

The Influence of Migration Sizes and Intervals on Island Models

Zbigniew Skolicki
Computer Science Department
George Mason University
4400 University Drive, MSN 4A5
Fairfax, VA USA
zskolick@gmu.edu

Kenneth De Jong
Computer Science Department
George Mason University
4400 University Drive, MSN 4A5
Fairfax, VA USA
kdejong@gmu.edu

ABSTRACT

A need for solving more and more complex problems drives the Evolutionary Computation community towards advanced models of Evolutionary Algorithms. One such model is the island model which, although the subject of a variety of studies, still needs additional fundamental research. In this paper we have experimentally studied the influence of various migrations sizes and intervals on island models using a set of special functions. One of the surprising observations from these experiments is that the migration interval seems to be a dominating factor, with migration size generally playing a minor role with regard to the best solution found. Additional experiments measuring genetic diversity show that too frequent migrations cause islands to dominate others and lose global diversity before they are able to exchange solutions to produce better results. Also, we observe that even small migrations already make a significant impact on the behavior of an island model and therefore the effects are comparable to those of bigger migrations. On the other hand rare migrations cause a degraded performance due to the slow convergence. Collectively, these observations provide useful guidance for island model applications.

Categories and Subject Descriptors: I.2.8 [Artificial Intelligence]: Problem Solving, Control Methods, and Search — heuristic methods; I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence; G.1.6 [Numerical Analysis]: Optimization — global optimization

General Terms: Algorithms, Performance, Experimentation

Keywords: evolutionary computation, island model, migration, migration size, migration interval, test suite

1. INTRODUCTION

The interest of the Evolutionary Computation community in more and more complex problems requires the use

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.

of advanced models of Evolutionary Algorithms (EAs), because standard models are not powerful enough. One family of such models are island models, in which individuals are split into subpopulations (*islands*), evolving on their own, and from time to time exchanging individuals by *migrations*. In addition to the obvious property of distributing computations, island models generally behave qualitatively different than simpler models and therefore are able to solve certain problems that standard models cannot [14, 17].

Although island models have been used many times in the past and there are different studies focusing on certain features, there is still no good understanding of the role of the basic parameters. Because there are so many possible parameter settings for an island model, it is important to understand what influence each parameter has and how it interacts with other parameters.

Two very basic and important parameters of island models are migration sizes and intervals, denoting how many individuals migrate and how frequently migrations occur. In this paper we study in a systematic experimental way how the behavior of an island model changes in reaction to changes in these two parameters, and how they are related.

We start with a brief introduction to island models in section 2. In section 3 we state the goals of our study. Then, in section 4 we describe the experiments we did and present the test suite of eight functions that was used. In section 5 we present the results of our experiments and make some analysis of them. Finally, in section 6 we summarize our work, try to answer the questions set in the goals section, and point out some possible future work.

2. BACKGROUND

Previous studies of island models have been both theoretical and experimental in nature. Strict mathematical analysis (using e.g. Markov chains) of island models is difficult due to their complexity. Consequently, in order to obtain any results, researchers are constrained to simplified models (for example takeover time was analyzed by Sprave [15] and Rudolph [12]). By contrast, experimental studies are generally performed on the island models themselves, but generalization of the results is difficult. As a result, general rules that are useful to the practitioner are difficult to create. To date our best understanding of island models comes from analyzing particular features such as the number of populations and their sizes analyzed by Whitley [18] or Cantú-Paz [3, 5].

The main difference between an island model and a single-population model is the separation of individuals into islands. The islands interact by means of migration, and this is the only means of communication between them. If there is no migration, an island model is nothing more than a set of separate runs. Therefore migration is very important. Migrations were also studied by Cantú-Paz [4].

The standard way of implementing migration assumes sending some number of individuals (*migration size*) every certain number of generations (*migration interval*). The migration size and interval are arguably the two most important migration parameters, controlling the quantitative aspect of migration (the total number of individuals sent in a run of the EA). There exist models where migrations are triggered when certain criteria are met (Branke [2]), however fixed schedules are much more popular and easier to analyze. The level of migration for specific context of agent individuals was studied by Krink [10]. Although different aspects of migration size and interval were studied, we are unaware of any work studying directly the influence on these parameters on the behavior of GA island models, as done in this paper. There exists a similar study done for GP [7, 9].

Another migration parameter defines what is being sent (*migration policy*). A common policy is to choose the best individuals from the source population and replace the worst individuals in the target population (we refer to it as *best-worst policy*). However, as shown by Cantú-Paz [4], this leads to an increased selection pressure in the system. Another possibility is to choose a random individuals and replace random individuals (*random-random policy*). Such policy should not increase the selection intensity, and thus seems appropriate for studying other migration parameters. Other policies are also possible.

The migration topology describes which islands send individuals to which islands. The most often used topologies are: an array, a ring, a grid, a star or a fully connected topology. Sparser connected topologies allow for slower spread of information in the system. Recently Sekaj performed a detailed study analyzing different island model parameters including topology [13].

Multi-machine implementations of island models allow the possibility of islands running asynchronously. In particular, researchers focusing on the use of island models for distributing computations over many machines have studied the properties of asynchronous models [1]. However, a synchronous model seems better for theoretical studies, because experiments are then not dependent on external factors and thus are repeatable. Also, in the case of identical islands, we don't expect to see big qualitative differences in behavior between the two models.

Island models have also been studied outside of computer science. In fact one of the first publications on subpopulations comes from Wright in 1932 [19]. An interesting theory is the *punctuated equilibria* theory [8] stating that subpopulations remain in equilibria, unless a change in environment causes a rapid evolutionary process. Migration can serve as such a trigger [6].

3. RESEARCH GOALS

Whereas it's obvious that the lack of migration results in no cooperation between islands, it is not so obvious how the behavior of an island model changes when we are increasing the size and/or changing the frequency of migrations.

The goal for the experiments presented in the paper was to clarify this matter, by answering as many of the following questions as possible:

1. Does the behavior of an island model (i.e. the average performance) depend only on the total number of migrants (no matter in how many, how big and how frequent migrations), or can we approximate it with such model? For example: do frequent and small migrations have a similar effect to rare and big migrations? Can we divide the migration size-interval plane into regions where there is too much of migration, good amount of migration and too little of migration?
2. If it's more complex, how complex is it with regard to migration size and interval? How do particular parameters influence the observed performance?
3. Can some setups of an island model achieve better results than the panmictic setup with one big population, or the setup with islands separated from each other (which is equivalent to restarting a single population as many times as there are islands)?
4. What settings of these parameters cause island models to succeed?
5. What settings of migration size and interval cause island models to fail?
6. What are the underlying causes that result in some settings being better than others? How is it related to the content of populations?

4. EXPERIMENTAL METHODOLOGY

To thoroughly study the influence of migration sizes and intervals on the behavior of island model EAs we fixed all other island model parameters and performed experiments for all combinations of these two parameters within certain ranges. This allowed us to draw charts representing average EA performance as a function of migration size and interval and see how they are related. The experiments were preceded by manually choosing other parameters in such a way that the problems used were semi-difficult. This allowed for observing improvements and degradations of the algorithm.

4.1 Experimental details

In the experiments we used a model with identical islands (the same parameters) evolving synchronously. We have experimented with setups of 2 and 5 islands. The total number of individuals in the system was set to 100. This matched the problems used (described in the next subsection) so that they were still relatively difficult. Using a smaller number would cause probably too high granularity of each island.

The EA used in islands was a standard one, so that the results could later be used as a guidance for using island models. A real valued representation was used (one gene for each argument of the given test function). We used one-point crossover at a rate of 0.7 (otherwise cloning). The mutation was Gaussian with a self-adapting sigma, and the mutation rate was $1/L$. The EA was generational (non-overlapping) with the exception of the best individual carried over (elitism). The ranked selection strategy was chosen as a good balance in terms of selection pressure.

Each island was run for for a maximum of 500 generations; however runs were stopped earlier in case of detecting a convergence, which was measured by computing the global level of diversity. In fact only around 1% of runs reached the limit.

Migrations occurred every *migration interval* number of generations for each island, and their size was *migration size*. The experiments were run with a random-random policy, in which a group of migrants is chosen at random from the source population (without repetition) and random individuals are replaced in the target population. This policy which was chosen because it doesn't change the selection intensity and thus a comparison between settings with different level of migrations is more meaningful. Some experiments with the best-worst policy (take the best, replace the worst individuals) suggested that the results described in later sections are also valid for other policies. However, we didn't perform the full set of experiments for the best-worst policy.

A random (dynamic) topology was used [7]. Each time a population was expecting a migration, the population was randomly paired with another, from which individuals would be migrated. Therefore the relations between any two populations were identical, as opposed to the topologies like a ring or torus, where some islands are direct neighbors and some are more distant from each other. In the same time the topology didn't "flood" the populations with migrants (as it would be with the fully connected topology).

We varied the migration interval using the values of 1, 2, 3, 4, 5, 10, 15, 20, 25, 30. The migration size was varied using the values of 1, 2, 4, 6, 8, 10, 12, 14, 16, 18, 19, 20. Because of the stochasticity of the process, for each combination of these two parameters we used at least 45 repetitions. To make it easier to compare, for different combinations the runs were starting from the same set of random seeds (meaning that for different combinations of parameters, corresponding runs were identical until the first migration occurred - which migration was naturally different for different combinations)

The performance of the EA was measured with the best-so-far value from all islands.

4.2 Test suite

A total number of eight functions was used. All of them operate on real numbers and return a real value. The first four of them (described below) are specially created functions to study properties of island models. The task is to maximize them. We tried to keep the functions as simple as possible so that it can be easily understood what is driving the algorithm. However, because island models are generally more suitable for functions that are complex and not easily solvable by standard EAs, we created the functions so that they possess some features characteristic for more difficult problems (e.g., IM2-IM4 require that islands cooperate to find a good solution).

The remaining four functions are standard multimodal test functions, known as Rosenbrock, Schwefel, Rastrigin and Griewangk functions. The task is to minimize them. Below we present the characterization of each function.

4.2.1 Function IM1

The goal of the first function was to see when too much migration worsens the behavior of an island model. It consisted of one wide low peak and one narrow and high peak.




Figure 1: function IM1




Figure 2: function IM2

If there is too much information exchange between the islands, one can expect populations would drag each other onto the wide deceptive peak. Because we are reporting the best-so-far value from all islands, migration cannot improve it — if one island finds the best solution, it will be reported.

The formula for the function is:

$$\text{IM1}(x_1, x_2) = \max(0, -(2x_1 - 1)^2 - (2x_2 - 1)^2 + 1, \\ -(5(x_1 - 1))^2 - (5(x_2 - 1))^2 + 1.25)$$

The domain for both arguments is $[0, 1]$.

4.2.2 Function IM2

This function consists of two wide and one narrow peak, reachable by a crossover of solutions from the wide peaks. It is unlikely that the global solution is found without performing this crossover, and thus without islands cooperating. Such cooperation was in fact observed in single preliminary runs of an EA on this problem.

The formula for the function is:

$$\text{IM2}(x_1, x_2) = \max(0, -x_1^2 - \left(\frac{x_2 - 1}{2}\right)^2 + 1, \\ -\left(\frac{x_1 - 1}{2}\right)^2 - x_2^2 + 1, \\ -(30(x_1 - 1))^2 - (30(x_2 - 1))^2 + 1.25)$$

The domain for both arguments is $[0, 1]$.

4.2.3 Function IM3

This function is an n-dimensional equivalent of function IM2 (we have used 10 dimensions). There are $2^n - 1$ peaks

in this function, located at points $(x_1, \dots, x_n) : x_i \in \{0, 1\}$ (omitting $(0, \dots, 0)$). The height of each peak depends on the number of 1s in its location, more 1s giving higher value. However, higher peaks are narrower and again it is unlikely to find them without crossing over solutions from lower peaks (function IM3 is an example of a hierarchical function). The global optimum consists of all 1s. Because of a bigger complexity than in function IM2, one would expect a bigger need for cooperation between islands to find the global solution.

The formula for the function is:

$$\text{IM3}(x) = \max_{p \in \{0,1\}^n} \text{peak}_p(x)$$

where $x = (x_1, \dots, x_n)$, and if $p \equiv 0$ then

$$\text{peak}_p(x) \equiv 0$$

if $\sum_i p_i = 1$ then

$$\text{peak}_p(x) = 1 - \sum_{i=1}^n \frac{(x_i - p_i)^2}{n}$$

and if $\sum_i p_i > 1$ then

$$\text{peak}_p(x) = \frac{n-1}{n} + \sum_{i=1}^n \left(\frac{p_i}{n} - k p_i (x_i - 1)^2 - \frac{(1-p_i)x_i^2}{n} \right)$$

We have set $n = 10$, $k = 100$. The domain for all arguments is $[0, 1]$.

4.2.4 Function IM4

For this function better solutions are achieved by a combination of not fully optimized local solutions. Overfitting to a local solution results later in an unproductive crossover. Therefore one would expect that a cooperation between islands is needed to find the global solution — and must happen relatively often, before islands converge to locally ideal solution. Such behavior is supposed to mimic the nature of more complex real problems, for which we expect that a repeated interaction between islands is required to find a good solution.

The formula for the function is:

$$\begin{aligned} \text{IM4}(x_1, x_2) = & \max(\text{decept}\left(\frac{2x_1}{3}, \frac{2x_2}{3}\right) \\ & + \sum_{i=2}^4 i \text{decept}(2^i(x_1 - 2 + 2^{(1-i)}), 2^i(x_2 - 2 + 2^{(1-i)})), \\ & - \frac{x_1^2 + x_2^2}{32} + 0.25) \end{aligned}$$

where

$$\begin{aligned} \text{decept}(x_1, x_2) = & \max(-64(x_1 + x_2)^2 - \frac{(x_1 - x_2 + 2)^2}{9} + 1, \\ & -64(x_1 + x_2)^2 - \frac{(x_1 - x_2 - 2)^2}{9} + 1, \\ & -(x_1 + x_2 + 1)^2 - 100(x_1 - x_2 + 2)^2 + 2, \\ & -(x_1 + x_2 + 1)^2 - 100(x_1 - x_2 - 2)^2 + 2) \end{aligned}$$

The domain for both arguments is $[-2, 2]$.




Figure 3: function IM4

4.2.5 Rosenbrock function

The formula for the function is:

$$f(x_1, \dots, x_n) = \sum_{i=2}^n (100(x_i - x_{i-1}^2)^2 + (1 - x_{i-1})^2)$$

The domain for all arguments is $[-2.048, 2.048]$.

4.2.6 Schwefel function

The formula for the function is:

$$f(x_1, \dots, x_n) = \sum_{i=1}^n -x_i \sin(\sqrt{|x_i|})$$

The domain for all arguments is $[-500, 500]$.

4.2.7 Rastrigin function

The formula for the function is:

$$f(x_1, \dots, x_n) = 10n + \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i))$$

The domain for all arguments is $[-5.12, 5.12]$.

4.2.8 Griewangk function

The formula for the function is:

$$f(x_1, \dots, x_n) = 1 + \sum_{i=1}^n \frac{x_i^2}{4000} - \prod_{i=1}^n \cos\left(\frac{x_i}{\sqrt{i}}\right)$$

The domain for all arguments is $[-600, 600]$.

5. EXPERIMENTAL RESULTS

For all setups the results looked quite similar, which suggested that there are certain types of dependencies between migration parameters and island model EA performance that are quite general.

5.1 Varying size and interval

Because of space limitations, on Fig.4 - Fig.11 we show the best-so-far values achieved only in experiments with 5 islands. It seems that the regularities observed are better visible for 5 island setups. In all charts the values are averaged over multiple runs and are the final values after the algorithm has converged.

The surprising observation from the experiments was that the migration interval was playing much bigger role than the




Figure 4: IM1, best-so-far fitness
(panmictic = 1.145, separated = 1.165)




Figure 8: Rosenbrock function, best-so-far fitness
(panmictic = 8.652, separated = 3.437)




Figure 5: IM2, best-so-far fitness
(panmictic = 1.105, separated = 1.075)




Figure 9: Schwefel function, best-so-far fitness
(panmictic = 259.7, separated = 738.9)




Figure 6: IM3, best-so-far fitness
(panmictic = 1.184, separated = 1.12)




Figure 10: Rastrigin function, best-so-far fitness
(panmictic = 6.607, separated = 4.397)




Figure 7: IM4, best-so-far fitness
(panmictic = 2.843, separated = 3.111)




Figure 11: Griewangk function, best-so-far fitness
(panmictic = 0.101, separated = 0.038)




Figure 12: IM3, best fitness in runs with different intervals , size = 10

migration size, especially for short intervals, when we could regularly observe a strong decrease in the performance. The reason for that was populations "overwriting" other ones before a successful exchange and recombination of solutions could be made. The behavior is deeper investigated in the next section.

The size of migrations didn't influence the behavior as much, with the exception for the sizes equal to the size of subpopulations (in our case equal to 20), when target populations are overwritten with source populations. One could expect a little bit more sensitivity for small migration sizes, which is however not observed.

The values plotted on the charts are the final ones, after all populations having converged to the same solutions. However for different setups the convergence takes different time and in particular it takes longer for infrequent migrations. Therefore if resources are limited, one should avoid using too rare migrations, because the island model will have no time to benefit from exchanging solutions and the best fitness will be worse than expected. Interestingly, for rare migrations the convergence usually occurred faster for medium sizes, than for big or small ones. Migrations close to subpopulation sizes probably make it difficult to mix the genetic material, since they swap the contents of whole populations. With very small migrations it is more probable that the migrant will die without affecting the target population and thus the longer average time to converge. Also it is possible that a recombination of the old and new genetic material may produce better results and extend the evolution.

Finally, in brackets under the graphs we show the best-so-far fitness for the panmictic case, and for the separated islands case. One can notice that the island model is very often achieving better results.

5.2 Measuring diversity for selected runs

To understand better why we observe such a strong influence of the migration interval on performance, we felt it was important to understand what was happening inside the populations. In particular, we did additional experiments measuring genetic diversity during selected runs. The diversity was computed as standard deviation on genomes [11].

First, we studied changes in diversity during runs with migration size fixed (equal to 10). We show the fitness and diversity plots for function IM3 in Fig. 12 - Fig. 14 (error bars plot for confidence level 95%). The graphs clearly show that frequent migrations lead to a fast drop of diversity. This




Figure 13: IM3, average local diversity inside populations in runs with different intervals , size = 10

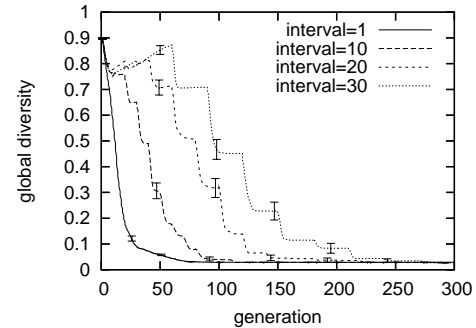


Figure 14: IM3, average global diversity in runs with different intervals , size = 10




Figure 15: IM3, average fitness in runs with different sizes , interval = 20




Figure 16: IM3, average local diversity inside populations in runs with different sizes , interval = 20

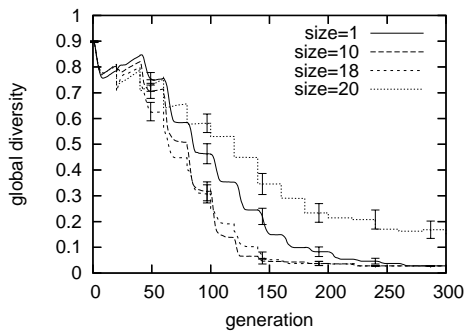


Figure 17: IM3, average global diversity in runs with different sizes , interval = 20

is caused by individuals from one island dominating individuals from another island. If one waits long enough before migrating, the global diversity remain large, even though individual island diversities drop. This means that different islands converge to different peaks. Exchange of individuals at this point results in better average fitness.

Later we studied the EA behavior with fixed interval (equal to 20) and variable migration size (Fig.15 - Fig.17 show fitness and diversity). The graphs show why migration size may not be influencing the results so much. We observe that a migration of even a single individual often quickly affects the whole target population, which results in a diversity drop and fitness increase comparable to the case in which a large number of individuals migrate (Fig.17). Moreover, for very small migrations we have sometimes observed an increase in diversity in generations following a migration, which suggests that a single individual is able to act as a seed for changes in a stagnated population. Big migrations tend to have an immediate effect on diversity simply because of replacing a larger number of individuals. However they diversity doesn't change much in following generations.

One of the questions we wanted to answer was whether setups with different migration intervals and sizes, but with the same total number of migrants (on average per generation) will show the same behavior. For the experiments we did, we observed that most of such setups converge to relatively similar fitnesses, with exception of the extreme setups. A plot for IM3 function is shown in Fig. 18. However, in our experiments we observed little sensitivity to changes in migration size and interval for moderate values of these parameters. Studying diversity plots (Fig. 19) shows that setup with bigger intervals and sizes need more time to converge.

6. CONCLUSIONS AND FUTURE WORK

Let us try to summarize the results of experiments and answer the questions posed in the section 3.

1. The experiments showed that for the chosen island model it is not simply the total number of exchanged individuals that is important, but rather the behavior is a result of several phenomena. However, it is difficult to generalize and give a definite answer for any island model.

2. The performance of an island model is affected by at least three factors, combined together (plus some domain-specific behavior):

- For frequent migrations, the effects of the migration interval are much stronger than the effects of migration size. If the migration interval is too small, performance drops drastically.

- There is a degradation in performance for migration sizes approaching the subpopulation size.

- There may be a degradation in performance for large migration intervals if the algorithm stops prematurely.

3. Comparing the best-so-far fitness from different setups we observed that many of them achieved better results than both panmictic and separated setup.

4. The best performance is achieved with moderate migration intervals and small migration sizes. Therefore island models' "sweet spot" is not connected in parameter space to regions that would correspond to one big population (frequent big migrations), nor to many separated ones (rare and very small migrations). This is in agreement with "rules of thumb" existing in literature (for example stating that optimal migration settings for island GP is 10% population every 5-10 generations [16]). However, we now know that the island model behavior is very sensitive to small migration interval and therefore we must be careful to keep it big enough. Moreover, our results suggest that a migration size of 10% is unnecessarily big.

5. Too frequent migrations should be avoided, because they make islands to exchange individuals so fast that all islands start to share the same individuals. This loss of diversity negatively affects performance. Although we noticed a negative influence only of artificially big migrations (equal to subpopulation sizes), for different policies (like best-worst) the influence may be visible earlier also as a loss of diversity. In addition, for more complex functions rare migrations may occur too few times to be useful and should be avoided.

6. Our experiments showed the necessity of maintaining diversity in the system to produce novel results and the necessity for enough information exchange between islands to combine partial results. Since these properties are not specific to EA or representation we used, we believe that for other island model EAs we should observe a similar dependency between migration size, interval and EA performance.

There is a lot of future work with regard to island models and migrations. The above observations together with the diversity charts suggest a hypothesis that migrations should occur only after the diversity in the system stabilizes (or reaches appropriate level). This hypothesis could be studied in the context of the *punctuated equilibria* theory [6].

We could get different results for finer-grained migration sizes. Therefore studying bigger populations, allowing for percent-wise smaller migrations could shed more light on this matter. One can expect that there should be a sudden drop in performance when shifting to tiny migrations that are ultimately too weak to have any effect on the island model.




Figure 18: IM3, average best fitness in different setups with the same average number of migrants per generation




Figure 19: IM3, average global diversity in different setups with the same average number of migrants per generation

Finally, the functions studied here were relatively simple. It would be worth seeing how the properties described in this paper extend for even more complex problems.

7. REFERENCES

- [1] E. Alba and J. M. Troya. An analysis of synchronous and asynchronous parallel distributed genetic algorithms with structured and panmictic islands. In *IPPS/SPDP Workshops*, pages 248–256, 1999.
- [2] J. Branke, A. Kamper, and H. Schmeck. Distribution of evolutionary algorithms in heterogeneous networks. In *Proceedings of Genetic and Evolutionary Computation Conference - GECCO 2004*. Springer-Verlag, 2004.
- [3] E. Cantú-Paz. *Designing Efficient and Accurate Parallel Genetic Algorithms*. PhD thesis, University of Illinois at Urbana-Champaign, 1999.
- [4] E. Cantú-Paz. Migration policies, selection pressure, and parallel evolutionary algorithms. *Journal of Heuristics*, 7(4):311 – 334, 2001.
- [5] E. Cantú-Paz and D. E. Goldberg. Are multiple runs of genetic algorithms better than one? In E. C.-P. et. al., editor, *Proceedings of the Genetic and Evolutionary Computation Conference 2003*, 2003.
- [6] J. Cohoon, S. Hedge, W. Martin, and D. Richards. Punctuated equilibria: A parallel genetic algorithm. In J. Grefenstette, editor, *Proceeding of the Second International Conference on Genetic Algorithms and their Applications*, pages 148–154, Hillsdale, NJ, 1987. Lawrence Erlbaum Associates.
- [7] F. F. de Vega, M. Tomassini, W. F. Punch III, and J. M. Sanchez-Prez. Experimental study of multipopulation parallel genetic programming. In *Proceedings of the European Conference on Genetic Programming*, volume 1802 of *Lecture Notes In Computer Science*, pages 283 – 293. Springer-Verlag, 2000.
- [8] N. Eldredge and S. Gould. Punctuated equilibria: an alternative to phyletic gradualism. In T. Schopf, editor, *Models of Paleobiology*, pages 82–115. Freeman, Cooper, San Francisco, 1972.
- [9] F. Fernandez, M. Tomassini, and L. Vanneschi. An empirical study of multipopulation genetic programming. *Genetic Programming and Evolvable Machines*, 4(1):21–51, March 2003.
- [10] T. Krink, B. H. Mayoh, and Z. Michalewicz. A PATCHWORK model for evolutionary algorithms with structured and variable size populations. In W. Banzhaf, J. Daida, A. E. Eiben, M. H. Garzon, V. Honavar, M. Jakiela, and R. E. Smith, editors, *Proceedings of the Genetic and Evolutionary Computation Conference*, volume 2, pages 1321–1328, Orlando, Florida, USA, 13-17 1999. Morgan Kaufmann.
- [11] R. W. Morrison. *Designing Evolutionary Algorithms for Dynamic Environments*. PhD thesis, George Mason University, Fairfax, VA, 2002.
- [12] G. Rudolph. On takeover times in spatially structured populations: Array and ring. In K. K. Lai, O. Katai, M. Gen, and B. Lin, editors, *Proceedings of the Second Asia-Pacific Conference on Genetic Algorithms and Applications (APGA '00)*, pages 144–151, Hong Kong, PR China, 2000. Global-Link Publishing Company.
- [13] I. Sekaj. Robust parallel genetic algorithms with re-initialisation. In X. Yao, E. K. Burke, J. A. Lozano, J. Smith, J. J. M. Guervós, J. A. Bullinaria, J. E. Rowe, P. Tiño, A. Kabán, and H.-P. Schwefel, editors, *Parallel Problem Solving from Nature - PPSN VIII, 8th International Conference, Birmingham, UK, September 18-22, 2004, Proceedings*, volume 3242 of *Lecture Notes in Computer Science*, pages 411–419. Springer, 2004.
- [14] Z. Skolicki and K. D. Jong. Improving evolutionary algorithms with multi-representation island models. In *Parallel Problem Solving from Nature - PPSN VIII 8th International Conference*. Springer-Verlag, 2004.
- [15] J. Sprave. A unified model of non-panmictic population structures in evolutionary algorithms. In P. J. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzalá, editors, *Proceedings of the Congress on Evolutionary Computation*, volume 2, pages 1384–1391, Mayflower Hotel, Washington D.C., USA, 6-9 1999. IEEE Press.
- [16] M. Tomassini. Spatially Structured EAs. In *GECCO'04 Tutorials*, pages xx–xx, June 2004.
- [17] D. Whitley, S. Rana, and R. B. Heckendorn. Island model genetic algorithms and linearly separable problems. In *Selected Papers from AISB Workshop on Evolutionary Computing*, volume 1305 of *Lecture Notes In Computer Science*, pages 109 – 125. Springer-Verlag, 1997.
- [18] D. Whitley, S. Rana, and R. B. Heckendorn. The island model genetic algorithm: On separability, population size and convergence. *Journal of Computing and Information Technology*, 7(1):33–47, 1999.
- [19] S. Wright. The roles of mutation, inbreeding, crossbreeding and selection in evolution. In D. F. Jones, editor, *Proceedings of the Sixth International Conference of Genetics*, pages 356–366, Brooklyn Botanic Garden, 1932.