function of this number for the initial randomly generated population and the expected average fitness of all generations before k. We can also derive a recursive relationship for the expected average fitness of the new generation in terms of the data from the parental generation:

$$\Phi_{k+1} = \frac{1}{N\Phi_k}\sum_{s=0}^{l} F^2(s)\,\varpi_k(s) \qquad (5)$$

Equations (4) and (5) describe the dynamics of interaction of our model. Theoretically for infinite population size and appropriate recombination rate the population should always tend to the optimum. If there is no recombination the population should fixate to the most fit individual initially present in the population. However we have observed that in numerous cases the **finite** population size in the simulations causes the model to fixate on local optima. Although we can not predict with certainty the direction of fixation of the population based on the initial conditions, we will describe a simple rule that lets us predict the concentration of individuals in a given generation based on the distribution of its parent generation. We will illustrate this rule for the case of the initial distribution in the simulation.

Let us denote with p_I the initial probability of 1's in the first randomly generated population. In our discussion so far we have assumed that every allele for each individual in the initial population has equal chances (0.5) of being a 1 or a 0. In general however we can choose different distribution of 1's and 0's and we can vary that using p_I. Let us also denote with X the brake point $p_I.l$ (i.e. $X = p_I.l$). Then we can form the following quantities, which we will call "left" and "right moment of F with respect to X":

$$\Phi_0^+ = \frac{1}{N}\sum_{s=0}^{\lfloor X\rfloor} F(s)\varpi_0(s) \qquad \text{and}$$

$$\Phi_0^- = \frac{1}{N}\sum_{s=\lceil X\rceil}^{l} F(s)\varpi_0(s) \qquad (6)$$

Using relationships (4) and (6) we can show that the condition $\Phi_0^+ > \Phi_0^-$ is equivalent to the condition:

$$\sum_{s=0}^{X} \varpi_1(s) > \sum_{s=X}^{l} \varpi_1(s) \qquad (7)$$

In other words if the left moment of F with respect to X" is larger than the right moment in the initial population, in the first generation there will be larger expected number of individuals whose fitness is to the left of X than those to the right. This property is also true for any two consecutive generations but it does not necessarily propagates throughout all generations. In particular we might encounter a situation, where even though the left moment with respect to a fixed point is larger than right one, these moments are actually smaller than the corresponding moments in the previous generation. In that case the fixation point is likely to be on the right side of X, contrary to what the first generation will suggest.

For completeness we will point out that for our model the expected number of individuals with fitness $F(s)$ in the initial population is:

$$\varpi_0(s) = N\binom{l}{s} p_I^s (1-p_I)^{l-s} \qquad (8)$$

4.2 Derivation with recombination

In this section, we will provide some formulas for the general case when the recombination rate $r \neq 0$. There are two ways to generate an individual with fitness $F(s)$ to the new generation. The first one is if as a result of a recombination, a new individual with fitness $F(s)$ is created and it is added to the new generation. The probability of a recombination event is r because we compare the recombination rate with a uniformly distributed random number. The second way of selecting an individual with fitness $F(s)$ corresponds to the case when there is no recombination and an individual is chosen as in Section 4.1. The probability of no recombination is $1-r$ thus in analogy to the previous section we can derive:

$$\varpi_{k+1}(s) = N[r.PR + (1-r).PN] \qquad (9)$$

Here PR is the probability of selecting an individual with fitness $F(s)$ with recombination and adding it to the new generation, while PN is the corresponding probability without recombination. If the individual that is selected is not added to the new generation, then we perform a second try to add new individual, then third try, etc. until a new individual is added. Thus using the same reasoning as in the previous section, we can show that the overall probability of adding an individual with fitness $F(s)$ is $\dfrac{F(s)}{\Phi_k}$. Using that and the reasoning about the probability of selecting individuals with fitness $F(s)$ without recombination, we can derive:

$$\varpi_{k+1}(s) = N[r.\mathrm{P_r}(s) + (1-r)\frac{1}{N}\varpi_k(s)]\frac{F(s)}{\Phi_k} \qquad (10)$$

where $P_r(s)$ is the probability of generating an individual with fitness $F(s)$ with recombination.

At this point if we had mutation in the model, the probability $P_r(s)$ would have been similarly subdivided to account for the effect of mutation on the individual genotype. However as we pointed out earlier, in that case there would be no fixation in the finite population model. Thus we assume that the Mutation Rate is set to zero.

We can derive the following formula for $P_r(s)$:

$$\frac{1}{(l+1)N^2}\sum_{x=0}^{s}\sum_{z=x}^{l}\sum_{p=s-x}^{l}\left[\sum_{y=0}^{l}\frac{\binom{y}{x}\binom{l-y}{z-x}\binom{l-y}{s-x}\binom{y}{p-s+x}}{\binom{l}{z}\binom{l}{p}}\right].\omega_k(z)\omega_k(p)$$

$$= \frac{1}{(l+1)N^2}\sum_{x=0}^{s}\sum_{z=x}^{l}\sum_{p=s-x}^{l} B(x,z,p,s)\omega_k(z)\omega_k(p) \qquad (11)$$

Here the coefficients $B(x,z,p,s)$ are numerical constants for given x, z, p and s. From (11) we can see that the probability $P_r(s)$ is a linear combination with constant coefficients of terms of the type $\omega_k(z)\omega_k(p)$ where the parameters z and p vary over the range $[0,l]$. We can add the coefficients for terms that have permuted indices and arrange the resulting coefficients in lower triangular matrices where each entry r_{ij} of the matrix $R_{k+1}(s)$ corresponds to the coefficient in (11) in front of the term $\omega_k(i)\omega_k(j)$.

We can see from (11) that with the increase of s we get less representation in the R-matrices of the $\omega_k(z)$ for $z < s$ and stronger representation of those with $z > s$. At the same time if $\Phi_0^+ > \Phi_0^-$ then higher weight is given to the individuals with fitness less than the mean of the initial population X. Because those individuals are more represented in the R-matrices to the left of X, that means that in the next generation those individuals will be even more represented and thus by induction the population in the next generation will be moving toward the left of X. By analogy if $\Phi_0^+ < \Phi_0^-$ the individuals to the right of X get more weight and representation.

For larger l we can clearly see the structure of this argument because for a given X, all the R-matrices for $s > X$ will have upper left rectangles of zeros. Accordingly the R-matrices for $s < X$ will have zeroes in there bottom right rows. These properties confirm the validity of our conclusion for the general case.

We should note however that this effect is not as strong as in the case of no recombination, where there was a clear tendency toward one direction or the other. In this case although the R-matrices have a lot of zeroes for say small s to the right of a given value, not all of the coefficients are 0, and those that are represented have a neutralizing influence. An interesting question that is beyond the scope of this paper, is to analyze the influence of r on the dynamics of the system.

Finally we want to point out that although the derivation referred to the "front" positions of the first parent and the "back" positions of the second, in fact the exact location or even adjacency of the positions is not important here. In fact our results will be valid with slight modifications for any single event recombination (not necessarily simple one-point ones).

5. Summary and Conclusions

With respect to fixation on suboptimal fitness in the simulations we can conclude that:

- The increase of the population size decreases the rate of "wrong" fixation due to less sampling error.

This result is in accord with the results in population genetics that are due to sampling error. As described in Chapter 5 of [11] selection strength has to be above certain value for given population size, in order for the system to overcome the "genetic drift". Genetic drift in that context refers to the sampling error by the gametes - an internal source of chance effects in evolution.

We have shown that:

- More rugged landscapes fix faster and with higher overall fitness to the local optimum that is closest to the mean value of the fitness of the initial population.

The extent to which this result holds depends on the recombination rate. As shown in Section 3 our general conclusion is that:

- The presence of recombination in the model leads to a decrease in the fixation on suboptimum fitness especially for a simple, Gaussian landscape due to introduction and maintenance of diversity. For a Gaussian distribution of the initial population, recombination reduces the variability in the population.

The results for rate of fixation reflect the speed for all fixations, correct or incorrect. We can summarize the results from Section 3 as follows:

- The more complicated the fitness landscape, the faster fixation occurs. Populations with complicated, rugged fitness terrain fixate faster with recombination rather than without it. Increased Recombination increases the speed of fixation.

- As expected larger populations take longer to fixate.

These results also follow from the equations (4), (10) and (11) derived in Section 4. Since the rate of selection $\frac{F(s)}{\Phi_k}$ does not depend on the population size, clearly larger populations will take longer to bring all except one of the $\varpi_k(s)$ to zero (i.e. to fixate). In addition when F is relatively flat on the top (larger s), for Gaussian initial distribution, there are several suboptimal values of s around the optimum for which both $F(s)$ and $\varpi_k(s)$ (and thus their product) are close to the optimal ones, i.e. selection pressure is weak. From (4) and (11) follows that in future generations the number of these suboptimal individuals will be close to the number of the optimal. Thus on average it will take longer to fixate on the optimal individual and the probability of wrong fixation will be higher.

In addition to the above results we emphasize the following new insight that is a direct result from this work:

It comes from the rule that we derived experimentally and proved theoretically regarding the expected distribution of the population at a given generation based on the distribution at the previous generation. We showed that if the average fitness (the moment) of the individuals at one side of the mean of the initial population is larger than the corresponding one at the other, then in the next generation there will be more individual on that side than on the other. As we pointed out, this property is only valid for one generation at a time and it does not necessarily propagate throughout the entire simulation. In addition the formulas in Section 4 describe theoretically the dynamics of interaction within our system and allow for further exploration of the influences of different parameters and conditions.

Finally we would like to mention that the ruggedness of the fitness function have similar effect as the strength of the selection pressure. If we compare the analysis in Section 3 with the results in [9] we can see that both the increase of the ruggedness (larger N) and the increase of the selection pressure lead to faster fixation and less "wrong" fixation.

Acknowledgments

I thank Dr. Aviv Bergman and Professor Marc Feldman for the detailed and useful comments and corrections on the paper.

References

[1] L. Altenberg, *The Schema Theorem and Price's Theorem,* FOGA-3 (eds. Whitley and Vose), 1995, Morgan Kaufmann, pp.23-49.

[2] A.Bergman and M.Feldman, *More on Selection for and against Recombination,* Journal of Theoretical Population Biology, Vol. 38, No.1, August 1990.

[3] A.Bergman and M.Feldman, *Recombination Dynamics and the Fitness Landscape*, Physica D 56, 1992, p. 57-67.

[4] W. J. Ewens, *Mathematical Population Genetics*, Springer-Verlag, 1979.

[5] D.Goldberg, *Sizing Populations for Serial and Parallel Genetic Algorithms,* Proceedings of the Third International Conference on Genetic Algorithms, 1989, p. 70-79.

[6] J. Grefenstette, *Conditions for Implicit Parallelism,* FOGA-1 (ed. Rawlins), 1991, Morgan-Kaufmann, p. 252-261.

[7] D.Hartl and A.Clark, *Principles of Population Genetics*, Sinauer Associates, Inc., 1989.

[8] J.Holland, *Adaptation in Natural and Artificial System*, MIT Press, 1992.

[9] K.Kolarov, *The Role of Selection in Evolutionary Aagorithms*, Proceedings of the 1995 IEEE International Conference on Evolutionary Computation, Perth, 1995, p. 86-91.

[10] B.Manderick, M. de Weger and P.Spiessens, *The Genetic Algorithm and the Structure of the Fitness Landscape*, Proceedings of the Fourth International Conference on Genetic Algorithms, 1991, p. 143-150.

[11] J.Roughgarden, *Theory of Population Genetics and Evolutionary Ecology: An Introduction*, Macmillan Publishing Co., 1979.

[12] G. Rudolph, , Proceedings of the 1996 IEEE International Conference on Evolutionary Computation, 1996, p. 50-54.

[13] M. Vose and C. Liepens, *Punctuated Equilibria in Genetic Search,* Complex Systems 5, 1991, p. 31-44