

Takeover Time Curves in Random and Small-World Structured Populations

Mario Giacobini
Information Systems
Department
University of Lausanne
Lausanne, Switzerland
mario.giacobini@unil.ch

Marco Tomassini
Information Systems
Department
University of Lausanne
Lausanne, Switzerland
marco.tomassini@unil.ch

Andrea Tettamanzi
Information Technologies
Department
University of Milano
Milano, Italy
andrea.tettamanzi@unimi.it

ABSTRACT

We present discrete stochastic mathematical models for the growth curves of synchronous and asynchronous evolutionary algorithms with populations structured according to a random graph. We show that, to a good approximation, randomly structured and panmictic populations have the same growth behavior. Furthermore, we show that global selection intensity depends on the update policy. The validity of the models is confirmed by a comparison with experimental results of simulations. We also present experimental results on small-world and scale-free population graph topologies. We show that they lead to qualitatively similar results. However, the different nature of the nodes can be exploited to obtain a more varied evolutionary behavior.

Categories and Subject Descriptors

I.2 [Artificial Intelligence]: Problem Solving, Control Methods, and Search

General Terms

Algorithms

Keywords

Evolutionary algorithms, Takeover times, Structured populations, Small-World networks

1. INTRODUCTION

Spatially structured populations have been proposed in evolutionary computation as a means for improving the search properties of EAs. In fact, it has often been experimentally shown that topology affects population dynamics and that spatially structured populations have distinctive advantages over panmictic (mixing) ones (two recent reviews of the field

are [1, 4]). Although many topologies are in principle possible, the simplest and those that are best known both experimentally and theoretically are the so-called *island* or *multipopulation* models, and the *cellular* models. Systematic studies, both experimental and theoretical, of the cellular model of EAs (cEAs) have been published in the last few years. In particular, models for the selection intensity in regular one-dimensional and two-dimensional grids have been proposed in [8, 7, 9]. These models correctly predict the experimental growth curves in synchronous and asynchronous cellular EAs. Being able to understand, and possibly control, selection intensity is important since selection strength is related to the exploration vs. exploitation trade-off of the algorithm.

In this paper we study irregular topologies for the cellular model, in order to understand the behavior of the global selection intensity with respect to the known bounding panmictic and regular lattice cases. The most general irregular topology is the random graph, which is why we concentrate on this kind of structure in the first part of the paper. While the random graph is a well-studied structure in mathematics, biology, and the social sciences, it has not, to our knowledge, been investigated in cEAs. For example, models of infection transmission over a population of individuals with random links among them have been known for years [12]. Likewise, information transmission in society has been sometimes modeled using a random graph structure [12]. Recently, families of graphs, collectively called *small world* that are not regular nor completely random have attracted a lot of interest [12] because of their striking topological and dynamical properties. In the second part of the paper we study selection intensity on populations structured according to small-world topologies for the first time, as far as we know.

Most propagation models are based on differential equations. Here we propose instead a discrete model, which seems to us more suitable for finite evolving spatial populations.

After a brief introduction to cEAs, we describe their synchronous and asynchronous variants and the concept of takeover time. Next, we present our models, and specialize them for the random graph cases, comparing them with simulation data. Then small-world population topologies are briefly introduced and experiments are conducted to investigate the behavior of the global selection pressure on these structures. Finally, we present our conclusions and ideas for future work.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

GECCO'05, June 25–29, 2005, Washington, DC, USA.
Copyright 2005 ACM 1-59593-010-8/05/0006 ...\$5.00.

2. TAKEOVER TIMES IN CEAS

Synchronous and Asynchronous cEAs. In the cellular model individuals are structured according to a given spatial topology, e.g. a ring or a grid, and each individual only interacts with a limited number of surrounding individuals called its *neighborhood*.

Updating a cell (individual) in a cEA 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. The replacement could also be probabilistic. 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 potentially many ways for asynchronously updating the cells of a cEA (see, for instance, [16]). The most general update scheme is independent random ordering in time, which consists of randomly choosing the cell to be updated next with replacement. This update policy will be called *uniform choice* (UC) in the following. We also consider another update method, *new random sweep* (NRS), in which the next cell to be updated is chosen with uniform probability without replacement; this will produce a certain update sequence $(c_1^j, c_2^k, \dots, c_n^m)$, where c_q^p means that cell number p is updated at time q and (j, k, \dots, m) is a permutation of the n cells. A new random cell permutation is generated 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 the synchronous and the asynchronous NRS, and possibly less than n different cells in the asynchronous UC method, since some cells might be updated more than once during a single time step.

Takeover Time. In order to study the induced selection pressure 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 an initial copy of the best individual over time [10, 5]. 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. The shape of the *growth curve*, representing the percentage of the best individual in the population with respect to time is also of interest. It has been shown that going from a panmictic population to one that is spatially structured as a lattice of the same size, the global selection pressure induced on the entire population is qualitatively similar but weaker [14]. Sarma and De Jong [15] 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. Gorges-Schleuter studied growth curves for local Evolution Strategies in [11] with similar results for populations structured as a ring or a torus. Quantitative models for the growth curves for the ring and the torus, both for synchronous as well as asynchronous updates, have recently been presented in [13, 8, 7]. In the following sections we present general models for the propagation of the best individual and then we examine the random graph case in detail. Small-world graphs will be dealt with in section 6.

3. DIFFERENCE EQUATION MODELS

Let us consider the random variables $V_i(t) \in \{0, 1\}$ indicating the presence in cell i ($1 \leq i \leq n$) of a copy of the best individual ($V_i(t) = 1$) or of a worse one ($V_i(t) = 0$) at time step t , where n is the the population size. The random variable

$$N(t) = \sum_{i=1}^n V_i(t)$$

denotes the number of copies of the best individual in the population at time step t . Initially $V_i(1) = 1$ for some individual i , and $V_j(1) = 0$ for all $j \neq i$.

Following Rudolph's definition [13], if the selection mechanism is non-extinctive, the expectation $E[T]$ with $T = \min\{t \geq 1 : N(t) = n\}$ is called the takeover time of the selection method. In the case of spatially structured populations the quantity $E_i[T]$, denoting the takeover time if cell i contains the best individual at time step 1, is termed the takeover time with initial cell i . Assuming a uniformly distributed emergence of the best individual among all cells, the takeover time is therefore given by

$$E[T] = \frac{1}{n} \sum_{i=1}^n E_i[T].$$

In the following sections we give the recurrences describing the growth of the random variable $N(t)$ in cEA evolving populations structured on graphs for the synchronous and the two asynchronous update policies described in Section 2. In general, such recurrences take the common form

$$E[N(t)] = \sum_{i=1}^n P[N(t-1) = i] (i + \Delta_i(t-1)), \quad (1)$$

where $\Delta_i(t-1) = N(t) - N(t-1)$, given $N(t-1) = i$, is a random variable as well. This random variable will depend on the update method. Let's suppose that at time step $t-1$ there are i copies of the best individuals. At the next time step each of the other $n-i$ individuals will contain a copy of the best with a probability depending on its number of neighbors, the number of its neighbors containing a copy of the best and the selection operator. The first two conditions can be seen as random variables, both depending on the topology of the population.

Therefore, in Equation 1 we have

$$\Delta_i(t-1) = \sum_{r=1}^{n-i} \sum_{j=1}^{n-1} P[K = j] \sum_{l=0}^j P[B_j = l] p_{\text{sel}}(j, l), \quad (2)$$

where K denotes the number of neighbors of an individual, B_j the number of copies of the best individual in a neighborhood of j individuals, and $p_{\text{sel}}(j, l)$ is the probability of selecting a copy of the best individual among the l best of j neighbors.

4. THE RANDOM GRAPH TOPOLOGY

A random graph with n vertices can be constructed by taking all possible pairs of vertices and connecting each pair with probability q (and thus not connecting it with probability $1-q$) [12]. In the general case of a cEA evolving a population structured as a random graph with probability $0 < q < 1$ of having an edge between any pair of vertices, the

random variables K and B_j of equation 2 have the following probability functions:

$$\begin{aligned} P[K = j] &= q^j (1 - q)^{n-1-j}, \\ P[B_j = l] &= \begin{cases} 1 & \text{if } l = \frac{j}{n-1}, \\ 0 & \text{otherwise,} \end{cases} \end{aligned}$$

since any of the j neighbors of an individual has a probability $i/(n-1)$ of containing a copy of the best individual.

Cellular evolutionary algorithms are good candidates for using selection methods that are easily extensible to small local pools such as ranking and tournament. The equations for individual growth in the case of those customary local selection policies are mathematically rather complicated, involving higher moments of the distribution. We therefore use a simplified selection policy, called Uniform Selection, which gives rise to a useful and interesting model.

This selection mechanism randomly selects an individual in the selection pool (i.e. the neighborhood of a given individual). The selected individual then replaces the considered individual if it has a better fitness. Such operator is similar to the *local parent selection* introduced by Gorges-Schleuter in [11], except that a $(\mu + \lambda/\mu, \nu)$ -LES is used instead of a $(\mu, \lambda/\mu, \nu)$ -LES, with $\kappa = \infty$ (κ being the upper limit for the life span) and $\rho = 1$ (ρ being the number of selected ancestors).

In a cEA whose population is structured as a random graph, since the number of edges incident on a given edge is a binomial random variable, we can use the *mean field hypothesis*, which consists in taking for all vertices the average number of neighbors $q(n-1)$. In this way, any vertex “sees” the same isotropic average environment. Under this hypothesis, the expected number of copies of the best individual in a neighborhood not containing a copy of the best at time step t is $E[B_{q(n-1)}] = qN(t)$. We will see that this approximation is good unless the probability q is very low. In the case of uniform selection, the probability that a copy of the best is selected at time step t is

$$p_{\text{sel}}^{\text{rnd}}(q(n-1), qN(t)) = \frac{qN(t)}{q(n-1)} = \frac{N(t)}{n-1}. \quad (3)$$

This approximation, valid when the mean field hypothesis can be used, gives a probability equal to the panmictic case, where the graph describing the topology of a population is a complete graph. In fact, for this structure, each individual has exactly $K = n-1$ neighbors, and the number of copies of the best individual in a neighborhood not containing a copy of the best at time step t is $B_{n-1} = N(t)$. In the case of uniform selection, the probability that a copy of the best is selected at time step t is

$$p_{\text{sel}}^{\text{pan}}(n-1, N(t)) = \frac{N(t)}{n-1}. \quad (4)$$

Therefore, the selection pressure for a randomly structured population is similar to that of a panmictic one, when the mean field hypothesis can be applied.

For synchronous update, using equations 1 and 4, the recurrence for $N(t)$ can be written as

$$\begin{cases} N(0) = 1 \\ E[N(t)] = E[N(t-1)] + (n - E[N(t-1)]) \frac{E[N(t-1)]}{n}, \end{cases}$$

which is a typical form of a discrete logistic recurrence.

In the asynchronous new random sweep case, the $n - E[N(t-1)]$ individuals not containing a copy of the best

in the previous time step will contain a copy of the best following the recurrence

$$\begin{cases} E[M(1)] = E[N(t-1)] \left(1 + \frac{1}{n}\right) \\ E[M(\tau)] = E[N(t-1)] \left(1 + \frac{1}{n}\right)^\tau, \end{cases}$$

where $M(\tau)$ is the probability that individual at position τ in the sweep, among those not containing the best, will be taken over. Therefore, the growth of $N(t)$ can be described by the following recurrence:

$$\begin{cases} N(0) = 1 \\ E[N(t)] = E[N(t-1)] + \left(1 + \frac{1}{n}\right)^{n-E[N(t-1)]}. \end{cases}$$

When employing an asynchronous uniform choice update policy, the growth of random variable $N(u)$ in terms of single update step u can be described by the recurrence:

$$\begin{cases} N(0) = 1 \\ E[N(u)] = E[N(u-1)] + \frac{n-E[N(u-1)]}{n} \frac{E[N(u-1)]}{n-1}. \end{cases}$$

5. EXPERIMENTAL RESULTS

We report experimental data averaged over 100 independent runs for the three update policies and for panmictic and randomly structured populations, using uniform selection described and modeled in Section 4. Although, as stated in that section, we do not provide mathematical models for linear ranking and binary tournament selection, experiments using those selection methods, not shown here for lack of space, give qualitatively similar results. Note also that, although the curves are represented as being continuous for the sake of clarity, they are obviously discrete. In all the curves of Figure 1 the population grows exponentially at first and then saturates, giving the usual sigmoidal shape for the growth curves. However, one can clearly distinguish the three update policies with NRS being faster than synchronous. The UC policy starts similar to NRS and then joins the synchronous case when saturation sets in.

Figure 1 (a) depicts the theoretical curves corresponding to the three update modes given in the previous section. Of course, according to the mean field approximation, the panmictic and the random graph cases are actually the same. This is clearly confirmed by Figures 1 (b) and (c), which show, respectively, the panmictic experimental curves and the random graph case with probability $q = 0.1$. Figure 1 (d) shows the experimental random graph case with $q = 0.01$: we observe that for low probabilities the mean field hypothesis gives a worse approximation of the experimental results. It should be noted that for low q values there is a non-negligible probability that the generated random graph is disconnected. To avoid these cases, we only consider connected graphs in our experiments, randomly sampled among the family $G_{n,q}$ of all possible random graphs with n vertices and edge probability q .

Table 1 gives the predicted and the experimental average takeover times for the three update modes for panmictic and randomly structured populations.

In Figures 2 and 3 we provide a direct comparison by superposing on the same graph the theoretical curves for the synchronous (a), asynchronous NRS (b), and asynchronous UC (c) updates, and ten randomly chosen corresponding experimental curves. This is more informative than a comparison with the average experimental curves since the theoretical result is an expectation curve.




Figure 1: Theoretical growth curves for the three update policies (a). Experimental growth curves for panmictic (b), random graph with $q = 0.05$ (c), and random graph with $q = 0.01$ (d), for the three updates policies. Population size is 1024. The experimental curves in (b), (c), and (d) are averages over 100 independent runs.

	Synchronous	Asynchronous NRS	Asynchronous UC
Predicted Takeover Time	14	11	16
Panmictic Mean Takeover Time	14.65 (1.43)	11.47 (1.28)	16.6 (1.73)
Random ($q=0.05$) Mean Takeover Time	15.49 (1.46)	12.2 (1.2)	17.07 (1.9)
Random ($q=0.01$) Mean Takeover Time	19.5 (2.55)	15.56 (2.01)	20.2 (2.09)

Table 1: Predicted takeover times and experimental mean takeover times (with corresponding standard deviations) for the three update methods. The experimental results are obtained over 100 independent runs. Population size is 1024 in all cases.




Figure 2: Theoretical curve (black) and ten randomly chosen experimental curves (gray) for a random graph population with $q = 0.05$ using synchronous update (a), asynchronous NRS (b), asynchronous UC (c).

Figure 2 reports results for the random graph with $q = 0.1$; it is clear that there is a very good agreement between the prediction of the models and experiments. As stated

above, we cannot hope that such an agreement also holds for random graphs with very low probability q . In fact, Figure 3 shows that the approximation is much worse. This




Figure 3: Theoretical curve (black) and ten randomly chosen experimental curves (gray) for a random graph population with $q = 0.01$ using synchronous update (a), asynchronous NRS (b), asynchronous UC (c).

is understandable qualitatively on the following grounds. In a random graph the node degree is binomially distributed by construction. Therefore, for 1024 vertices, the average number of neighbors of an individual is about 50 for $q = 0.05$ and about 10 for $q = 0.01$. The standard deviation in the latter case is around 3, while it is around 7 in the former in absolute terms; however, in terms of relative deviation (i.e., σ/μ), we have around $1/3$ for $q = 0.01$, ca. 228% more than around $1/7$ for $q = 0.05$. This means that a significant number of nodes will have much fewer edges than the average in the $q = 0.01$ case, which will slow down the propagation rate of the best individual.

6. SMALL-WORLD GRAPH TOPOLOGIES

It has been shown in recent years that graphs that occur in many social, biological, and man-made systems are often neither completely regular, such as lattices, nor completely random [18, 17]. They have instead what has been called a *small-world* topology, in which nodes are highly clustered yet the path length between them is small. This behavior is due to the presence of *shortcuts* i.e., a few direct links between nodes that would otherwise be far removed. Following Watts' and Strogatz's discovery, Barabasi *et al.* [3] found that several important networks such as the World Wide Web, Internet, author citation networks, and metabolic networks among others, also have the small world property but their degree distribution differs: they have more nodes of high degree that are likely in a random graph of the same size and edge density. These graphs have been called *scale-free* because the degree probability distribution function follows a power law. In the next sections we briefly describe how small-world and scale-free graphs can be constructed, more details can be found in [2, 12].

6.1 The Watts-Strogatz Model

Although this model has been a real breakthrough in the technical sense when it appeared, today it is clear that it is not a good representation of real networks as it retains many features of the random graph model. In fact, scale-free and other types of graphs have been successively proposed as more faithful description of the kind of big technological, human, and biological networks we observe. In spite of this, the Watts-Strogatz model, because of its simplicity of construction and the richness of behavior, is still an interesting topology in artificial systems where there is no "natural" constraint on the type of connectivity.

According to Watts and Strogatz [18], a small-world graph

can be constructed starting from a regular ring of nodes in which each node has k neighbors ($k \ll N$) by simply systematically going through successive nodes and "rewiring" a link with a certain probability β . When the edge is deleted, it is replaced with an edge to a randomly chosen node. If rewiring an edge would lead to a duplicate edge, it is left unchanged. This procedure will create a number of *shortcuts* that join distant parts of the lattice. Figure 4 schematically depicts this process for a small ring with $k = 4$.




Figure 4: (a) regular one-dimensional lattice with $k = 4$. (b) a small-world graph obtained by randomly rewiring some of the nearest neighbor links.

Shortcuts are the hallmark of small worlds. While the average path length¹ between nodes scales logarithmically in the number of nodes for a random graph, in Watts-Strogatz graphs it scales approximately linearly for low rewiring probability but goes down very quickly and tends to the random graph limit as β increases. This is due to the progressive appearance of shortcut edges between distant parts of the graph, which obviously contract the path lengths between many vertices. However, small worlds typically have a higher clustering coefficient² than random graphs. Small-world networks have a degree distribution $P(k)$ i.e., the probability that a randomly selected vertex has degree k , that is close to binomial for intermediate and large values of the rewiring probability β as for a random graph. $P(k)$

¹The average path length L of a graph is the average value of all pair shortest paths.

²The clustering coefficient C of a node is a measure of the probability that two nodes that are its neighbors are also neighbors among themselves. The average $\langle C \rangle$ is the average of the C s of all nodes in the graph.

tends to a delta function for $\beta \rightarrow 0$ since in this case we recover the regular lattice.

6.2 The Barabási-Albert Model

Albert and Barabási were the first to realize that real networks grow incrementally and that their evolving topology is determined by the way in which new nodes are added to the network. They proposed an extremely simple model that is still useful based on these ideas [3]. At the beginning one starts with a small clique (a completely connected graph) of m_0 nodes. At each successive time step a new node is added such that its $m \leq m_0$ edges link it to m nodes already in the graph. When choosing the nodes to which the new nodes connects, it is assumed that the probability π that a new node will be connected to node i depends on the degree k_i of i such that nodes that have already many links are more likely to be chosen over those that have few. This is called *preferential attachment* and is an effect that can be observed in real networks. The probability π is given by:

$$\pi(k_i) = \frac{k_i}{\sum_j k_j},$$

where the sum is over all nodes already in the graph. Barabási and Albert have shown that the model evolves into a stationary scale-free network with a power-law probability distribution for the vertex degree $P(k) \sim k^{-\gamma}$, with $\gamma \sim 3$. There exist other more general and more refined models that are capable of producing graphs with a power law degree distribution (see, e.g. [6]). However, the basic Barabási-Albert model is enough for our initial investigation.

6.3 Experimental Results on Small Worlds

In all the experiments described below we have used a population size of 1024, and a total number of edges of the same order of that of a random graph with $q = 0.01$. The selection mechanism employed has been uniform selection in all cases. All the curves are averages of 100 independent runs.

6.3.1 Watts-Strogatz Model

We have used Watts' small world construction with 1024 individuals starting from a ring with $k = 10$ neighbors i.e., a regular radius-five one-dimensional lattice. In this way the mean number of neighbors (10.24) is equal to that of random graphs with $q = 0.01$.

Figure 5 shows the growth curves with synchronous update for different values of the rewiring probability β and for the ring. The trend is clear: increasing β from 0 (the ring case), to 0.8 (topologies approaching that of a random graph), the selection pressure increases slowly at first and then very quickly around $\beta = 0.005$. This can be easily understood if one takes into account how the mean path length and the clustering coefficient vary in a small-world graph. From Figure 6 (redrawn from [17]) one can see that for β around 0.005 there is a sudden drop in the average path length from values that pertain to the lattice to values that are close to those of a random graph. This means that, suddenly, many short paths become available through the network between most nodes, which explains the higher growth rate.

Figure 7 depicts the growth curves for synchronous update, and for the two asynchronous policies in small worlds with $\beta = 0.005$. As for panmictic populations and random

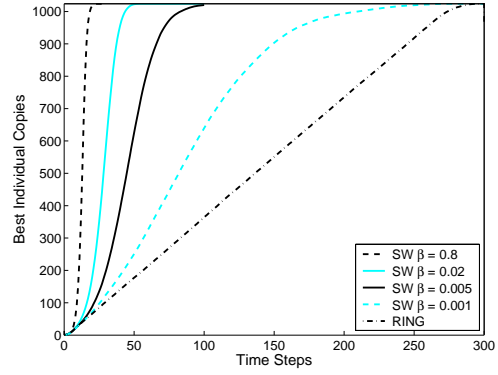


Figure 5: Growth curves for synchronous update for different values of the rewiring probability β . Rightmost curve is for a ring ($\beta = 0$). Leftmost curve corresponds to $\beta = 0.8$, which is almost in the random graph region.




Figure 6: Average path length and clustering coefficient for small-world graphs as a function of the rewiring probability β . Note the log scale on the abscissae.

graphs, new random sweep is faster than uniform choice, which is in turn faster than synchronous. The experimental takeover time values are to be compared with those of random graphs with the same average number of edges i.e., figure 1 (d). Clearly, the corresponding small-world graphs induce a lower global selection pressure.

6.3.2 Barabási-Albert Model

Scale-free graphs have been generated according to the Barabási-Albert model of section 6.2. We start from a clique of $m_0 = 14$ nodes, and we add $1024 - 14 = 1010$ individuals, each creating $m = 10$ edges with preferential attachment, following the algorithm.

Figure 8 shows the behavior of the growth curves for the three update policies used here. For the position of the initial best individual any vertex is equally likely. The order of the curves is the same as that observed for random graphs and Watts-Strogatz small-world networks. The inversion in the last part of the uniform choice curve is due to the fact that cells are chosen with replacement, and thus the last few non-conquered individuals are increasingly unlikely to be chosen. Apart from this effect, the takeover times are very close to those observed in the corresponding random graphs. This confirms that scale-free graphs are a topology

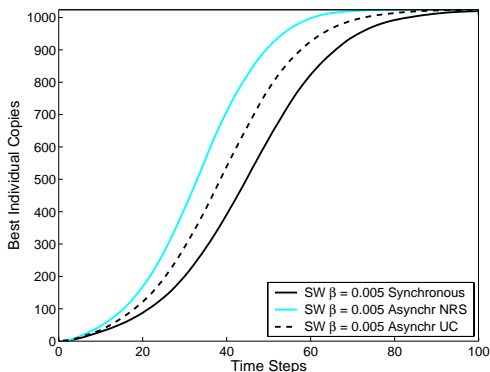


Figure 7: Growth curves for synchronous, new random sweep, and uniform choice asynchronous update. The rewiring probability is $\beta = 0.05$. Please note the change of horizontal axis scale with respect to Figure 5.

in which propagation is at least as fast as for random graphs which, for example, has important consequences in infection rates [12].




Figure 8: Growth curves for synchronous, new random sweep, and uniform choice asynchronous update in scale-free graphs. Initial best individual uniformly distributed at random among the nodes.

But scale-free graphs have other surprising properties. In particular, those networks are extremely tolerant to attacks to randomly chosen target nodes, which is due to the fact that there are few important (highly connected) nodes and many unimportant (sparsely connected) ones. On the other hand, deliberate suppression of highly connected nodes is likely to produce a lot of damage [2]. The different status of highly connected nodes is demonstrated in the following experiment, where the initial best individual has always been placed in a “hub” node (see Figure 9).

The takeover time is very short; shorter than the random graph case (see Figure 1 and Table 1). This is also known to happen in infectious processes, where scale-free communication patterns have the effect of eliminating the so-called infection threshold [12].

Table 2 summarizes numerical results of the takeover times in Watts-Strogatz and scale-free topologies for the synchronous and the two asynchronous updates.

One could probably exploit the effect of small-world topologies on the dynamical properties of evolutionary computa-




Figure 9: Growth curves for synchronous evolution in scale-free graphs when the initial best individual is placed on a highly connected node.

tion processes by letting the topology dynamically adapt or self-organize in order to control the selection pressure, and thus the explorative or exploitative characteristics of the algorithm.

7. CONCLUSIONS

In the first part of this study we presented general stochastic models for individual growth that are valid for any graph-structured population, including the classical panmictic EA as a limiting case. We have given discrete difference equations for the synchronous and two typical asynchronous cell update mechanisms, which are new, as far as we can tell, at least for the asynchronous cases. These discrete logistic recurrence equations seem to us more adequate for finite populations than the usual continuous logistic fitting. Although finding closed forms for the recurrences appears to be mathematically very difficult, the comparison of predicted growth curves with experimental ones confirms that the models are a good description of the propagation phenomena. We have used a selection method that lends itself to an approximate mathematical treatment in terms of the mean-field hypothesis. Under this hypothesis, we found that the panmictic and random graph cases have the same behavior, unless the edge probability of the random graph is very small. A practical consequence of our results is that it appears unnecessary to use the whole population as a selection pool, given that the random graph case using a fraction of the population shows the same behavior as the panmictic one. The models and the experiments also confirm previous results for regular lattices [8, 7] i.e. that the selection pressure can be varied using different updating schemes. In the future we intend to pursue the study of propagation under classical selection schemes such as linear ranking and tournament.

In the second part of the paper we presented an experimental investigation of individual growth in a couple of families of networks that are neither regular nor random: the Watts-Strogatz model and the scale-free Barabási-Albert model. We have empirically shown that the individual propagation properties and the global induced selection pressure are qualitatively similar to those found in panmictic populations and random graphs. However, the inhomogeneity of these small-world networks opens-up new possibilities for evolutionary computation when the nature of a given node is taken into account. Thus, for example, hubs in scale-

	Synchronous	Asynchronous NRS	Asynchronous UC
Ring ($\beta = 0$)	279.64 (9.35)	201.51 (9.03)	229.30 (10.54)
WS ($\beta = 0.001$)	168.89 (45.93)	116.50 (27.12)	135.15 (35.10)
WS ($\beta = 0.005$)	80.16 (14.47)	60.18 (8.75)	70.43 (12.08)
WS ($\beta = 0.02$)	45.96 (4.32)	36.58 (4.15)	41.31 (4.27)
WS ($\beta = 0.8$)	19.16 (1.69)	15.43 (1.76)	20.50 (2.46)
BA	17.94 (2.59)	14.60 (2.66)	19.77 (3.55)

Table 2: Experimental mean takeover times (with corresponding standard deviations) for the three update methods and for the small-world topologies discussed in the text. WS stands for Watts-Strogatz and BA stands for Barabási-Albert. The experimental results are obtained over 100 independent runs. Population size is 1024 in all cases.

free networks largely determine the dynamical properties of the population. By controlling these features, or letting the population network to self-organize, it is possible to change the selection pressure and thus the algorithm characteristics within a wide range. Future work along these lines includes using small-world structured populations together with selection and variation operators on benchmarks and real-life problems.

Acknowledgements

M. Giacobini and M. Tomassini gratefully acknowledge financial support by the Fonds National Suisse pour la Recherche Scientifique under contract 200021-103732/1.

8. REFERENCES

- [1] E. Alba and M. Tomassini. Parallelism and evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 6(5):443–462, October 2002.
- [2] R. Albert and A.-L. Barabasi. Statistical mechanics of complex networks. *Reviews of Modern Physics*, 74:47–97, 2002.
- [3] Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286:509–512, 1999.
- [4] E. Cantú-Paz. *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Press, 2000.
- [5] U. K. Chakraborty, K. Deb, and M. Chakraborty. Analysis of selection algorithms: A Markov chain approach. *Evolutionary Computation*, 4(2):133–167, 1996.
- [6] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of Networks: From Biological Nets to The Internet and WWW*. Oxford University Press, Oxford, UK, 2003.
- [7] M. Giacobini, E. Alba, A. Tettamanzi, and M. Tomassini. Modeling selection intensity for toroidal cellular evolutionary algorithms. In K. D. et al., editor, *Proceedings of the genetic and evolutionary computation conference GECCO'04*, pages 1138–1149. Springer Verlag, Berlin, 2004.
- [8] M. Giacobini, M. Tomassini, and A. Tettamanzi. Modelling selection intensity for linear cellular evolutionary algorithms. In P. L. et al., editor, *6th International Conference on Artificial Evolution, EA 2003*, volume 2936 of *Lecture Notes in Computer Science*, pages 345–356. Springer Verlag, Berlin, 2004.
- [9] M. Giacobini, M. Tomassini, A. Tettamanzi, and E. Alba. Selection intensity for cellular evolutionary algorithms for regular lattices. *IEEE Transactions on Evolutionary Computation*. To appear.
- [10] 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.
- [11] 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.
- [12] M. E. J. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.
- [13] G. Rudolph. On takeover times in spatially structured populations: Array and ring. In K. K. L. et al., editor, *Proceedings of the Second Asia-Pacific Conference on Genetic Algorithms and Applications*, pages 144–151. Global-Link Publishing Company, 2000.
- [14] J. Sarma and K. A. D. 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.
- [15] J. Sarma and K. A. D. 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.
- [16] B. Schönfisch and A. de Roos. Synchronous and asynchronous updating in cellular automata. *BioSystems*, 51:123–143, 1999.
- [17] D. J. Watts. *Small worlds: The Dynamics of Networks between Order and Randomness*. Princeton University Press, Princeton NJ, 1999.
- [18] D. J. Watts and S. H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.