

Analysis of the Simple Genetic Algorithm on the Single-peak and Double-peak Landscapes

Alden H. Wright *
Computer Science Department
University of Montana, USA
wright@cs.umt.edu

Jonathan E. Rowe
School of Computer Science
The University of Birmingham, UK
j.e.rowe@cs.bham.ac.uk

James R. Neil
School of Computer Science
The University of Birmingham, UK
j.r.neil@cs.bham.ac.uk

December 2, 2001

Abstract

We compare the behavior of a GA with and without crossover. A simple GA with crossover can have two stable fixed points (bistability) on the single-peak landscapes for string lengths at least 8. There are catastrophic transitions from one to two and back to one stable fixed point as the mutation rate increases. Crossover also causes the fixed-point population distribution to change. For a fixed point near the peak, there are fewer copies of the optimum string, more copies of near-optimum strings, and less copies of far-from-optimum strings. An example is given where average population fitness decreases to the minimum possible population fitness for the with-crossover GA, while the average population fitness increases for the without-crossover GA. The primary tool in obtaining these results is the Vose dynamical system model.

1 Introduction

A major goal of GA theory is to understand the role of crossover in a GA. Much of the early work on GA theory concentrated on the role of crossover in combining “building blocks” from different individuals into a single individual. This paper shows how crossover can have other kinds of effects. In comparison to a GA without crossover, it can prevent the GA from converging to a population with many copies of the optimum string.

The NEEDLE (needle-in-the haystack, single-peak) fitness function has fitness $1 + a$ where $a > 0$ for the string of all zeros, and fitness 1 for every other point. (This paper assumes a fixed-length binary string representation.)

*This paper was written while Alden Wright was visiting the School of Computer Science, University of Birmingham, UK. His visit was supported by supported by EPSRC grant GR/R47394.

For string lengths of at least 8, if the mutation rate is less than a lower critical value, then there is a single stable fixed point near the uniform optimum population. As the mutation increases through the lower critical value, a new stable fixed point which is close to a random population occurs. For mutation between the critical values, there are two stable fixed points (bistability), one near the needle and one corresponding to a random population. As the mutation rate increases through the upper critical mutation rate, the fixed point near the needle disappears, and the only stable fixed point is close to a random population. Thus, there is a “catastrophic” transition where the fixed point close to the optimum is lost. A mutation/selection GA has only a single fixed point for each mutation rate. It is well known from work on the Eigen model that there is an “error threshold” mutation rate above which the fixed point corresponds to an approximately random population. However, for a finite string length this transition is gradual, whereas with crossover the transition is abrupt and catastrophic.

The population distribution at a fixed-point population is quite different with and without crossover. Crossover decreases the frequency of the optimum string, increases the frequency of the near-neighbors of the optimum string, and decreases the frequency of the strings far from the optimum string.

These results on the NEEDLE fitness function confirm earlier results by Boerlijst et. al. [3] and others (see section 2). Our modeling methodology is quite different: we use the Vose dynamical system model which is exact in the infinite population limit.

The BINEEDLE (double needle-in-the haystack, double-peak) fitness function has fitness $1 + a$ for the string of all zeros and the string of all ones, and fitness 1 for every other point. We show that the a GA using crossover can exhibit tristability on this fitness function for a range of mutation rates. There can be a fixed point near each needle point plus a fixed point which is close to a random population. These results are confirmed by running a finite population GA.

We give an example using the BINEEDLE fitness function that shows that the with-crossover GA, when started from some initial populations, will follow a trajectory that decreases average fitness to a population with close to minimum fitness, whereas the without-crossover GA in the same situation will follow a trajectory that increases fitness.

2 Literature Review

2.1 The Boerlijst models and results

In [3] Boerlijst et. al. model the development of virus populations¹ on isolated peak and plateau landscapes. Their mathematical model, which assumes infinite populations, incorporates both mutation and crossover (an addition to many previous models that assume only mutation). Crossover is an important operator for many viruses, notably retroviruses such as HIV are able to recombine their genomes. Virus crossover occurs when an already infected cell is super-infected by another strain. Super-infected cells may be of homozygous or heterozygous types.

The Boerlijst model consists of a system of differential equations that attempt to capture the dynamics of cell infection. There are three levels of model, selection only, selection/mutation, and selection/mutation/recombination. The models keep track of the number of uninfected cells, infected cells and free viruses. In addition, the recombination model keeps track of super-infected cells. A virus strain

¹In the biological literature an ensemble of genetically distinct variants of a virus is known as a quasi-species

corresponds to a binary string, and since a super-infected cell has two virus strains, it corresponds to a pair of binary strings. Recombination corresponds to uniform crossover on binary strings. Boerlijst et. al. give very little information about how they go from their model to the results that they present in their paper, so it is hard to see what simplifying assumptions are used in presenting results where all strains in the same error (united) class have seem to have the same frequency.

Their specific results are for a string length of 15. When recombination is considered on an isolated peak landscape it is found that recombination leads to more compact populations - that is there are less rare mutants but also less of the fittest strain in a fixed-point population. They also find that there is a region of bistability where the population can maintain itself on the fitness peak, but fails to adapt to it when not started near the peak. Bistability arises if the recombination rate is large compared with the selection coefficient.

From extensive work on the Eigen model (see [5] for example) it is known that there is an “error threshold” for the selection/mutation model. For mutation above the error threshold, the population cannot maintain itself on the landscape peak and reverts to an approximately random population. Boerlijst et. al. find that the presence of recombination shifts the error threshold to lower mutation rates. They find that on the single-peak landscape, recombination is disadvantageous for the virus.

They also consider a plateau landscape where error classes close to the single-peak class have increased fitness. They also find bistability in this case, and they find that recombination can be selected for and can be advantageous for the virus for small mutation rates.

Boerlijst et al. also use a crude approximation to the isolated peak landscape called the error tail model. The error tail model considers only the number of copies of the optimal genotype $X(t)$ (the master sequence) and the number of non-optimal genotypes $Y(t)$ (the error tail) at a given time t . The error tail reproduces with a reduced rate to that of the master sequence and only mutation and recombination that take the optimal genotype to the sub-optimal genotype are considered. This model is bistable for higher recombination rates, and recombination shifts the error threshold to lower mutation rates.

2.2 Other models and results

Ochoa and Harvey [7] adjust the error threshold model (without recombination) to account for finite populations. Then they do evolutionary computation simulations that give population distributions that are qualitatively similar to that shown by Boerlijst et. al. [3] for infinite populations. Ochoa and Harvey use a generational GA with fitness proportionate selection and uniform crossover, they set the crossover rate to 0.5 and use string lengths of 15 to make the results more comparable to [3]. Ochoa and Harvey do not mention the bistability phenomenon.

In [1] Barnett treats finite population dynamics analytically using the error tail simplification cast in a Moran birth death model. Mutation is again considered only in the optimal to sub-optimal direction but crossover is considered in both directions. The probability of producing the optimal strain by crossover is dependent on the distribution of the error tail $Y(t)$, this means that the Markov property is not conserved by the model. To rectify this it is assumed that the distribution in the tail is uniform.

The model exhibits stochastic bistability with one stationary distribution bunched around the master sequence and the other in the error tail. Stochastic bistability means that over time a simulation will move between the stationary distributions (modes).

In [2] Barnett also analyzes isolated peak landscapes with a model that assumes an infinite population

and an infinite string length. In this model an incidence vector of the classes of unitation is created and then expressed in terms of its generating function. Since he cannot solve the model explicitly, he gives two empirical approximations, one where the height of the peak is low, and one where the height of the peak is high. He also evaluates the stability of fixed points by calculating the principal eigenvalues of the differential. However, since his model assumes that all elements of each unitation class have equal representations, this does not show stability in the space of all populations. He finds that below the error threshold (the upper critical mutation rate in this paper), there are two stable and one unstable fixed points, and above the error threshold there is one stable fixed point.

3 The Vose Dynamical System Model

Michael Vose [9] has developed a deterministic discrete-time dynamical system model of the simple genetic algorithm. This model is exact in the infinite population limit. In addition, it gives the exact probability distribution that is used in selecting individuals for the next generation in a finite population GA. Thus, it can also be used as the basis for a Markov chain model which is exact for the finite population simple GA.

In this model, populations are represented as vectors indexed over the search space Ω . If $k \in \Omega$ is a binary string, and if x represents a population, then x_k is the proportion of string k in the population. For example, if the string length is 2, then $\Omega = \{00, 01, 10, 11\}$. The multiset $\{00, 01, 01, 11, 11\}$ is a population that is represented by the vector $\langle 1/5, 0, 2/5, 2/5 \rangle$. A population vector can also be interpreted as a probability distribution over the search space Ω .

Let ℓ be the string length, and let $n = 2^\ell$. We identify Ω with the integers in the range from 0 to $2^\ell - 1$ by identifying a string with the integer given by its binary representation.

A *function of unitation* is a real-valued function on binary strings that depends only on the number of ones in the string. The NEEDLE and BINEEDLE functions are functions of unitation. A *unitation class* of strings is the set of strings with the same number of ones. This is also known as an “error class”.

Population vectors are points in the set $\Lambda = \{x : \sum x_i = 1 \text{ and } x_k \geq 0 \text{ for all } k \in \Omega\}$. This set is the unit simplex in \mathbb{R}^n . The Vose model is given by a function $\mathcal{G} : \Lambda \rightarrow \Lambda$. $\mathcal{G} = \mathcal{M} \circ \mathcal{F}$, where \mathcal{M} describes mixing (crossover and mutation) and \mathcal{F} describes proportional selection. The trajectory of the dynamical system starting at $x \in \Lambda$ is the sequence $x, \mathcal{G}(x), \mathcal{G}^2(x), \dots$

A fixed point of the dynamical system defined by \mathcal{G} is a point x such that $\mathcal{G}(x) = x$. A fixed point x is asymptotically stable if there is a neighborhood of x such that the dynamical system, when started in this neighborhood, will converge to x . It is well known from dynamical system theory that if all of the eigenvalues of the differential $d\mathcal{G}_x$ have modulus less than 1, then x is asymptotically stable. The differential of \mathcal{G} can be calculated by the chain rule from the differential of \mathcal{M} and the differential of \mathcal{F} , and these are calculated in Theorems 6.13, 6.14, and 7.1 of [9].

4 A comparison of models

The Boerlijst differential equations model assumes an infinite population and apparently keeps track of the frequencies of all strings and pairs of strings. As mentioned earlier, it is not clear how they reduce the number of variables, and there may be unstated simplifying assumptions. (In general, the evolution

of a model that includes recombination depends on the full distribution of sequence types within the population.) They do not address stability in their model, although they discuss the stability of fixed points in their results section. Their model is not an evolutionary computation model.

The Boerlijst error-tail model is a very approximate infinite population model. Barnett [1] uses the Moran model to adapt the error-tail model to finite populations.

The Vose model is an evolutionary computation model and is well-developed and documented. The model is exact in the infinite population limit. Since the model keeps track of the frequencies of all strings, no assumption is made in this regard. Thus, the model makes fewer assumptions than any of the other infinite population models. The Vose model allows for the computation of the principal eigenvalue of the differential at the fixed point, thus allowing us to specifically evaluate the stability of fixed points relative to the space of all populations.

With the Vose model, we are limited to relatively short string lengths (8 in this paper). We are also unable to solve the fixed point equations, and so must rely on numerical techniques to find fixed points.

5 Results for the NEEDLE fitness

The Vose model was implemented as a computer program. For efficiency, the Walsh basis was used (chapters 6 and 7 of [9]). The practical limit of our current implementation of this program is string length 8, and this string length was used for all experiments with the Vose model. The needle height parameter a was set to 1 for all runs. When the GA included crossover, uniform crossover with a crossover rate of 1 was used.

Fixed points were found by running the Vose model for enough iterations so that the frequency of the needle string did not change by more than 10^{-6} from one iteration to the next. Except as otherwise noted, the initial population was either a population consisting entirely of the needle string, or a population corresponding to a uniform distribution over all strings.

Figure 1 shows the fixed points for mutation rates from 0.030 to 0.12 when the with-crossover was started in the center of the simplex and at the needle, and the fixed points for the no-crossover GA. All fixed points shown are stable as verified by computing the principal eigenvalue of the differential.

When crossover is included, for a mutation rate of 0.039 and below there is a single stable fixed point near the needle. As the mutation rate increases above 0.039 a fixed point that is close to a random population becomes stable. For mutation rate between 0.40 and 0.49, both fixed points are stable. For mutation rates at 0.50 and above, only the close-to-random fixed point is stable.

When crossover is not included, there is only a single stable fixed point which moves from the needle to the center of the simplex as the mutation rate increases. The Eigen model [5] shows that there is an “error threshold” mutation rate above which the population is close to random. However, this “phase transition” is only sharp in the infinite-string-length limit. In our example, this transition appears to be at a mutation rate of approximately 0.1. The phase transition value predicted by an Eigen infinite population model formula given in Ochoa and Harvey [7] is $1/8$.

The transition at the both critical mutation rates appears to be what is known as a “catastrophe” in mathematical catastrophe theory [4]. For mutation rates just below the lower critical rate, there is no stable fixed point near the random population. As the the mutation increases through the lower critical mutation rate, a stable fixed point and an unstable fixed point come into existence near the center of the simplex (a random

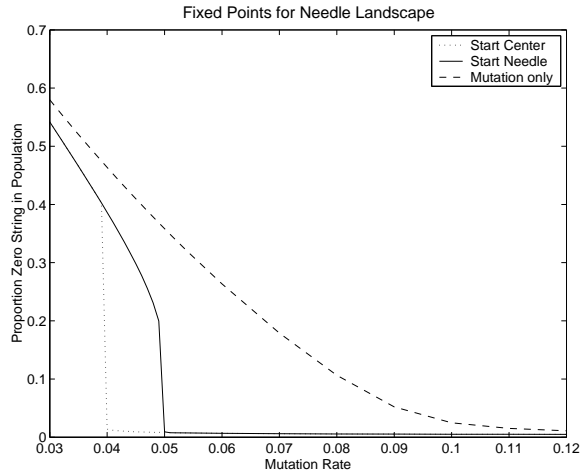


Figure 1: Proportion of the zero string in the fixed-point population of a GA with and without mutation.

population). When the mutation rate increases through the upper critical mutation rate, the unstable fixed point and the stable fixed point near the needle come together and then disappear. Above the upper critical mutation rate, there is only the random population fixed point.

The whole system exhibits what is known as hysteresis. In a system with varying parameters, the state of the system depends on the history of how that state was reached in parameter space. For example, suppose that the system is started with a random population and a mutation rate between the upper and lower mutation rates. It will then converge to the random fixed point. If the mutation is lowered below the lower critical mutation rate, the system will jump to the needle fixed point, and if the mutation is then raised back to its original value, the system will stay on the needle fixed point.

6 The distribution of a population at a fixed point

As pointed out by [3], the distribution of the population at a fixed point is quite different with recombination than with no recombination. For example, the fixed point distributions over the unitation classes for a mutation rate of 0.048 are shown in figure 2. The with-crossover fixed point used is the one close to the needle. The frequency of the all-zeros string is considerably higher when there is no recombination. The unitation classes with 1, 2, 3, and 4 ones are larger under recombination, and the unitation classes with 5, 6, 7, and 8 ones are larger under mutation.

The same pattern appeared for all situations where there was a with-crossover fixed point that corresponded to a without-crossover fixed point. The frequency of the all-zeros string was smaller with recombination, and the frequency of the first unitation class was larger.




Figure 2: Unitation classes for the needle fitness function with a mutation rate of 0.048

7 Bineedle results

The with-crossover GA model on the BINEEDLE fitness, there are again catastrophic transitions at lower and upper critical mutation rates. Below the lower critical mutation rate of approximately 0.033, there are stable fixed points near each needle and there is an unstable fixed point at the center of the simplex. As the mutation increases through the lower critical mutation rate, the fixed point at the center of the simplex becomes stable, and the fixed points near the needle remain stable. There is an upper critical mutation rate between 0.052 and 0.054. When the mutation rate increases through the upper critical mutation rate, the fixed points near the needles become unstable and the only stable fixed point is at the center of the simplex. Figure 3 shows the unitation distribution at the two stable crossover fixed points and the no-crossover fixed point for a mutation rate of 0.048.

For the without-crossover model it is known [8] that there is a unique fixed point for non-zero functions of unitation. The fixed point has an eigenvalue which is very close to 1. This means that the unitation distribution is not very meaningful since there are other distributions which are skewed to one or the other needle and which are almost stable. Individuals in the population interact only through selection, so an individual of a given fitness which is close to one needle is completely equivalent to an individual of the same fitness which is close to the other needle. Thus, there is very little selective pressure to move to one needle or the other. This means that the GA can maintain a population with representatives near both needles.

In contrast, when there is recombination, crossing an individual which is at one needle with an individual which is at the other needle will have a high probability of having low fitness. Thus, if the mutation rate is not too high (less than the lower critical rate), the GA will converge to a fixed point which is close to the needle of higher initial frequency.

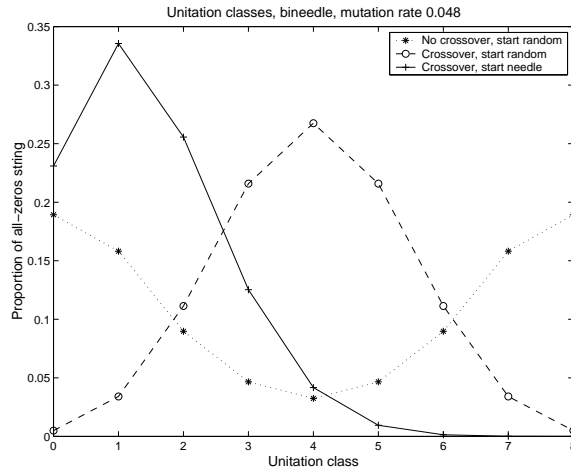


Figure 3: Unitation classes for the bineedle fitness function with a mutation rate of 0.048

8 Examples where crossover causes fitness to decrease

Using the results of the last section it is not difficult to get examples where the presence of crossover causes fitness to decrease.

As a first example, we take a mutation rate of 0.048 and string length 8 with the BINEEDLE fitness. The initial population is in linkage equilibrium with a zero allele probability of 0.66. Figure 4 gives the frequency of the all-zeros string for typical runs for various population sizes. The infinite population model with crossover converged to the random population fixed point which is stable. For population size 500, 99 out of 100 with-crossover runs “converged” to a population corresponding to random population fixed point.

We can construct even an more interesting example by increasing the string length. For string length 20 we initialized the population with a zero allele probability of 0.95 and used a mutation rate of 0.018 and a population size of 500. The initial expected number of copies of the all-zeros string is 172. In 100 runs with crossover, after 100 generations 98 had no copies of an optimum string, one run has 1 copy, and 1 run had 4 copies. In 100 runs without crossover, the average number of copies of an optimum string was 200, and the minimum number of copies was 169.

Thus, the addition of crossover to the GA causes the GA to converge to a stable fixed-point population of minimum or near minimum fitness. These results are similar to those of Rowe [8] where he constructs a fitness function for which a finite population mutation/selection GA usually converges to a population distribution centered around the global fitness minimum.

9 Discussion

Crossover by itself does not change the frequency of the alleles in the population. Geiringer’s theorem ([6], [10], [9]) gives the limiting population of a crossover-only infinite-population GA model on a flat fitness landscape. The limiting population has the same frequency of each allele as the initial population,

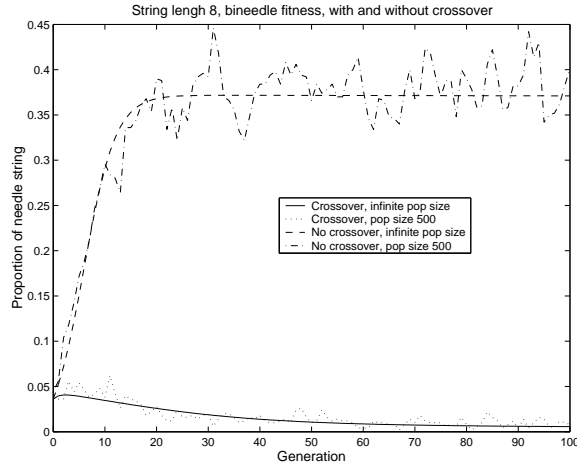


Figure 4: Sample runs for bineedle fitness, mutation rate 0.048, string length 8

and the frequency of a string is the product of the frequency of the alleles. Such a population is said to be in linkage equilibrium or in Robbins’ proportions. When recombination is strong (such as the uniform crossover with crossover rate 1 using in this paper), the infinite-population dynamics of a GA can be approximated by looking at the effect of mutation and selection on the allele frequencies.

Gene pool crossover takes the population to linkage equilibrium in one step. A paper [11] which completely analyzes the fixed points for a gene pool GA on the NEEDLE and BINEEDLE fitness functions is in preparation.

A mutation-only GA infinite-population model converges to the center of the simplex (Theorem 10.8 of [9]). Thus, the effect of mutation is to move the allele frequencies towards $1/2$.

The effect of selection on the allele frequencies for the NEEDLE fitness function is to increase the frequency of the all-zeros string. Thus, it has the effect of increasing the frequency of the zero alleles.

It appears that for the NEEDLE fitness function and for a mutation rate below the lower critical rate, the effect of selection is larger than the effect of mutation, so the GA model converges to a fixed point near the needle. For a mutation rate between the critical values, the effect of mutation is larger near the center of the simplex, and the effect of selection is larger near the needle. For a mutation rate larger than the upper critical value, the effect of mutation is larger than the effects of selection.

10 Conclusion

We have used the Vose dynamical system model to show that the simple GA with uniform crossover and fitness proportionate selection can exhibit bistability on single-peak and double-peak fitness landscapes. Crossover creates catastrophic error threshold transitions as mutation increases. Stable fixed points come into and out of existence at these transitions. The presence of crossover decreases the mutation rate at which the error threshold occurs. In addition, crossover causes the fixed-point population to have fewer copies of the optimum string, more copies of strings close to the optimum string, and less copies of strings far from the optimum string. The Vose model rigorously demonstrates these phenomena.

We have also shown that crossover can cause the GA to move to a population of minimum fitness while the same GA without crossover goes to a population containing many copies of the optimum string.

References

- [1] L. Barnett. Recombination and bistability in finite populations. Technical report, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK, BN1 9QH, 1998. <http://www.cogs.susx.ac.uk/users/lionelb/>.
- [2] L. Barnett. The effects of recombination on a haploid quasispecies evolving on a single peak landscape. Technical report, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK, BN1 9QH, 2000. <http://www.cogs.susx.ac.uk/users/lionelb/>.
- [3] M. C. Boerlijst, S. Bonhoeffer, and M. A. Nowak. Viral quasi-species and recombination. *Proc. Royal Society London B*, 263:1577–1584, 1966.
- [4] D. P. L. Castiglano and S. A. Hayes. *Catastrophe Theory*. Addison Wesley, Reading, Massachusetts, USA, 1993.
- [5] M. Eigen, J. McCaskill, and P. Schuster. Molecular quasi-species. *J. Phys. Chem.*, 92:6881–6891, 1988.
- [6] H. Geiringer. On the probability of linkage in mendelian heredity. *Annals of Mathematical Statistics*, 15:25–57, 1944.
- [7] G. Ochoa and I. Harvey. Recombination and error thresholds in finite populations. In *Foundations of Genetic Algorithms 5*, pages 245–264, San Mateo, 1997. Morgan Kaufmann.
- [8] J. E. Rowe. Population fixed-points for functions of unitation. In Wolfgang Banzhaf and Colin Reeves, editors, *Foundations of genetic algorithms (FOGA-5)*, pages 60–84, San Mateo, 1999. Morgan Kaufmann.
- [9] M. D. Vose. *The Simple Genetic Algorithm: Foundations and Theory*. MIT Press, Cambridge, MA, 1999.
- [10] M. D. Vose and A. H. Wright. The simple genetic algorithm and the walsh transform: Part I, theory. *Evolutionary Computation*, 6(3):253–273, 1998.
- [11] A. H. Wright, J. E. Rowe, R. Poli, and C. R. Stephens. A fixed point analysis of a gene pool ga with mutation. 2001. in preparation.