

Studying Fault-Tolerance in Island-Based Evolutionary and Multimemetic Algorithms

Rafael Nogueras · Carlos Cotta

Received: date / Accepted: date

Abstract The use of parallel and distributed models of evolutionary algorithms (EAs) is widespread nowadays as a means to improve solution quality and reduce computational times when solving hard optimization problems. For this purpose, emergent computational environments such as P2P networks and desktop grids are offering a plethora of new opportunities but also bring new challenges: functioning on a computational network whose resources are volatile requires fault tolerance and resilience to churn. In this work we analyze these issues from the point of view of island-based EAs. We consider two EA variants, genetic and multimemetic algorithms, and analyze the impact on them of design decisions regarding the logical interconnection topology among islands and the particular fault-management policy used. To be precise, we have conducted an extensive empirical evaluation of five topologies (ring, von Neumann grid, hypercube and two kinds of scale-free networks) and four policies (including checkpoint creation and population reinitialization variants) on four benchmark problems, considering three different scenarios of increasing resource volatility. The statistical analysis of the results underlines the inherent fault-tolerance of these EAs and indicates that, while checkpointing is the most robust strategy and is superior in the most fragile topologies, a seemingly simpler guided reinitialization strategy provide statistically comparable results on the top-performing topologies, namely von Neumann grids and hypercubes.

Keywords Genetic Algorithms · Memetic Algorithm · Island Model · Fault Tolerance

This work is partially supported by Ministerio de Ciencia e Innovación project ANYSELF (TIN2011-28627-C04-01), by Junta de Andalucía project DNEMESIS (P10-TIC-6083) and by Universidad de Málaga, Campus de Excelencia Internacional Andalucía Tech.

R. Nogueras, C. Cotta
Dept. Lenguajes y Ciencias de la Computación, Universidad de Málaga, ETSI Informática,
Campus de Teatinos, 29071 Málaga, Spain
Tel.: +34 952 137158
Fax: +34 952 131397
E-mail: ccottap@lcc.uma.es

1 Introduction

Metaheuristics have steadily become during the last decades one of the weapons of choice for solving hard optimization problems. Among others, one of the salient features of these techniques is their flexibility and amenability to run on different computational environments. This is specifically true for parallel environments, on which metaheuristics in general –and population-based techniques in particular– can provide notably results in shorter times, cf. [1,3]. Not surprisingly, this has been an active field of research since the early years of the paradigm. Indeed, the deployment of population-based metaheuristics –such as, e.g., evolutionary algorithms (EAs) – on parallel systems has been around since the late 80s and early 90s, e.g., [15,37,45], and the design factors influencing their performance have been extensively studied – see, e.g., [4,8,9,41].

The last decade has witnessed the emergence of massively parallel environments such as peer-to-peer networks [30] and volunteer computing networks [39], offering vast possibilities for these techniques but also new challenges. Among the latter we can cite the inherent volatility of computational resources in these networks. The term *churn* has been coined to denote the collective effect of a plethora of peers entering or leaving the system independently along time. A parallel EA running on such an environment has to be able to cope with this effect, whereby the ideal functioning the algorithm would exhibit in a stable environment is disrupted by communication failures or by information losses resulting from computing nodes going out of the system. Fortunately, EAs have been shown to be fairly resilient to this phenomenon on a fine-grained scale such as, for example, master-slave models in which a reliable node takes care of the algorithm logic and individual-wise operations (e.g., evaluation, operator application, etc.) are distributed in the system, or even in which the latter are run in a fully decentralized way – see [21,25–27] – suggesting these techniques are inherently fault-tolerant at this level. Indeed, as shown in [20], the computing resources can be diminished due to failures in up to one order of magnitude without major impact on the quality of the results. This fact notwithstanding, different fault-management policies have been used in this fine-grained context, such as redundancy [12] or epidemic algorithms [13] to cite a few – see [17] and references therein for a recap of some of these approaches as well as frameworks supporting them. The situation admits different perspectives when a more coarse level of information distribution is considered. In particular, we can consider the distribution of islands or population demes across the system. This situation has been tackled for example in [29], wherein a checkpointing policy (namely saving periodically a snapshot of the state of a task to avoid having to restart computation from scratch) is proposed. More precisely, they propose a mixed architecture in which islands run on stable client nodes and costly individual operations such as local search are done on volatile workers, and emphasize the need for an effective management of the fault-tolerance issue. A related take on this is provided precisely in [17]. Therein, the robustness of an island-based EA with ring topology is studied: whenever a fault takes place a whole island is disconnected from the ring, which is re-wired to close the gap. Even with such a simple fault management policy, the EA only exhibits a minor degradation of performance, providing results of similar quality to a non-faulty EA.

Spurred on by these results, we approach the fault-tolerance issue in this work from different perspectives. On one hand, we intend to study the influence that

factors such as the underlying topology of the island-model or the actual fault-tolerance policy used have on the performance of the algorithm, thus shedding some light on the impact that these design decisions have. On the other hand, we consider two different population-based algorithms, a genetic algorithm and a multimemetic algorithm, so as to analyze the inherent resilience of each of these techniques. Next section provides a detailed account of the algorithmic setting and the experimental framework we have considered for this purpose.

2 Materials and Methods

2.1 Algorithmic Setting

One of the goals of this work is to provide a comparative account of the fault tolerance of two different evolutionary algorithms. In particular, we have considered genetic algorithms (GAs) and multimemetic algorithms (MMAs) [19]. Let us start by describing the latter. MMAs are a subclass of memetic algorithms [16, 31, 32] which explicitly handle *memes* (representing in this case a local search operator), having them evolve alongside solutions. They thus provide a self-adaptive search approach. The precise MMA considered is inspired by the work of Smith [42, 43] wherein each individual in the population carries a binary genotype and a single meme representing a rewriting rule $A \rightarrow C$, where both A and C are patterns of a certain length. A ternary alphabet $\Sigma = \{0, 1, \#\}$ is used to express these patterns. Here, ‘#’ is used as a wildcard symbol. Thus, given a genotype $b_1 b_2 \dots b_n$, a meme $a_1 \dots a_r \rightarrow c_1 \dots c_r$ could be applied on any segment of the former matching the antecedent (i.e., on any position i for which $b_i b_{i+1} \dots b_{i+r-1} = a_1 \dots a_r$). This application would result on the substitution of this genotypic segment by the consequent (i.e., letting $b_i b_{i+1} \dots b_{i+r-1} \leftarrow c_1 \dots c_r$). For these purposes, the wildcard symbol ‘#’ is taken as a don’t-care symbol in the antecedent (hence it can match both ‘0’ and ‘1’) and as a don’t-change symbol in the consequent (hence leaving the corresponding position unchanged). For instance, let a genotype be 11011101, and let a rule be $10\# \rightarrow 1\#0$. A possible application of the rule could be the following:

$$1 \overbrace{101}^A 1101 \xrightarrow{\text{rule}} 1 \underbrace{100}_C 1101$$

The meme could also be applied on the sixth position, resulting in the genotype 11011100. To avoid positional bias, the order in which application points are sought is randomized. Upon finding a match, the resulting genotype is generated and evaluated according to the objective function under consideration. The total cost of the process is kept under control by using a parameter w , which determines the maximal number of meme applications. The best neighbor generated (if better than the current solution) is kept. Note also that the length of each meme is not fixed but evolves itself, increasing or decreasing by one with probability p_r within a certain length range $[l_{\min}, l_{\max}]$.

Apart from the use of memes embedded within individuals, this MMA otherwise resembles a generational memetic algorithm in which parents are selected using binary tournament, and recombination, mutation and local-search (conducted

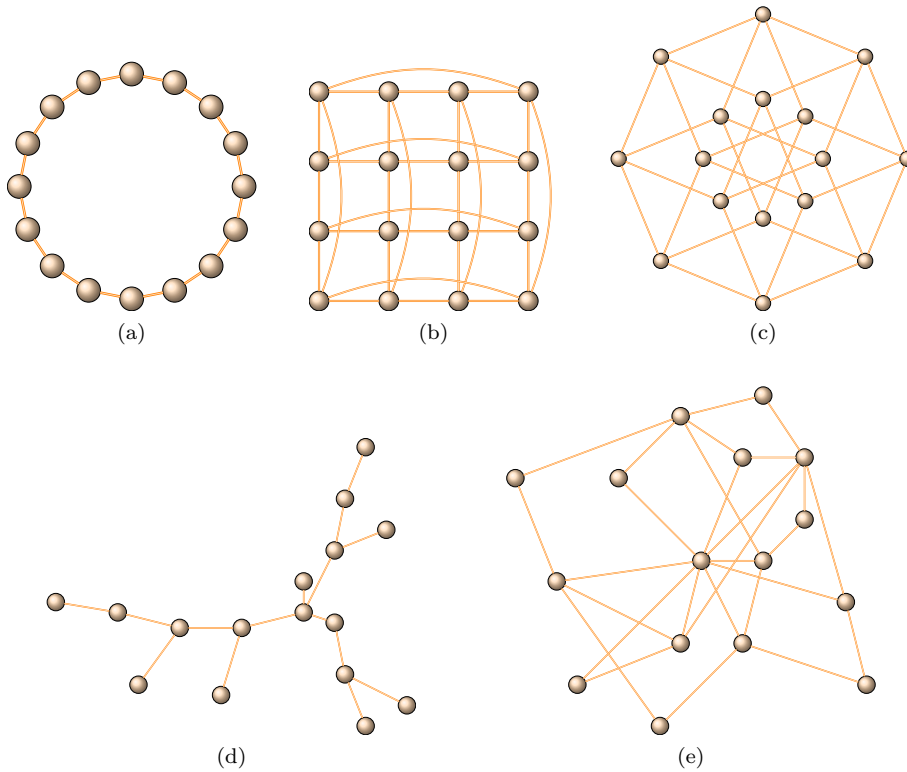


Fig. 1: Topologies considered for the interconnection of islands. (a) Ring (b) Toroidal with von Neumann neighborhood (c) Hypercube (d) Scale-free ($m = 1$) (e) Scale-free ($m = 2$)

using the meme carried out by the individual as illustrated before) are used to generate the offspring, which replaces the worst parent, following the model presented in [33]. Notice that by disabling meme evolution and the use of local search, the algorithm reduces to a GA. This GA model is also included in the experimentation.

Both EAs have been deployed on an island-based architecture, whereby multiple independent populations (islands) are initially created and subsequently evolve in partial isolation. The definition of this architecture involves some design decisions regarding (1) and interconnection topology among islands and (2) a collection of migration parameters determining how and when information is exchanged among islands [4,9,41]. Regarding the former, we have considered the following topologies:

- *Ring*: islands are arranged in a logical ring so that island i can communicate with island $i - 1$ and island $i + 1$ (all operations module the number of islands so as to close the ring).
- *Von Neumann* (VN): islands are arranged in a toroidal $p \times q$ grid, and connections are assumed between each island and the islands located at a Manhattan distance of ρ (the neighborhood radius) at most. For $\rho = 1$ this implies each

island communicates with the islands located north, south, east and west from its position.

- *Hypercube* (HC): islands are arranged on the vertices of a κ -dimensional hypercube (hence, 2^κ islands are assumed to exist), and connections are assumed along the edges of the hypercube, namely between vertices whose indices expressed in binary differ in just one bit.
- *Scale-Free* (SF): this is a non-regular complex topology commonly observed in many natural and social processes, in which node degrees exhibit a power-law distribution. We generate this kind of topology using the Barabási-Albert model [5], whereby a network is grown by adding a new node at a time. This node is connected to m existing nodes, selected with a probability proportional to their current degree (the model is thus driven by preferential attachment [7]).

The aforementioned topologies are illustrated in Fig. 1. Regarding the migration strategy, we follow previous results in the context of GAs and MMAs [4,34] indicating that selecting a random individual as migrant and having it replace the worst individual in the target population is a robust strategy. Note finally that migration is performed synchronously on active islands. While this could not correspond with real environments, it is simpler and helps as a proof-of-concept.

2.2 Fault Simulation and Management

The island-based model described in the previous section is deployed on a simulated distributed system composed of n nodes whose availability changes dynamically. More precisely, we assume all n nodes are initially available but any of them can abandon the system during the run (i.e., a fault), probably returning at a later point (i.e., a recovery) or maybe not becoming available again at all. There are several options to model these faults/recoveries. The easiest approach is considering independent faults/recoveries, namely assuming they happen with a certain probability p_f per time unit. In this case, availability times are exponentially distributed. In a more general (and realistic) situation, faults/recoveries are not the result of a memoryless –i.e., time-independent– process and hence availability times follow a different distribution such as, for example, the Weibull distribution [48]. This distribution is often utilized in survival analysis to model individual lifetimes or time-to-failure in mechanical devices [23], and provides a nice generalization of the exponential distribution: hazard rates (the rate at which faults take place) need not be constant anymore; indeed, these rates can increase or decrease with time according to a certain shape parameter η . For $\eta = 1$ we obtain the exponential distribution mentioned before, whereas for $\eta > 1$ (resp. $\eta < 1$) the hazard rate increases (resp. decreases) with time – see Fig. 2 (left). We actually consider $\eta > 1$, which implies that the longer a node has been available, the more likely it is that it goes down in the near future. In mathematical terms, the probability of a node being available up to time t_1 given that it was available up to time t_0 is

$$p(t_0, t_1, \eta, \beta) = e^{-[(t_1/\beta)^\eta - (t_0/\beta)^\eta]} \quad (1)$$

where η and β are respectively the shape and scale parameters of the distribution. Under this model, the mean availability stint is given by $\beta\Gamma(1+1/\eta)$, where $\Gamma(\cdot)$ is




Fig. 2: (a) Weibull-distributed node availability probabilities for some configurations used in the experimentation. Values of $k = 5, 10, 20$ correspond to scale parameter $\beta \simeq 79.5, 159.5, 319.5$ – see Section 3. (b) Example of the evolution of the number of available nodes in a simulation of the system.

the gamma function. For simplicity, we assume this same model for node recoveries. Note that this fits for example the case of volunteering computing systems, in which computer nodes are contributed when idle, since it is known that the duration of some human tasks on the computer follows a Weibull distribution, e.g., see [24, 44]. An example of how the number of available nodes fluctuates using this model is shown in Fig. 2 (right).

When a fault takes place, the corresponding node leaves the system. Unlike [17], we do not re-wire the interconnections among islands in this case. On one hand, while how to perform such a re-wiring is conceptually straightforward in cases such as the ring topology, it is not clear how it should be done in general, preserving the properties of the underlying topology (i.e., without having this topology reducing to something completely different after a few reconnections, thus introducing noise in the analysis). On the other hand, keeping the topology static allows for a more focused study on the inherent robustness of each interconnection scheme and is a conceptually simpler strategy. Our management policies thus manifest themselves when a node that was down becomes available again (in some cases using information gathered during the previous functioning of the node). To be precise, we have considered the following possibilities:

- **no action:** any node which becomes inactive is subsequently ignored, even if it re-enters the network. This is the simplest option and can be used as a baseline to gauge the remaining strategies.
- **checkpoint:** the functioning of each node is monitored and checkpoints are periodically created to store its state. When a node becomes available again, its activity resumes from the last checkpoint saved. This is the most classical

approach to fault management in a general sense. In our experimentation we consider that such a checkpoint is created in each generation of the EA.

- **reinitialization**: islands located at nodes which become available again are reinitialized in some way. We consider the following two reinitialization strategies:
 - **random**: the island is created from scratch, much like it was done at the beginning of the run.
 - **probabilistic**: once a node is reactivated, its active neighbors –according to the topology used– are used to re-build the new island (if there were no adjacent active populations, a random reinitialization would be done). More precisely, a probabilistic model is built using the adjacent populations, and subsequently sampled to create the new population, in the spirit of estimation of distribution algorithms (EDAs) [22, 28, 36]. In our experimentation we have considered the use of COMIT [6] for this purpose. This is an EDA that creates a probabilistic distribution using bivariate dependencies, and was shown to provide good results in the context of multimemetic optimization [35]. See Appendix A for details. The rationale of this strategy is to re-create a population state which places the island at an analogous level of neighboring demes in the current search stage, hence avoiding re-starting from scratch.

The main advantage of these reinitialization strategies is the fact that they do not require active monitoring for checkpoint creation. Hence, they can be arguably simpler to deploy and may introduce less overhead.

Once the algorithms and fault-management strategies considered have been presented, next section describes the test suite used in the experimentation in order to evaluate these strategies.

2.3 Test Suite

In order to test the robustness of the different topologies and the performance of the fault-management strategies, we have used a test suite comprising four different problems defined on binary strings, namely Deb’s trap (TRAP) function [10], the Massively Multimodal Deceptive Problem (MMDP) [14], and Watson et al.’s Hierarchical-if-and-only-if (HIFF) and Hierarchical-Exclusive-OR (HXOR) functions [46]. These are described below.

Fully-deceptive trap function Deb’s 4-bit fully deceptive function (TRAP henceforth) has a single global optimum surrounded by low-fitness solutions and a local optimum surrounded by increasingly good solutions. Hence, gradient-based methods are deceived to follow the path towards this local optimum. In mathematical terms, TRAP is defined as:

$$f(b_1 \cdots b_4) = \begin{cases} 0.6 - 0.2 \cdot u(b_1 \cdots b_4) & \text{if } u(b_1 \cdots b_4) < 4 \\ 1 & \text{if } u(b_1 \cdots b_4) = 4 \end{cases} \quad (2)$$

where $u(s_1 \cdots s_i) = \sum_j s_j$ is the unitation (number of 1s) of the binary string. A higher-order problem is built by concatenating k 4-bits blocks, and defining the fitness of this $4k$ -bit string as the sum of the function value for all blocks/subproblems. In our experiments we have considered $k = 32$ subproblems (and hence $opt = 32$).

Massively Multimodal Deceptive Problem The MMDP is a bipolar deceptive function with two global optima located at extreme unitation values (and hence far apart from each other), and with a local deceptive attractor halfway between them. This location of the deceptive attractor results in massively more local optima than global optima (i.e., $\binom{L}{L/2}$ local vs 2 global). The basic MMDP is defined for 6-bit blocks as follows:

$$f(b_1 \cdots b_6) = \begin{cases} 1 & u(b_1 \cdots b_6) \in \{0, 6\} \\ 0 & u(b_1 \cdots b_6) \in \{1, 5\} \\ 0.360384 & u(b_1 \cdots b_6) \in \{2, 4\} \\ 0.640576 & u(b_1 \cdots b_6) = 3 \end{cases} \quad (3)$$

We concatenate k copies of this basic block to create a harder problem. We have considered $k = 24$ (thus, $opt = 24$).

Hierarchically consistent functions The hierarchically consistent test problems are recursive epistatic functions defined for 2^k -bit strings. They use two auxiliary functions, namely $f : \{0, 1, \times\} \rightarrow \{0, 1\}$ and $t : \{0, 1, \times\} \rightarrow \{0, 1, \bullet\}$, the first one being used to score the contribution of building blocks, and the second one to capture their interaction. In the case of the Hierarchical if-and-only-if (HIFF) function f and t are defined as:

$$f(a, b) = \begin{cases} 1 & a = b \neq \bullet \\ 0 & \text{otherwise} \end{cases} \quad t(a, b) = \begin{cases} a & a = b \\ \bullet & \text{otherwise} \end{cases}$$

These two functions are used as follows:

$$\text{HIFF}_k(b_1 \cdots b_n) = \sum_{i=1}^{n/2} f(b_{2i-1}, b_{2i}) + 2 \cdot \text{HIFF}_{k-1}(b'_1, \dots, b'_{n/2}) \quad (4)$$

where $b'_i = t(b_{2i-1}, b_{2i})$ and $\text{HIFF}_0(\cdot) = 1$. The Hierarchical-XOR (HXOR) works similarly but changing f so as to provide a fitness contribution of 1 when $a = 1$ and $b = 0$ or vice versa, and having in that case $t(a, b) = a$ (and $t(a, b) = \times$ otherwise). We have considered $k = 7$ (i.e., 128-bit strings, $opt = 576$).

3 Experimental Results

The experiments have been realized with island-based GAs and MMAs comprising $n = 16$ islands interconnected by each of the topologies considered. In the case of the SF topology, we consider values $m = 1$ (SF₁) and $m = 2$ (SF₂) when generating the network (generated anew in each run of the algorithm). As to the VN topology, we consider a 4×4 toroidal grid. Each island has a population size of $\mu = 16$ individuals, and follows a generational reproductive plan with binary tournament for parent selection, one-point crossover ($p_X = 1.0$), bit-flip mutation ($p_M = 1/\ell$, where ℓ is the number of bits, 144 for MMDP and 128 for the remaining functions), local-search if required (conducted using the meme linked to the individual, $w = 1$) and replacement of the worst parent [33]. Offspring inherit the meme of the best parent, which is subsequently subject to length adaptation with probability p_r and mutation with probability p_M . We consider meme lengths bounded by $l_{\min} = 3$ and

$l_{\max} = 9$, and use $p_r = 1/l_{\max}$ for length self-adaptation. Migration (performed as indicated in Section 2.1) is done every 20 generations, thus allowing for a reasonable lapse of isolated evolution in each deme. Aiming to compare strategies/topologies, the computational time to reach the optimum is used in some related works with distributed EAs, cf. [2] but in this case the hardness of solving to optimality some of the test functions precludes this approach. For this reason, we have opted for measuring performance for a fixed number of 50,000 evaluations.

Regarding fault simulation, we use the model described in Section 2.2, using $\eta = 1.5$ as shape parameter to induce an increasing fault rate. The scale parameter β is used to modulate the frequency of faults. The value of this parameter has been chosen so as to have on average one node failure/recovery per k generations under an hypothetical time-independent fault distribution (the actual distribution is obviously not time-independent since $\eta > 1$, but this analogy provides an anchor to interpret the value of the parameter, if only in an approximate way). Since, there are n nodes, this implies that the survival probability per generation in this hypothetical scenario would be $p = 1 - (kn)^{-1}$, which translates into a scale parameter $\beta = -1/\log(p)$ (to a good approximation, $\beta \simeq kn - 1/2$). We consider three different scenarios of increasing volatility: $k = 20$ (low volatility), $k = 10$ (moderate volatility) and $k = 5$ (high volatility) – see Fig. 2 (right). The case $k = \infty$ would represent the fault-less execution of the algorithms. We take the best individual in the last population to measure the performance of each run. Twenty runs are performed for each problem, algorithm, configuration (island topology and fault-management strategy) and fault scenario, for a total of 10,400 runs.

Full numerical data is provided in Tables 3-10 in Appendix B. A summary of these data is provided in Fig. 3-6. Therein, fitness has been normalized using the mean fitness for $k = \infty$ (no faults) as reference, in order to get relative degradation figures. This allows aggregating the results for all four problems in the test suite to obtain a degradation curve for a specific configuration (topology, strategy) (Fig. 3 and 4). As expected, the steepest degradation is exhibited by **no action**. On the contrary, **checkpoint** offers the most graceful degradation for all topologies –Fig. 3 (b) and 4 (b)– with a loss in the most severe scenario of about 5% for the GA and just about 2% for the MMA. The two reinitialization strategies offer slightly worse results (the statistical significance of this difference shall be analyzed later on) and most interestingly, they seem to be more sensitive to the interconnection topology, in particular in the case of the MMA – see Fig. 4 (c)-(d). Specifically in the case of **probabilistic reinitialization**, it can be seen how the performance is fully competitive with that of **checkpoint** when using VN or HC topology. We attribute this effect to their richer connectivity, which allows for a more effective probabilistic modeling of neighboring islands when a node is reactivated. If we analyze the data from the complementary perspective of topologies, we can observe that in general the degradation profile is going to be more dependent on the fault-management strategy used than the other way around. This is true both for the GA –Fig. 3 (e)-(i)– and the MMA –Fig. 4 (e)-(i)– although notice that in the latter case the difference between **checkpoint** and the two reinitialization strategies is less marked than for the GA with the same topology. This can be due to the enhanced search capabilities of the MMA that allow reactivated demes to catch-up more easily with their neighbors. Moreover, while VN, HC and SF₂ seem to provide robust performance just in combination with **checkpoint** for the GA, they also provide competitive results in

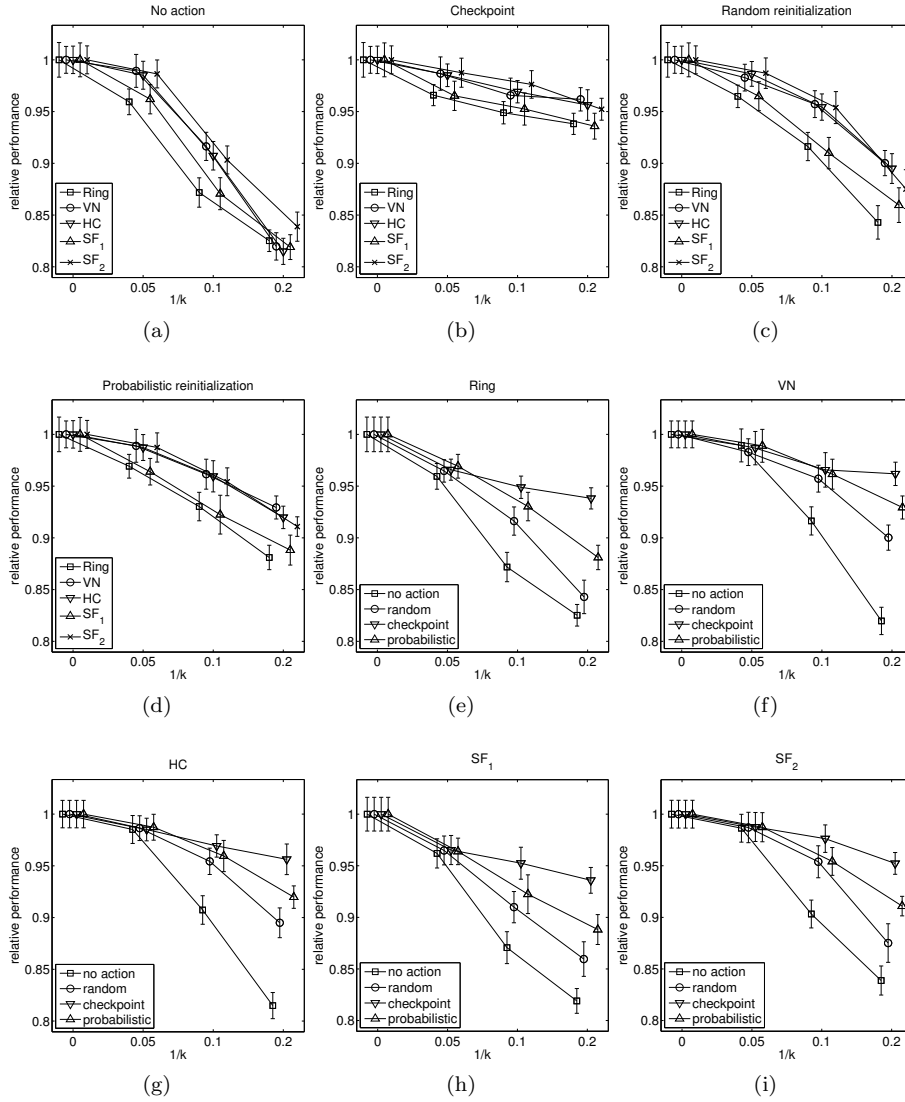


Fig. 3: Fitness degradation in island-based GAs for specific fault-management policies (a)-(d) and topologies (e)-(i).

the case of the MMA in combination with **probabilistic reinitialization** – Fig. 4 (f)(g) and (i). Let us also note en passant that there are a couple of instances in which the degradation profile is slightly non-monotonic. This is due to the stochastic nature of these techniques and the disturbance introduced by the volatility of the environment and the fault-management policy introduced (as pointed out in [40], faults can also play a role in increasing diversity in the GA search process, which may have beneficial consequences). At any rate, note that these non-monotonic

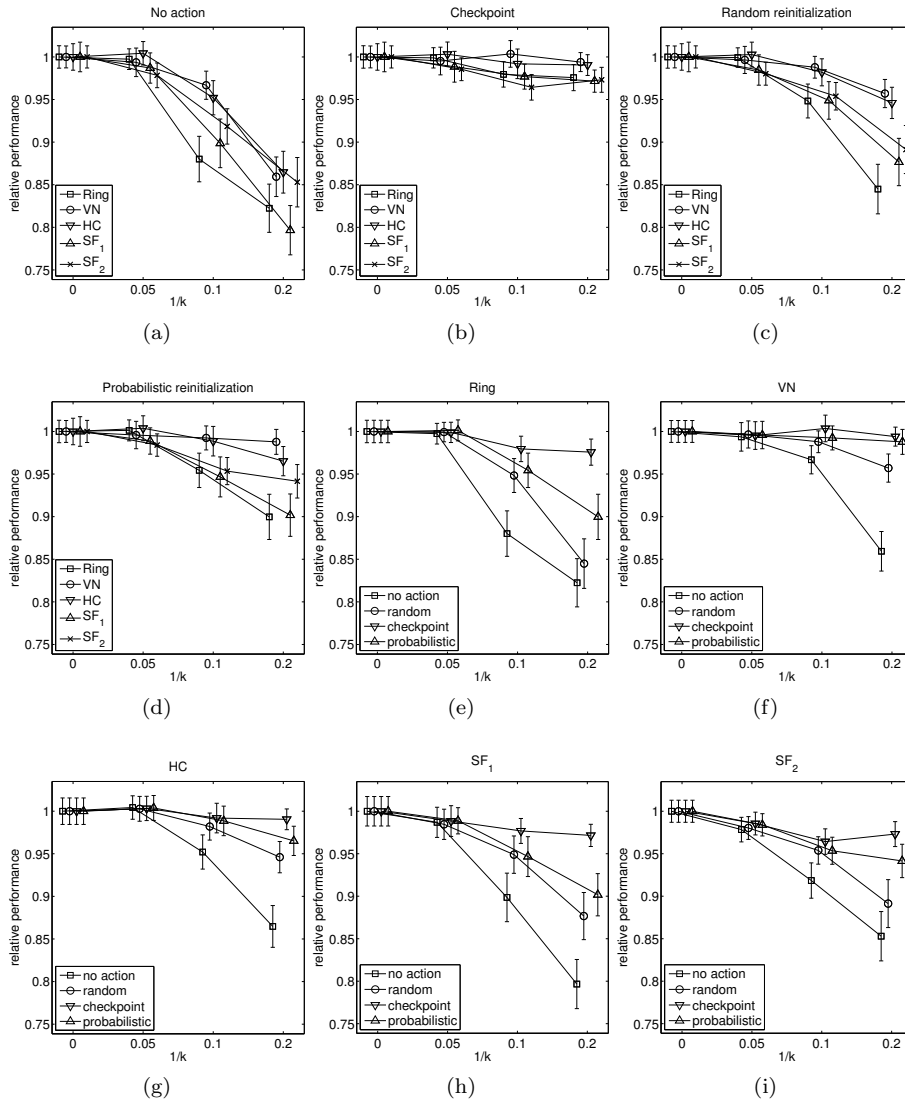


Fig. 4: Fitness degradation in island-based MMAs for specific fault-management policies (a)-(d) and topologies (e)-(i).

fluctuations are not statistically significant (see Tables 3-10). From a general perspective, Fig. 5 and 6 provide a global picture of the behavior of both algorithmic models. As advanced before, the MMA provides better performance and lower degradation than the GA for a given topology or fault-management policy. It is also not surprising to have ring and SF₁ as the topologies with the highest sensitivity to the instability of the network, since they are more prone to get partitioned in disconnected clusters as a result of node faults. This said, merely augmenting

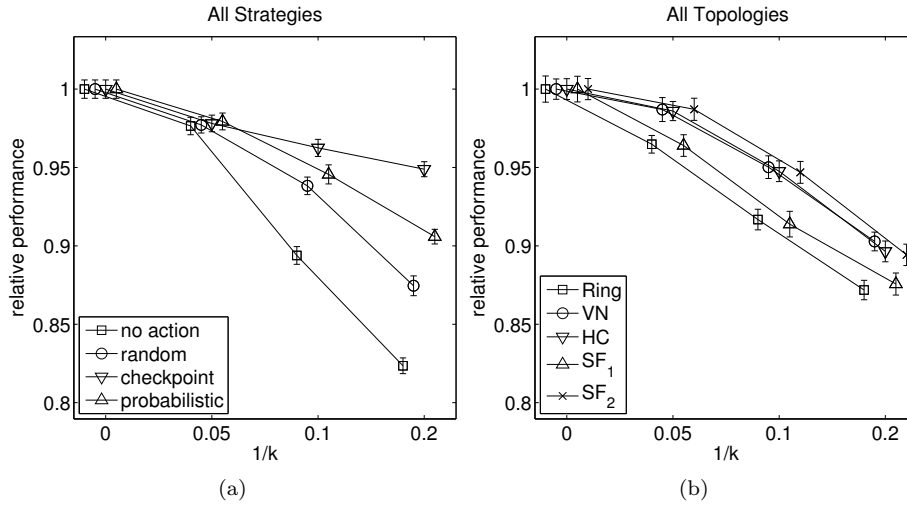


Fig. 5: Summary of fitness degradation in island-based GAs. (a) According to fault-management policy. (b) According to topology.

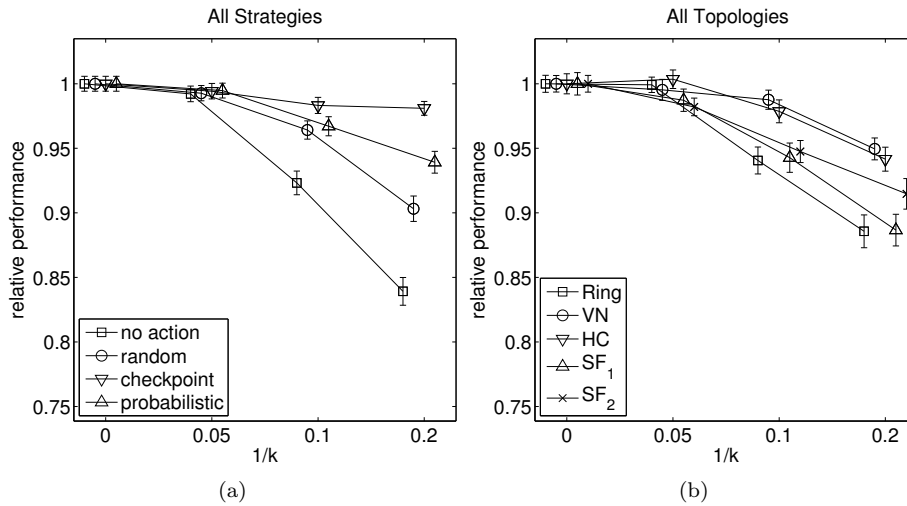


Fig. 6: Summary of fitness degradation in island-based MMAs. (a) According to fault-management policy. (b) According to topology.

the connectivity does not tell the whole story, since there are marked performance differences between the remaining topologies. This is tackled next.

Having analyzed the degradation trends, let us now focus on the statistical significance of the data and how they compare on a head-to-head basis. To do so, we perform a rank analysis, computing the relative ordering of the different

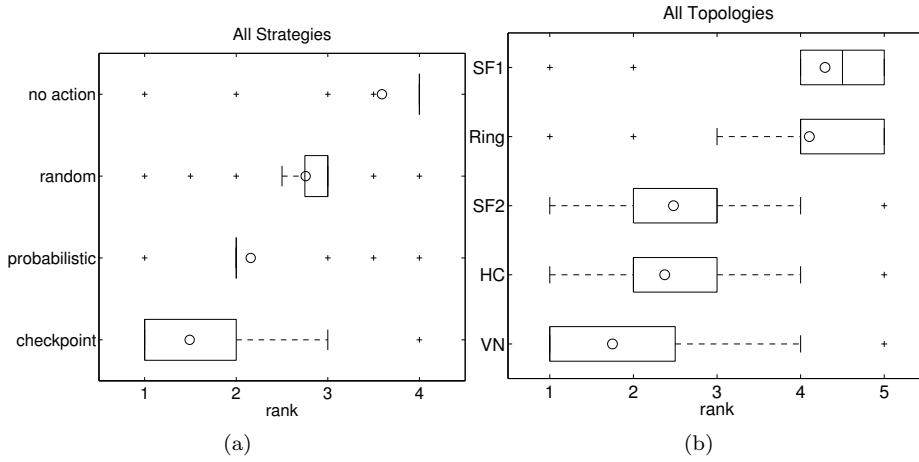


Fig. 7: Global rank distribution of performance degradation indices in island-based GAs. (a) According to fault-management policy. (b) According to topology.

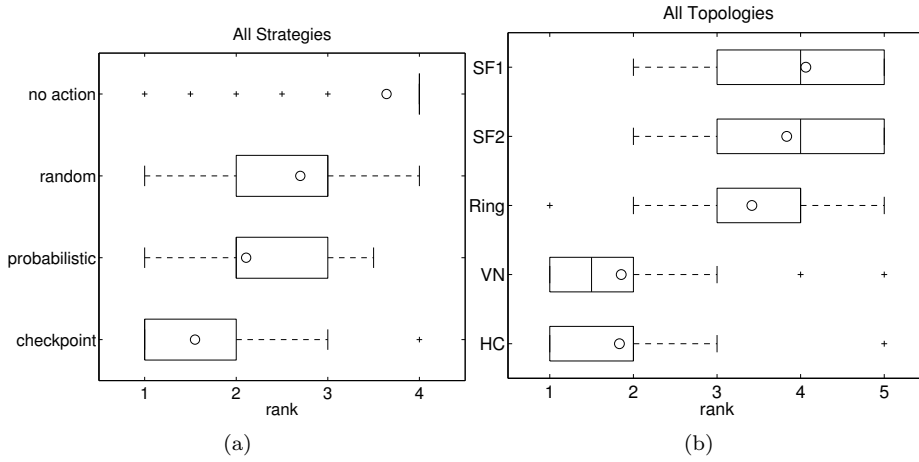


Fig. 8: Global rank distribution of performance degradation indices in island-based MMAs (a) According to fault-management policy. (b) According to topology.

strategies for a each problem, fault scenario and topology (and conversely, the relative ordering of the different topologies for each problem, fault scenario and strategy). We assign rank 1 to the strategy (resp. topology) with the best mean performance in each configuration and rank 4 (resp. rank 5) to the worst one (in case of ties the average rank of the tied positions is awarded). The distribution of ranks is shown in Fig. 7-10. Then, we have used Quade test [38] –a non-parametric test more sensitive than Friedman test, see [11]– to determine whether there are significant differences between these ranks. Table 1 shows the outcome of the test




Fig. 9: Rank distribution of performance degradation indices in island-based GAs for specific fault-management policies (a)-(d) and topologies (e)-(i).

for each of the comparisons: strategy-wise for a given topology, topology-wise for a given strategy, or global (strategy-wise for all topologies or the other way around). As can be seen, the p -values are very small in all cases, supporting the significance of the differences. Thus, a post-hoc test is in order. We have used Holm test for this purpose [18] and the results are shown in Table 2.

Let us firstly consider the global picture. Both for the GA and the MMA checkpoint is the strategy with the best rank, and this position is statistically significant ($\alpha = 0.05$), as seen in Table 2 (first column, upper half). Topology-wise, there is not a clear winner since the test is not passed for HC, and SF₂ with respect to VN in the GA, and for VN with respect to HC in the MMA. This would indicate in first instance that checkpoint is clearly the best strategy and VN/HC are the

Table 1: Results (p -values) of Quade statistical test on island-based GA/MMA performance degradation data. The upper half (resp. bottom half) corresponds to the comparison of different fault-management policies (resp. topologies) for a given topology (resp. strategy) and for all of them.

		Topology				
All	ring	VN	HC	SF ₁	SF ₂	
GA	0.00e-00	1.58e-08	1.89e-05	6.42e-06	7.81e-07	3.75e-06
MMA	0.00e-00	2.73e-06	1.46e-04	2.65e-05	9.16e-07	6.66e-07

		Fault-management policy			
All	no action	checkpoint	random	probabilistic	
GA	0.00e-00	2.29e-04	2.19e-04	2.45e-08	5.11e-08
MMA	0.00e-00	8.65e-05	2.13e-04	4.44e-06	1.17e-06

Table 2: Results of Holm test ($\alpha = 0.05$). In each entry of the table we firstly indicate in boldface the control topology/strategy. Then, we list the remaining strategies/topologies in rank order, using italics to denote items which do not pass the test with respect to the control subject (the item with the best rank) and regular face for those which do pass the test (a horizontal segment separates the former from the latter). The upper half (resp. bottom half) corresponds to different fault-management policies (resp. topologies) for a given topology (resp. strategy) and for all of them.

		Topology				
All	ring	VN	HC	SF ₁	SF ₂	
GA	checkpoint	checkpoint	checkpoint	checkpoint	checkpoint	checkpoint
	<i>probabilistic</i>	<i>probabilistic</i>	<i>probabilistic</i>	<i>probabilistic</i>	<i>probabilistic</i>	<i>probabilistic</i>
	random	random	random	random	random	random
	no action	no action	no action	no action	no action	no action
MMA	checkpoint	checkpoint	probabilistic	checkpoint	checkpoint	checkpoint
	<i>probabilistic</i>	<i>probabilistic</i>	<i>checkpoint</i>	<i>probabilistic</i>	probabilistic	<i>probabilistic</i>
	random	random	random	random	random	random
	no action	no action	no action	no action	no action	no action

		Fault-management policy			
All	no action	checkpoint	random	probabilistic	
GA	VN	VN	VN	VN	VN
	HC	<i>SF₂</i>	<i>HC</i>	<i>HC</i>	<i>HC</i>
	SF ₂	<i>HC</i>	<i>SF₂</i>	<i>SF₂</i>	<i>SF₂</i>
	SF ₁	ring	ring	SF ₁	ring
	ring	SF ₁	SF ₁	ring	SF ₁
MMA	HC	VN	HC	HC	VN
	VN	<i>HC</i>	<i>VN</i>	<i>VN</i>	<i>HC</i>
	ring	<i>ring</i>	ring	SF ₂	ring
	SF ₂	SF ₂	SF ₁	ring	SF ₂
	SF ₁	SF ₁	SF ₂	SF ₁	SF ₁

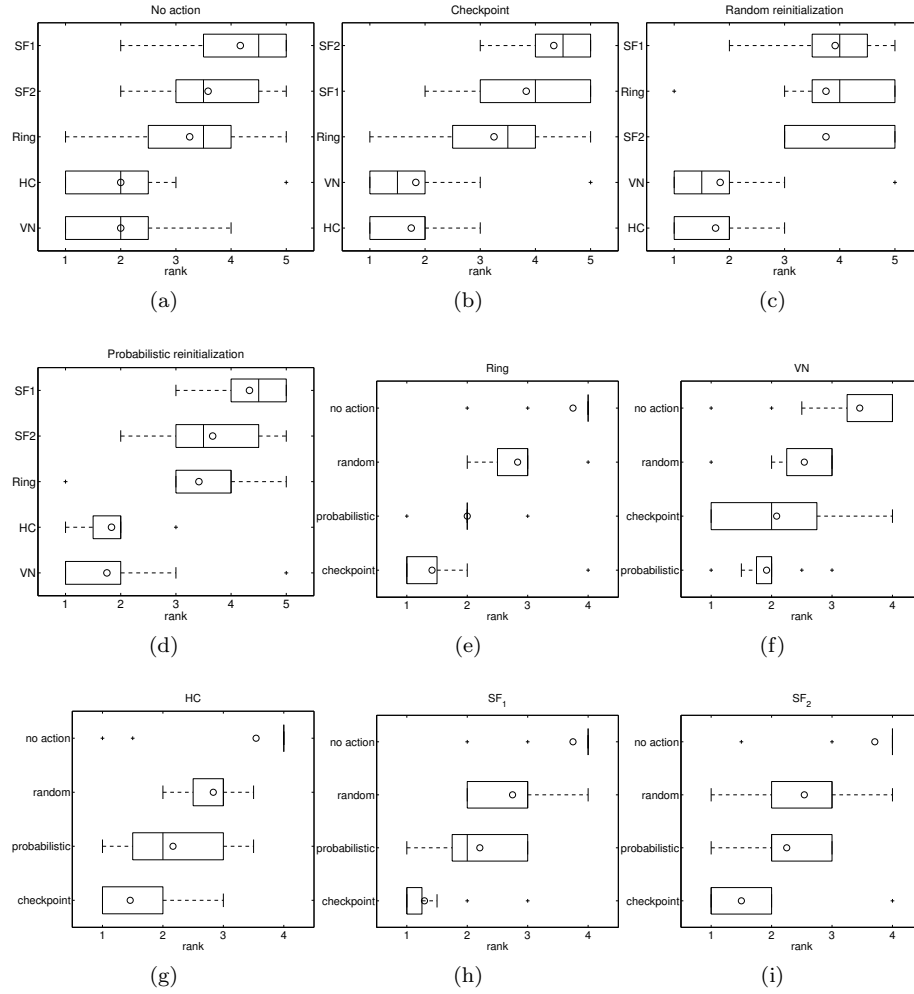


Fig. 10: Rank distribution of performance degradation indices in island-based MMAs for specific fault-management policies (a)-(d) and topologies (e)-(i).

most robust topologies. Indeed, **checkpoint** does provide top performance for most topologies and so do HC and VN for all fault-management strategies. However, if we look closely in the intersection of these two winners, the situation admits nuances: both for the GA and the MMA with VN and HC topologies, **checkpoint** does not achieve statistically significant differences with **probabilistic reinitialization** (and in fact, the latter ranks first in MMA with VN topology). Notice also how both HC and VN are not significantly different when either **checkpoint** or **probabilistic reinitialization** are used as fault-management strategy. Overall, this suggests that **checkpoint** is certainly a powerful strategy, but its global superiority is dominated by its robustness when dealing with poor performing or fragile topologies. This is certainly a strong point of this strategy, but notice that when the topology is

appropriately selected to be one of the top performing ones, **probabilistic reinitialization** offers comparable performance. This is interesting if we take into account that this latter strategy may be arguably simpler to implement in practice, since it does not require saving periodically the state of the island in some stable data store.

4 Conclusions

Fault tolerance and resilience to churn are undoubtedly important issues in island-based EAs given the increasing relevance of emerging environments for parallel computation, such as for example P2P networks or desktop grids. In this sense, we have studied here how two EA variants behave in the presence of volatile resources, taking into account design factors such as the interconnection topology and the fault-management strategy used. The results firstly indicate that this coarse-grained model is seemingly robust (thus confirming from a different perspective some results anticipated in the literature, e.g. [17]), tolerating well scenarios of low and even moderate volatility, and being amenable to the incorporation of fault-management strategies for the most unstable scenarios. Related to this, structured topologies such as von Neumann grids and hypercubes seem to provide the most graceful degradation profiles. From the point of view of the fault-management policy, a checkpointing strategy appears to be a robust choice, capable of providing top performance in most scenarios. However, a reinitialization strategy based on checking the active neighbors when a node is reactivated and using this information to build the new island has shown to be very competitive, providing results comparable to those of the checkpoint strategy without requiring to create periodic checkpoints in external stable data stores. More work is required to confirm these findings in scaled-up scenarios, but this work paves the way for conducting such studies in a more focused way. It would be interesting to analyze the impact of using other complex network topologies aside from scale-free networks, such as for instance small-world networks [47], to study comparatively their robustness (note in this sense that we cannot generalize the performance of SF networks to the whole class of complex networks). Another interesting line for future developments is deepening in the self-adaptiveness of the EA, trying to endow it with the capability to grasp the characteristics of the underlying computing environment and react accordingly, potentially introducing dynamic changes of topology or fault-management strategy.

Acknowledgements

Thanks are due to the anonymous reviewers for their helpful comments to improve the paper.

References

1. Enrique Alba. *Parallel Metaheuristics: A New Class of Algorithms*. Wiley-Interscience, 2005.

2. Enrique Alba and Gabriel Luque. Evaluation of parallel metaheuristics. In Luís Paquete, Marco Chiarandini, and Dario Basso, editors, *Workshop on Empirical Methods for the Analysis of Algorithms – EMAA 2006*, pages 9–14, Reykjavik, Iceland, 2006.
3. Enrique Alba and Marco Tomassini. Parallelism and evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 6(5):443–462, Oct 2002.
4. Enrique Alba and José M. Troya. Influence of the migration policy in parallel distributed GAs with structured and panmictic populations. *Appl. Intell.*, 12(3):163–181, 2000.
5. Réka Albert and Albert-László Barabási. Statistical mechanics of complex networks. *Review of Modern Physics*, 74(1):47–97, Jan 2002.
6. Shumeet Baluja and Scott Davies. Using optimal dependency-trees for combinatorial optimization: Learning the structure of the search space. In *14th International Conference on Machine Learning*, pages 30–38. Morgan Kaufmann, 1997.
7. Albert-László Barabási and Réka Albert. Emergence of Scaling in Random Networks. *Science*, 286(5439):509–512, 1999.
8. Erick Cantu-Paz. *Efficient and Accurate Parallel Genetic Algorithms*. Kluwer Academic Publishers, Norwell, MA, USA, 2000.
9. Erick Cantu-Paz. Migration policies, selection pressure, and parallel evolutionary algorithms. *Journal of Heuristics*, 7(4):311–334, July 2001.
10. Kalyanmoy Deb and David E. Goldberg. Analyzing deception in trap functions. In L. Darrell Whitley, editor, *Second Workshop on Foundations of Genetic Algorithms*, pages 93–108, Vail, Colorado, USA, 1993. Morgan Kaufmann.
11. Joaquín Derrac, Salvador García, Daniel Molina, and Francisco Herrera. A practical tutorial on the use of nonparametric statistical tests as a methodology for comparing evolutionary and swarm intelligence algorithms. *Swarm and Evolutionary Computation*, 1(1):3–18, 2011.
12. Christian Gagné, Marc Parizeau, and Marc Dubreuil. Distributed beagle: An environment for parallel and distributed evolutionary computations. In *Proceedings of the 17th Annual International Symposium on High Performance Computing Systems and Applications (HPSCS) 2003*, pages 201–208, Sherbrooke (QC), May 11-14 2003.
13. Maribel García Arenas, Pierre Collet, A. E. Eiben, Márk Jelasity, Juan J. Merelo Guervós, Ben Paechter, Mike Preuß, and Marc Schoenauer. A framework for distributed evolutionary algorithms. In Juan J. Merelo Guervós, Panagiotis Adamidis, Hans-Georg Beyer, José Luis Fernández-Villacañas Martín, and Hans-Paul Schwefel, editors, *PPSN*, volume 2439 of *Lecture Notes in Computer Science*, pages 665–675. Springer, 2002.
14. David E. Goldberg, Kalyanmoy Deb, and Jeffrey Horn. Massive multimodality, deception, and genetic algorithms. In *Parallel Problem Solving from Nature – PPSN II*, pages 37–48, Brussels, Belgium, 1992. Elsevier.
15. Martina Gorges-Schleuter. ASPARAGOS: an asynchronous parallel genetic optimization strategy. In J. David Schaffer, editor, *Third International Conference on Genetic Algorithms*, pages 422–427, San Francisco, CA, 1989. Morgan Kaufmann.
16. W.E. Hart, N. Krasnogor, and J.E. Smith. *Recent Advances in Memetic Algorithms*, volume 166 of *Studies in Fuzziness and Soft Computing*, chapter Memetic Evolutionary Algorithms, pages 3–27. Springer-Verlag, Berlin Heidelberg, 2005.
17. J. Ignacio Hidalgo, Juan Lanchares, Francisco Fernández de Vega, and Daniel Lombrana. Is the island model fault tolerant? In *Proceedings of the 9th Annual Conference Companion on Genetic and Evolutionary Computation*, GECCO '07, pages 2737–2744, New York, NY, USA, 2007. ACM.
18. Sture Holm. A simple sequentially rejective multiple test procedure. *Scandinavian Journal of Statistics*, 6:65–70, 1979.
19. Natalio Krasnogor, B.P. Blackburne, Edmund K. Burke, and Jonathan D. Hirst. Multi-meme algorithms for protein structure prediction. In J.J. Merelo et al., editors, *Parallel Problem Solving From Nature VII*, volume 2439 of *Lecture Notes in Computer Science*, pages 769–778. Springer-Verlag, Berlin, 2002.
20. Juan Luis Jiménez Laredo, Pascal Bouvry, Diego Lombrana González, Francisco Fernández de Vega, Maribel García Arenas, Juan Julián Merelo Guervós, and Carlos M. Fernandes. Designing robust volunteer-based evolutionary algorithms. *Genetic Programming and Evolvable Machines*, 15(3):221–244, 2014.
21. Juan Luis Jiménez Laredo, Pedro A. Castillo, Antonio Miguel Mora, Juan Julián Merelo Guervós, and Carlos M. Fernandes. Resilience to churn of a peer-to-peer evolutionary algorithm. *IJHPSA*, 1(4):260–268, 2008.

22. Pedro Larrañaga and Jose A. Lozano. *Estimation of Distribution Algorithms*, volume 2 of *Genetic Algorithms and Evolutionary Computation*. Springer US, 2002.
23. Elisa T. Lee and John W. Wang. *Statistical Methods for Survival Data Analysis*. John Wiley & Sons, Inc., Hoboken, NJ, 2003.
24. Chao Liu, Ryan W. White, and Susan Dumais. Understanding web browsing behaviors through weibull analysis of dwell time. In *Proceedings of the 33rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '10, pages 379–386, New York, NY, USA, 2010. ACM.
25. Daniel Lombrana González, Francisco Fernández de Vega, and Henri Casanova. Characterizing fault tolerance in genetic programming. *Future Generation Computer Systems*, 26(6):847 – 856, 2010.
26. Daniel Lombrana González, Juan Luis Jiménez Laredo, Francisco Fernández de Vega, and Juan Julián Merelo Guervós. Characterizing fault-tolerance of genetic algorithms in desktop grid systems. In Peter Cowling and Peter Merz, editors, *Evolutionary Computation in Combinatorial Optimization*, volume 6022 of *Lecture Notes in Computer Science*, pages 131–142. Springer Berlin Heidelberg, 2010.
27. Daniel Lombrana González, Juan Luis Jiménez Laredo, Francisco Fernández de Vega, and Juan Julián Merelo Guervós. Characterizing fault-tolerance in evolutionary algorithms. In Francisco Fernández de Vega, José Ignacio Hidalgo Pérez, and Juan Lanchares, editors, *Parallel Architectures and Bioinspired Algorithms*, volume 415 of *Studies in Computational Intelligence*, pages 77–99. Springer, 2012.
28. Jose A. Lozano, Pedro Larrañaga, Iñaki Inza, and Endika Bengoetxea. *Towards a New Evolutionary Computation*, volume 192 of *Studies in Fuzziness and Soft Computing*. Springer Berlin Heidelberg, 2006.
29. Nouredine Melab, Mohand Mezmaz, and El-Ghazali Talbi. Parallel hybrid multi-objective island model in peer-to-peer environment. In *Proceedings of the 19th IEEE International Parallel and Distributed Processing Symposium (IPDPS'05) - Workshop 6 - Volume 07*, IPDPS '05, pages 190.2–, Washington, DC, USA, 2005. IEEE Computer Society.
30. Dejan S. Milošević, Vana Kalogeraki, Rajan Lukose, Kiran Nagaraja, Jim Pruyne, Bruno Richard, Sami Rollins, and Zhichen Xu. Peer-to-peer computing. Technical Report HPL-2002-57, Hewlett-Packard Labs, 2002.
31. P. Moscato. On Evolution, Search, Optimization, Genetic Algorithms and Martial Arts: Towards Memetic Algorithms. Technical Report Caltech Concurrent Computation Program, Report. 826, California Institute of Technology, Pasadena, California, USA, 1989.
32. Ferrante Neri, Carlos Cotta, and Pablo Moscato. *Handbook of Memetic Algorithms*, volume 379 of *Studies in Computational Intelligence*. Springer-Verlag, Berlin Heidelberg, 2012.
33. Rafael Nogueras and Carlos Cotta. Analyzing meme propagation in multimemetic algorithms: Initial investigations. In *Proceedings of the 2013 Federated Conference on Computer Science and Information Systems*, pages 1013–1019, Cracow (Poland), 2013. IEEE Press.
34. Rafael Nogueras and Carlos Cotta. An analysis of migration strategies in island-based multimemetic algorithms. In T. Bartz-Beielstein, J. Branke, B. Filipič, and J. Smith, editors, *Parallel Problem Solving From Nature – PPSN XIII*, number 8672 in Lecture Notes in Computer Science, pages 731–740, Berlin Heidelberg, 2014. Springer.
35. Rafael Nogueras and Carlos Cotta. A study on multimemetic estimation of distribution algorithms. In T. Bartz-Beielstein, J. Branke, B. Filipič, and J. Smith, editors, *Parallel Problem Solving From Nature – PPSN XIII*, number 8672 in Lecture Notes in Computer Science, pages 322–331, Berlin Heidelberg, 2014. Springer.
36. Martin Pelikan, Kumara Sastry, and Erick Cantú-Paz. *Scalable Optimization via Probabilistic Modeling*, volume 33 of *Studies in Computational Intelligence*. Springer Berlin Heidelberg, 2006.
37. Chrisila B. Pettey, Michael R. Leuze, and John J. Grefenstette. A parallel genetic algorithm. In *Proceedings of the Second International Conference on Genetic Algorithms on Genetic Algorithms and Their Application*, pages 155–161, Hillsdale, NJ, USA, 1987. L. Erlbaum Associates Inc.
38. Dana Quade. Using weighted rankings in the analysis of complete blocks with additive block effects. *Journal of the American Statistical Association*, 74:680–683, 1979.
39. Luis F.G. Sarmenta. Bayanihan: Web-based volunteer computing using java. In Yoshifumi Masunaga, Takuya Katayama, and Michiharu Tsukamoto, editors, *Worldwide Computing and Its Applications – WWCA '98*, volume 1368 of *Lecture Notes in Computer Science*, pages 444–461. Springer Berlin Heidelberg, 1998.

40. Yuji Sato and Mikiko Sato. Parallelization and fault-tolerance of evolutionary computation on many-core processors. In *IEEE Congress on Evolutionary Computation*, pages 2602–2609. IEEE, 2013.
41. Zbigniew Skolicki and Kenneth De Jong. The influence of migration sizes and intervals on island models. In *Genetic and Evolutionary Computation Conference 2005*, pages 1295–1302, New York, NY, 2005. ACM.
42. James E. Smith. Self-adaptation in evolutionary algorithms for combinatorial optimisation. In Carlos Cotta, Marc Sevaux, and Kenneth Sörensen, editors, *Adaptive and Multilevel Metaheuristics*, volume 136 of *Studies in Computational Intelligence*, pages 31–57. Springer Berlin Heidelberg, 2008.
43. James E. Smith. Self-adaptative and coevolving memetic algorithms. In Ferrante Neri, Carlos Cotta, and Pablo Moscato, editors, *Handbook of Memetic Algorithms*, volume 379 of *Studies in Computational Intelligence*, pages 167–188. Springer-Verlag, Berlin Heidelberg, 2012.
44. Daniel Stutzbach and Reza Rejaie. Understanding churn in peer-to-peer networks. In *Proceedings of the 6th ACM SIGCOMM Conference on Internet Measurement*, IMC '06, pages 189–202, New York, NY, USA, 2006. ACM.
45. Reiko Tanese. Distributed genetic algorithms. In *3rd International Conference on Genetic Algorithms*, pages 434–439, San Francisco, CA, USA, 1989. Morgan Kaufmann Publishers Inc.
46. Richard A. Watson, Gregory S. Hornby, and Jordan B. Pollack. Modeling building-block interdependency. In Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors, *Parallel Problem Solving from Nature – PPSN V*, volume 1498 of *Lecture Notes in Computer Science*, pages 97–106. Springer-Verlag, Berlin Heidelberg, 1998.
47. Duncan J. Watts and Steven H. Strogatz. Collective dynamics of 'small-world' networks. *Nature*, 393:440–442, 1998.
48. Waloddi Weibull. A statistical distribution function of wide applicability. *Journal of Applied Mechanics*, 18(3):293–297, 1951.

A Probabilistic Reinitialization Model

Given m individuals $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$, each of them comprising n variables $\mathbf{x}^{(i)} = x_1^{(i)} \dots x_n^{(i)}$, the COMIT-based model used factorizes their joint probability distribution $p(\mathbf{x})$ as

$$p(\mathbf{x}) = p(x_{i_1}) \prod_{j=2}^n p(x_{i_j} | x_{i_{a(j)}}),$$

where $i_1 \dots i_n$ is a permutation of the indices $1 \dots n$, and $a(j) < j$ is the permutation index of the variable which x_{i_j} depends on. We assume i_1 is the variable with the lowest entropy $H(X_k)$ in the selected sample $\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(m)}$, and then we pick i_j ($j > 1$) as the variable that minimizes the conditional entropy $H(X_k | X_{i_s}, s < j)$. Thus, we have a tree dependence structure.

Once the dependency tree has been computed, variables are sorted in topological order and randomly sampled using the conditional probability distributions $p(x_{i_j} | x_{i_{a(j)}})$ associated to each variable.

B Numerical Data

Table 3: Results (averaged for 20 runs) of the different island-based GAs on the four problems considered (**no action**). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated. In this table and subsequent ones, a bullet (\bullet) indicates the corresponding results are significantly different (using a Wilcoxon test, $\alpha = 0.05$) to the results for $k = \infty$.

		TRAP			HIFF		
topology	k	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	28.4	28.4 ± 0.1	1	397.0	407.9 ± 10.1
	20	0	28.0	28.1 ± 0.2	0	371.0	377.2 ± 5.5
	10	0	26.2	26.3 ± 0.3	\bullet	0	325.0 ± 5.0
	5	0	25.1	25.3 ± 0.3	\bullet	0	298.0 ± 3.7
VN	∞	0	30.0	30.0 ± 0.1	0	436.0	433.0 ± 6.5
	20	0	30.0	30.2 ± 0.2	0	414.0	420.2 ± 6.7
	10	0	28.8	28.8 ± 0.2	\bullet	0	384.0 ± 5.9
	5	0	26.7	26.5 ± 0.2	\bullet	0	326.5 ± 7.2
HC	∞	0	30.3	30.3 ± 0.1	0	436.0	433.2 ± 6.6
	20	0	30.3	30.1 ± 0.1	0	428.0	429.8 ± 7.4
	10	0	28.8	28.6 ± 0.1	\bullet	0	368.0 ± 7.9
	5	0	26.6	26.3 ± 0.3	\bullet	0	323.0 ± 5.0
SF ₁	∞	0	29.0	28.9 ± 0.1	0	397.0	403.6 ± 7.6
	20	0	28.4	28.5 ± 0.2	0	380.0	383.6 ± 6.9
	10	0	26.8	26.5 ± 0.4	\bullet	0	326.0 ± 5.8
	5	0	25.0	25.0 ± 0.3	\bullet	0	309.0 ± 5.8
SF ₂	∞	0	29.8	29.8 ± 0.1	0	402.0	413.4 ± 8.1
	20	0	29.9	29.7 ± 0.2	0	400.0	411.0 ± 7.7
	10	0	28.0	27.9 ± 0.3	\bullet	0	362.0 ± 5.6
	5	0	26.3	26.5 ± 0.3	\bullet	0	340.0 ± 6.9

		HXOR			MMDP		
topology	k	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	389.0	394.8 ± 5.8	0	20.0	20.2 ± 0.1
	20	0	364.0	370.3 ± 5.7	\bullet	0	19.9 ± 0.2
	10	0	319.0	332.8 ± 6.9	\bullet	0	18.5 ± 0.2
	5	0	310.0	309.7 ± 4.2	\bullet	0	17.9 ± 0.2
VN	∞	0	436.0	435.9 ± 7.0	0	21.2	21.4 ± 0.1
	20	1	422.0	426.4 ± 9.8	0	21.5	21.5 ± 0.1
	10	0	371.0	380.9 ± 7.7	\bullet	0	20.3 ± 0.1
	5	0	332.0	329.6 ± 5.1	\bullet	0	19.0 ± 0.2
HC	∞	0	420.0	425.8 ± 7.0	0	21.5	21.6 ± 0.1
	20	0	408.0	410.9 ± 6.7	0	21.5	21.3 ± 0.1
	10	0	379.0	378.9 ± 5.9	\bullet	0	20.2 ± 0.2
	5	0	320.0	326.4 ± 5.3	\bullet	0	18.8 ± 0.2
SF ₁	∞	0	397.0	407.7 ± 8.7	0	20.0	20.2 ± 0.1
	20	0	368.0	377.7 ± 6.2	\bullet	0	19.9 ± 0.2
	10	0	342.0	341.3 ± 6.3	\bullet	0	18.6 ± 0.3
	5	0	307.0	307.4 ± 4.2	\bullet	0	18.2 ± 0.2
SF ₂	∞	0	414.0	417.9 ± 5.2	0	21.1	21.3 ± 0.1
	20	0	398.0	402.4 ± 4.5	\bullet	0	21.1 ± 0.2
	10	0	349.0	358.2 ± 6.0	\bullet	0	20.1 ± 0.2
	5	0	327.0	328.6 ± 5.1	\bullet	0	18.4 ± 0.2

Table 4: Results (averaged for 20 runs) of the different island-based GAs on the four problems considered (checkpoint). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP			HIFF				
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$		
Ring	∞	0	28.4	28.4 ± 0.1	1	397.0	407.9 ± 10.1		
	20	0	28.4	28.4 ± 0.1	0	379.0	381.1 ± 5.3	•	
	10	0	28.0	27.9 ± 0.2	•	0	365.5	368.0 ± 5.0	•
	5	0	27.4	27.5 ± 0.2	•	0	362.5	361.1 ± 4.4	•
VN	∞	0	30.0	30.0 ± 0.1	0	436.0	433.0 ± 6.5		
	20	0	30.0	30.2 ± 0.1	0	412.0	416.8 ± 7.0		
	10	0	30.0	29.7 ± 0.2	0	408.0	417.6 ± 6.3		
	5	0	29.6	29.6 ± 0.1	0	402.0	403.5 ± 5.7	•	
HC	∞	0	30.3	30.3 ± 0.1	0	436.0	433.2 ± 6.6		
	20	0	30.4	30.1 ± 0.1	0	426.0	428.0 ± 6.1		
	10	0	30.0	30.0 ± 0.1	0	406.0	411.8 ± 5.6	•	
	5	0	29.2	29.3 ± 0.1	•	1	391.0	403.9 ± 10.2	•
SF ₁	∞	0	29.0	28.9 ± 0.1	0	397.0	403.6 ± 7.6		
	20	0	28.4	28.5 ± 0.2	•	0	381.5	387.7 ± 6.9	
	10	0	28.0	28.1 ± 0.2	•	0	368.0	373.9 ± 7.8	•
	5	0	27.6	27.6 ± 0.1	•	0	359.0	365.8 ± 5.6	•
SF ₂	∞	0	29.8	29.8 ± 0.1	0	402.0	413.4 ± 8.1		
	20	0	29.6	29.7 ± 0.2	0	405.0	414.7 ± 8.0		
	10	0	29.6	29.3 ± 0.2	0	404.0	408.1 ± 6.2		
	5	0	29.0	29.1 ± 0.1	•	0	392.5	386.5 ± 5.1	•

topology	k	HXOR			MMDP				
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$		
Ring	∞	0	389.0	394.8 ± 5.8	0	20.0	20.2 ± 0.1		
	20	0	368.0	371.3 ± 3.5	•	0	20.1	20.0 ± 0.2	
	10	0	369.0	371.6 ± 4.9	•	0	19.7	19.6 ± 0.1	•
	5	0	368.0	368.9 ± 5.0	•	0	19.3	19.5 ± 0.1	•
VN	∞	0	436.0	435.9 ± 7.0	0	21.2	21.4 ± 0.1		
	20	1	418.0	425.1 ± 10.0	0	21.5	21.5 ± 0.1		
	10	1	400.0	403.3 ± 10.8	•	0	21.1	21.1 ± 0.1	
	5	0	408.0	415.1 ± 5.8	•	0	20.8	20.9 ± 0.1	•
HC	∞	0	420.0	425.8 ± 7.0	0	21.5	21.6 ± 0.1		
	20	0	408.0	412.0 ± 5.0	0	21.5	21.4 ± 0.1		
	10	0	408.0	407.1 ± 5.1	0	21.1	21.1 ± 0.1		
	5	0	404.0	406.1 ± 4.2	•	0	20.8	20.9 ± 0.1	•
SF ₁	∞	0	397.0	407.7 ± 8.7	0	20.0	20.2 ± 0.1		
	20	0	374.5	379.1 ± 6.6	•	0	19.9	19.9 ± 0.1	
	10	0	368.0	378.4 ± 6.5	•	0	20.0	19.8 ± 0.2	
	5	0	370.5	371.1 ± 6.2	•	0	19.3	19.6 ± 0.1	•
SF ₂	∞	0	414.0	417.9 ± 5.2	0	21.1	21.3 ± 0.1		
	20	0	392.0	400.7 ± 5.0	•	0	21.5	21.1 ± 0.2	
	10	0	389.0	398.6 ± 6.7	•	0	20.8	20.8 ± 0.2	•
	5	0	395.0	390.3 ± 4.8	•	0	20.4	20.5 ± 0.1	•

Table 5: Results (averaged for 20 runs) of the different island-based GAs on the four problems considered (random reinitialization). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP			HIFF		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	28.4	28.4 ± 0.1	1	397.0	407.9 ± 10.1
	20	0	28.3	28.4 ± 0.1	0	380.0	381.1 ± 4.6
	10	0	27.3	27.2 ± 0.2	0	353.0	351.1 ± 6.0
	5	0	24.9	25.0 ± 0.4	0	315.0	318.4 ± 7.2
VN	∞	0	30.0	30.0 ± 0.1	0	436.0	433.0 ± 6.5
	20	0	30.1	30.2 ± 0.1	0	406.0	414.7 ± 7.1
	10	0	29.7	29.7 ± 0.2	0	412.0	414.7 ± 8.0
	5	0	28.2	28.4 ± 0.2	0	369.0	374.5 ± 7.9
HC	∞	0	30.3	30.3 ± 0.1	0	436.0	433.2 ± 6.6
	20	0	30.2	30.1 ± 0.1	0	416.0	425.8 ± 6.7
	10	0	29.6	29.5 ± 0.2	0	402.0	407.3 ± 6.6
	5	0	28.4	28.4 ± 0.3	0	370.0	366.9 ± 7.2
SF ₁	∞	0	29.0	28.9 ± 0.1	0	397.0	403.6 ± 7.6
	20	0	28.4	28.5 ± 0.2	0	380.0	383.1 ± 6.0
	10	0	27.6	27.3 ± 0.3	0	343.5	347.1 ± 6.4
	5	0	26.2	25.7 ± 0.4	0	332.0	324.4 ± 7.0
SF ₂	∞	0	29.8	29.8 ± 0.1	0	402.0	413.4 ± 8.1
	20	0	29.8	29.7 ± 0.2	0	400.0	412.6 ± 8.6
	10	0	28.6	28.9 ± 0.2	0	388.0	392.4 ± 6.4
	5	0	27.2	27.0 ± 0.4	0	350.0	347.3 ± 8.5
topology	k	HXOR			MMDP		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	389.0	394.8 ± 5.8	0	20.0	20.2 ± 0.1
	20	0	365.0	370.9 ± 5.3	0	20.0	19.9 ± 0.1
	10	0	345.0	350.9 ± 6.5	0	19.3	19.3 ± 0.1
	5	0	319.0	317.9 ± 6.4	0	18.5	18.3 ± 0.2
VN	∞	0	436.0	435.9 ± 7.0	0	21.2	21.4 ± 0.1
	20	0	422.0	422.4 ± 6.4	0	21.5	21.4 ± 0.1
	10	0	390.0	392.6 ± 5.3	0	21.1	21.0 ± 0.1
	5	0	370.0	367.2 ± 3.0	0	20.1	20.3 ± 0.1
HC	∞	0	420.0	425.8 ± 7.0	0	21.5	21.6 ± 0.1
	20	0	408.0	417.2 ± 5.6	0	21.5	21.4 ± 0.1
	10	0	396.0	395.8 ± 5.8	0	21.0	21.0 ± 0.2
	5	0	371.0	368.4 ± 6.4	0	20.4	20.1 ± 0.2
SF ₁	∞	0	397.0	407.7 ± 8.7	0	20.0	20.2 ± 0.1
	20	0	376.0	382.2 ± 7.2	0	19.9	19.9 ± 0.1
	10	0	357.5	357.6 ± 6.2	0	19.3	19.4 ± 0.2
	5	0	337.5	337.9 ± 7.3	0	18.3	18.5 ± 0.2
SF ₂	∞	0	414.0	417.9 ± 5.2	0	21.1	21.3 ± 0.1
	20	0	392.0	400.8 ± 5.5	0	21.3	21.2 ± 0.2
	10	0	382.0	386.4 ± 7.6	0	20.9	20.7 ± 0.3
	5	0	357.0	349.6 ± 8.0	0	19.7	19.5 ± 0.2

Table 6: Results (averaged for 20 runs) of the different island-based GAs on the four problems considered (probabilistic reinitialization). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP			HIFF		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	28.4	28.4 ± 0.1	1	397.0	407.9 ± 10.1
	20	0	28.3	28.4 ± 0.2	0	379.0	380.4 ± 4.8
	10	0	27.6	27.5 ± 0.2	•	358.0	361.6 ± 5.4
	5	0	26.0	26.1 ± 0.3	•	332.0	341.4 ± 5.1
VN	∞	0	30.0	30.0 ± 0.1	0	436.0	433.0 ± 6.5
	20	0	30.1	30.2 ± 0.1	0	408.0	414.5 ± 6.8
	10	0	30.0	29.8 ± 0.2	0	406.0	414.4 ± 8.5
	5	0	29.0	28.8 ± 0.1	•	395.0	390.4 ± 6.7
HC	∞	0	30.3	30.3 ± 0.1	0	436.0	433.2 ± 6.6
	20	0	30.0	30.1 ± 0.1	0	424.0	431.0 ± 7.4
	10	0	29.8	29.7 ± 0.2	•	401.0	415.7 ± 9.0
	5	0	29.2	29.1 ± 0.2	•	377.0	380.6 ± 4.5
SF ₁	∞	0	29.0	28.9 ± 0.1	0	397.0	403.6 ± 7.6
	20	0	28.4	28.5 ± 0.2	0	382.0	382.4 ± 5.4
	10	0	27.6	27.5 ± 0.3	•	354.0	355.8 ± 8.2
	5	0	26.9	26.4 ± 0.3	•	343.0	343.3 ± 6.4
SF ₂	∞	0	29.8	29.8 ± 0.1	0	402.0	413.4 ± 8.1
	20	0	29.7	29.6 ± 0.2	0	404.0	415.9 ± 8.3
	10	0	29.2	29.0 ± 0.2	•	386.0	392.8 ± 6.4
	5	0	28.0	28.1 ± 0.1	•	370.0	368.7 ± 4.8

topology	k	HXOR			MMDP		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	389.0	394.8 ± 5.8	0	20.0	20.2 ± 0.1
	20	0	369.0	377.6 ± 5.6	•	20.0	20.0 ± 0.1
	10	0	362.5	358.0 ± 6.4	•	19.3	19.4 ± 0.2
	5	0	337.0	336.9 ± 3.8	•	18.6	18.5 ± 0.2
VN	∞	0	436.0	435.9 ± 7.0	0	21.2	21.4 ± 0.1
	20	1	428.0	430.4 ± 9.7	0	21.5	21.5 ± 0.1
	10	0	402.0	400.1 ± 6.4	•	20.8	21.0 ± 0.1
	5	0	380.0	387.2 ± 4.3	•	20.8	20.7 ± 0.1
HC	∞	0	420.0	425.8 ± 7.0	0	21.5	21.6 ± 0.1
	20	0	408.0	414.4 ± 5.5	0	21.3	21.3 ± 0.1
	10	0	390.0	393.6 ± 6.2	•	21.1	21.0 ± 0.2
	5	0	384.0	380.1 ± 5.6	•	20.4	20.4 ± 0.1
SF ₁	∞	0	397.0	407.7 ± 8.7	0	20.0	20.2 ± 0.1
	20	0	371.5	381.7 ± 6.4	•	20.0	19.9 ± 0.2
	10	0	360.0	364.4 ± 8.7	•	19.7	19.4 ± 0.2
	5	0	358.5	351.3 ± 6.0	•	18.4	18.7 ± 0.2
SF ₂	∞	0	414.0	417.9 ± 5.2	0	21.1	21.3 ± 0.1
	20	0	392.0	399.8 ± 4.6	•	21.5	21.1 ± 0.2
	10	0	384.0	385.7 ± 5.9	•	20.8	20.7 ± 0.2
	5	0	362.0	361.5 ± 3.9	•	20.1	20.1 ± 0.1

Table 7: Results (averaged for 20 runs) of the different island-based MMAs on the four problems considered (no action). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP				HIFF		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	17	32.0	31.8 ± 0.1		19	576.0	570.4 ± 5.6
	20	19	32.0	31.8 ± 0.2		19	576.0	570.0 ± 6.0
	10	17	32.0	30.0 ± 0.6	•	17	456.0	451.6 ± 20.9
	5	12	28.1	28.6 ± 0.7	•	12	387.0	411.1 ± 22.0
VN	∞	17	32.0	31.8 ± 0.1		18	576.0	565.6 ± 7.2
	20	19	32.0	31.9 ± 0.1		15	576.0	550.8 ± 10.4
	10	17	32.0	31.5 ± 0.2		18	576.0	556.6 ± 11.9
	5	13	30.8	29.9 ± 0.5	•	12	456.0	470.0 ± 18.2
HC	∞	16	32.0	31.6 ± 0.2		16	576.0	555.2 ± 9.6
	20	19	32.0	31.9 ± 0.1		18	576.0	565.2 ± 7.5
	10	17	32.0	31.7 ± 0.2		17	576.0	542.8 ± 15.6
	5	14	32.0	30.6 ± 0.5		11	456.0	472.3 ± 20.2
SF ₁	∞	19	32.0	31.9 ± 0.1		19	576.0	568.8 ± 7.2
	20	17	32.0	31.7 ± 0.2		17	576.0	554.4 ± 12.5
	10	17	32.0	30.3 ± 0.6	•	14	576.0	499.7 ± 23.0
	5	6	27.9	28.0 ± 0.7	•	11	374.5	396.9 ± 21.7
SF ₂	∞	19	32.0	32.0 ± 0.0		19	576.0	570.0 ± 6.0
	20	18	32.0	31.8 ± 0.2		17	576.0	553.0 ± 10.7
	10	16	32.0	31.1 ± 0.4		14	496.0	506.6 ± 15.5
	5	13	32.0	29.6 ± 0.7	•	13	472.0	479.9 ± 22.4

topology	k	HXOR				MMDP		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	402.5	408.1 ± 7.0		7	23.6	23.0 ± 0.3
	20	0	395.0	397.2 ± 6.5		12	24.0	23.4 ± 0.2
	10	0	351.0	354.9 ± 5.5	•	8	20.8	21.0 ± 0.5
	5	0	319.0	322.9 ± 5.4	•	2	19.0	20.2 ± 0.5
VN	∞	0	446.0	444.0 ± 7.5		12	24.0	23.5 ± 0.2
	20	1	436.0	440.3 ± 9.9		15	24.0	23.6 ± 0.2
	10	0	400.0	403.1 ± 7.0	•	11	23.6	23.1 ± 0.3
	5	0	332.0	336.9 ± 6.1	•	2	21.1	21.3 ± 0.4
HC	∞	1	440.5	447.9 ± 8.6		13	24.0	23.5 ± 0.2
	20	1	436.0	440.7 ± 8.6		14	24.0	23.6 ± 0.2
	10	0	392.0	391.1 ± 7.1	•	7	22.9	22.4 ± 0.3
	5	0	339.5	338.8 ± 4.8	•	3	20.4	20.7 ± 0.4
SF ₁	∞	1	396.0	412.9 ± 10.7		11	24.0	23.2 ± 0.3
	20	0	405.0	406.6 ± 7.8		8	23.5	23.0 ± 0.3
	10	0	355.0	356.6 ± 5.7	•	5	20.8	20.9 ± 0.4
	5	0	318.0	321.7 ± 8.0	•	0	18.6	19.2 ± 0.4
SF ₂	∞	1	447.0	442.5 ± 8.7		14	24.0	23.6 ± 0.1
	20	0	418.0	421.6 ± 6.7		13	24.0	23.6 ± 0.2
	10	0	380.0	379.5 ± 6.7	•	8	23.5	22.5 ± 0.4
	5	0	344.5	346.9 ± 7.5	•	3	19.9	20.3 ± 0.5

Table 9: Results (averaged for 20 runs) of the different island-based MMAs on the four problems considered (random). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP			HIFF			
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	
Ring	∞	17	32.0	31.8 ± 0.1	19	576.0	570.4 ± 5.6	
	20	19	32.0	31.8 ± 0.1	19	576.0	570.0 ± 6.0	
	10	17	32.0	31.1 ± 0.4	16	576.0	536.6 ± 14.1	•
	5	13	28.1	28.1 ± 0.6	14	440.5	441.0 ± 22.9	•
VN	∞	17	32.0	31.8 ± 0.1	18	576.0	565.6 ± 7.2	
	20	19	32.0	31.9 ± 0.1	15	576.0	550.8 ± 10.4	
	10	19	32.0	31.9 ± 0.1	18	576.0	565.7 ± 7.6	
	5	17	32.0	31.6 ± 0.2	17	576.0	557.0 ± 10.5	
HC	∞	16	32.0	31.6 ± 0.2	16	576.0	555.2 ± 9.6	
	20	19	32.0	31.9 ± 0.1	16	576.0	554.8 ± 9.8	
	10	17	32.0	31.8 ± 0.1	17	576.0	555.3 ± 11.5	
	5	15	32.0	31.0 ± 0.5	18	576.0	555.6 ± 11.3	
SF ₁	∞	19	32.0	31.9 ± 0.1	19	576.0	568.8 ± 7.2	
	20	17	32.0	31.7 ± 0.2	17	576.0	555.6 ± 11.5	
	10	19	32.0	31.2 ± 0.4	14	576.0	522.2 ± 15.9	•
	5	12	29.2	28.6 ± 0.7	14	576.0	500.6 ± 19.9	•
SF ₂	∞	19	32.0	32.0 ± 0.0	19	576.0	570.0 ± 6.0	
	20	18	32.0	31.8 ± 0.2	17	576.0	558.6 ± 9.6	
	10	17	32.0	31.7 ± 0.2	15	576.0	545.5 ± 13.2	
	5	14	31.4	29.7 ± 0.6	17	576.0	503.9 ± 22.4	•

topology	k	HXOR			MMDP			
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	
Ring	∞	0	402.5	408.1 ± 7.0	7	23.6	23.0 ± 0.3	
	20	0	396.0	398.7 ± 6.2	11	24.0	23.4 ± 0.2	
	10	0	371.0	371.4 ± 5.8	6	22.0	22.1 ± 0.4	•
	5	0	338.0	338.1 ± 6.4	6	20.0	20.5 ± 0.5	•
VN	∞	0	446.0	444.0 ± 7.5	12	24.0	23.5 ± 0.2	
	20	1	450.0	445.9 ± 9.0	14	24.0	23.6 ± 0.2	
	10	0	428.0	422.2 ± 7.0	12	23.8	23.4 ± 0.2	
	5	0	388.5	394.9 ± 8.1	6	22.6	22.6 ± 0.3	•
HC	∞	1	440.5	447.9 ± 8.6	13	24.0	23.5 ± 0.2	
	20	1	436.0	445.3 ± 8.0	15	24.0	23.7 ± 0.2	
	10	0	408.0	416.1 ± 6.9	12	24.0	23.3 ± 0.3	•
	5	0	380.0	378.8 ± 6.1	8	23.1	22.5 ± 0.4	•
SF ₁	∞	1	396.0	412.9 ± 10.7	11	24.0	23.2 ± 0.3	
	20	0	392.0	404.1 ± 7.9	8	23.5	22.9 ± 0.3	
	10	0	378.0	387.3 ± 7.5	7	22.7	22.3 ± 0.4	•
	5	0	337.5	344.4 ± 7.8	3	20.3	20.8 ± 0.4	•
SF ₂	∞	1	447.0	442.5 ± 8.7	14	24.0	23.6 ± 0.1	
	20	0	412.0	419.6 ± 6.5	14	24.0	23.6 ± 0.2	
	10	0	393.0	394.7 ± 4.7	11	23.6	23.1 ± 0.3	•
	5	0	360.0	357.4 ± 8.2	6	22.7	22.4 ± 0.4	•

Table 10: Results (averaged for 20 runs) of the different island-based MMAs on the four problems considered (probabilistic reinitialization). The number of times the optimum is found (n_{opt}), the median (\tilde{x}), the mean (\bar{x}) and the standard error of the mean (σ_x) are indicated.

topology	k	TRAP			HIFF		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	17	32.0	31.8 ± 0.1	19	576.0	570.4 ± 5.6
	20	19	32.0	31.8 ± 0.1	19	576.0	570.0 ± 6.0
	10	16	32.0	31.1 ± 0.5	16	576.0	528.3 ± 15.6 •
	5	16	32.0	30.6 ± 0.5 •	17	460.0	470.0 ± 22.1 •
VN	∞	17	32.0	31.8 ± 0.1	18	576.0	565.6 ± 7.2
	20	19	32.0	31.9 ± 0.1	15	576.0	550.8 ± 10.4
	10	19	32.0	31.9 ± 0.1	17	576.0	560.3 ± 9.2
	5	19	32.0	32.0 ± 0.0	18	576.0	564.0 ± 8.3
HC	∞	16	32.0	31.6 ± 0.2	16	576.0	555.2 ± 9.6
	20	19	32.0	31.9 ± 0.1	16	576.0	554.8 ± 9.8
	10	18	32.0	31.9 ± 0.1	16	576.0	552.8 ± 10.7
	5	18	32.0	31.7 ± 0.2	13	576.0	537.6 ± 12.1
SF ₁	∞	19	32.0	31.9 ± 0.1	19	576.0	568.8 ± 7.2
	20	17	32.0	31.7 ± 0.2	17	576.0	558.8 ± 9.4
	10	17	32.0	30.9 ± 0.5	13	576.0	521.0 ± 17.3 •
	5	11	31.6	30.1 ± 0.5 •	14	520.0	498.9 ± 19.4 •
SF ₂	∞	19	32.0	32.0 ± 0.0	19	576.0	570.0 ± 6.0
	20	18	32.0	31.8 ± 0.2	17	576.0	558.6 ± 9.6
	10	17	32.0	31.6 ± 0.2	14	576.0	544.2 ± 12.2 •
	5	17	32.0	31.4 ± 0.3	15	576.0	529.4 ± 16.1 •
topology	k	HXOR			MMDP		
		n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$	n_{opt}	\tilde{x}	$\bar{x} \pm \sigma_x$
Ring	∞	0	402.5	408.1 ± 7.0	7	23.6	23.0 ± 0.3
	20	0	396.0	400.6 ± 7.1	12	24.0	23.5 ± 0.2
	10	0	377.0	377.2 ± 4.7 •	8	22.9	22.7 ± 0.3
	5	0	355.0	354.4 ± 5.0 •	3	22.0	21.7 ± 0.4 •
VN	∞	0	446.0	444.0 ± 7.5	12	24.0	23.5 ± 0.2
	20	1	449.0	444.4 ± 9.3	14	24.0	23.6 ± 0.2
	10	0	422.0	429.1 ± 8.1	15	24.0	23.7 ± 0.1
	5	0	400.0	416.2 ± 9.3 •	14	24.0	23.7 ± 0.1
HC	∞	1	440.5	447.9 ± 8.6	13	24.0	23.5 ± 0.2
	20	1	442.0	445.7 ± 8.2	15	24.0	23.7 ± 0.1
	10	1	413.0	424.8 ± 10.3 •	11	24.0	23.5 ± 0.2
	5	0	416.0	416.1 ± 6.8 •	6	22.7	22.5 ± 0.3 •
SF ₁	∞	1	396.0	412.9 ± 10.7	11	24.0	23.2 ± 0.3
	20	0	402.5	404.9 ± 7.6	9	23.5	23.1 ± 0.3
	10	0	384.0	388.9 ± 6.9 •	8	22.7	22.3 ± 0.4
	5	0	364.0	365.2 ± 5.8 •	5	20.4	20.9 ± 0.4 •
SF ₂	∞	1	447.0	442.5 ± 8.7	14	24.0	23.6 ± 0.1
	20	0	420.0	425.1 ± 6.3	14	24.0	23.7 ± 0.1
	10	0	388.0	393.9 ± 5.7 •	10	23.8	23.2 ± 0.3
	5	0	380.0	391.6 ± 6.3 •	9	23.1	22.9 ± 0.3 •