Selection Intensity in Asynchronous Cellular Evolutionary Algorithms

Mario Giacobini[1], Enrique Alba[2], and Marco Tomassini[1]

[1] Computer Science Institute, University of Lausanne, Lausanne, Switzerland
[2] Department of Computer Science, University of Málaga, Málaga, Spain

Abstract. This paper presents a theoretical study of the selection pressure in asynchronous cellular evolutionary algorithms (cEAs). This work is motivated by the search for a general model for asynchronous update of the individuals in a cellular EA, and by the necessity of better accuracy beyond what existing models of selection intensity can provide. Therefore, we investigate the differences between the expected and actual values of the selection pressure induced by several asynchronous update policies, and formally characterize the update dynamics of each variant of the algorithm. New models for these two issues are proposed, and are shown to be more accurate (lower fit error) than previous ones.

1 Introduction

Cellular evolutionary algorithms (cEAs), also called diffusion or fine-grained models, have been popularized, among others, by early work of Gorges-Schleuter [4] and Manderick and Spiessen [6]. These models are based on a spatially distributed population in which genetic operations may only take place in a small neighborhood of each individual. Usually, individuals are arranged on a regular grid of dimensions $d = 1, 2$ or 3. Cellular EAs are a kind of decentralized EA model. They are not just a parallel implementation of an EA; in fact, although parallelism could be used to speed up the search, we do not address it in this work.

Although fundamental theory is still an open research line for cEAs, they have been empirically reported as being useful in maintaining diversity and promoting slow diffusion of solutions through the grid (exploration). Part of their behavior is due to a lower selection pressure compared to that of panmictic EAs (here *panmictic* means that any one chromosome may mate with any other in the population). The influence of the neighborhood, grid topology, and high efficiency in comparison to other EAs have all been investigated in detail in [2, 5,7,8], and tested on different applications such as combinatorial and numerical optimization.

Cellular EAs can be seen as stochastic cellular automata (CAs) [11,13] where the cardinality of the set of states is equal to the number of points in the search space. CAs, as well as cEAs, usually assume a *synchronous* or "parallel" update policy (reproduction at a time), in which all the cells are formally updated

E. Cantú-Paz et al. (Eds.): GECCO 2003, LNCS 2723, pp. 955–966, 2003.

simultaneously. However, this is not the only option available. Indeed, several works on *asynchronous* CAs have shown that sequential update policies have a marked effect on their dynamics (see e.g. [9,10]). Thus, it would be interesting to investigate asynchronous cEAs and their problem solving capabilities. A first step in that direction was made in [1], where a set of standard problems were studied under several asynchronous update policies in a 2-d cGA environment. The main observation was that, although asynchronous update is not always the best choice in terms of solution quality, it is numerically faster, and the speed of convergence can be varied by changing the updating scheme. Thus, since convergence and diversity in EAs are related to selection, we would like to get a better understanding of the behavior of selection in asynchronous cEAs as compared to synchronous update and to the panmictic case. We present here an extension of the works [2,5,7,8] on selection pressure to asynchronous cEAs. For reasons of space, we limit ourselves to the two-dimensional grid case, the most common in practice.

The paper is organized as follows. The next section contains some background on asynchronous cEAs. Section 3 describes the results of our experiments on selection pressure in asynchronous cEAs. Section 4 analyzes the current logistic model of selection pressure and presents an improved characterization of asynchronous algorithms, leading to a new model proposal. Finally, section 5 offers our conclusions, as well as some comments on future work.

2 Asynchronous cEAs

Updating a cell (individual) in a cellular EA means selecting two parents in the individual's neighborhood (including the individual itself), applying genetic operators to them, and finally replacing the individual with the best offspring. In a conventional synchronous cEA, all the individuals in the grid are updated simultaneously. This step makes up a generation, and the process is repeated until a termination condition is reached.

There exist many ways for sequentially updating the cells of a 2-d cEA . Here we employ *step-driven* updates and ignore the so-called *time-driven* methods, in which (real) time is explicit. Time-driven methods are more realistic for physical simulation but are not needed in the EA case (an excellent discussion of asynchronous update in CAs is available in [9]). The most general update scheme is independent random ordering of updates in time, which consists of randomly choosing the cell to be updated next with replacement. This corresponds to a binomial distribution for the update probability. This update policy will be called *uniform choice* (UC) in the following and it is similar to the time-driven Poisson update in the limit of large n, n being the population size.

In our study we also consider three other update methods: *fixed line sweep*, *fixed random sweep*, and *new random sweep* (we employ the same terminology as in [9]).

- In *fixed line sweep* (LS), the simplest method, the n grid cells are updated sequentially $(1, 2 \ldots n)$, line by line of the 2-d grid.

- In the *fixed random sweep* update (FRS), the next cell to be updated is chosen with uniform probability without replacement; this will produce a certain update sequence $(c_1^j, c_2^k, \ldots, c_n^m)$, where c_q^p means that cell number p is updated at time q and $(j, k, \ldots, m)$ is a permutation of the n cells. The same permutation is then used for all update cycles.
- The *new random sweep* method (NRS) works like FRS, except that a new random cell permutation is chosen anew for each sweep through the array.

A *time step* is defined as updating n times sequentially, which corresponds to updating *all* the n cells in the grid for LS, FRS and NRS, and possibly less than n different cells in the uniform choice method, since some cells might be updated more than once. It should be noted that, with the exception of fixed line sweep, the other asynchronous updating policies are stochastic, representing an additional source of non-determinism besides that of the genetic operators.

3 Takeover Times

In order to study the induced selection pressure by itself (without introducing the perturbing effect of recombination or mutation operators) a standard technique is to let selection be the only active operator, and then monitor the growth rate of the best individual in the initial population [3]. The takeover time is the time it takes for the single best individual to conquer the whole population. A shorter takeover time thus means a higher selection pressure. It has been shown that when we move from a panmictic population to a spatially structured one of the same size with synchronous updating of the cells, the global selection pressure induced on the entire population is qualitatively similar but weaker (Sarma and De Jong [7]).

Three standard selection algorithms were used in [7], namely fitness proportionate, linear ranking, and binary tournament. The cellular EA structure was a two-dimensional toroidal grid of size 32 × 32 with three different neighborhood shapes with 5, 9 and 13 neighbors respectively, which are the most common in practice. In the spatially distributed case it was observed that, for all three mentioned neighborhoods, the global selection pressure induced by fitness-proportionate selection was smaller than the pressures induced by linear ranking and binary tournament, with binary tournament being roughly equivalent to ranking as the neighborhood size increases, as it can be inferred from well-known existing theoretical considerations on selection pressure.

Sarma and De Jong [8] performed a more detailed empirical analysis of the effects of the neighborhood's size and shape on the local selection algorithms. They were able to show that propagation times are closely related to the neighborhood size, with larger neighborhoods giving rise to stronger selection pressures.

In the following, we report results (for three different neighborhoods) on the selection pressure for three selection methods in the case of asynchronous update. The neighborhoods used are Von Neumann (5 neighbors along the NWSE directions and the center cell, also called Linear 5), Moore (9 neighbors, including the central cell and its eight nearest neighbors, also called Compact 9), and

Compact 13 (like Moore, with four neighbors added along the N,W,S, and E directions -like a diamond-).

For the sake of comparison, we also include the curves corresponding to the panmictic case and to the synchronously updated grid. Since the results with the asynchronous Fixed Random Sweep are very similar to those using the New Random Sweep policy, only the latter curves are reported. Each of these results is the average of 100 independent runs.

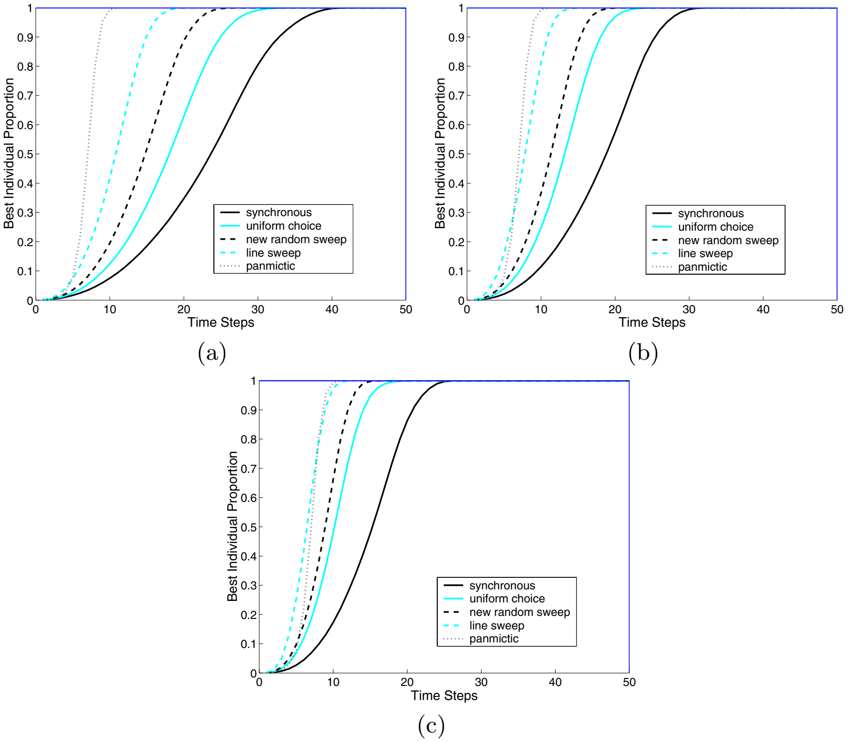

(a) (b) (c)

Fig. 1. Takeover times with rank selection. Linear 5 neighborhood (a); Compact 9 neighborhood (b); Compact 13 neighborhood (c). Mean values over 100 runs. The vertical axis represents the proportion of population consisting of best individual as a function of the time step.

Figure 1 shows the mean growth curves of a cEA using rank selection in a Linear 5, Compact 9 and Compact 13 neighborhood, respectively. Figure 2 depicts graphs of mean growth curves for a cEA using binary tournament selection respectively in the Linear 5, Compact 9 and Compact 13 neighborhood. The mean takeover times results for a cEA using a roulette wheel selection respectively in a Linear 5, Compact 9, and Compact 13 neighborhood are reported in Figure 3.

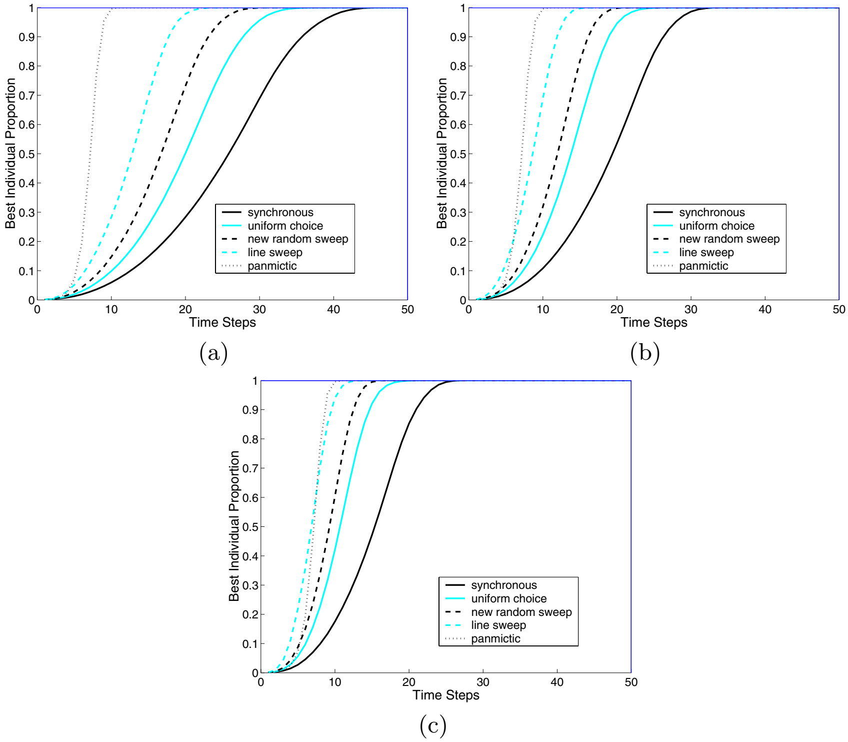

Fig. 2. Takeover times with binary tournament selection. Linear 5 neighborhood (a); Compact 9 neighborhood (b); Compact 13 neighborhood (c). Mean values over 100 runs. The vertical axis represents the proportion of population consisting of copies of the best individual as a function of the time step.

These results largely confirm the findings of Sarma and De Jong as far as synchronous and panmictic cEAs are concerned. Indeed, binary tournament and ranking induce very similar global selection pressure, while proportional selection exhibits less pressure. Moreover, for a given selection policy, larger neighborhoods induce a stronger selection intensity.

What is new in this paper (our contribution) is the behavior of the asynchronous models. Generally speaking, it can be observed that the asynchronous models give an emergent selection pressure that is between the panmictic upper bound and the synchronous lower bound. All graphs show that the global selection intensity grows going from uniform choice update to line sweep, with new random sweep and fixed random sweep in between, although an analysis of the variances should be conducted to quantitatively confirm the trend. This suggests that, by choosing the appropriate asynchronous update policy, one is able to control the selection pressure without using ad hoc numerical parameters. This opens new possibilities for dynamical EAs in which the selection pressure is under the control of the modeler even during the run (work in progress).

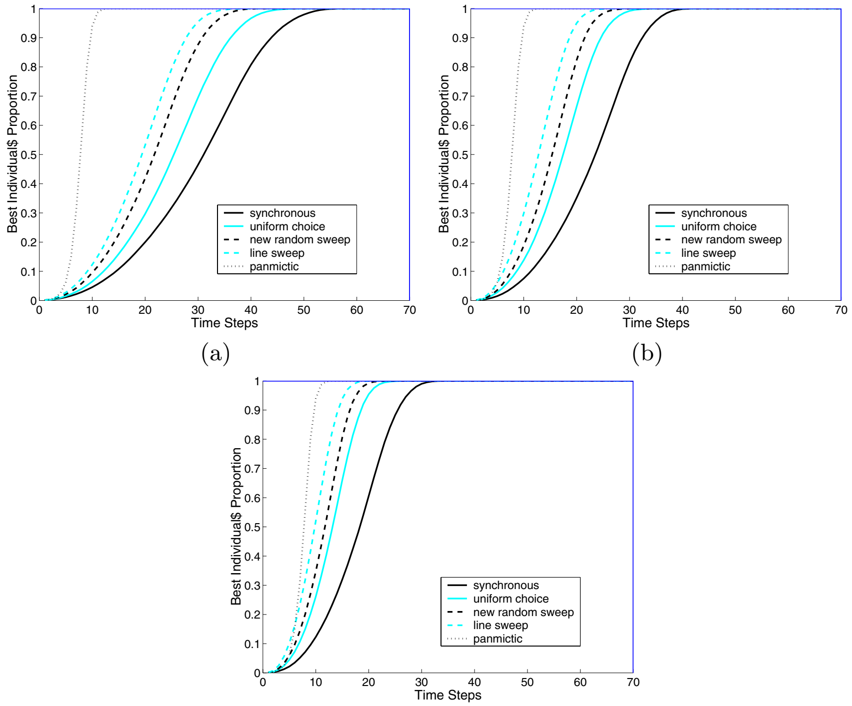

(a) (b)

(c)

Fig. 3. Takeover times with fitness proportional selection. Linear 5 neighborhood (a); Compact 9 neighborhood (b); Compact 13 neighborhood (c). Mean values over 100 runs. The vertical axis represents the proportion of population consisting of copies of the best individual as a function of the time step. Note the change of scale on the horizontal axis to avoid cutting off the curves in figure (a).

The impression is confirmed by Table 1, where mean takeover times of all the update methods for each selection mechanism and each of the three considered neighborhoods are reported with their standard deviations.

4 Modelling Individual Growth

In this section, quantitative models for the individual growth (and thus for the different global selection pressures induced in the population) are presented using asynchronous update. We first give some statistical results valid for all finite 2-d cEAs discrete lattices. Next, we offer a quantitative analysis of the takeover time, and finally we hint at some possible improvements in the existing logistic model.

Table 1. Mean takeover time for the three selection mechanisms (vertically) and the five update methods (horizontally). Upper part: Linear 5 neighborhood. Middle: Compact 9 neighborhood. Lower part: Compact 13 neighborhood. The last column refers to the classical panmictic case. Standard deviations in parentheses.

LINEAR 5	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	52 (3.7)	34 (2.8)	37 (3.0)	37 (3.0)	44 (3.7)	12 (1.0)
Tournament	42 (2.7)	21 (1.9)	26 (2.0)	28 (1.9)	33 (3.7)	10 (0.7)
Ranking	39 (2.1)	18 (1.6)	24 (1.5)	24 (1.7)	30 (3.2)	10 (0.8)

COMPACT 9	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	38 (2.6)	23 (2.5)	26 (1.8)	26 (2.2)	29 (4.2)	12 (1.0)
Tournament	31 (1.8)	15 (1.4)	19 (1.4)	19 (1.5)	23 (2.9)	10 (0.7)
Ranking	30 (1.7)	13 (1.4)	18 (1.3)	18 (1.4)	22 (2.7)	10 (0.8)

COMPACT 13	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	31 (1.7)	18 (1.9)	20 (1.9)	20 (1.8)	23 (3.3)	12 (1.0)
Tournament	25 (1.4)	12 (1.2)	15 (1.0)	15 (1.2)	18 (2.9)	10 (0.7)
Ranking	25 (1.2)	11 (1.1)	14 (1.1)	15 (1.0)	18 (2.5)	10 (0.8)

4.1 Statistical Results on Information Propagation

Schönfish and de Roos [9] derived the expected value $E(Z)$ and the variance $V(Z)$ of the number of single steps between an update of a cell x and the next update of a cell $y \neq x$ in the neighborhood of x, $U(x)$, in order to compare the different asynchronous CA updating methods. These results can also be applied to asynchronous cEAs, where the local transition function f_0 does not deterministically determine the next state of the cell, but describes a probabilistic rule for such an update [11,13]. This rule is generally determined by the different selection, crossover and mutation mechanisms used in the EA. In our study of the takeover times induced by different update methods the local function f_0 only depends on the selection mechanism used.

Table 2. Values of the expected value $E(Z)$ and of the variance $V(Z)$ of the number of time steps between an update of a cell x and the next update of a cell $y \neq x$ in the neighborhood $U(x)$ for a cEA on a square grid. For Asynchronous Line Sweep, which depends from the chosen neighborhood, the Linear 5 neighborhood result is shown.

	Synchro	LS	FRS	NRS	UC
E(Z)	1	$\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{12}\left(7+\frac{1}{n}\right)$	1
V(Z)	0	$\frac{1}{2}\left(\frac{1}{n}-\sqrt{n}+\frac{1}{2}n\right)$	$\frac{1}{12}(n-2)$	$\frac{1}{12}\left(\frac{23}{12}n-\frac{13}{6}-\frac{13}{12n}\right)$	n-1

Schönfish and de Roos used the number of cell updates in their statistics. In order to be able to extend their results to the synchronous cEA, their statistics need to be translated using the number of time steps between an update of a cell x and the next update of a cell $y \neq x$ in the neighborhood $U(x)$ of x. In fact, for synchronous cEAs it is not possible to count the single cells updating, but a generational synchronous step can be compared to a time step of an asynchronous cEA. Table 2 contains the values of $E(Z)$ and $V(Z)$ for the four asynchronous and the synchronous updating policies. If in a cEA we keep the selection mechanism fixed and we vary the updating method, the results of Table 2 explain the ranking of the observed takeover times; notice that some experimental values are slightly different while theoretically they should be identical on the average. This is the case for Line Sweep and Fixed Random Sweep, as well as for Uniform Choice and Synchronous. However, these results do not explain the actual shapes of the selection pressure curves. We are currently working on the analytical study of the curves and some preliminary results are reported in the next two sections.

4.2 Fitting the Selection Pressure Curves

Sarma and De Jong [7] proposed a simple quantitative model for the study of the selection pressure curves for cEAs. They assumed that the diffusion of the best individual in the artificial evolution of a structured population would follow a logistic curve. Let us analyze their result, shown in Equation 1.

$$P_b(t) = \frac{1}{1 + \left(\frac{1}{P_b(0)} - 1\right) e^{-\alpha t}} \tag{1}$$

This equation, where $P_b(t)$ represents the proportion of the best individual in the population at time t, was proposed for synchronous cEAs, and therefore we wondered whether it holds for asynchronous ones. Consequently, we proceeded to analyze the error (mean squared error) between an actual average selection pressure and the theoretically predicted values, for all the update methods considered in this work. The steps were (1) to compute the theoretical value of α, (2) to generate the predicted curve by using one point of the average observed performance curve, and (3) to compare it against the whole set of points of this observed curve.

To derive the α parameter we selected a mid point (with $P_b(t)$ around 0.5) from the experimental curves. Then, we generated the corresponding curve, and computed the squared error. Table 3 shows our measurements.

This table also shows that, although the fitting is satisfactory, it is not that good, since there exists a gap between the fitted curves and the experimental points. This claim is confirmed by Figure 4, in which the panmictic case is particularly good, while the other fittings could clearly be improved. This lead us to think that there could exist a better fitting for cellular EAs than the logistic one, whose main advantage is its similitude to the theoretical results existing for panmictic algorithms [3].

Table 3. Mean squared error between predicted and actual logistic fittings for Linear 5, Compact 9, and Compact 13 neighborhoods for all the update modes.

LINEAR 5	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	0.00406	0.00358	0.00372	0.00427	0.00358	0.00069
Tournament	0.00309	0.00270	0.00281	0.00274	0.00248	0.01053
Ranking	0.00366	0.00282	0.00257	0.00234	0.00290	0.00181

COMPACT 9	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	0.00349	0.00376	0.00284	0.00280	0.00332	0.00069
Tournament	0.00287	0.00194	0.00193	0.00178	0.00202	0.01053
Ranking	0.00311	0.00190	0.00197	0.00184	0.00209	0.00181

COMPACT 13	Synchro	LS	FRS	NRS	UC	Panmictic
Roulette	0.00339	0.00328	0.00199	0.00242	0.00281	0.00069
Tournament	0.00231	0.00148	0.00142	0.00149	0.00145	0.01053
Ranking	0.00225	0.00166	0.00153	0.00150	0.00122	0.00181

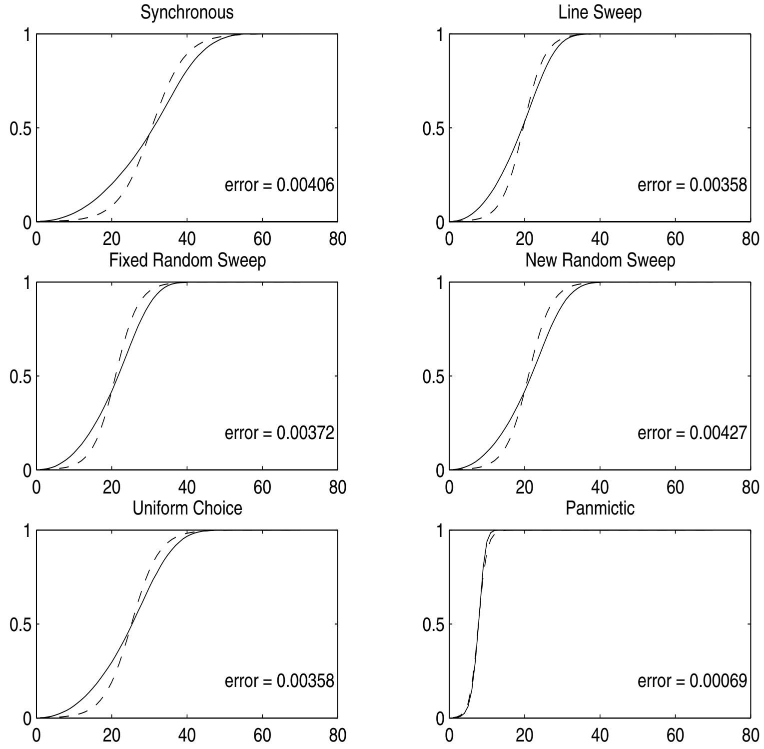

Fig. 4. Fitting of the experimental takeover time curves (full) with the logistic model (dashed) for the various update modes. Results refer to fitness proportional selection with Linear 5 neighborhood.

We have shown in this section that the logistic fitting should be improved for decentralized algorithms. In the next section some steps toward a more accurate model are described.

4.3 An Improved Model

It is well known since the work of Verhulst [12], that the assumption of logistic growth is true for biological populations within bounded resources. It is easy to see that this behavior also holds for the best individual growth in the artificial evolution of a finite panmictic population [3]. In fact, if we consider a population of size n, the number $N(t)$ of copies of the best individual in the population at time step t is given by the following recurrence:

$$\begin{cases} N(0) = 1 \\ N(t) = N(t-1) + p_s N(t-1)(n - N(t-1)) \end{cases} \tag{2}$$

where p_s is the probability that the best individual is chosen. This recurrence can be easily transformed into one that describes a logistic population growth in discrete time:

$$\begin{cases} N(0) = 1 \\ N(t) = N(t-1) + (p_s n) N(t-1) \left(1 - \frac{1}{n} N(t-1)\right) \end{cases} \tag{3}$$

Such a recurrence can be expressed in analytical form by the logistic equation:

$$N(t) = \frac{n}{1 + \left(\frac{n}{N(0)} - 1\right) e^{-\alpha t}} \tag{4}$$

where the growth coefficient α depends on the probability p_s and the population size n. This happens to be the approach taken in [8] for synchronous cEAs.

As suggested by Gorges-Schleuter in [5], in the artificial evolution of locally interacting, spatially structured populations, the assumption of a logistic growth doesn't hold anymore. In fact, in the case of a ring or a torus structure we have respectively a linear and a quadratic growth. We complete here her analysis which holds for unrestricted growth, extending it to bounded synchronously updated spatial populations.

For a structured population, let us consider the limiting case, which represents an upper bound on growth rate, in which the selection mechanism is deterministic, and a cell always chooses its best neighbor for updating. If we consider a population of size n with a ring structure (like that of 1-d cellular automata in which the two cells on the borders are linked) and a neighborhood radius of k (i.e. a neighborhood of a cell contains $2k + 1$ cells), the following recurrence describes the growth of the number of copies of the best individual:

$$\begin{cases} N(0) = 1 \\ N(t) = N(t-1) + 2k \end{cases} \tag{5}$$

This recurrence can be described by the closed equation $N(t) = N(0) + 2kt$, that clearly shows the linear character of the growth rate.

In the case of a population of size n disposed on a toroidal grid of size $\sqrt{n}\times\sqrt{n}$ (assuming $\sqrt{n}$ odd) and the Linear 5 neighborhood structure, the number of copies of the best individual can be described by the following recurrence:

$$\begin{cases} N(0) = 1 & \\ N(t) = N(t-1) + 4t & , \quad for\ 0 \leq t \leq \frac{\sqrt{n}-1}{2} \\ N(t) = N(t-1) + 4(\sqrt{n} - t) \ , & \quad for\ t \geq \frac{\sqrt{n}-1}{2} \end{cases} \tag{6}$$

This growth is described by a convex quadratic equation followed by a concave one, as the two closed forms of the recurrence clearly show:

$$\begin{cases} N(t) = 2t^2 + 2t + 1 & , \quad for\ 0 \leq t \leq \frac{\sqrt{n}-1}{2} \\ N(t) = -2t^2 + 2(2\sqrt{n} - 1)t + 2\sqrt{n} - n \ , & \quad for\ t \geq \frac{\sqrt{n}-1}{2} \end{cases} \tag{7}$$

Thus, a more accurate fitting should take into account the non-exponential growth followed by saturation (crowding effect).

5 Conclusions

In this work we have presented different update policies to deal with cellular EAs in a search for a more efficient algorithm with respect to its canonical form. We have investigated the induced selection pressure of such policies with respect to the widespread synchronous update and panmictic algorithms. Our results indicate that these two algorithms represent the smaller and higher (respectively) bounds to selection intensity, and that the asynchronous update methods represent intermediate values of pressure (which can even be ranked for such methods).

We have studied the existing proposals dealing with a logistic fitting that, although very similar to that existing for panmictic EAs and globally valid, are susceptible to further improvement for cellular EAs. We provide such an improved model of the best individual growth and additionally characterize the expected time between the update of two individuals residing in the same neighborhood for all the update methods considered in the paper.

Our aim in this paper has been to advance in the study of selection pressure in cellular EAs. Future work will consider more suitable functions for fitting the experimental curves (maybe a quadratic one). We will consider as well further extensions of the results to other aspects influencing the selection pressure of cellular EAs, such as the relationship existing between the neighborhood and the grid topology. Also, extensions to dimensions larger than two are being considered, as well as the application of cellular EAs to several hard problems. Our global aim is to obtain a body of knowledge of these algorithms, especially with regard to the numerical efficiency of the search.

Acknowledgments. The authors would like to acknowledge useful discussions on the models with A. Tettamanzi. The second author wishes to thank founding from the Spanish "Ministerio de Ciencia y Tecnología" and FEDER through contract TIC2002-04498-C05-02 (the TRACER project).

References

1. E. Alba, M. Giacobini, and M. Tomassini. Comparing synchronous and asynchronous cellular genetic algorithms. In J. J. Merelo et al., editor, *Parallel Problem Solving from Nature - PPSN VII*, volume 2439 of *Lecture Notes in Computer Science*, pages 601–610. Springer-Verlag, Heidelberg, 2002.
2. E. Alba and J. M. Troya. Cellular evolutionary algorithms: Evaluating the influence of ratio. In M. Schoenauer et al., editor, *Parallel Problem Solving from Nature, PPSN VI*, volume 1917 of *Lecture Notes in Computer Science*, pages 29–38. Springer-Verlag, 2000.
3. D. E. Goldberg and K. Deb. A comparative analysis of selection schemes used in genetic algorithms. In G. J. E. Rawlins, editor, *Foundations of Genetic Algorithms*, pages 69–93. Morgan Kaufmann, 1991.
4. M. Gorges-Schleuter. ASPARAGOS an asynchronous parallel genetic optimisation strategy. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 422–427. Morgan Kaufmann, 1989.
5. M. Gorges-Schleuter. An analysis of local selection in evolution strategies. In *Genetic and evolutionary conference, GECCO99*, volume 1, pages 847–854. Morgan Kaufmann, San Francisco, CA, 1999.
6. B. Manderick and P. Spiessens. Fine-grained parallel genetic algorithms. In J. D. Schaffer, editor, *Proceedings of the Third International Conference on Genetic Algorithms*, pages 428–433. Morgan Kaufmann, 1989.
7. J. Sarma and K. A. De Jong. An analysis of the effect of the neighborhood size and shape on local selection algorithms. In H. M. Voigt, W. Ebeling, I. Rechenberg, and H. P. Schwefel, editors, *Parallel Problem Solving from Nature (PPSN IV)*, volume 1141 of *Lecture Notes in Computer Science*, pages 236–244. Springer-Verlag, Heidelberg, 1996.
8. J. Sarma and K. A. De Jong. An analysis of local selection algorithms in a spatially structured evolutionary algorithm. In T. Bäck, editor, *Proceedings of the Seventh International Conference on Genetic Algorithms*, pages 181–186. Morgan Kaufmann, 1997.
9. B. Schönfisch and A. de Roos. Synchronous and asynchronous updating in cellular automata. *BioSystems*, 51:123–143, 1999.
10. M. Sipper, M. Tomassini, and M. S. Capcarrere. Evolving asynchronous and scalable non-uniform cellular automata. In *Proceedings of International Conference on Artificial Neural Networks and Genetic Algorithms (ICANNGA97)*, pages 67–71. Springer-Verlag KG, Vienna, 1998.
11. M. Tomassini. The parallel genetic cellular automata: Application to global function optimization. In R. F. Albrecht, C. R. Reeves, and N. C. Steele, editors, *Proceedings of the International Conference on Artificial Neural Networks and Genetic Algorithms*, pages 385–391. Springer-Verlag, 1993.
12. P. F. Verhulst. *Mem. Acad. Roy. Bruxelles*, 28(1), 1844.
13. D. Whitley. Cellular genetic algorithms. In S. Forrest, editor, *Proceedings of the Fifth International Conference on Genetic Algorithms*, page 658. Morgan Kaufmann Publishers, San Mateo, California, 1993.