

Migrants Selection and Replacement in Distributed Evolutionary Algorithms for Dynamic Optimization

Yesnier Bravo, Gabriel Luque, and Enrique Alba

Departamento de Lenguajes y Ciencias de la Computación
Universidad de Málaga, Málaga, España.
{yesnier, gabriel, eat}@lcc.uma.es

Abstract. Many distributed systems (task scheduling, moving priorities, changing mobile environments, ...) can be linked as Dynamic Optimization Problems (DOPs), since they require to pursue an optimal value that changes over time. Consequently, we have focused on the utilization of Distributed Genetic Algorithms (dGAs), one of the domains still to be investigated for DOPs. A dGA essentially decentralizes the population in islands which cooperate through *migrations* of individuals. In this article, we analyze the effect of the migrants selection and replacement on the performance of the dGA for DOPs. Quality and distance based criteria are tested using a comprehensive set of benchmarks. Results show the benefits and drawbacks of each setting in dynamic optimization.

1 Introduction

Dynamic Optimization Problems (DOPs) are important research challenges appearing in real life applications, many of them in fact linked to distributed systems like task scheduling, moving priorities, changing mobile environments, etc. A DOP is actually a problem where the definition changes as the solving algorithms is progressing. This forces the constant research in new techniques for tracking the moving optima over time.

Several authors have proposed the use of multiple populations for solving Dynamic Optimization Problems (DOPs) using Genetic Algorithms (GA) [4–7], with the aim of tracking the optimum changes by specializing and pursuing promising regions of the search space. However, one of the domains still to be investigated is the utilization of Distributed Genetic Algorithms (dGAs) [2, 3], characterized by decentralizing the population in demes, named *islands*, independently evolving and communicating through *migrations* of individuals. Conversely, their use should be valuable for DOPs because of the natural diversity enhancement and speciation-like features [2].

In this article, we study distinct migrants selection and replacement strategies, based on both quality (fitness) and distance based criteria. Specifically, we analyze how they affect the performance of a physically parallel dGA in a comprehensive set of DOP benchmarks. Results show the benefits and drawbacks of each strategy for addressing distinct DOP features.

The remainder of this paper is organized as follows. Section 2 provides a brief background on DOP. Section 3 introduces the dGA model. Section 4 summarizes the experimental design, and the obtained results are analyzed in Sections 5 and Sections 6. Finally, conclusions and future works are stated in Section 7.

2 Dynamic Optimization Problems

An optimization problem in which environmental data, constraints and/or objectives change over time, is named Non-stationary, Time-variant or Dynamic Optimization Problem [1]. In contrast to other types of uncertainty in real-world optimization problems, in this domain the fitness function is deterministic at any point in time, but is dependent on time t , i.e.,

$$F(X) = f_t(X). \quad (1)$$

Consequently, the algorithm for DOPs should be able to track the optimum changes over time. The following features are commonly used in literature to characterize DOPs:

- **Change frequency:** It determines how often the changes occurs. It is usually measured in number of generations between two consecutive changes.
- **Change severity:** It defines how different the fitness landscape is after a change, the higher this value the more abrupt is the change.
- **Cycle length, cycle accuracy (noise):** They characterize cyclic environments, where a finite set of states recur over time. They denote how often it takes to return to a previous state, and how close this return is to it, respectively.

These features are used later in this article to build the set of DOP benchmarks for analyzing the influence of the migration policies of dGAs for DOPs.

A few authors have previously studied the influence of the migration policies in stationary environments. Cantu-Paz [8] and Alba et al. [9] showed the benefits of sending a random individual instead of the best individual. Current multi-population approaches for DOPs have used migration policies. For instance, Oppacher and Wineberg in [5], send the elite (best) individuals from *colonies* subpopulations to a *core* subpopulation. Other policies used in literature involve a global knowledge of the entire population, like Ursem in [6], by applying the *hill-valley detection* mechanism among the best individuals of each subpopulation, named *nation*. Recently, Park et al. [7] have used two populations with different evolutionary objectives and, given the inconvenience of normal migrations, they applied crossbreeding as a means of information exchange.

However, to the best of our knowledge, no coherent and comprehensive study has been accomplished regarding migrants selection and replacement in DOPs. Our work is a contribution in this direction, with the aim at supporting the design of new approaches for DOP based on dGA models.

3 Distributed GA for DOPs

As aforementioned, the dGA [2], also known as island model, structures the population in *islands*. Each island independently evolves, usually in parallel, and communicates with the other ones through migration of individuals. The pseudo-code for each island is shown in Alg. 1, which basically consists of a standard GA with an additional communication step for migrations.

Algorithm 1 Pseudocode for an island evolution in a dGA model.

```

Initialize & evaluate the subpopulation
while not stop_condition do
  Select parents for reproduction
  Apply crossover, mutation, and evaluation operators
  Select new parents and replace the old subpopulation
  if migration_period is met then
    Select outgoing migrants & send them to the neighbor islands
    Select individuals to replace & replace them by the incoming migrants
  end if
end while

```

As you can notice from Alg. 1, two essential parameters in the specification of a dGA [2] are the criteria for selecting emigrants and for replacing existing individuals in the target subpopulation by incoming migrants, often referred to in the literature as the *migration policy*. The importance of understanding their effects on the performance of the dGA has been previously reported for static problems [8], but it gets even more important for DOPs. The reason is that they control the convergence rate of the dGA. An excessively fast convergence removes the ability to react after a change, while a very low one does not allow for adaptation to the new environmental conditions.

In this paper, we will focus on the criteria for selecting outgoing migrants and selecting existing individuals, at the target subpopulation, for replacing with incoming migrants. We test both quality (fitness) and distance based selection methods, thus influencing the behavior of the dGA from the phenotype and genotype, respectively. Table 1 shows the most representative strategies.

In particular, the most *distant* individuals are determined in the Hamming space from the centroid of each deme. In the next section we describe the experimental design used to accomplish this study.

4 Experimental Setup

Aiming to analyse the influence of the chosen migrants selection and replacement criteria (see Table 1 from previous section) we use a canonical dGA consisting of eight islands evolving homogenously. In every island, we use a simple generational GA with 64 individuals, binary-tournament selection, one-point crossover (with

Table 1: Migrants selection and replacement criteria

<i>Migration policy</i>	<i>Description</i>
<i>best – worst</i>	Good migrants replace poor individuals
<i>best – random</i>	Good migrants replace random individuals
<i>best – distant</i>	Good migrants replace distant individuals
<i>random – worst</i>	Random migrants replace poor individuals
<i>random – distant</i>	Random migrants replace distant individuals
<i>worst – worst</i>	Poor migrants replace poor individuals
<i>distance – distant</i>	Distant migrants replace the distant individuals

$pX = 1.0$), without mutation (to remove possible bias in the results). Migrations occur synchronously on a unidirectional ring topology after every generation, and the migration rate is 10%.

The behavior of algorithms is tested using four dynamic functions (Onemax, Royal-Road, P-Peaks, and MMDP) built with the XOR-DOP benchmark generator [10], thus addressing different difficulties: epistasis, multimodality, and deception. We use binary strings of 100 bits, separated in 25 contiguous building blocks (BBs) of 4 bits for the Royal-Road Problem, $P = 50$ peaks for the P-Peaks Problem instance, and a MMDP with $k = 16$ deceptive subproblems. OneMax and Royal-Road are *unimodal* DOPs (only one suboptimal solution), while P-Peaks and MMDP are *multimodal* DOPs (multiple suboptimal solutions). For each problem instance, we also test distinct change modes (cyclic, cyclic with noise, and random) and change severities ($\rho \in \{0.05, 0.1, 0.2, 0.5, 1.0\}$). The higher the ρ value the severer the change; $\rho = 1.0$ means a random severity in the range $[0.01, 0.99]$. For all DOP instances the change frequency is $\tau = 10$ generations, the cycle length is 5 changes, and the noise sums a severity of 0.05.

Algorithms and benchmarks were implemented in C++, using the MALLBA library¹. All experiments were performed in a PC with an Intel Core i7-720QM processor at 1.60GHz, 4GB of RAM, and running GNU/Linux Ubuntu 12.4.

To describe the performance of algorithms we compute the mean *accuracy* or *relative error*, i.e.:

$$Acc = \frac{1}{N} \sum_{i=1}^N \frac{f(\text{generation_best}_i) - Min_i}{Max_i - Min_i} \quad (2)$$

where N is the total number of generations, and Max_i and Min_i are the current maximum and the minimum fitness values, respectively. High values of this metric indicate a better adaptation of the algorithm to the changing optimum along the run. In addition, we use the standard deviation (STD) in the Hamming space of the solutions to determine the global (inter-deme) and local (intra-deme) genotypic diversity. In both cases, we average the results of over 100 independent runs and evaluate the statistical significance. First, we use the Kolmogorov-Smirnov test to check whether the data follow a normal distribution or not. If so, then we do an ANOVA test to compare the means; otherwise Kruskal-Wallis test is

¹ Online available at <http://neo.lcc.uma.es/mallba/easy-mallba>

used to compare the medians. In each case, a level of confidence of 95 % is used. In the next section we summarize and discuss the obtained results.

5 Influence of the Migration Policy on Accuracy

We begin by analyzing the influence of migrants selection and replacement on the performance of the dGA model for dynamic environments. Our goal is to show how different migration policies can help to reach a good balance between convergence (quickly adapt to the optimum movements) and the natural speciation-like behavior (search multiple sub-optima at the same time). This conclusion confirms the importance of tracking multiple suboptimal values on dynamic optimization, since they are candidate optima after a change in the environment, thereby underlining our interest on the dGA model for DOPs.

Table 2 summarizes the results for all DOPs, change modes (**Cyclic**, **Cyclic** with **Noise**, and **Random**), and the five change severities tested. Since we are interested in the effect of the migration policies, we have grouped results with different change severities, accounting the number of experiments in which the migration policy is statistically better than the rest (values range from 0 to 5). A better value of this metric indicates a better adaptation to a wider range of change severities.

Table 2: Number of experiments in which the migration policy is statistically better than the rest. The best values are boldfaced.

<i>Migration policy</i>	<i>Onemax</i>			<i>Royal-Road</i>			<i>P-Peaks</i>			<i>MMDP</i>			Total
	Cy	CyN	Ra	Cy	CyN	Ra	Cy	CyN	Ra	Cy	CyN	Ra	
better-worst	3	1	1	3	0	0	2	0	0	0	0	0	10
better-rand	0	0	0	0	0	0	1	0	0	0	0	0	1
better-distant	0	4	4	0	4	4	0	0	0	0	0	0	16
rand-worst	4	3	2	3	2	2	3	1	1	0	0	0	19
rand-distant	1	4	4	0	4	4	0	0	0	0	0	0	17
worst-worst	5	4	2	5	4	2	5	5	5	5	5	5	52
distant-distant	0	3	4	0	4	4	0	0	0	0	0	0	15

You can notice from Table 2 that the *worst-worst* strategy (poor migrants replace poor individuals) reaches the best accuracy level in a wide range of DOP instances (52 of 60 problem instances), most notably in multimodal DOPs (*P-Peaks* and *MMDP*) where the results are significantly better than the rest in all cases. These problems consist of a large number of suboptimal solutions and require not only to search for multiple candidate optima in the search space, but also to pursue their movements. In this scenario, the *worst-worst* strategy promotes an isolated evolution among islands, which becomes basins of attraction for the low-quality migrants and, at the same time, prevents the best solution from dominating the whole population.

However, we also notice that the fitness-based replacement (including the *worst-worst* criteria) quickly degrades the accuracy of the algorithm when larger




Fig. 1: Effect of the different migration policies in the performance of the dGA for Onemax with Cyclic with Noise (a) and Random (b) change modes, and severities ($\rho = \{0.05, 0.1, 0.2, 0.5, 1.0\}$).

movements of the optimum are considered (e.g., increasing the severity, including noise, doing random changes). In these scenarios, replacing the most distant individuals is many times a better migration policy, as you can observe in (see Fig. 1) for the Onemax.

This good performance of distance-based replacement is due the enhancement in the global diversity (low coupling among islands) and the speciation-like behavior of the dGA. The diversity enables the algorithm to react and adapt to abrupt and discontinuous changes, while the speciation helps to store old optima solutions that are useful to bias the search process in the future.

Therefore, in addition to the benefits of the *worst-worst* migration policy, we underline that the replacement criterion is statistically significant and must be taken into account when dealing with DOPs. This is an interesting finding, since such parameter is often referred to as less important in the literature (see, for example, [8] or [9]). In the next section we will deeply examine this issue by analyzing the effect of the migration strategies on diversity.

6 Influence of the Migration Policy on Diversity

In this section, we analyze the influence of distinct migration policies on the population diversity. For this purpose, we compute the standard deviation (STD) among the individuals inside every island, and average the obtained values (intra-deme diversity). We also compute the STD among the centroid individuals of the islands (inter-deme diversity). Fig. 2 shows the obtained results for the MMDP with random change mode and severity $\rho = 0.5$ (the same behavior has been observed for the other problem instances).

You can notice from Fig. 2a that the distance-based replacement ensure a high global (inter-deme) diversity. However, the frequent replacement of good




Fig. 2: Effect on the population diversity for the MMDP with random (Non-cyclic) change mode and severity $\rho = 0.5$.

individuals, since they are far from the remainder members of the subpopulation, reduce the local (intra-deme) diversity (see Fig. 2b). Consequently, each island quickly converges to the local suboptima, and it is unable to track their movements after a change.

The good performance previously noticed with this strategy (see Section 5) was obtained for unimodal DOPs with high change severity, which is a hard scenario to converge, and a rapid adaptation to the first stationary periods used to succeed.

In addition, you can notice from Fig. 2b that the (*worst-worst*) migration policy provides the best overall diversity. It also exhibits the loss of diversity over time; however, this issue can be solved by considering mutation, or incorporating common dynamic optimization techniques like *random-immigrants*, *fitness sharing*, etc.

7 Conclusions

In this paper, we have analyzed the influence of the migration selection and replacement criteria, two important parameters in the specification of dGA models, for DOPs. We used a canonical version of the algorithm without mutation to remove possible bias in the results and tested combinations of both quality and distance-based migration policies. Finally, we used a comprehensive test environment based on real-world difficulties, with several change modes and severities.

On the one hand, results showed the benefits when both migrant selection and replacement are chosen the least-fit individuals of each subpopulation. The performance was notably better when addressing unimodal DOPs with small changes. However, this migration policy also exhibits the loss of global diversity, denoting inability to track the moving optima for a long time.

On the other hand, distance-based replacement strategies showed to be more robust to track larger number of local optima, improving the speciation-feature

of the population. However, it is not good to explore the changing landscape efficiently, since it produces the early loss of local diversity.

In conclusion, we have shown how migration policies allow us to balance between search of optima solutions and track of environmental changes for adapting to distinct dynamic optimization scenarios. In future works, we aim at developing adaptive or self-adaptive dGAs that exploit the main findings of this work with respect to the migration policy, thus enhancing the local behavior of the sub-populations with other techniques (random immigrants, fitness sharing, ...).

Acknowledgment

Authors acknowledge funds from the Spanish Ministry of Sciences and Innovation European FEDER, under contract TIN2011-28194 (RoadMe <http://roadme.lcc.uma.es>), and AUIP as sponsors of the Scholarship Program Academic Mobility.

References

1. T. T. Nguyen, S. Yang, and J. Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and Evolutionary Computation*, 2012.
2. G. Luque and E. Alba. Parallel Genetic Algorithms: Theory and Real World Applications. In *Studies in Computational Intelligence*, Springer, ISBN 978-3-642-22083-8, July 2011.
3. H. Homayounfar, S. Areibi and F. Wang. An Island based GA for static/dynamic optimization problems, In *3rd International DCDIS Conference on Engineering Applications and Computational Algorithms*, Guelph, Ontario, May, 2003.
4. J. Branke, T. Kaussler, C. Schmidt, and H. Schmeck. A multi-population approach to dynamic optimization problems. In *4th International Conference on Adaptive Computing in Design and Manufacture*, 299–308. Springer, 2000.
5. F. Oppacher and M. Wineberg. The shifting balance genetic algorithm: Improving the GA in a dynamic environment. In *Proc. of the Genetic and Evolutionary Computation Conference*, 504–510, San Francisco, 1999. Morgan Kaufman.
6. R. K. Ursem. Multinational GAs: Multimodal optimization techniques in dynamic environments. In *Proc. of the GECCO*, 19–26, Las Vegas, 2000. Morgan Kaufmann.
7. T. Park, R. Choe, and K. R. Ryu. Dual-population genetic algorithm for non-stationary optimization. In *Proc. of the GECCO*, 1025–1032, New York, 2008. ACM.
8. E. Cantú-Paz. Migration policies and takeover times in parallel genetic algorithms. In *Proc. of the GECCO*, 1999. Morgan Kaufmann.
9. E. Alba and J. M. Troya. Influence of the migration policy in parallel distributed GAs with structured and panmictic populations. *Applied Intelligence*, 12:163–181, 2000.
10. S. Yang. and X. Yao. Experimental study on population-based incremental learning algorithms for dynamic optimization problems. *Soft Computing*, 9(11): 815-834, November 2005.