

The Importance of a Two-level Perspective for Island Model Design

Z. Skolicki, K. De Jong, *George Mason University*

Abstract—Our theoretical understanding of island models (IMs) is much worse than of single-population evolutionary algorithms (EAs). As a consequence there is relatively little guidance available to a practitioner for even the most basic aspects of IM design such as choosing the size and number of the islands. In this paper we improve on this situation by showing how a particular two-level perspective can in fact provide guidance for IM design.

I. INTRODUCTION

Island models (IMs) are evolutionary algorithms (EAs), in which individuals are split into sub-populations (islands), each one evolving independently. Islands interact by means of occasional migrations. One of the reasons for using IMs is to take advantage of the parallelism obtained by distributing the computation over multiple machines. However, in addition to a possible speedup, IMs also have very different dynamics than standard single-population EAs, resulting in different problem solving behavior. Unfortunately, our theory for IMs is much less developed than for single population EAs.

In this paper we show how our understanding of IMs can be improved by taking a two-level perspective of their structure and behavior. On one hand each island has its own local EA. On the other hand interactions among islands can be viewed as forming a higher-level structure. This duality makes it possible to identify two interacting evolutionary processes - one at the local level, and another at the inter-island level. Such perspective makes it easier to understand the dynamics of IMs and provides guidance for setting IMs parameters.

The paper is organized as follows. After this introduction, a short background on IMs is given in section II. Section III describes experimental details. In section IV we discuss the two levels of evolution. After we get a better understanding of this perspective, in section V we show how this viewpoint may be used for making various design decisions. In particular we explain how to set up the number and size of islands and how those decisions interact with the choice of local EA. The paper finishes with a discussion and conclusions.

II. BACKGROUND

Island models were initially studied in evolutionary biology. One of the first publications on the subject was a Shifting Balance Theory by Wright [1]. In computer science distributed EAs were both studied in theory and verified in practice. A good background introduction was given in

Z. Skolicki is a Ph.D. candidate (phone 703-993-3919; fax 703-993-1710; zskolick@gmu.edu) and K. De Jong is a Professor (phone 703-993-1553; fax 703-993-1710; kdejong@gmu.edu) at the Department of Computer Science, George Mason University, 4400 University Drive, Fairfax, VA, USA

[2]. A short review was also included in our previous paper [3]. Below we mention a few studies relevant to the design decision guidelines developed in this paper.

Cantú-Paz developed a model for population sizing based on an analysis of building blocks supply and decision [4]. Cantú-Paz and Goldberg also compared isolated runs of small populations with single runs of a big population in [5], generally concluding the advantage of the latter, but with some exceptions.

Whitley et al. analyzed the behavior of IMs on linearly separable problems and concluded that IMs are useful if simply increasing population size in an EA does not help [6].

Fernández et al. studied the number and size of sub-populations using a genetic programming approach [7]. They found that comparing to isolated setups, interaction between islands improves the results considerably and smaller islands (in bigger quantity) can be used. The total number of individuals required were highly dependent on a particular domain.

III. EXPERIMENTAL SETUP

In this section we describe a certain IM implementation used in various experiments later. Unless explicitly stated otherwise, we use the parameter values given below.

A default model with $N = 10$ identical islands of size $M = 10$ individuals (described as $N \times M$ model) is used. Generations are synchronized and migrations are occurring every migration interval $i = 10$. Instead of specifying a constant number of migrants we use a migration probability $\alpha = 0.1$ applied to each individual. Such approach allows us to experiment in general with migration ratios different than fractions resulting from an island size.

We use a dynamic full topology, which means that each time a migration is about to occur, a single target island for migrants is chosen randomly out of all islands. A similar approach was used in the literature [7]. Dynamic topologies let us compare setups with different number of islands and constant α in a fair way, because the total number of migrants remain constant. We use a random migration policy, that is, we choose emigrants randomly and replace random individuals in the target islands.

Two different EA setups are used, called EA-1 and EA-2 [8]. EA-1 has weaker selection using a binary tournament parent selection, no survival selection and non-overlapping generations (similarly to GA setups). EA-2 has stronger selection using a uniform stochastic parent selection, a truncation survival selection, overlapping generations, and a brood ratio of 1.0 (similarly to $(\mu + \lambda)$ ES setups, where $\mu = \lambda$). Both recombination and mutation are used. For

recombination, uniform crossover is always used. Mutation is performed at a $1/L$ ratio. It is either bit-flip mutation for a binary representation, or a non-adaptive Gaussian mutation with $\sigma = 0.01\%$ of each appropriate domain range (quite small) for a real-valued representation.

Maximal fitness is the best fitness value in any island in all generations. To report the maximal fitness, simulations are run until the whole system converges, if needed until 2000 generations. For all measures an average of 60 runs (with the same configuration) is reported. This number of experiments results in small confidence intervals, and therefore we do not plot them in charts to avoid visual cluttering.

IV. TWO LEVELS OF EVOLUTION

In this section we focus on the two levels of evolution in IMs, and suggest that they are to some extent independent. Since the local-level evolution inside islands is better understood, we mainly discuss the importance of global-level mechanisms.

Islands often converge quickly and represent areas around single individuals. With N islands this results in a “population” of N “individuals”. Local EAs play a function of (directed) mutation at the global scale. Migrations between two islands result in better individuals surviving, so they correspond to selection. A recombination of individuals from two islands (after migration) that creates a novel solution may be treated as a recombination of islands.

A. Diversity measures

We can measure diversity of individuals at both levels of evolution. Interaction between these levels will be reflected in changes to appropriate diversities.

1) *Local diversity*: Local diversity is computed in a way analogous to a standard EA approach. We compute standard deviation of genotypes (they are fixed length) [9] and divide by \sqrt{L} for normalization. For the i -th island, the centroid c_i can be computed by the following formula (please note that we operate on vectors)

$$c_i = \frac{\sum_{j=1}^M x_{ij}}{M}$$

and the diversity is computed by

$$d_i = \sqrt{\frac{\sum_{j=1}^M \|x_{ij} - c_i\|^2}{LM}}$$

where M is the number of individuals, L is the length of genotypes, x_{ij} is the j -th individual in the i -th island. For binary genotypes, we treat their alleles as 1.0 and 0.0 (and one can speed up computation by appropriately regrouping ones and zeros). Diversity inside each island may differ, but we generally report an average local diversity $d_L = \frac{\sum_i d_i}{N}$.

2) *Inter-island diversity*: Analogously to the local diversity, we could compute a global one, d_G , using individuals from all islands. However, instead of looking at the total variation of all individuals in the system, we can analyze only the diversity between populations, or more precisely between their centroids. This approach reflects the diversity of the global-level evolutionary process. Such *inter-island* diversity is defined by

$$d_P = \sqrt{\frac{\sum_{i=1}^N \|c_i - c\|^2}{LN}}$$

A standard analysis of variance lets us divide the total variance (global diversity) into “within islands” (local diversity) and “between islands” parts (inter-island diversity)

$$\sum_{i=1}^N \sum_{j=1}^M \|x_{ij} - c\|^2 = \sum_{i=1}^N \sum_{j=1}^M \|x_{ij} - c_i\|^2 + M \sum_{i=1}^N \|c_i - c\|^2$$

The left hand side is from definition equal to $LN M d_G^2$ and the first component of the right hand side is equal to $LM \sum_i d_i^2$. Therefore the inter-island diversity d_P is equal to

$$d_P = \sqrt{\frac{LN M d_G^2 - LM \sum_i d_i^2}{LN M}} = \sqrt{d_G^2 - \frac{\sum_i d_i^2}{N}}$$

Assuming that islands on average have similar local diversities (as it should be with homogenous islands), we can approximate $\frac{\sum_i d_i^2}{N} \approx d_L^2$, and so $d_P \approx \sqrt{d_G^2 - d_L^2}$.

B. A growth of inter-island diversity

Similarly to a local evolution, diversity is also important at the global level, i.e., islands should differ from each other. The initial source of both local and global diversity is stochasticity. All islands are usually initialized using the same procedure and use the same EA, but due to their finite sizes, random initialization and stochastic sampling they follow different evolutionary paths. The initial random state makes inter-island diversity relatively low. However, due to local evolution focusing on different suboptimal peaks, inter-island diversity grows quickly (please compare diversity charts later).

C. Spreading and selection of solutions

Shifting Balance Theory (SBT) explained a benefit of distributed search and spreading good solutions among islands. This behavior corresponds to a selection of better islands and is to much degree independent from the local search, which may but does not have to improve solutions. In fact SBT suggests random drift as the main lower-level mechanism. Spreading solutions is useful in finding peaks of multi-modal functions, specially if they are located in an irregular way.

Let us create a function IM1, which consists of one wide low peak and one narrow and high peak (see Fig. 1). The function is defined over $[0, 1]^2$ and the formula is




Fig. 1. Selected functions used in this paper

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

A single population running on this landscape would most of the time converge to the suboptimal (wide) peak. However, if we used a set of populations, a few of them would converge to the high peak. If in addition migration was used, this optimum would spread over in the whole system. In Fig. 2 we show the mean best fitness achieved with a constant island size M , and a growing number of islands N . Bigger N gives a better search ability and is a reason for fitness improvement. For comparison, results from experiments with a panmictic (single-population) model are also drawn.

A more exploitative EA-2 obtains better results which may mean a focus put on global evolution, at the cost of local exploration. A more explorative EA-1 means more local exploration, from which however the global process does not benefit and it is more difficult to spread solutions.

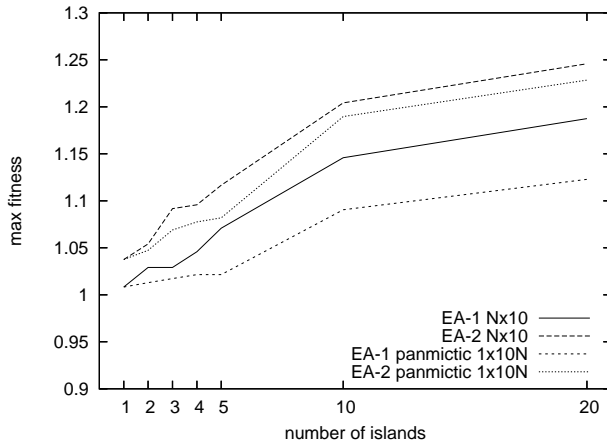


Fig. 2. Spreading solutions becomes easier with bigger N . IM1 function.

D. Recombination of solutions

Due to inter-island diversity various individuals survive in their niches. Many engineering problems have modular solutions and such individuals can optimize different aspects of a solution. This in turn may occur useful later after recombination. Local interaction of individuals, due to an earlier convergence of the islands and evolution after migration may effectively cause an island to converge to a combination of previously known solutions. Such process can be treated as a higher-level “recombination” of islands.

Let us define a function IM2 (see in Fig. 1) over $[0, 1]^2$ by the following formula

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

It is highly unlikely that the highest peak is found randomly, so individuals representing the two lower peaks need to be recombined. Since a single island will most likely converge to either one or another, the probability of a successful recombination grows with a higher number of islands, as shown in Fig. 3.

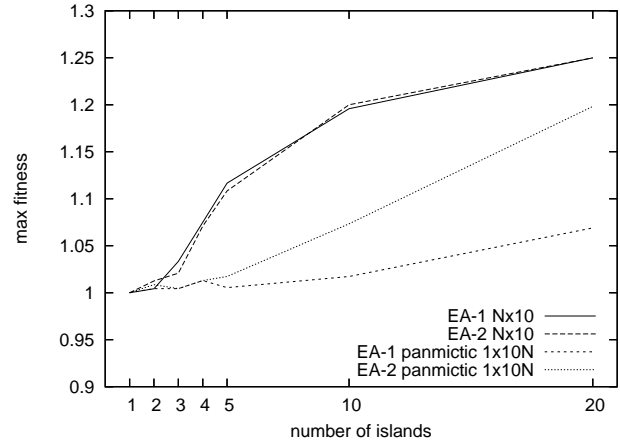


Fig. 3. Recombining solutions becomes easier with bigger N . IM2 function.

E. Interaction between global and local evolution

It is well understood that migrations increase local diversity. However, such increase is very different from the effects of simply increasing exploration with a higher mutation or recombination rate. This is because it is not random and occurs at a cost of the inter-island diversity, as will be seen in diversity charts. This process keeps increasing the local diversity until the “reservoir” of the inter-island diversity finally runs out.

The diversity increase is “selective”, that is it affects only those loci which had different alleles in migrants and locals. In this way the global-level evolution serves as a first stage of search, fixing certain loci. The local evolution, by mixing

migrants and local individuals focuses on details and parts of solutions that still need optimization.

While the above mechanism is the most expected, there exist problems deceptive for IMs, for which the interaction between the two levels of evolution is harmful. In the rest of this section we will show examples of such functions.

One such possibility are functions/problems which have various “paths” to reach the global solution, and for which mixing solutions from different paths creates stable but sub-optimal solutions (a “trap”). With bigger N the probability that islands choose different paths and get trapped after migrations is higher.

Let us define a function IM-Trap, which operates on binary genomes and returns the greater of two numbers: the number of consecutive zeros starting from left or the number of consecutive ones starting from the right. In addition there is also a “trap” in the form of a sub-optimal, high value returned if there is simultaneously a block of zeros from the left and a block of ones from the right, of certain lengths. The IM-Trap formula is

$$\begin{aligned} \text{left}(a) &= \max_{i=0\dots L} \{i : \forall j \quad 1 \leq j \leq i \quad a_i = 0\} \\ \text{right}(a) &= \max_{i=0\dots L} \{i : \forall j \quad 0 \leq j \leq i-1 \quad a_{L-i} = 0\} \\ \text{IM-Trap}(a) &= \begin{cases} 0.9 \cdot L & \text{if } \text{left}(a) > \frac{L}{4} \\ & \text{and } \text{right}(a) > \frac{L}{4} \\ \max\{\text{left}(a), \text{right}(a)\} & \text{otherwise} \end{cases} \end{aligned}$$

This is an example of a situation when the local evolution might solve a problem, but the global process disturbs achieving this goal. If mixing is very low (and there is hardly any inter-island evolution), a single island may discover the global optimum before islands get trapped due to an exchange of individuals. If mixing is very high (and inter-island evolution converges fast), islands may unify and follow the same path to the global optimum. If however mixing is moderate (and inter-island evolution is present), building blocks for the trap will be found in different islands, and then mixed together.

In Fig. 4 we show the results of simulation for various N , using EA-2 (for EA-1 different parameter settings would have to be used). As expected, using more islands results in a worse performance, even though we do not change the size of islands, so more computational resources is used. There is again a slight improvement for $N > 5$, as a result of either effectively sparser interaction between islands (as we use a dynamic topology) or simply much more individuals in the system.

Another example of a problem difficult for an IM is one in which local-level evolution diminishes rather than enhances inter-island diversity. With $M \rightarrow \infty$ islands start to behave more similarly to each other, following a theoretical evolutionary path. This results in a low inter-island diversity creating similar partial solutions, which cannot be recombined effectively. With smaller M stochastic effects

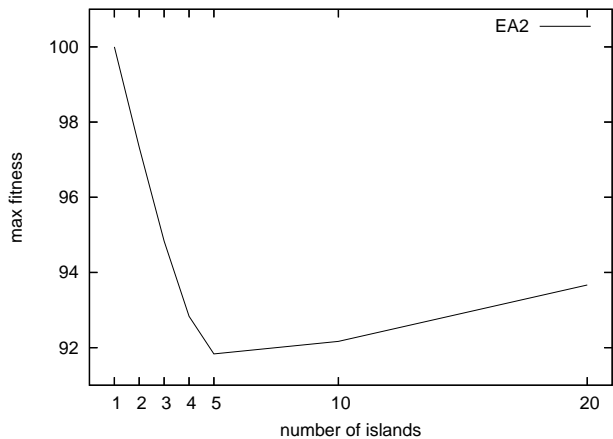


Fig. 4. Increasing N for IM-Trap function worsens the performance

play much bigger role. In IMs drift may be desired, specially that it may last only until next migration.

Let us define a function IM2' (see in Fig. 1) over $[0, 1]^2$ by a formula given below. It is qualitatively similar to IM2, but one of the wide peaks is higher than another.

$$\text{IM2}'(x, y) = \begin{cases} 3 & \text{if } x, y > 0.999 \\ \max(x^2(y-1)^2, 2(x-1)^2y^2) & \text{otherwise} \end{cases}$$

It is not likely that the global optimum will be found by random sampling. Islands will more likely converge to the higher of the sub-optimal peaks and the tendency becomes stronger with the growing size of the island. The best-in-run fitness as a function of island size is shown in Fig. 5 and confirms our predictions. Additionally, we see that for smaller islands more exploitative EA-2 achieves better results but for bigger islands a higher exploration of EA-1 results in better results.




Fig. 5. Too big islands cannot converge to different optima to recombine and create the optimal solution. IM2' function

V. PARAMETER SETTINGS IN IMs

The understanding of the dual nature of IMs has implications on setting up IMs parameters and we will illustrate it

with two parameters: the number of islands and island size.

We will perform some more experiments with a function IM6, defined over $[0, 1]^n$ and given by the following formula

$$\text{IM6}(x) = \prod_{i=1}^n \frac{(x_i^2 - 0.45)^2}{0.3025}$$

This function has local optima in each corner of the $[0, 1]^n$ cube, and the global optimum in $(1, \dots, 1)$ with a value of 1. IM6 can be understood as a more complex version of IM2 function. A two-dimensional version is shown in Fig. 1. Multiplication used in the formula causes deep valleys in the landscape and therefore mutation-based improvement is highly unlikely. The easiest way for an EA to reach the global optimum is through a combination of partial solutions. Thanks to a real-valued representation, parents' genomes are mixed easily, and each gene can be treated as a module. Such setup lets us study the behavior of both local and global processes of IMs.

A. The number of islands

As already discussed, the number of islands N decide about how many areas the system converges to (or the size of "population" at the global level). Assuming that each island quickly converges, N is obviously the upper limit on the number of areas *simultaneously* analyzed in the system. If the number of islands is too small, it may be difficult for an EA to find solutions for modular or hierarchical functions that require a combination of multiple partial solutions together.

Not surprisingly experiments with the IM6 function and different number of N show that increasing the number of islands increases the maximal fitness in the run. Increasing N is useful until it reaches a level when the global optimum is usually found. Increasing N further may speed up to some degree the convergence (higher chance that one of them will find the optimum and spread it), but generally seems questionable. This is shown in Fig. 6. We have used $M = 50$.

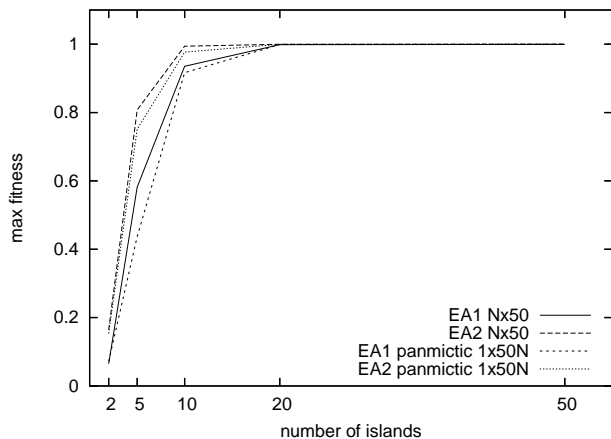


Fig. 6. Increasing N for IM6 function

Looking at diversity plots gives us some more insight. In Fig. 7(a), 7(c), 7(e) and 7(c) we see that bigger N results in a higher inter-island diversity. The inter-island diversity is then

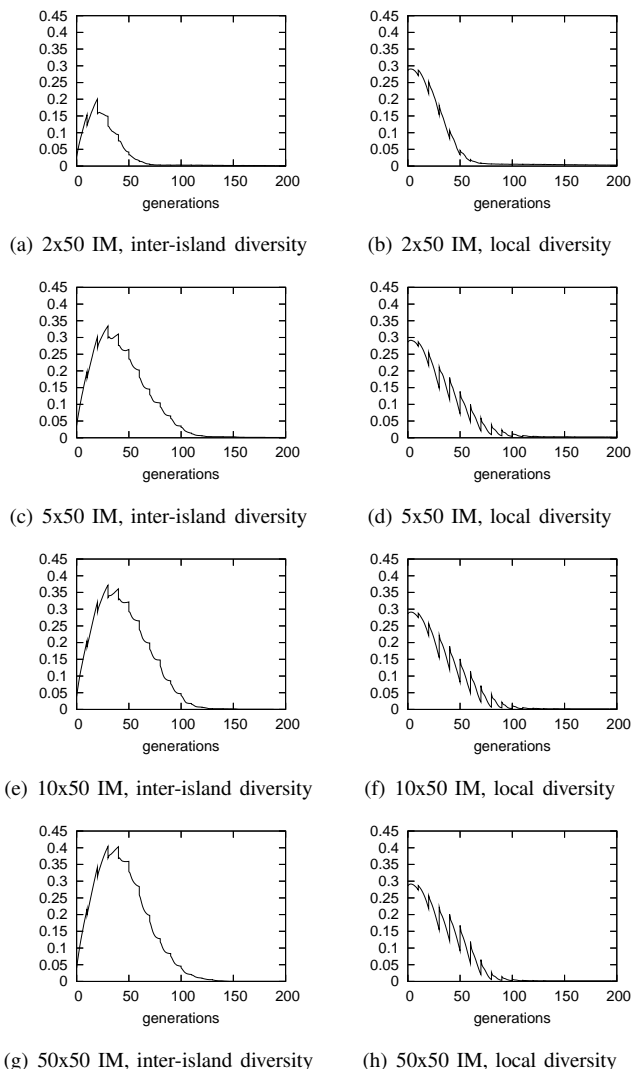


Fig. 7. Inter-island and local diversity with various N , IM6 function, EA-1

gradually "transformed" by migrations into local diversity, as seen by drops in its value occurring every migration interval. In Fig. 7(b), 7(d), 7(f) and 7(h) we see an expected repeated increase in local diversity due to migrations. We also see that for bigger N the local diversity is maintained longer. This may be counter-intuitive, because the size of islands did not change. However, a higher inter-island diversity for bigger N "helps" islands maintain their local diversity longer.

B. Island size

We have already seen that bigger M may have a negative impact on the global-level evolution. In the same time, it makes the local-level evolution produce better results, mainly due to better local exploration and a smaller impact of drift (similarly to a standard single-population EAs).

Small islands converge fast and are responsible for the more local aspect of search, whereas the inter-island evolution operates at the global level. If the domain requires a lot of local improvements however, or is deceptive with regard to details, bigger islands may be needed. Increasing

the island size, there is more space and time until convergence for storing and mixing genes coming from other islands and individuals. Such situation seems favorable for modular/hierarchical functions, where mixing sub-solutions is important.




Fig. 8. Increasing M for IM6 function

In Fig. 8 we see that for IM6 function increasing the island size has positive effect on the performance. Inter-island diversity chart is shown in Fig. 9. As expected, in all cases migrations cause a temporary decrease in this diversity. What is interesting is that for bigger islands we observe a further decrease in periods between migrations, and for smaller islands we observe an increase, which behavior ultimately results in a longer period of high diversity. This suggests that in one case the effects of migration are enhanced (inter-island diversity is further diminished, presumably by either migrants or locals dominating the other), and in the other the effects are diminished (inter-island diversity grows back, which suggests that the islands remain different, so migrants genes must have not survived). A faster drop in diversity is correlated with a higher maximal fitness values achieved for big islands. Smaller islands resulted in worse fitness of individuals, and smaller differences in fitness between individuals might caused that migrants were rarely dominating the target islands. Even though smaller islands seem to maintain diversity longer, it is probably just an artifact resulting from no fitness improvement.

Nevertheless, again we observe a transition of diversity from the global to the local level. Initially, smaller islands converge faster, and bigger ones maintain diversity. After a few migration intervals though, in all cases local diversity drops considerably. It is then small islands for which a high inter-island diversity keeps increasing the local diversities much longer.

C. The interaction between the number and size of islands

If the total number of individuals $N \cdot M$ is limited, island size is coupled with the number of islands. How should one split individuals in such case to maximize the solving ability of an IM?

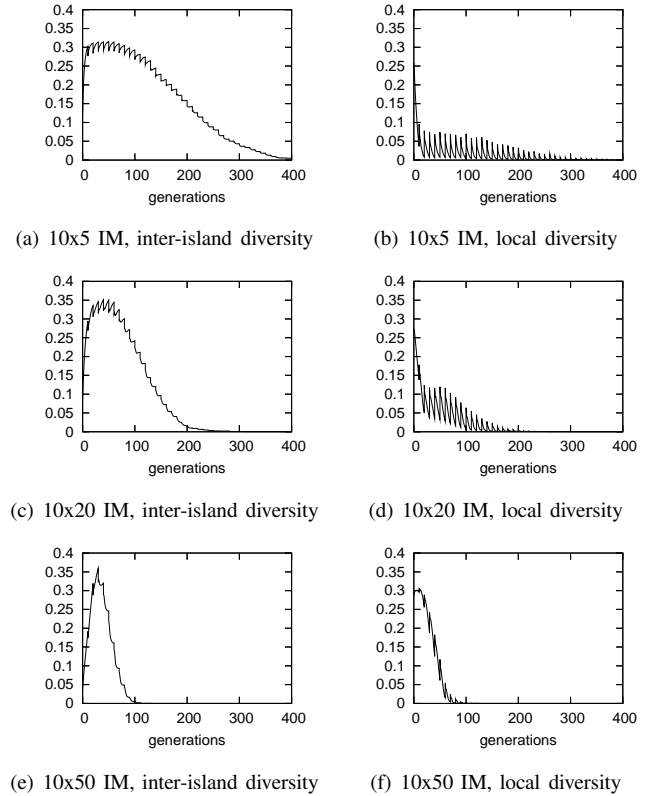


Fig. 9. Inter-island and local diversity with various M , IM6 function, EA-1

Previous analysis suggests that increasing N favors the inter-island evolution and increasing M favors the local evolution. Therefore, to keep a balance, islands cannot be too small, but also there cannot be too few of them. There is some middle ground, where the optimum is reached. Where exactly it is, it depends both on a particular problem (the difficulty at finding and combining sub-solutions) as well as the EAs used on islands.

In Fig. 10 we show some results for the IM6 function. For both EA-1 and EA-2 setups we see that a large number of relatively small islands yield the best results. Comparing between EA-1 and EA-2, the more exploitative EA-2 is able to effectively evolve solutions with smaller islands. With stronger selection small islands converge very fast anyway, so a minimal M is probably a better choice, allowing for maximizing N and focusing on inter-island evolution. A more explorative EA-1 apparently requires bigger islands for local evolution. Note that for the islands of size one local evolution becomes a random walk for EA-1 (non-overlapping generations) and a stochastic hill-climbing for EA-2.

Hierarchical if-and-only-if (H-IFF) is the original function used by Watson [10]. It operates on binary genomes. Evaluation procedure checks whether gene blocks of length 2, 4, 8, \dots , $2^{\log_2(L)}$ consist of all 1s or all 0s. Only blocks starting at positions resulting from composing proper two subblocks of lower level are checked. Each such block may contribute to fitness. Longer blocks can contribute only if their sub-blocks also contributed (were homogeneous). Each next level has




Fig. 10. Maximal fitness when keeping $N \cdot M$ constant, IM6 function.

twice less potential blocks, but its weight is doubled, so it may have the same contribution to the general fitness. H-IFF is a hierarchical function and requires improvement of solutions through combining modules at multiple levels of generality.

In Fig. 11 we show results with a H-IFF function [10]. We used 64 bits and a 0.015 mutation rate. Again EA-2 shows better results, with an exception for setups with a few big islands – which matches previous observations. A stronger selection does not harm small islands that are about to converge fast anyway. However, in the case of a few big islands selection must be weaker so that local evolution takes over the task of finding good solutions.




Fig. 11. Maximal fitness when keeping $N \cdot M$ constant, H-IFF function.

Finally, in Figure 12 we show results with four other multi-modal functions often used as test functions in literature: Rastrigin, Rosenbrock, Schwefel and Griewangk functions. Please note that these are minimization problems. Their formulas are given below.

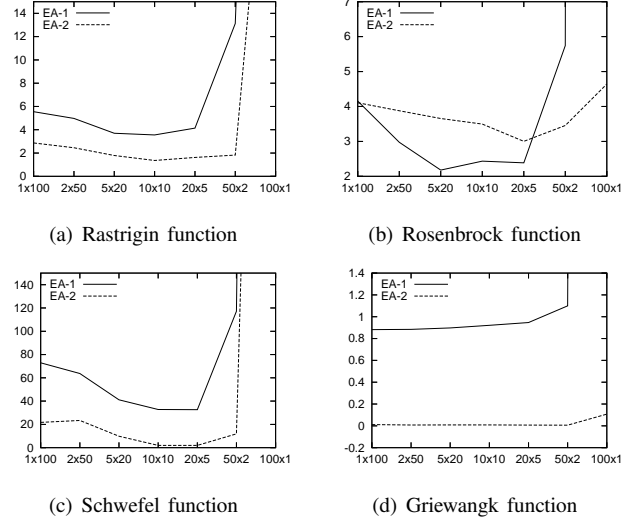


Fig. 12. Minimal fitness when keeping $N \cdot M$ constant

$$\text{Rosenbrock}(x) = \sum_{i=2}^n (100(x_i - x_{i-1}^2)^2 + (1 - x_{i-1})^2)$$

$$\text{Schwefel}(x) = \sum_{i=1}^n -x_i \sin(\sqrt{|x_i|})$$

$$\text{Rastrigin}(x) = 10n + \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i))$$

$$\text{Griewangk}(x) = 1 + \sum_{i=1}^n \frac{x_i^2}{4000} - \prod_{i=1}^n \cos\left(\frac{x_i}{\sqrt{i}}\right)$$

These experiments confirm that the best fitness is usually achieved for many small islands. One must be however careful of a sudden drop in performance when islands become too small. Not surprisingly, results for extreme 100x1 setups often lay far outside the charts' boundaries.

VI. DISCUSSION AND CONCLUSIONS

In this paper we have shown that adopting a two-level perspective on IMs allows for a better understanding of their behavior and provides guidance for setting IMs parameters. By understanding the interaction between local processes operating inside islands and global evolutionary process at the inter-island level, a better understanding of the overall behavior of IMs is obtained.

An important aspect of this interaction was obtained by analyzing diversity changes, which illustrated how IMs first “accumulate” inter-island diversity, and then slowly “release” it to help the local-level evolution. Because the whole process occurs gradually, each individual has a better chance to contribute to evolution, unlike in single-population models where a fast decrease of diversity occurs. Examples of when the two levels of evolution disturb each other were also shown.

Both the global and the local processes, while generally complementary, may contribute with different strength

toward the overall result. A proper balance between the two levels allows for improving performance. Understanding their dynamics for a given problem, helps us correctly set IMs parameters. Below we summarize concrete guidelines resulting from our experiments.

The number of islands affects evolution mainly at the global level, and the size of islands affects local evolution. With a fixed budget for evaluations these parameters are coupled. Contrary to a standard approach of setting up a few big islands, we saw that many small islands may be more beneficial. An $N = M$ balance seems reasonable, giving equal chances to both levels of evolution. Fixation in islands is not a big problem any more, because migration brings local diversity up again, and in fact it is the inter-island evolution that may be more sensitive, because there is no standard mechanism effectively increasing its diversity back. For simpler functions even a $N > M$ setup may occur the best. Note that technically nothing prevents from setting up more than one island on a single machine.

Selection pressure plays an important role too, and the $N \times M$ setup should be correlated with a proper choice of the island EA. We have shown that more exploitative EA requires using more islands which can be smaller. In the view of the two-level model, this can be explained by global evolution becoming responsible for a major part of the exploration task. A more explorative EA requires bigger islands, but possibly fewer of them. In such case, because local evolution is carried out by a more balanced EA, global evolution only serves to exchange bigger parts of solutions.

Although not studied directly in this paper, the two-level evolution approach helps us also predict a relation between migration level and required island EA. A higher level of migration (bigger migrations, smaller migration interval, denser topology, stronger migration policy) results usually in a faster global convergence. To keep the evolvability at the global scale, this convergence should be opposed by an increased global exploration. This in turn means that populations should either converge to new individuals or simply evolve into new regions. Thus local evolution should be more explorative, which can be achieved by a weaker selection pressure in islands.

A smaller level of migration (smaller migrations, longer migration intervals, sparser topology, weaker migration policy) means a slower convergence at the global scale. This keeps inter-island diversity higher and thus more exploitation is allowed at the global scale. Islands may dominate others in the result of migration and we may use a stronger selection pressure in islands.

Finally, one must remember that we have experimented with very simple functions, and for more complex domains significant amount of resources for “local” improvement may be required – and the global scale may have a less significant function of mixing a few partial solutions. The methodology shown in this paper was to first understand the dynamics of the two levels of evolution, and then appropriately set up parameters. This methodology should help in deciding proper

values even for quite different domains.

REFERENCES

- [1] S. Wright, “The roles of mutation, inbreeding, crossbreeding and selection in evolution,” in *Proceedings of the Sixth International Conference of Genetics*, D. F. Jones, Ed., Brooklyn Botanic Garden, 1932, pp. 356–366.
- [2] S. M. Gustafson, “An analysis of diversity in genetic programming,” Ph.D. dissertation, The University of Nottingham, 2004.
- [3] Z. Skolicki and K. De Jong, “The influence of migration sizes and intervals on island models,” in *Proceedings of Genetic and Evolutionary Computation Conference - GECCO 2005*. ACM, 2005.
- [4] E. Cantú-Paz, “Designing efficient and accurate parallel genetic algorithms,” Ph.D. dissertation, University of Illinois at Urbana-Champaign, 1999.
- [5] E. Cantú-Paz and D. E. Goldberg, “Are multiple runs of genetic algorithms better than one?” in *Proceedings of the Genetic and Evolutionary Computation Conference 2003*, E. C.-P. et. al., Ed., 2003.
- [6] 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*, vol. 7, no. 1, pp. 33–47, 1999.
- [7] F. Fernández, M. Tomassini, and L. Vanneschi, “An empirical study of multipopulation genetic programming,” *Genetic Programming and Evolvable Machines*, vol. 4, no. 1, pp. 21–51, March 2003.
- [8] K. A. De Jong, *Evolutionary Computation: A Unified Approach*. Cambridge, MA: MIT Press, 2006.
- [9] R. W. Morrison, “Designing evolutionary algorithms for dynamic environments,” Ph.D. dissertation, George Mason University, 2002.
- [10] R. A. Watson, *Compositional Evolution: The impact of Sex, Symbiosis and Modularity on the Gradualist Framework of Evolution*, ser. Vienna Series in Theoretical Biology. Cambridge, MA: MIT Press, 2006.