J Glob Optim (2015) 63:677–707
DOI 10.1007/s10898-015-0302-7

Island models for cluster geometry optimization: how design options impact effectiveness and diversity

António Leitão[1] · **Francisco Baptista Pereira**[1,2] · **Penousal Machado**[1]

Received: 5 September 2014 / Accepted: 4 April 2015 / Published online: 7 May 2015

Abstract Designing island models is a challenging task for researchers. A number of decisions is required regarding the structure of the islands, how they are connected, how many individuals are migrated, which ones and how often. The impact of these choices is yet to be fully understood, specially since it may change between different problems and contexts. Cluster geometry optimization is a widely known and complex problem that provides a set of hard instances to assess and test optimization algorithms. The analysis presented in this paper reveals how design options for island models impact search effectiveness and population diversity, when seeking for the global optima of short-ranged Morse clusters. These outcomes support the definition of a robust and scalable island-based framework for cluster geometry optimization problems.

Keywords Cluster geometry optimization · Island models · Diversity

1 Introduction

Chemical clusters, which are aggregates of particles ranging from a few to thousands, have been intensively studied in the past decades due to their importance across different fields of study such as Physics, Chemistry, Biochemistry and Nanoscience [86]. Determining the geometry of atomic or molecular clusters that minimize the potential energy held is impor-

✉ António Leitão
apleitao@dei.uc.pt

Francisco Baptista Pereira
xico@dei.uc.pt

Penousal Machado
machado@dei.uc.pt

[1] Department of Informatics Engineering, CISUC, University of Coimbra, 3030-290 Coimbra, Portugal

[2] Instituto Politécnico de Coimbra, ISEC, Rua Pedro Nunes, 3030-199 Coimbra, Portugal

tant as it corresponds to the arrangements that maximize their stability. Such a problem is equivalent to finding the global minimum on the potential energy surface (PES) [23].

PES are multidimensional functions that describe the interactions between the particles that compose a cluster (*3N* dimensional when a cluster of N atoms is modeled on the three dimensional space). PES are regarded as very complex functions that result in highly roughed landscapes, rigged with a large number of local minima and a deep multiple-funnel character [22,75,80]. In more detail, PES can be thought of as a collection of basins of attraction, each corresponding to a local minimum. The multiple funnels correspond each to a cluster motif (icosahedral, decahedral, close packed, etc.) and contain a large number of basins under the same motif. Moreover, they may contain smaller funnels that share some of their characteristics. The number of basins increases exponentially with the cluster size [13].

Cluster Geometry optimization (CGO) has become relevant as a test problem to benchmark the effectiveness of global optimization algorithms. Finding the lowest potential energy of a cluster and the corresponding geometry is a NP-hard task that provides difficult instances to challenge new algorithms [13]. Lennard-Jones [42,47] clusters and Morse clusters [56] are two of the most relevant models as their functions depend solely on the distance between every pair of particles that compose the cluster. Between these pair-wise models, Morse clusters are regarded as being particularly tough. While choosing Lennard-Jones clusters as a test system offers a number of advantages, the fact that its instances are mostly structurally uniform and have strong funneling properties makes Lennard-Jones clusters less challenging. Morse clusters overcome these setbacks, being regarded as a challenging, more adequate test system and have been adopted by many research groups. The Morse function can be used to model both long range and short range interactions, the latter being considered particularly tough [22].

Evolutionary algorithms were first applied to CGO problems in the early 1990s [34,84] and are currently considered state-of-the-art approaches. The most successful approaches use Cartesian coordinates to represent particles in space and problem specific operators. Local-search methods that use first order derivative information to guide individuals into the nearest local optimum have later been integrated and represent a major breakthrough. Such hybrid approaches were first applied by Deaven and Ho [20] and have been widely used, contributing to better results [31,34,62].

The first population-based approach to discover all putative optima for Morse clusters up to 80 atoms was the population basin-hopping (PBH) proposed by Grosso et al. [31] which relies on a two-phase local optimization method designed to enhance its search efficiency. The approach shows however some weaknesses. The fact that it requires the specification of problem dependent parameters may induce bias and makes it difficult to extend to instances where the optimal geometry is unknown. In order to overcome this limitation, Cassioli et al. [10] proposed an approach where each individual encodes different settings, achieving enhanced performance on Lennard-Jones clusters without the need of preliminary parameter definition. Pereira et al. [60] proposed a self-adaptive EA that evolves the needed parameters along each run, thus removing the need to specify them, avoiding the incurring bias. The approach was tested on Morse clusters ranging up to 80 atoms, successfully finding all putative global optima. They also suggested that the enhanced results were in most part the outcome of the increased diversity caused by the adaption of the settings. Later Pereira et al. [61] introduced an unbiased hybrid steady-state EA combined with a single phase local optimization method based on a general quasi-Newton method. This study includes mechanisms designed to ensure the maintenance of diversity and demonstrates that its preservation is a key point to efficiently tackle CGO problems.

Island models (IM) are a possible extension to standard EAs that distribute individuals and restrict interactions between them. Other methods for the distribution of individuals are available [2]. Such models are often referred to as non-panmictic models, as opposed to single-population models, known as panmictic. Non-panmictic models result in behaviours that are distinct from standard EAs [71]. IMs, in particular, divide the global population in smaller sub-populations (islands) that undergo periods of isolated evolution punctuated by stages of migration of individuals between islands. The creation of such borders causes individuals to be at larger and unequal distances from each other thus causing the spread of genes in the population to naturally slow down. While on the one hand this may prevent the mixing of successful individuals, on the other hand it helps prevent individuals from prematurely dominating the population. Such delay contributes to the maintenance of diversity of genetic material which may help achieve enhanced results [71].

Since they depend mostly on local interactions and need little or no synchronization, IMs are appealing distributed models that have established themselves among the EA community. Considering the discussed characteristics of the PES, IMs seem an adequate paradigm to tackle CGO problems. Designing an appropriate IM is however a difficult task as it requires a wide number of decisions such as how many islands are used, how are they connected, how many individuals are migrated, which ones and how often. These decisions are generally problem dependent as search spaces with distinct characteristics will allow algorithms to benefit differently from the particularities of IMs. This study focuses on making appropriate IM design choices for tackling CGO problems and the corresponding PES functions by including an extensive discussion on the impact of different parameters and operators on the behaviour of the algorithm. Performance and behavioural differences between IMs and a sequential model are also addressed. Overall, this study intends to provide extensive knowledge on pertinent decisions regarding IMs on this particular context and hopefully in others.

This paper is structured as follows: Sect. 2 introduces the problem of Morse clusters optimization, related work and a hybrid steady-state approach to tackle the problem; Sect. 3 covers Island Models, their characteristics, related work, and a distributed approach to the optimization of Morse clusters; Sect. 4 describes the experimental setup, includes a discussion on the effect of different design choices on the behaviour of IMs and finally compares the obtained results with those of a panmictic approach; Sect. 5 draws conclusions on the study findings and addresses future work.

2 Morse clusters optimization

The following subsections cover optimization of Morse clusters. The problem is introduced in Sect. 2.1, followed by an overview of EA approaches and key achievements in Sect. 2.2. Finally, Sect. 2.3 describes a state-of-the-art panmictic approach applied to the optimization of Morse clusters.

2.1 Morse clusters

The Morse potential is the N-particle additive potential defined by Eq. 1 which depends on the Cartesian distance between particles i and j, r_{ij}, the bond dissociation energy, ϵ, the equilibrium bond length, r_0, and the range exponent of the potential, β [56]. As in various previous studies [13,21,61,73], we have adopted a simplified scaled version of the Morse potential by setting both r_0 and ϵ to 1.

$$V_{Morse} = \epsilon \sum_{i}^{N-1} \sum_{j>i}^{N} \{exp[-2\beta(r_{ij} - r_0)] - 2exp[-\beta(r_{ij} - r_0)]\} \quad (1)$$

Regarding the range exponent of the potential, β, it has a strong influence on the characteristics of the function, thus allowing it to model aspects of particle interactions for a wide range of materials. As the value of β is increased, the number of local minima in the PES increases rapidly, changing the landscape from smooth to rough [21,22]. Also, as β increases, the PES is more likely to assume a multiple-funnel character, therefore increasing the challenge for global optimization algorithms. While for $\beta = 6$ the structural behaviour of the Morse potential is very similar to that of Lennard-Jones clusters, the global optimization of Morse clusters for $\beta = 10$ or $\beta = 14$, which correspond to short range interactions, is regarded as a specially tough task, one that is orders of magnitude harder than Lennard-Jones clusters [21,22]. In this study, following previous works [13,14,60–62], we have set $\beta = 14$ leaving r_{ij} as the only variable in Eq. 1, resulting on a Morse potential model that depends only on the interactions between particles. A database of putative Morse cluster minima up to 80 atoms is available for different values of β [1].

2.2 Related work

A general overview of the most relevant EA approaches to CGO can be found in [36]. This section covers important milestones, followed by a brief review of key achievements of the application of EAs to Morse clusters in particular.

Early EA approaches relied on binary strings to represent atoms. Floating point representation was first introduced in CGO by Zeiri [85] and soon became a widely preferred alternative [17,20,33,40,58], mostly because it allows non-discrete continuous positioning of particles. Advantages include increased precision and the ability to apply genetic operators that take into account properties of cluster geometry in a more convenient way and without the need of converting genes to atomic coordinates.

Various genetic operators that use context specific knowledge have been proposed to tackle CGO problems. These operators have been designed to be sensitive to the semantic properties of clusters. Regarding mutation, the most relevant operators include particle permutation, where the position of two or more particles are switched, particle displacement, where one or several particles are slightly displaced in a random fashion, piece rotation, where part of the cluster is rotated around a chosen axis, and, finally, piece reflection, where part of the cluster is replaced by its own reflection or the reflection of a different part [41,64,66].

Concerning crossover, the Cut and Splice (C&S) operator has been most widely used. It works by defining a random cutting plane in one parent and merging one of the resulting sub-clusters with a complementary sub-cluster from the second parent also obtained through a cutting plane [20]. It therefore ensures that each parent contributes with particles that are close together, meaning that they will likely have a low potential energy level. This feature increases the likelihood of mixing useful blocks of genes and is regarded as beneficial to the evolutionary process [36,67]. The efficiency of this operator may be increased by using knowledge of the geometry of each individual to determine the position of the cutting planes [35].

Apart from specific operators, Lamarckian local optimization of individuals has also played an important role. Gregurick et al. [30] have proposed the use of conjugate gradient type local relaxation which has also been applied by Deaven and Ho [20] and Rata et al. [66]. Alternative methods such as the Nelder–Mead simplex local search has been applied by Iwamatso [40] and also by Sastry and Xiao [68] while a limited memory Broyden–Fletcher–

Goldfarb–Shanno (L-BFGS) [49] has been applied by Hobday and Smith [37], Zhao and Xie [86] and Pereira et al. [61]. The combination of local optimization and EAs into hybrid approaches is nowadays regarded as having striking effects on search efficacy.

Optimization of Morse clusters is considered nowadays a very relevant benchmark problem which has established itself in the scientific community. Doyes and Wales [21] were the first to tackle it by applying a basin-hopping algorithm, similar in many ways to a Monte Carlo method aided by local optimization. They report finding all but 12 putative optima in clusters ranging up to 80 atoms. Roberts et al. [67] were the first to suggest using an EA on this problem. Their approach combined a real-valued representation [85] with a Cut and Splice crossover operator [20] and Lamarckian local optimization [20]. In their study they focused on medium and short-range Morse clusters with 19 to 50 atoms achieving nearly all putative optima. A revised version of their algorithm, which successfully found all the presumed solutions in the same range, was published in 2003 by the same research group [41].

A two-phase local search algorithm was designed by Locatelli and Schoen [50] with the specific purpose of tackling multi-funneled PESs. The idea behind the algorithm is to use a modified potential function which combines knowledge from the true PES as well as from the structure of the putative solution. By performing local optimization, the algorithm biases the search toward local optima with specific properties that aid the search for a global minimum when searching on the true PES during the second step. This method was combined with a basin-hopping approach and applied first to Lennard-Jones clusters [50,51] and latter to Morse clusters [22]. In the later, it was able to find almost all putative global optima for instances between 41 and 80 atoms. The two-phase local optimization has also been coupled with a population-based algorithm in a posterior study showing a high rate of success in finding the assumed global minima on Morse clusters ranging from 41 to 80 atoms [31].

While these approaches have been applied successfully to Morse clusters, the proposed two-phase local optimization method has an important weakness. In order for the first phase to work properly, a number of parameters need to be specified a priori. These parameters depend on the geometric structure of the optimal solution corresponding to each cluster size. This property raises questions regarding bias, since the parameters will guide the search to solutions with a given geometric structure, as well as regarding what happens in instances where we lack that knowledge.

Two self-adaptive approaches have been proposed to overcome this weakness. Cassioli et al. [10] have presented a PBH algorithm where the required parameters are encoded in each individual as weight factors that are inherited from parents to children. Evolutionary pressure ensures that only the most adequate spread throughout the population. On Lennard-Jones clusters, this approach has confirmed that the existence of different weights is helpful. Pereira and Marques [60] proposed an extension to the two-phase local optimization where each individual encodes the parameters to be applied on its own local relaxation method. In their study, each individual starts with a random set of parameters that undergo evolution by online adjustment, thus searching for the most suitable configuration for each instance. Results on the optimization of Morse clusters show that although the number of successful runs decreases when compared to Locatelli's two-phase approach, the algorithm is able to find the putative global optima of Morse clusters up to 80 atoms without the need of setting instance-specific parameters beforehand.

Pereira and Marques [60] argue that the success of their approach was largely the result of an increase in diversity induced by the adaption of parameters. A later study from the same research group confirmed the hypothesis [61]. The study relied on a hybrid steady-state EA coupled with an unbiased single-phase local optimization procedure. Their analysis shows

that maintenance of diversity is a key factor when tackling CGO problems. They report finding all putative global optima on Morse clusters ranging from 41 to 80 atoms, showing how an unbiased hybrid EA can be used to tackle such problems efficiently. Many other studies support that diversity is a key factor for the success of the optimization algorithm [10,11,31].

Although not directly related to EAs, it is worth mentioning dynamic lattice searching (DLS) [70]. This method was applied to the optimization of short-ranged Morse clusters and obtained good quality results, as it established several new upper bounds on the potential energy for instances until 240 atoms [13,14,25]. DLS creates a dynamic lattice that is formed by all possible locations (in general, sites on the surface of the cluster) for atoms around an initial random minimum structure. Then it relies on a greedy strategy to iteratively optimize the cluster surface by moving atoms between positions in the lattice. The method is restarted several times from different random starting geometries to improve the likelihood of obtaining the local minimum. Notwithstanding its effectiveness, DLS construction strategy slightly biases the search towards specific cluster structures, thus compromising its robustness (e.g., the dynamic lattice construction strategy adopted by Shao et al. [70] favors the appearance of close-packed motifs).

2.3 A hybrid steady-state EA for Morse clusters optimization

This section describes the main components that compose the algorithm used in this study. They have been previously used in different approaches and have been extensively discussed in Pereira's work on diversity in CGO problems [61]. The work compared different operators and parameters and discusses the observed behaviours. It focuses on mechanisms responsible for the maintenance of diversity. These are linked to the success of the approaches which was taken into account on our choices for the present study.

2.3.1 Local optimization

In this study, local optimization relies on the L-BFGS method. This technique attempts to combine both the reduced storage needs of conjugate gradient methods and the superlinear convergence that characterizes full memory quasi-Newton methods. It needs to be supplied with the function to be minimized as well as first order derivative information. The generic element n of the Morse cluster potential gradient is obtained through Eq. 2 where x_{ni} represents the difference between the Cartesian coordinates of particles n and i, $x_{ni} = x_n - x_i$ Similar equations apply for the y and z axis.

$$g_n = -2\beta\epsilon \sum_{i \neq n}^{N} \left(\frac{x_{ni}}{r_{ni}} \right) \{exp[-2\beta(r_{ni} - r_0)] - exp[-\beta(r_{ni} - r_0)]\} \tag{2}$$

L-BFGS is applied to every individual generated prior evaluation, guiding the solution to the nearest local optimum. The accuracy of the local optimization method is set to 1.0E−8.

2.3.2 Representation and genetic operators

Individuals represent clusters of N atoms on the *3D* space as a $3 \times N$ real-valued array that represents the Cartesian coordinates of each particle. The coordinates range between 0 and $N^{1/3}$ which, despite the lack of formal proof, is widely accepted to allow the correct scaling of the cluster volume with N [41,61,85]. Restrictions on the distance between each

pair of atoms are enforced since two atoms coming too close to each other will result on an excessively repulsive potential. During initialization and when genetic operators are applied, inter-atomic distances are verified and only solutions where they are larger than the pre-specified parameter δ are allowed. In this study $\delta = 0.5$.

A generalization of the C&S operator is used for crossover. This operator (GenC&S) has been proposed by Pereira et al. [63] and operates in a very similar way to the original C&S but differs in the way the exchanged sub-clusters are determined. Their study shows that the GenC&S exhibits a more appropriate behaviour than the original C&S on cluster geometry optimization. Instead of using cutting planes and exchanging the resulting sub-clusters, GenC&S tries to ensure that each parent contributes with atoms that are close together by Euclidean distance. In detail, GenC&S builds *D1* from parents *P1* and *P2* by applying Algorithm 1 where N is the number of atoms that compose the clusters. *D2* is obtained by swapping the roles of the parents. A different variation of the original C&S is given by Pullan [65].

Algorithm 1 Generalized Cut and Splice operator

1: Select a random atom *CP* from *P1*
2: Select a random number $X \in [1, N-2]$
3: Copy *CP* and the X closest atoms from *P1* to *D1*
4: Select the $N-(X+1)$ atoms from P2 that are closest to *CP*
5: **for** each atom selected from *P2* **do**
6: **if** distance constraints between the atom and those in *D1* are met **then**
7: Include it in *D1*
8: **else**
9: Skip it
10: **end if**
11: **end for**
12: **if** atoms in $D1 \leq N$ **then**
13: Complete *D1* with randomly placed atoms
14: **end if**

Sigma mutation [63] is applied to every individual after crossover. This particle displacement operator acts on one atom at each time by slightly moving it on the *3D* space. The new location is obtained by perturbing each coordinate with a random value sampled from a Normal distribution with mean 0 and standard deviation σ. Local optimization is then applied before evaluation takes place.

2.3.3 Population model

Steady State EAs are behaviourally different from generational EAs. In these models only one or two offspring are generated at each step and they have to compete with their parents and remaining individuals for a place in the population. Also, if they are allowed in the population, slots need to be vacant, meaning that individuals need to be selected to forcibly perish. This procedure is repeated until a stop criterion is met.

The replacement strategy that controls this process is arguably the most important operator on steady state EAs with a major impact on its performance and behaviour [19,52,74]. In general, replacement strategies can be divided into fitness-based, where the quality of the individuals weights on the decision, or age-based, where the oldest individuals are removed from the population. These criteria can however be combined with more complex mechanisms that aim to maintain diversity.

A widely-known straightforward fitness-based approach is the GENITOR replacement strategy [83] that works by replacing the worst individuals in the population with emerging offspring, given that the later are better, otherwise they are discarded. Various studies show however that such a strategy induces a high selective pressure that is bound to guide the population to premature convergence even when parents are selected randomly [28,74]. GENITOR is no more than a simple strategy in a wide set that includes plenty of strategies that work in more appropriate ways regarding selection pressure and diversity.

When considering a problem such as CGO, where the search landscape is rigged with local minima which increases exponentially with cluster size, it is imperative to provide the EA with mechanisms that prevent or at least slow down takeover time. Many replacement strategies that try to enforce population diversity derive from crowding methods [18,53]. They work by having new offspring replace individuals that are most similar to themselves based on genotype similarity. This way, the population is less likely to be flooded by similar solutions. Such crowding methods result on the formation of niches.

A different approach to promote diversity has been proposed by Krasnogor [43]. His study presents an adaptive memetic algorithm where the local search component plays a double role. When there is high diversity, it acts as a regular local search algorithm, however, when diversity starts to fall, it changes behaviour and accepts individuals that deteriorate fitness in order to avoid premature convergence. The study also considers different local search methods that are able to complement each other as well as a self-adaptive approach where each individual encodes its own local search strategy that evolves along the run. Such approaches that learn along the run which strategy to apply are known as multimeme algorithms (MMA) [43,44,59].

The replacement strategy used in this study was proposed by Lee et al. [45] and was later incorporated in the PBH algorithm [31]. It is a fitness-based strategy that includes mechanisms that aim at maintaining diversity. Algorithm 2 shows the steps that determine if a new offspring D is allowed into the population Pop and what individual is replaced. Briefly, in each iteration, two parents are selected and two offspring built. Considering an individual D, already locally optimized and evaluated, if Pop contains a solution that is close to D, then only the best of the two is kept. Otherwise, D is different from all the solutions in the population and will replace the worst one, given that D is better. In order to perceive if individuals are close, a distance measure $d(X, Y)$ that captures the similarity between individuals X and Y is required. Moreover, a minimum allowed distance between individuals in the population d_{min} needs to be specified.

Algorithm 2 Replacement strategy adopted by the EA

1: Find $X \in Pop$ such that $d(X, D)$ is minimum
2: **if** $[d(X, D) \leq d_{min}$ and $V_{Morse}(D) < V_{Morse}(X)]$ **then**
3: D replaces X in Pop
4: **else if** $[d(X, D) > d_{min}]$ **then**
5: Select $Y \in Pop$, such that $V_{Morse}(Y)$ is maximum
6: **if** $[V_{Morse}(D) < V_{Morse}(Y)]$ **then**
7: D replaces Y in Pop
8: **end if**
9: **else**
10: D is discarded
11: **end if**

2.3.4 Diversity measures

A large number of measures have been designed for the comparison of individuals in CGO problems and extensively discussed in the literature. Broad reviews and classifications of different metrics can be found in various publications [11,54,61]. The proposed metrics can be as simple as the fitness distance between individuals [20] or much more complex and aware of properties of cluster geometry and the semantics of individuals [12,31,35,45,62].

This study relies on a metric based on the distances of atoms to the clusters center of mass, first proposed by Grosso et al. [31]. This choice was made based on the results and analysis presented by Pereira et al. in a study where the metric is compared with several others [61]. Briefly, this measure works as follows: First, for each cluster, the distances of the atoms to the center of mass are organized in a vector in non-decreasing order. Then, given two vectors and the respective ordered sets Ord_A and Ord_B, the dissimilarity between the clusters is given by Eq. 3, where K represents de size of the sets.

$$d^3_{ord}(A, B) = \sum_{i=1}^{K} |(Ord_A(i) - Ord_B(i))|^3 \tag{3}$$

To estimate d_{min} we rely on information provided by the initial randomly generated population using Eq. 4 [31] where N represents the size of the population, $d(A, B)$ the distance between individuals A and B and ζ specifies the proportion among the average distance calculated with a set of individuals randomly generated and the minimum allowed distance allowed between every pair of solutions that belong to the population. Following the discussion in the study by Pereira et al. [61], ζ has been set to 0.25.

$$d_{min} = \zeta \times \frac{\sum_{A=1}^{N-1} \sum_{B=A+1}^{N} d^3_{ord}(A, B)}{\frac{1}{2}(N^2 - N)} \tag{4}$$

3 Island models for cluster geometry optimization

This section introduces IMs through a brief discussion of their characteristics in Sect. 3.1, and relevant related work in Sect. 3.2. Finally, an IM approach to hybrid steady-state EAs is described in Sect. 3.3

3.1 Island models

Parallel EAs can be subdivided into three main groups depending on their characteristics: global single-population master-slave EAs, single-population fine-grained EAs and multiple-population coarse-grained EAs [6]. While a single population is used on the first two approaches, the third one is particularly interesting due to the fact that it relies on multiple sub-populations that undergo independent evolution processes with the occasional exchange of individuals between populations. Such an approach has become quite popular in the last couple of decades despite being the parallel EA model that is least understood, mostly due to complex behaviour induced by the exchange of individuals and all the design choices implied. Many different approaches and studies have been presented on different applications and the model has been referred to by many names such as Multi-population EAS, Multiple-deme EAs, Distributed EAs or Island models (IM). We chose to use the IM designation, a name

that derives from its resemblance with a model used to describe natural populations that are isolated by geographical barriers, such as in islands.

More importantly, apart from the advantages brought by the distribution of computational effort, IMs show different behaviours from standard EAs (panmictic) [71] and therefore may be able to perform better than panmictic approaches on certain problems [3]. The most distinctive aspect of IMs is that individuals are restricted to interact with others from the same island. In detail, during periods of isolation, selection operates on each island independently and parents are able to mate only with others from the same island. Moreover, on steady-state approaches, offspring compete only with individuals from the same island through the replacement strategy. These restrictions can therefore be seen as borders that separate individuals, promoting independent evolutionary processes and slowing down the flow of genes in the global population. The existence of independent processes makes IMs very different from standard EAs while using migration separates IMs from isolated runs of panmictic EAs [71].

The effects of the mentioned features on the global population are striking. On the one hand, such an approach may prevent the successful mixing of individuals but on the other hand it prevents temporarily better individuals from dominating the whole population. Instead, each island exploits its own genetic material, enabling the subpopulations to occupy different regions of the search space. During these periods of exploitation, progress is slow and based on small, almost insignificant changes in the genetic pool of the islands. This equilibrium is however perturbed by migration from time to time which triggers rapid evolutionary changes through the constant intra-island competition. This effect has been noticed to have strong similarities with the theory of punctuated equilibria [24]. Cohen et al. [15] were the first to notice the resemblance, showing in their experiments with parallel EAs that while relatively little changes occurred during isolation periods, new solutions were often found soon after migration.

The aforementioned characteristics are bound to promote diversity on the global population as well as provide weaker individuals with ecological opportunities to contribute with their genes and hopefully help to reach better solutions [55]. Regarding CGO problems, such features may be helpful in covering a large part of the search space. Considering the PES, while intra-island evolution pushes individuals towards different local optima, migration introduces new genetic material into the islands, ultimately helping the population escape and progress.

IMs require a number of design choices as well as the specification of various operators and parameters: the number and size of islands, a topology that defines how islands are connected, operators for the selection of emigrants and the inclusion of immigrants (migration policies), migration sizes and frequency. All these components are crucial to IMs and need to be balanced in order to achieve the desired effect.

3.2 Related work

The idea of using parallel computation to simulate the evolution of natural populations goes back to the 1950s, when Holland proposed a parallelization architecture aiming various applications [38,39]. While the project was never built, some of the ideas remained. In the 1970s, Bethke conducted a study on the parallelization of two different GA implementations [4] and analysed their efficiency. He showed that parallel approaches may reach close to 100 % efficiency on single instruction, multiple data class computers. He also identified some of the possible bottlenecks on parallel architectures. Early in the 1980s, Grefenstette proposed four possible approaches to parallel GAs [29]. One of them was an IM where the best individual

from each island is broadcast to all others every generation. Grefenstette also focused on the effects of migration on the maintenance of diversity and raised important questions regarding frequency and topology.

In the mid 1980s, Grosso performed what might have been the first systematic study on IMs [32]. He divided a population in five islands and experimented with different migration rates. He discusses that in isolated islands average fitness improvements occurred faster but stopped at worst fitness values than in the panmictic population. With low migration rates the results did not improve but with intermediate rates the results were as good as those achieved by the panmictic population. These results suggest that there is a lower bound on migration rates, below which the IM performs worse than the panmictic EAs. A few years later Tanese conducted several studies regarding the effects of migration sizes and intervals on IM efficiency as well as on maintenance of gene diversity [77–79]. It is shown that the effect of the parameters are connected and that both need to be balanced in order to achieve efficiency, suggesting that an upper bound on migration rate exists, above which the IM performs worst than a panmictic approach. In a late 1990s survey of parallel distributed GAs by Alba and Troya it is proposed that migration rates should fall between 1 and 10 % of the population size with the best value to be considered problem dependent [3].

In one of Tanese's studies [77], three different experiments are conducted on a 4D hypercube topology. On the first one, migration occurred with intervals of 5 generations and two different migration rates were experimented on a varying number of islands. The results obtained were as good as those obtained by a panmictic approach. The second one used different crossover and mutation rates in each island in an attempt to balance exploration and exploitation. The last study experimented with varying frequencies of migration showing that too short or too long intervals between migrations resulted on a decrease of performance of the IM. Skolicki and De Jong have also conducted a study on migration sizes and intervals on IMs using a GA [72]. They conclude that the performance is affected by three main factors. Firstly, if migrations are frequent, then the interval length has a stronger effect than migration size, meaning that if the interval is too small there is a drop in performance. Secondly, as the migration size approaches the islands size, performance degrades. Lastly, performance may also drop for large migration intervals, an effect that may be reduced if the algorithm is ensured not to stop prematurely. Experiments on GP have been conducted by Fernández et al. [26,27]. They report similar conclusions, obtaining best results on a number of problems by migrating 10 % of the subpopulation size every 5 or 10 generations. They also report that if few individuals are migrated, the interval between migrations should be small whereas if many are migrated the interval should be large in order for each subpopulation to have enough time to improve and exploit the new genetic material. Moreover, they discuss that if too many individuals are exchanged, even with large intervals, the IM performance will drop due to the risk of homogenization of the islands. Also on the subject of migration intervals, Grosso [32] proposed that migration occurred only when the population was near convergence. Braun applied the same principle in a study where migration occurred with completely converged islands to restore diversity [5]. Munetomo et al. [57] have also applied the same idea and Cantú-Paz and Goldberg [8] have presented theoretical models that predict the quality of solutions using a fully connected topology.

Migration policies define what is being exchanged during migration. Cantú-Paz explored two different emigrant selection strategies, best or random, and two immigrant replacement strategies, worst and random in an IM where migration occurs at every generation and the exchanged individuals are copied (not moved) from the original island to the target one [7]. His study shows that choosing either emigrants or individuals to be replaced by immigrants, or

both, with a fitness based method highly increases the selection pressure, driving the islands to premature convergence. On the other hand, using random selection on both operations reportedly has no effect in takeover time. In [3], Alba and Troya study various selection operators for selection emigrants, among which random selection is reported as able to prevent the "conquest" effect on the target population.

Topologies determine how information flows from island to island. They play an important role since they control how fast solutions spread. The denser the topology is, the faster solutions reach all the islands in the model, which may lead to a few solutions taking over the population. The sparser the topology is, the more isolated islands will be from each other thus allowing the emergence of different solutions. The mixing of better solutions is more likely to occur on the first scenario while the preservation of diversity is more likely on the second. The work by Cohoon et al. [15] used a mesh topology as they defend that using highly connected topologies with a small radius works best for an adequate mixing of individuals through the evolution process. A later extension on their work using a 4-D hypercube used similar ideas [16]. Cantú-Paz has later experimented with a set of commonly used topologies such as fully connected, 4-D hypercubes, 4×4 torus, unidirectional and bidirectional rings also suggests that high density allows IMs to reach global solutions within fewer evaluations [9]. However, a more recent study [69] has concluded that highly dense topologies are adequate for the optimization of smooth cost functions, while sparser islands are more adequate for non-smooth cost functions as a mean to preserve a higher population independence and diversity. Dynamic approaches have also been studied using different strategies. In this case, an island is not enforced to communicate with a fixed set of neighbours but instead chooses which neighbours will receive its emigrants at each migration step. The choice may be made based on the effect the migration will bring to the target island according to various criteria such as population diversity [57] or genotype distance between the islands populations [48]. On a simpler approach, target islands may be chosen randomly [71]. Skolicki [71] has discussed that dynamic topologies are an adequate way to maintain the number of migrations fixed when experimenting with different topologies as well as to further reduce the speed of information spread and the risk of over flooding islands with new individuals.

Island size variation has been regarded by some researchers as a key factor in IM performance despite the fact that extensive parametric studies are scarce. Nonetheless, Whitley et al. have conducted such a study on both non-deceptive and deceptive functions [81]. They tested with populations of 5000 individuals divided into 1, 5, 10, 50 and 100 islands and populations of 500 individuals split into 1, 5 and 10 islands. In very large populations, it is shown that increasing the number of islands has a positive effect on the models efficacy on the optimization of deceptive functions, which is explained by the benefit of having multiple independent populations exploring potentially different areas of the search space that may contain different subsolutions. On nonseparable functions, the panmictic approach seems to gain advantage on most of the IM configurations, however certain configurations are able to perform as good as or even better than the single population approach. On the tests with small populations, IMs perform similarly or better than panmictic approaches on both deceptive or non-deceptive functions. Martin et al. [55] argue that increasing the number of islands is beneficial as long as the island size doesn't fall under a given threshold, which depends on the characteristics of the function being optimized. Whitley et al. [81,82] have also applied a distributed version of a GENITOR steady-state EA where immigrants replace the worst individuals in target islands.

3.3 An approach to hybrid steady-state EAs

IMs require a population of P individuals to be divided into N islands with M individuals each so that $P = N \times M$. Also, since we are using a steady state approach, the stop criterion needs to be adapted. In this case the evolutionary process stops at E evaluations in the panmictic approach, while on the IM each island stops evolving the population when I evaluations have been performed so that $E = N \times I$. As an example, a panmictic model where 100 individuals are evolved until 5,000,000 evaluations are reached may be divided into 5 islands consisting of 20 individuals each, that undergo 1,000,000 evaluations.

Existing studies suggest that increasing the number of islands used, at least to a certain point, is beneficial to IM performance, which may be specially relevant for CGO problems given their multiple-funnel character and high number of local optima. In this study we address the effects of IMs formed by 5, 10 or 20 islands and compare the resulting behaviours. Also, different topologies are regarded as playing a key role on the maintenance of diversity and exploration of the search space. Considering the importance of diversity maintenance on CGO problems, its mandatory to understand the impact of both sparse and dense topologies as well as their benefits and weaknesses. In order to do so, fully connected, ring, and torus topologies are tested. In the later two, unidirectional and bidirectional approaches are tested. Moreover, dynamic approaches are regarded as a mean to further reduce the speed of information spreading through the islands. An approach where each island chooses the target of each migration from its neighbourhood in a random fashion is also tested for the each topology.

Regarding migration, the widely popular *best-worst* and *random-random* policies are tested. The second one reportedly causes the less impact on the algorithms convergence. A third policy which we have named *random-rest* is also tested where, similarly to the *random-random* approach, individuals are selected randomly from their source island and will replace a random individual at the target island. However, restrictions were imposed so that individuals arriving at a given island are discarded if an equal individual is already part of the population as shown by Algorithm 3. Comparisons between individuals rely both on their fitness values and on the distance between them as described by Eq. 3. This metric has been designed to better fit the steady state approach as well as help maintain diversity by keeping migration from causing duplication of individuals on the same island. The best individual from each island is protected from being replaced. Different migration sizes and intervals are tested as they need to be balanced in order to obtain a better performance from the IM. Migration sizes should be large enough to promote innovation and a good mixing of individuals but also small enough to prevent flooding and fast duplication and deletion of genes on the global population. Rates of 1, 2 and 4 individuals were tested. Migrations intervals of 5 and 10 generations were also tested. The interval should be long enough to allow the exploitation of new genetic material in each island and the exchange of suboptimized individuals. However, too large intervals will result on the islands remaining on stasis for longer periods of time, wasting resources and reducing productivity. For migration and analysis purposes, a generation counter is incremented when P individuals are generated in the panmictic model or M in in each subpopulation of the IM. A study covering parts of the proposed approach and preliminary results has been published in 2012 [46].

Algorithm 3 Replacement strategy used by the migration policy

```
Consider a candidate immigrant (I) to a subpopulation (Pop)
for each individual (P) in Pop do
  if d(I, P) == 0 and V_Morse(I) == V_Morse(P) then
    Discard I
    Return
  end if
end for
sort Pop according to quality
Replace a random individual except Pop[0]
```

4 Experimental results

This section describes the experiments conducted on both the panmictic and the IM approaches. A study on the influence of various experimental setups on the IM follows and finally the results of the best setup are compared to those of the panmictic approach and discussed.

4.1 Experimental settings

The experimental settings are separately described on the following subsections.

4.1.1 Panmictic approach

Each experiment consists on the evolution of a population of 100 individuals until a limit of 5,000,000 evaluations is reached. Each iteration of the L-BFGS algorithm counts as 1 evaluation and a maximum of 1000 iterations per individual are allowed. Selection relies on tournament of 5 individuals and is followed by crossover at the rate of 0.7 and mutation at the rate of 0.05.σ has been set to $0.05 \times N^{1/3}$. Experiments were repeated 30 times for Morse clusters ranging from 41 to 80 atoms.

4.1.2 IM approach

Design choices have been thoroughly discussed in Sect. 3.3. In summary, approaches combining the following parameters were tested: 5, 10 or 20 islands; fully connected, ring and torus topologies; unidirectional, bidirectional and dynamic connections; *best-worst*, *random-random* and *random-rest* policies; migration rates of 1, 2 or 4 individuals; migration intervals of 5 and 10 generations. Both the size of the islands and the stop criterion were enforced so that the rules set in Sect. 3.3 are complied with.

Other parameter and design choices were kept the same as in the panmictic approach, except for selection which was downsized to a tournament of two individuals. While a tournament of 5 individuals may be well suited for a population of 100 individuals, in our IM experiments we deal with much smaller sub-populations. The operator ultimately prevents the 4 worst individuals from contributing with their genes on their present island which on a global level represents a great percentage of the gene pool (20, 40 and 80 % of the global population in IMs with 5, 10 and 20 islands respectively). Even if migration grants these individuals the chance to contribute in other islands, there is a high probability that they are replaced before being able to do so. Downsizing the size of the tournament helps relieve this effect and represents a better suited choice.

For each setup, experiments were repeated 30 times for a sub-set of instances. The discussion in Sect. 4.3 suggests that the best approach consists of a dynamic ring of 10 islands each migrating one individual every five generations using the *random-rest* policy. For this setup, experiments were repeated 30 times for Morse clusters sizing from 41 to 80 atoms.

4.2 Performance metrics

The performance of the IM on tackling the target problem is evaluated based on its ability to reach the putative optima. We have therefore registered the number of runs where the model was able to find the putative optima on the experimented range of clusters. This measure is widely accepted when assessing and comparing state-of-the-art approaches for CGO [14,22,31,61].

Other metrics were also registered along the run to aid in comparing and discussing the impact of different settings on the behaviour of the model. For each generation, the following elements were tracked in each island: fitness of best individual, average fitness of individuals, number of substitutions, average dissimilarity between individuals and fitness of incoming immigrants. On a global level, the following metrics were calculated for each generation: global best fitness, global average fitness, total number of substitutions and average inner diversity (AID) of the islands. Some of these metrics have been previously used on the analysis of island models [71] whereas others have been used for studying the behaviour of steady-state approaches to CGO [61]

4.3 The influence of IM parameters

This section details behavioural differences between varying design choices. The analysis focuses on the optimization of Morse clusters composed by 61 atoms as it would be impossible to cover the full range of clusters. All tested setups found the putative optima at most once on Morse clusters of 61 atoms as this instance is considered a particularly tough one. Interestingly, despite achieving very similar hit rates, the resulting behaviours are quite different as discussed on the following subsections. Each aspect of the design process is addressed individually. While comparing different choices for a given parameter all others keep the value of the best setup.

4.3.1 Number of islands

When designing an IM, one of the first questions that comes to mind is that of how many islands are going to be used. While the question seems trivial, the number of islands that is chosen will have a deep impact on the workings of the model. We want to take as much profit from parallelization as possible, however, splitting the population into very small islands may compromise their inner processes. If they are larger, we may get a better mixing of individuals in each island, but that will result on a model with only a few islands. Such a model will therefore share more resemblances with a panmictic approach, putting possible benefits from IMs at risk. Three configurations are tested, using either 5, 10 or 20 islands with the corresponding subpopulations sizing 20, 10 and 5 individuals.

Figure 1 shows the evolution of the model's AID averaged over 30 runs, using the aforementioned configurations. Differences between them are easily noticeable and reflect how they can produce different behaviours. When dividing the population of 100 individuals into 20 subpopulations of five individuals, one should consider that the selection method applied may harshly constrain individuals from sharing and contributing with their genes. In this

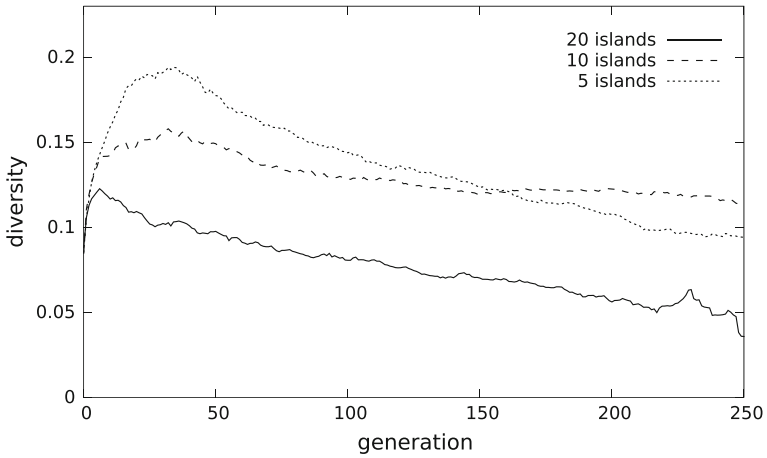

Fig. 1 Average inner island diversity obtained by different island sizes on the optimization of Morse clusters consisting of 61 atoms

case, even using a tournament of two individuals will keep the worst individual from each island from being selected, which results on an overall exclusion of 20 % of the genetic pool. Unless these individuals are selected for migration and have a better luck on their target islands, they will end up being substituted by offspring or by incoming migrants and that will impact diversity. Moreover, unless alternative methods are used, migration has a fixed rate per island, meaning that a higher number of individuals will be exchanged and duplicated on each migration step. In addition to the small size of the islands, this behaviour will keep the AID down and gradually diminishing toward the end of the run, eventually reaching a level lower than initial values. This is kept under control by the migration policy and the higher number of islands resulting on a larger distance between subpopulations, which help slow down the spreading of genes in the model and the AID from dropping fast. As seen in Fig. 1, the average AID continuously decreases, showing that a model with a high number of islands is unable to maintain an adequate AID level on this particular context.

The opposite configuration, where individuals are divided into five islands addresses the problems affecting the inner workings of each island. By including 20 individuals each, they are able to maintain a higher AID. This is more noticeable in the beginning of the run, but prevails throughout the entire runs. The AID is, in this case, more dependent on the steady-state model and its replacement strategy as the model has more resemblances to a panmictic approach. The impact of migration in this configuration is quite different from that discussed on the model with 20 islands. Due to the size of the islands, migration is less likely to highly disrupt the islands but migrating individuals are also less likely to contribute with their genes on their target island's evolution process. Moreover, while the number of simultaneous migrations drops, so does the distance between islands. Along the runs this will contribute to a higher homogenisation as less barriers separate the islands, making it easier for the genetic material from one island to reach all the others. Figure 1 backs up this analysis, showing that while islands profit from migration in the beginning, migration will push the average AID down due to its replicative and deleterious nature.

The third approach tested divides the population into 10 islands of 10 individuals each. While this may not necessarily be the optimal setting, it is a convenient and straightforward one. Its behaviour shows a balance between number of islands and their size which allowed it

to reach better results than its peers. Although the average AID shown in Fig. 1 is not as high in the beginning as in the setting with 5 islands, it remains above initial values for the whole run and evolves in a more stable manner. This effect is likely to be the result of an increased distance between islands, which while being smaller than on the 20 island configuration, seems large enough to slow down the spreading of genes. When comparing island sizes, subpopulations of 10 individuals seem large enough to allow the islands to nicely integrate new individuals and give them a chance to contribute with their genes without being badly disrupted. This contributes to a broader exploration of the search space by the IM.

Overall, this discussion shows the delicate balance between the number of islands and their size, suggesting that increasing the number of islands and therefore the barriers between individuals is beneficial both regarding behaviour and parallelization, but only to a given threshold from which the inner workings of the islands are rendered unproductive and unstable.

4.3.2 Topology

Two unidirectional topologies are initially addressed: ring and torus (fully connected is bidirectional by definition). Bidirectional approaches to the ring and torus as well as the fully connected topologies are then discussed and finally dynamic versions of these three topologies are considered.

Figure 2 shows the average AID of the IM observed along 30 runs, using unidirectional topologies. Despite the sudden peaks which are the result of migrations, the AID remains above initial values on both approaches. However, it also shows a decline toward the end of the run, an effect that is more intense on the torus topology in comparison to the ring.

Migration plays a delicate double role, on the one hand injecting new genetic material into the islands thus helping progress, and on the other hand causing the duplication of some genes and the deletion of others on a global level. Aside from the obvious differences in sparsity, the topologies also differ in the global number of exchanges at each migration step. Therefore it is unclear if the differences in AID are the result of divergent topologies or simply due to a higher number of migrated individuals in the torus topology.

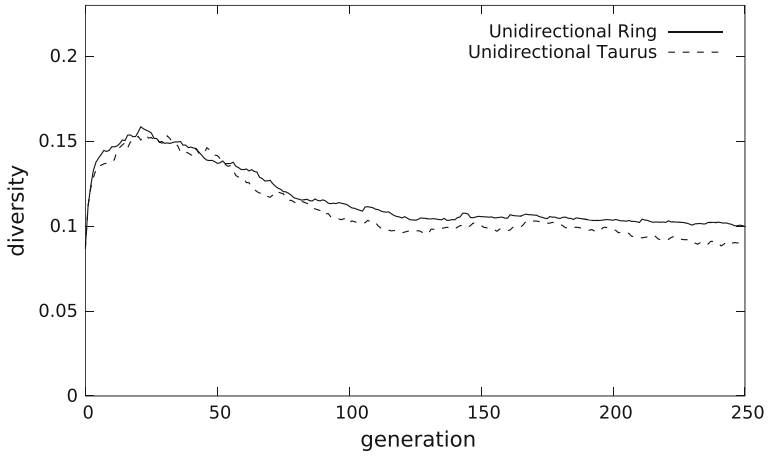

Fig. 2 Average inner island diversity obtained by unidirectional topologies on the optimization of Morse clusters consisting of 61 atoms

A closer look into the workings of unidirectional topologies uncovers nonetheless some particularities worth addressing. Figure 3 shows the fitness evolution of one island along one run as well as the immigrants accepted by its *random-rest* migration policy. As the run is closing to the end, we can see that less immigrants are allowed into the island's population and even those who are accepted have fitness values very close to the island's mean fitness value. This behaviour suggests that incoming immigrants are increasingly similar to those of their target island. While we only present one island, this effect expands to the other islands in the model and to other runs as well. We would hope for a bigger discrepancy between neighbour islands so that migrating individuals could impact and allow target islands to progress.

A likely reason for this effect is the unidirectional character of the model causing each island to gradually adapt to their sources, towards uniformity, and in turn causing their targets to adapt. Such an affinity between islands is likely to occur in a chain reaction that cycles through the topology. An illustration of such a reaction can be found in Fig. 4 which shows a

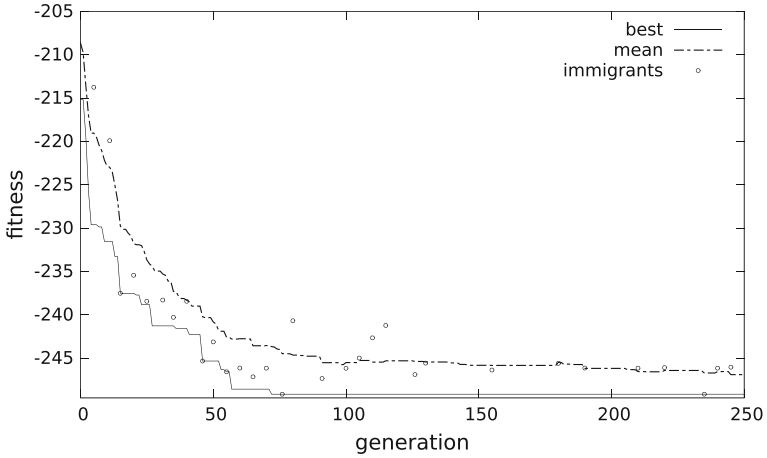

Fig. 3 Fitness evolution and included immigrants on an island using a unidirectional ring topology on the optimization of Morse clusters composed of 61 atoms

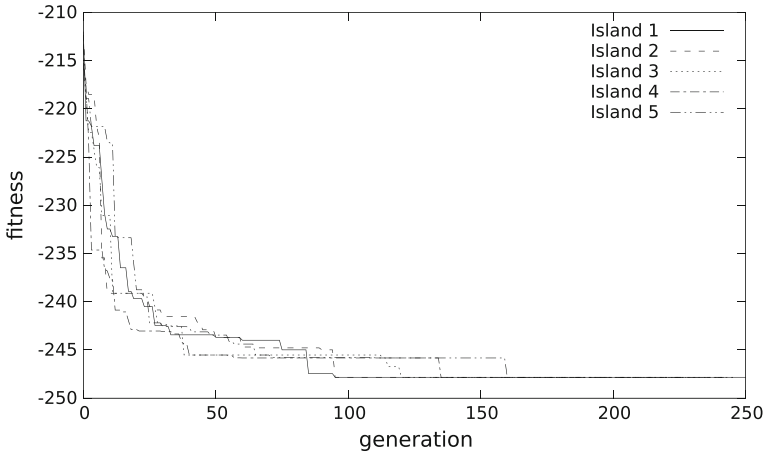

Fig. 4 Fitness evolution of the best individual on a subset of the unidirectional ring topology on the optimization of Morse clusters including 61 atoms

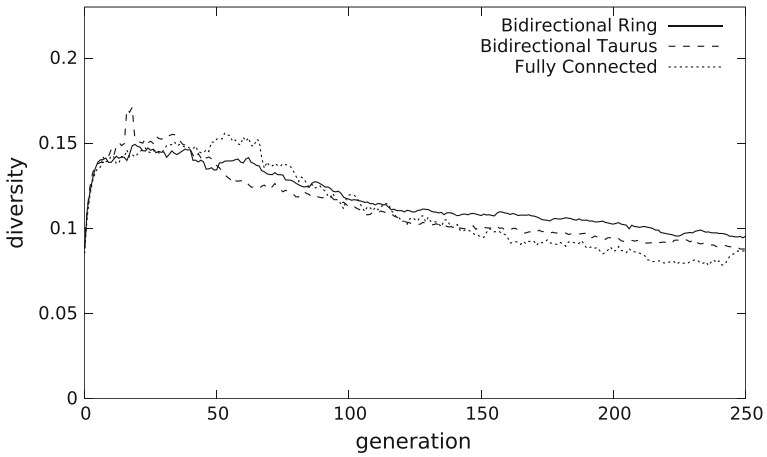

Fig. 5 Average inner island diversity obtained by bidirectional topologies on the optimization of Morse clusters consisting of 61 atoms

subset of islands from an unidirectional ring topology that consecutively progress towards a new best individual in an ordered fashion and as a result of unidirectional migrations. While these examples refer to the ring topology, as the effects are more evident, they can also be found on the torus topology.

Experiments with bidirectional approaches have shown to help mitigate the chain reaction effect found on unidirectional topologies. By giving islands the ability to move individuals to neighbours standing in either direction of the topology, a less predictable and less uniform mixing of individuals is promoted. As a result, migrations are more likely to impact the AID of target islands.

Figure 5 suggests that the ring topology has been able to maintain a slightly higher average AID than with unidirectional connections during most of the run, probably on account of avoiding chain reactions. Nonetheless, the number of exchanges increases as a result of an higher number of connections, which ultimately takes its toll towards the end of the run, causing the average AID to gradually decrease and eventually reach levels just below those achieved by the unidirectional approaches. On the grounds of being a denser topology, the differences between the torus approaches are less evident. Overall, bidirectional connections seem to produce a better mixing of genes and to prevent uniformity during the length of the runs and, even towards the end, to balance the increase in number of exchanges, maintaining a competitive average AID when compared to the unidirectional models. Finally, the fully connected topology obtains a lower AID than the other topologies along most of the run, eventually reaching values lower than initial ones. This extremely dense topology results on many more exchanges than any other, resulting in the fast deletion and duplication of genes that eventually leads to a decreasing AID.

While the previous experiments were useful for understanding certain characteristics of different topology choices, further investigation was needed in order to assess the influence of topologies without the inherent differences in number of migrations. Dynamic topologies are helpful in doing so since they keep the number of exchanges on a fixed value. Figure 6 shows how AID evolves, averaged over 30 runs, for dynamic versions of the three considered topologies. When compared with their static counterparts in Fig. 5, each dynamic approach registered a higher AID for most of the run. Such an observation suggests that the model

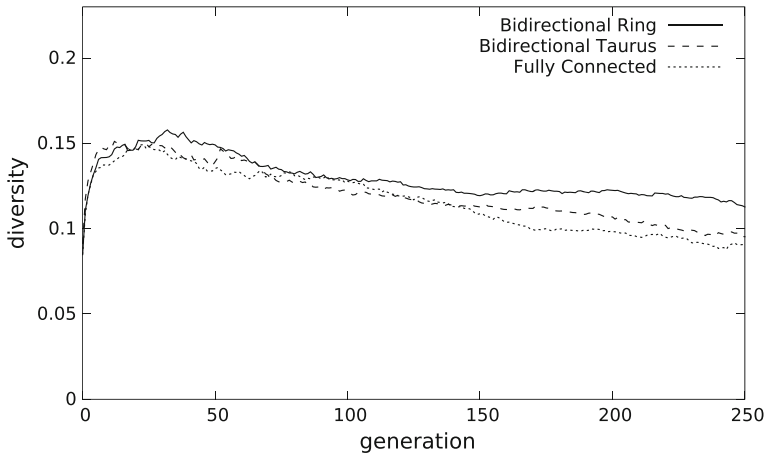

Fig. 6 Average inner island diversity obtained by dynamic topologies on the optimization of Morse clusters consisting of 61 atoms

benefits the most by taking advantage of moving individuals to multiple islands as in bidirectional topologies while preventing the resulting increase in simultaneous migrations. The random selection of target islands also introduces stochasticity which contributes to a more dynamic model, thus reducing uniformity and potentially increasing the impact of immigrant individuals.

Despite the general increase in average AID values, there are still differences in diversity handling between the three topologies. The density of the topologies seems to have a clear relation with how IMs handle diversity among its islands. Figure 6 shows a clear advantage on sparser topologies, with the ring topology being able to maintain AID values higher than the other two. Considering the multi-funneled character of the PES and its high number of local minima, maintaining a sparser topology may be beneficial since it presents the strongest barriers between islands, successfully slowing down the spreading of genes, their duplication as well as deletion, specially when combined with an appropriate migration policy. A balance between local and global approaches is achieved. While the replacement strategy applied by the steady-state model ensures that diversity is maintained inside the islands, the sparsity of the topology as well as the migration policy ensures that genetic material spreads slowly through the model helping maintain diversity on a global level.

4.3.3 Policy

Which individuals are selected for migration is an important decision. We need to make sure that migrants are able to impact target islands in meaningful ways. Also, it's important that the individuals that are replaced are chosen in a way that harms islands as little as possible. Three migration policies are addresses: *best-worst*, *random-random* and *random-rest*.

Figure 7 shows how AID evolves for each of the tested policies, averaged over 30 runs. We can see that the *best-worst* approach has a very different impact on diversity when compared to the other two. This behaviour was expected, as propagating only the best individuals is known to result on a conquering effect. Still, the idea behind this approach is worth addressing. It is expected that by migrating the best individuals of each island while replacing the worst will allow each island to benefit from new high performing genes to further evolve into better

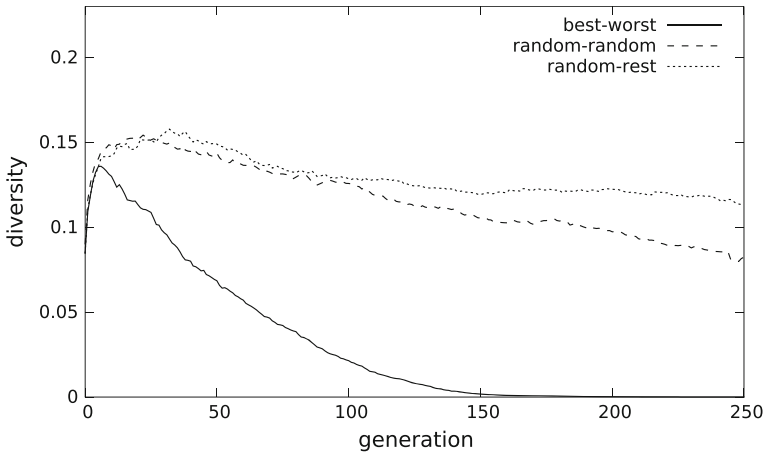

Fig. 7 Average inner island diversity obtained by different migration policies on the optimization of Morse clusters of 61 atoms

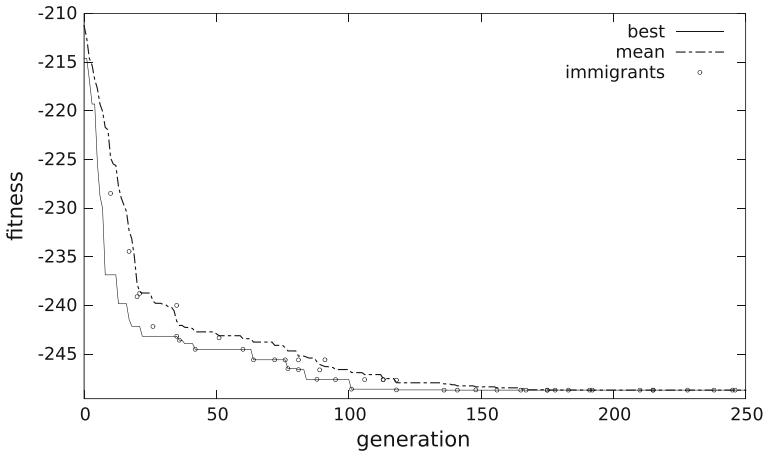

Fig. 8 Fitness evolution and included immigrants on an island using a *best-worst* migration policy on the optimization of Morse clusters composed of 61 atoms. Each island migrates one individual every five generations

candidate solutions. However, probably due the characteristics of the PES and the reduced size of the island, the migration of best individuals is likely to push them into local optima and undermine the exploration of the search space, which is vital when tackling CGO. The replacement strategy applied by the steady-state model will keep islands from converging, therefore, the behaviour observed in both Figs. 7 and 8 is the direct result of new immigrants conquering their target islands and eventually propagating one or a small set of individuals through the entire model.

The *random-random* approach has a much more desirable behaviour as seen in Fig. 7. By selecting individuals for migration in a random fashion we reduce the conquering effect since any individual can be picked. While the genetic quality of these individuals is not necessarily high, they will still be able to provide islands with fresh genes that will help them produce innovative solutions, escaping from local optima or stagnation. On a global level this

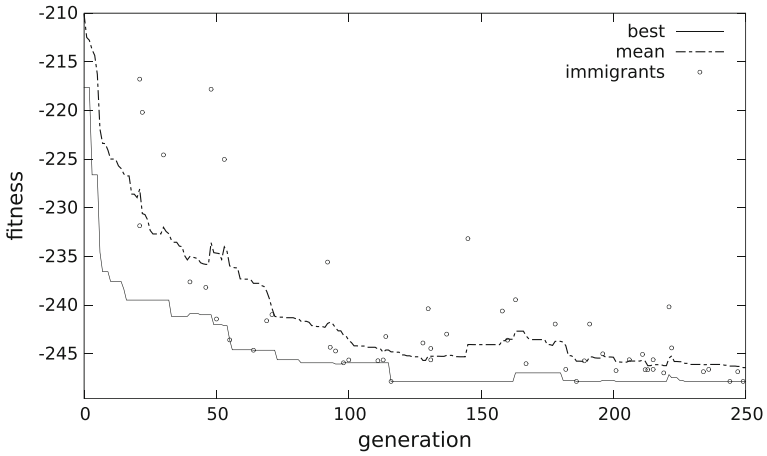

Fig. 9 Fitness evolution and included immigrants on an island using a *random-random* migration policy on the optimization of Morse clusters composed of 61 atoms. Each island migrates one individual every five generations

will allow the model to better explore the PES and search for a global optima. At the same time, incoming immigrants will replace random individuals, therefore worst individuals are not forcefully removed from the evolution process and may still be able to contribute, if not on their current islands, then in others through migration. There are nonetheless a few points at issue that may be specific of the steady state model applied. Firstly, the approach is able to migrate individuals that will bring no genetic diversity to target islands when they are structurally similar to one of those already in the subpopulation. Secondly, the best individual in each island is unprotected from being replaced which can cause drops in some island's best fitness values as well as keep those individuals from further contributing in the search process. Figure 7 shows how allowing the replicating individuals into the islands causes a gradual drop on the average AID. Also, Fig. 9 illustrates how new immigrants can cause fitness drops on target islands by replacing the best individual present.

The proposed *random-rest* approach is in its essence a *random-random* approach where the replacement of individuals in target islands by incoming migrants is restricted to address the previously mentioned issues. Figure 7 shows that this policy has an impact on the observed average AID, allowing it to keep at higher values than on both of the other tested approaches. Also when comparing Figs. 9 and 10, we can see that situations where the best individual in target islands was being replaced by immigrants no longer occur. The price paid to accomplish this is however that some migrations may be rendered useless as the migrated individuals are discarded which produces nevertheless a more beneficial behaviour and performance. This approach was designed for CGO alone and relies on the previously discussed dissimilarity measure. As long as dissimilarity measures are available to compare between two individuals, then this policy could be adapted and applied to other contexts as well.

4.3.4 Number of migrants

Three setups were tested, either migrating 1, 2 or 4 individuals simultaneously at each migration step. Figure 10 suggests that migrating 1 individual at each step is enough to help the islands progress both by the inclusion of new best individuals but also, and more interest-

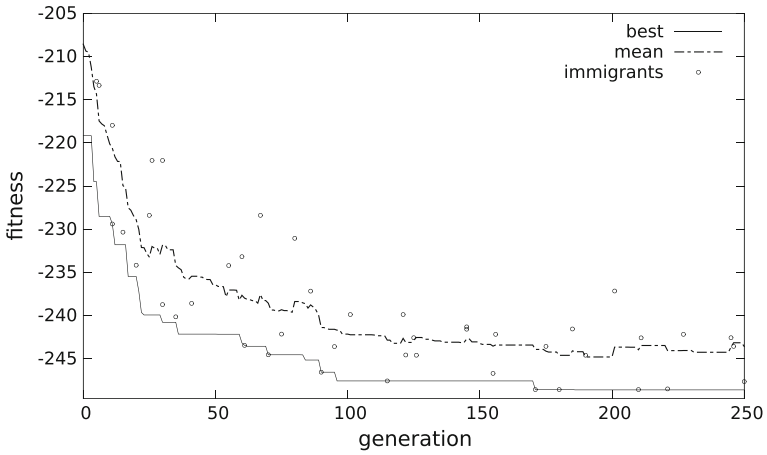

Fig. 10 Fitness evolution and included immigrants on an island using a *random-rest* migration policy on the optimization of Morse clusters composed of 61 atoms. Each island migrates one individual every five generations

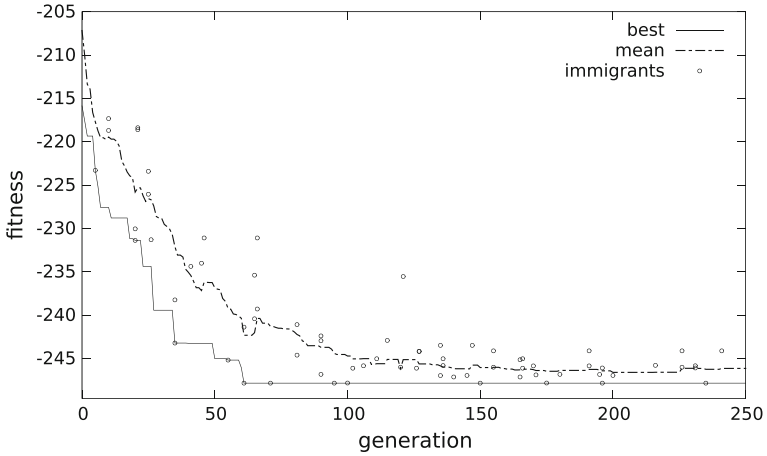

Fig. 11 Fitness evolution and included immigrants on an island, migrating 2 individuals every 5th generation on the optimization of Morse clusters composed of 61 atoms

ingly, the inclusion of genes that allow the reproduction of new best individuals. Considering the migration policy applied, one may think that by increasing the number of individuals migrated simultaneously, their impact will result on performance gains. The reasoning is that islands may have a better chance of having valuable genetic material injected at each migration step while the migration policy rejects duplicative individuals and therefore helps islands maintain their inner diversity. However, Fig. 11, which refers to the simultaneous migration of 2 individuals, shows no evidence of behavioural benefits, much less Fig. 12, referent to the concurrent migration of four individuals. In both cases, the islands seem to achieve better fitness values earlier in the runs but also seem to stagnate earlier. Furthermore, they seem to benefit less from new genetic material, even though more migrants are exchanged, which renders migration as a deleterious and duplicative operation rather than a useful tool.

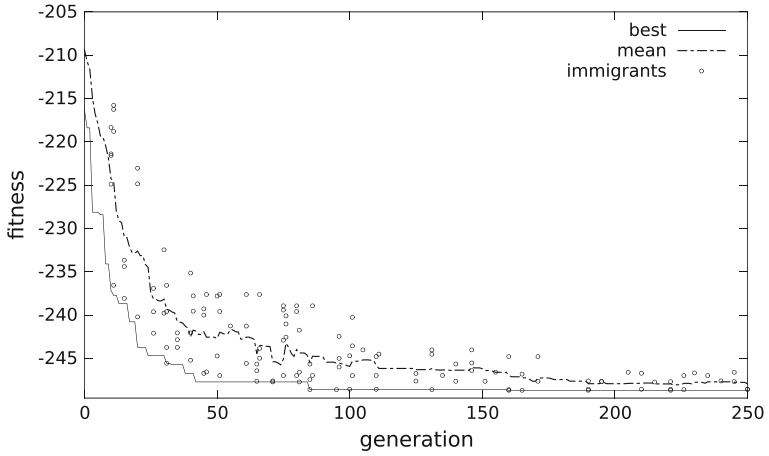

Fig. 12 Fitness evolution and included immigrants on an island, migrating four individuals every 5th generation on the optimization of Morse clusters composed of 61 atoms

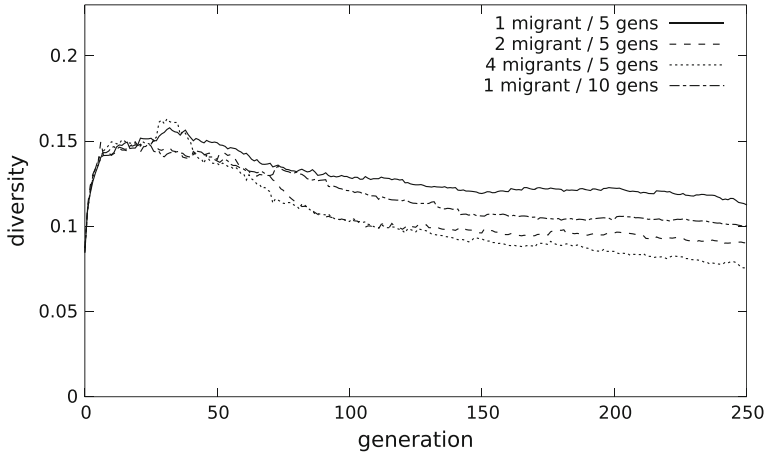

Fig. 13 Average inner island diversity obtained by simultaneously migrating a different number of individuals at each step on the optimization of Morse clusters of 61 atoms

Figure 13 shows that increasing the number of migrants causes the average AID to slowly degrade. The higher the number of individuals migrating at the same time gets, the lower the AID gets, specially toward the end. This is, once again, due to the nature of the migration operations. When migrating 1 individual at each step, the system seems able to incorporate new genes while maintaining a balanced AID for most of the run and, at the same time, to progress and gradually find better solutions as shown by Fig. 10. However, when that number increases, not only do the islands seem to make a poorer use of new immigrants, the possible benefits that higher migration rates may bring don't seem to counterbalance the inherent loss of AID.

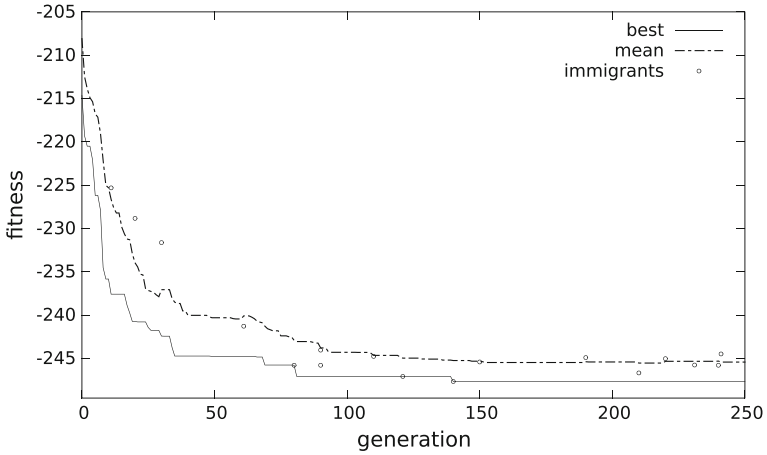

Fig. 14 Fitness evolution and included immigrants on an island, migrating 1 individual every 10th generation on the optimization of Morse clusters composed of 61 atoms

4.3.5 Migration intervals

Two intervals are addressed, either migrating individuals every 5th or 10th generation. Figure 13 suggests that in either case, the model is able to maintain average AID values above initial levels, despite the scenario with longer intervals showing a lower average AID . Furthermore, both Figs. 10 and 14 show that in both scenarios, migration is able to promote progress in target islands. However, Fig. 14 also shows that intervals of 10 generations causes islands to undergo long period of stasis where diversity is maintained in each island by the steady-state model but productive reproduction is very unlikely to occur.

While longer intervals could be appropriate for systems with larger islands, intervals of five generations seem appropriate to the size of the islands used. In this case, such periods of isolation seem long enough for each island to exploit incoming genetic material and reach new solutions. Arguably, if longer experiments were to be run with longer migration intervals, competitive results could be achieved. However, given the stop criterion applied, periods of stasis are no more than a waste of resources and evaluations.

4.4 Experimental results

Table 1 presents an overview of the results obtained by our final setup: 10 islands with 10 individuals each, connected through a dynamic bidirectional ring topology, where each island migrates 1 individual every 5 generations, abiding to a *random-rest* migration policy. It includes, for each instance N of the problem, the number of runs where the putative optimum was found. Lines labeled *Panmictic* present the results reported by Pereira et al. [61] using the same hybrid steady-state EA as applied here but on a panmictic system. Lines labeled *IM* correspond to the results obtained by our system, using the best set of parameters found. The symbol "-" was used to mark instances where no runs successfully reached the putative optima.

A first analysis of the results shows that our approach was able to find all the putative optima on clusters ranging from 41 to 80 atoms. Moreover, the IM seems to perform slightly better on a global level than the panmictic approach. It was able to yield a higher success rate

Table 1 Hit rate obtained by the panmictic and IM approaches on clusters ranging from 41 to 80 particles

N	41	42	43	44	45	46	47	48	49	50
IM	18	19	10	5	6	3	6	20	24	7
Panmictic	15	12	14	7	5	9	2	14	18	5
N	51	52	53	54	55	56	57	58	59	60
IM	4	11	9	10	4	9	8	**15**	5	**7**
Panmictic	7	6	5	12	6	11	8	4	2	–
N	61	62	63	64	65	66	67	68	69	70
IM	1	8	**14**	11	15	10	5	3	9	4
Panmictic	1	12	6	11	8	8	5	4	6	8
N	71	72	73	74	75	76	77	78	79	80
IM	7	10	6	**7**	3	5	3	5	2	1
Panmictic	6	7	6	1	2	4	7	3	6	5

on 22 instances, whereas the panmictic approach has a higher rate on 13 instances and ties were found on 5 other. A pairwise proportions test [76] was performed with a significance level of 0.01 and instances where significant differences were found were highlighted with bold in Table 1. No significant differences were found except on four instances where the IM performed significantly better, ergo backing up our assumption.

A more in-depth analysis of the results exposes behavioural differences between the panmictic and the IM approaches. While the IM performed well on some instances of the problem where the panmictic model performed poorly, the opposite happened as well. The explanation may lie on the putative solutions of the tested instances corresponding to both different geometric motifs and PES with different properties, which suggests that the two approaches may be better suited for instances with different characteristics. A reflection on the workings of both approaches further highlights behavioural differences between them. In IMs, both the selection and replacement strategies act differently from the original model. Having the population divided into subpopulations and reducing the size of the tournament will give individuals a higher chance of being selected for reproduction and the replacement strategy will give their offspring a higher chance of being accepted as they will only be measured against the individuals on their current island, facing therefore less competition. Furthermore, the migration policy will contribute as well by giving weak individuals the chance to migrate to islands where they will be more valuable. Finally, d_{min} is set independently on each island which means that their replacement strategies will ultimately behave differently along the runs. These characteristics help promote inter-island diversity but the lack of a global replacement strategy will have an impact as well and, considering the duplicative and deleterious nature of migration, is bound to cause some loss of diversity even if the migration policy helps lessen the consequences.

This discussion exposes how important it is to correctly design an IM in order to make it a valuable model. The final settings chosen for the IM were the result of the study in Sect. 4.3. The analysis shows that the average AID is maintained above initial values and more or less stable along the runs despite the effects of migration and the breeding of dissimilar offspring. The discussion also supports that this is the result of a careful balance between all aspects

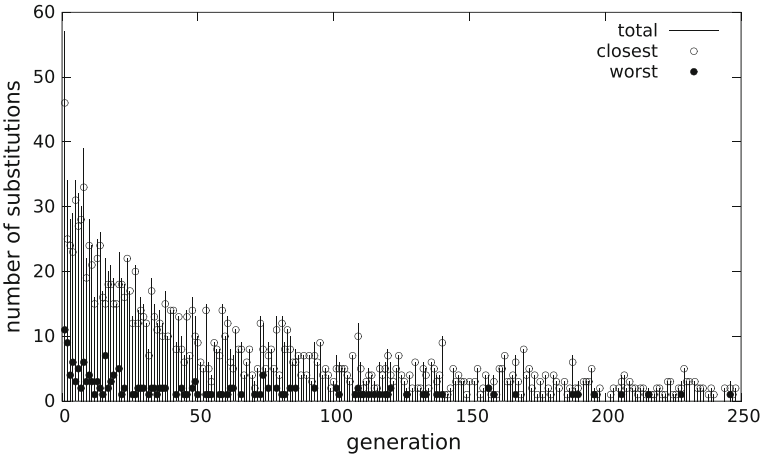

Fig. 15 Acceptance rate of the replacement strategy adopted by the EA on the optimization of Morse clusters composed of 61 atoms

of IMs and how volatile they are to changes in their design. In this particular context and for the reasons mentioned above, it seems important to keep migration and connections between islands to a minimum in order to keep AID high and duplication of genes low. Figure 10 shows that even though migration is kept to a minimum, the system is able to take advantage of such a setup and promote progress. Figure 15, which corresponds to one run, backs up this assumption on a global level, showing that newly bred individuals are able to integrate their subpopulations consistently, suggesting that both diversity and migration are maintained at suitable rates along the run.

When designing an IM there are plenty of design choices that should be taken into consideration and whose behaviour may well be problem dependent. The parametric study in Sect. 4.3 gives us insight of good design choices for the particular context of CGO and its characteristic PES. Firstly, it's on our best interest to split the population into as many islands as possible to take profit from parallelization. However, the size of the islands should be kept in mind so that their inner workings is not disrupted due to them being too small. Secondly, islands should be kept as sparse as possible and communication routes should be kept to a minimum, so that stronger barriers are maintained between individuals and the spreading of genes is kept under control in order to maintain a higher diversity. Still, stochasticity plays an important role and unidirectional topologies should be avoided since they can produced chained reactions. Thirdly, applying an appropriate migration policy that gives any individual an equal chance to participate in other islands while keeping duplicative individuals out is crucial to avoid conquering effects. Fourthly, a small migration rate on short intervals is enough to promote progress in target islands and allows them to exploit new genes while avoiding moments of stasis which are a waste of resources and evaluations.

5 Conclusions

IMs are distributed EAs that divide the population into islands that undergo isolated evolution punctuated by the migration of individuals between them. Apart from allowing the distribution of computational effort, IMs result on a distinct behaviour from panmictic models. While

isolation makes them different from standard EAs, migration makes them different from separate runs of Standard EAs.

The characteristics of IMs support that they can help maintain diversity on a global population and at the same time provide ecological opportunities for weaker individuals to share their genetic material. Given our current knowledge on CGO and the characteristics of the PES, IMs may be an appropriate approach to tackle such problems, hopefully bringing parallelization as well as behavioural and performance benefits. Designing an IM requires however a number of choices regarding the number of islands, topology, migration policies, rates and intervals. Those decisions are not straightforward and may have different outcomes in different contexts.

This study presents an extensive discussion on the impact of different design choices on the behaviour of IMs when tackling CGO problems. It shows how volatile IMs are to changes in their design and how important it is to balance these choices in order to achieve the expected behaviour and performance. In this particular context, it suggests a setup that helps make the model an asset and keep migration a valuable operation. By using this setup, the model was able to find all putative optima on Morse clusters ranging from 41 to 80 atoms. The overall results are promising and back up that IMs achieve a slightly more robust behaviour than a panmictic approach, being therefore a competitive parallel alternative to standard sequential EAs.

There are several possible avenues for future research. On the one hand, it would be important to apply IMs to other cluster geometry optimization variants, e.g., molecular clusters, in order to confirm if the general guidelines suggested in this study generalize well. On the other hand, it would also be relevant to investigate the possibility of granting the IMs the ability to self-adapt to different optimization scenarios. This would not only alleviate the burden of carefully making the necessary design choices but would also allow the islands to diverge from each other and adopt different behaviours, possibly enhancing their global performance.

References

1. http://physchem.ox.ac.uk/doye/jon/structures/Morse.html
2. Alba, E., Tomassini, M.: Parallelism and evolutionary algorithms. Evol. Comput. IEEE Trans. **6**(5), 443–462 (2002)
3. Alba, E., Troya, J.: A survey of parallel distributed genetic algorithms. Complexity **4**(4), 31–52 (1999)
4. Bethke, A.: Comparison of genetic algorithms and gradient-based optimizers on parallel processors: efficiency of use of processing capacity. Technical Report No. 197, University of Michigan, Logic of Computers Group, Ann Arbor (1976)
5. Braun, H.: On solving travelling salesman problems by genetic algorithms. In: Schwefel, H.P., Männer, R. (eds.) Parallel Problem Solving From Nature, pp. 129–133. Springer, Berlin (1990)
6. Cantú-Paz, E.: A survey of parallel genetic algorithms. Calculateurs paralleles, reseaux et systems repartis **10**(2), 141–171 (1998)
7. Cantú-Paz, E.: Migration policies, selection pressure, and parallel evolutionary algorithms. J. Heuristics **7**, 311–334 (2001)
8. Cantú-Paz, E., Goldberg, D.E.: Modeling idealized bounding cases of parallel genetic algorithms. In: Koza, J., Deb, K., Dorigo, M., Fogel, D., Garzon, M., Iba, H., Riolo, R. (eds.) Proceedings of the Second Annual Conference on Genetic Programming. Morgan Kaufmann, San Francisco (1997)
9. Cantú-Paz, E., Mejia-Olvera, M.: Experimental results in distributed genetic algorithms. In: International Symposium on Applied Corporate Computing, pp. 99–108 (1994)
10. Cassioli, A., Locatelli, M., Schoen, F.: Global optimization of binary Lennard–Jones clusters. Optim. Methods Softw. **24**, 819–835 (2009)
11. Cassioli, A., Locatelli, M., Schoen, F.: Dissimilarity measures for population-based global optimization algorithms. Comput. Optim. Appl. **45**, 257–281 (2010)

12. Cheng, L.: A connectivity table for cluster similarity checking in the evolutionary optimization method. Chem. Phys. Lett. **389**, 309–314 (2004)
13. Cheng, L., Feng, Y., Yang, J., Yang, J.: Funnel hopping: searching the cluster potential energy surface over the funnels. J. Chem. Phys. **130**(21), 214,112 (2009)
14. Cheng, L., Yang, J.: Global minimum structures of morse clusters as a function of the range of the potential: $81 \leq n \leq 160$. J. Phys. Chem. A **111**(24), 5287–5293 (2007)
15. Cohoon, J., Hegde, S., Martin, W., Richards, D.: Punctuated equilibria: a parallel genetic algorithm. In: Proceedings of the Second International Conference on Genetic Algorithms and their applications, pp. 148–154. L. Erlbaum Associates Inc., Hillsdale (1987)
16. Cohoon, J.P., Martin, W.N., Richards, D.S.: A multi-population genetic algorithm for solving the k-partition problem on hyper-cubes. In: Proceedings of the Fourth International Conference on Genetic Algorithms, vol. 91, pp. 244–248 (1991)
17. Darby, S., Mortimer-Jones, T.V., Johnston, R.L., Roberts, C.: Theoretical study of Cu–Au nanoalloy clusters using a genetic algorithm. J. Chem. Phys. **116**, 1536–1550 (2002)
18. De Jong, K.: An analysis of the behavior of a class of genetic adaptive systems. Ph.D. thesis (1975)
19. De Jong, K.A., Sarma, J.: Generation gaps revisited. In: Foundations of Genetic Algorihtms 2, pp. 19–28 (1992)
20. Deaven, D.M., Ho, K.M.: Molecular geometry optimization with a genetic algorithm. Phys. Rev. Lett. **75**, 288–291 (1995)
21. Doye, J.P., Wales, D.J.: Structural consequences of the range of the interatomic potential a menagerie of clusters. J. Chem. Soc. Faraday Trans. **93**(24), 4233–4243 (1997)
22. Doye, J.P.K., Leary, R.H., Locatelli, M., Schoen, F.: Global optimization of Morse clusters by potential energy transformations. INFORMS J. Comput. **16**, 371–379 (2004)
23. Dugan, N., Erkoc, S.: Genetic algorithm—Monte Carlo hybrid geometry optimization method for atomic clusters. Comput. Mater. Sci. **45**(1), 127–132 (2009)
24. Eldredge, N., Gould, S.J.: Punctuated equilibria: an alternative to phyletic gradualism. In: Schopf, T.J.M. (ed.) Models in Paleobiology, pp. 82–115. Freeman Cooper, San Francisco (1972)
25. Feng, Y., Cheng, L., Liu, H.: Putative global minimum structures of morse clusters as a function of the range of the potential: $161 \leq n \leq 240$. J. Phys. Chem. A **113**(49), 13651–13655 (2009)
26. Fernández, F., Tomassini, M., Punch, W., Sánchez, J.: Experimental study of multipopulation parallel genetic programming. In: Poli, R., Banzhaf, W., Langdon, W.B., Miller, J.F., Nordin, P., Fogarty, T.C. (eds.) Genetic Programming, vol. 1802, pp. 283–293. Springer, Berlin (2000)
27. Fernández, F., Tomassini, M., Vanneschi, L.: An empirical study of multipopulation genetic programming. Genet. Program. Evol. Mach. **4**, 21–51 (2003)
28. Goldberg, D.E., Deb, K.: A comparative analysis of selection schemes used in genetic algorithms. In: Rawlins, G. (ed.) Foundations of Genetic Algorithms, pp. 69–93. Morgan Kaufmann, San Mateo (1991)
29. Grefenstette, J.J.: Parallel Adaptive Algorithms for Function Optimization: (preliminary Report). Vanderbilt University, Computer Science Department (1981)
30. Gregurick, S.K., Alexander, M.H., Hartke, B.: Global geometry optimization of $(Ar)_n$ and $B(Ar)_n$ clusters using a modified genetic algorithm. J. Chem. Phys. **104**, 2684–2691 (1996)
31. Grosso, A., Locatelli, M., Schoen, F.: A population-based approach for hard global optimization problems based on dissimilarity measures. Math. Program. **110**, 373–404 (2007)
32. Grosso, P.B.: Computer simulations of genetic adaptation: parallel subcomponent interaction in a multilocus model. Ph.D. thesis (1985)
33. Guimarães, F.F., Belchior, J.C., Johnston, R.L., Roberts, C.: Global optimization analysis of water clusters $(h_2o)_n$ $(11 \leq n \leq 13)$ through a genetic evolutionary approach. J. Chem. Phys. **116**, 8327–8333 (2002)
34. Hartke, B.: Global geometry optimization of clusters using genetic algorithms. J. Phys. Chem. **97**(39), 9973–9976 (1993)
35. Hartke, B.: Global cluster geometry optimization by a phenotype algorithm with niches: location of elusive minima, and low-order scaling with cluster size. J. Comput. Chem. **20**(16), 1752–1759 (1999)
36. Hartke, B.: Application of evolutionary algorithms to global cluster geometry optimization. In: Johnston, R.L. (ed.) Applications of Evolutionary Computation in Chemistry, vol. 110, pp. 33–53. Springer, Berlin (2004)
37. Hobday, S., Smith, R.: Optimisation of carbon cluster geometry using a genetic algorithm. J. Chem. Soc. Faraday Trans. **93**(22), 3919–3926 (1997)
38. Holland, J.: A universal computer capable of executing an arbitrary number of sub-programs simultaneously. In: Eastern Joint IRE-AIEE-ACM Computer Conference, pp. 108–113. ACM (1959)
39. Holland, J.: Iterative circuit computers. In: Western Joint IRE-AIEE-ACM Computer Conference, pp. 259–265. ACM (1960)

40. Iwamatsu, M.: Global geometry optimization of silicon clusters using the space-fixed genetic algorithm. J. Chem. Phys. **112**, 10976–10983 (2000)
41. Johnston, R.L.: Evolving better nanoparticles: genetic algorithms for optimising cluster geometries. Dalton Trans. **22**, 4193–4207 (2003)
42. Jones, J.E.: On the determination of molecular fields. II. From the equation of state of a gas. R. Soc. Lond. Proc. Ser. A **106**, 463–477 (1924)
43. Krasnogor, N.: Towards robust memetic algorithms. In: Hart, W.E., Krasnogor, N., Smith, J.E. (eds.) Recent Advances in Memetic Algorithms, pp. 185–207. Springer, Berlin (2004)
44. Krasnogor, N., Blackburne, B., Burke, E.K., Hirst, J.D.: Multimeme algorithms for protein structure prediction. In: Merelo, J.J., Adamidis, P., Beyer, H.G. (eds.) Parallel Problem Solving from Nature–PPSN VII, pp. 769–778. Springer, Berlin (2002)
45. Lee, J., Lee, I.H., Lee, J.: Unbiased global optimization of Lennard-Jones clusters for $n \leq 201$ using the conformational space annealing method. Phys. Rev. Lett. **91**(8), 080,201 (2003)
46. Leitão, A., Pereira, F.B., Machado, P.: Enhancing cluster geometry optimization with island models. In: Proceedings of the IEEE Congress on Evolutionary Computation, CEC 2012, Brisbane, Australia, June 10–15, 2012, pp. 1–8. IEEE (2012)
47. Lennard-Jones, J.E.: Cohesion. Proc. Phys. Soc. **43**, 461–482 (1931)
48. Lin, S., Punch III, W., Goodman, E.: Coarse-grain parallel genetic algorithms: categorization and new approach. In: Proceedings of the 6th IEEE Symposium on Parallel and Distributed Processing, pp. 28–37 (1994)
49. Liu, D.C., Nocedal, J.: On the limited memory bfgs method for large scale optimization. Math. Program. **45**, 503–528 (1989)
50. Locatelli, M., Schoen, F.: Fast global optimization of difficult Lennard-Jones clusters. Comput. Optim. Appl. **21**, 55–70 (2002)
51. Locatelli, M., Schoen, F.: Efficient algorithms for large scale global optimization: Lennard-Jones clusters. Comput. Optim. Appl. **26**, 173–190 (2003)
52. Lozano, M., Herrera, F., Krasnogor, N., Molina, D.: Real-coded memetic algorithms with crossover hill-climbing. Evol. Comput. **12**(3), 273–302 (2004)
53. Mahfoud, S.W.: Niching methods for genetic algorithms. Ph.D. thesis (1995)
54. Marques, J.M.C., Llanio-Trujillo, J.L., Abreu, P.E., Pereira, F.B.: How different are two chemical structures? J. Chem. Inf. Model. **50**(12), 2129–2140 (2010)
55. Martin, W.N., Lienig, J., Cohoon, J.P.: Island (migration) Models: Evolutionary Algorithms Based on Punctuated Equilibria. Oxford University Press Inc, New York (1997)
56. Morse, P.M.: Diatomic molecules according to the wave mechanics. ii. vibrational levels. Phys. Rev. **34**(1), 57–64 (1929)
57. Munetomo, M., Takai, Y., Sato, Y.: An efficient migration scheme for subpopulation-based asynchronously parallel genetic algorithms. In: Proceedings of the 5th International Conference on Genetic Algorithms, p. 649. Morgan Kaufmann Publishers Inc. (1993)
58. Niesse, J.A., Mayne, H.R.: Global geometry optimization of atomic clusters using a modified genetic algorithm in space-fixed coordinates. J. Chem. Phys. **105**, 4700–4706 (1996)
59. Pelta, D.A., Krasnogor, N.: Multimeme algorithms using fuzzy logic based memes for protein structure prediction. In: Recent Advances in Memetic Algorithms, pp. 49–64. Springer, Berlin (2005)
60. Pereira, F., Marques, J.: A self-adaptive evolutionary algorithm for cluster geometry optimization. In: Proceedings of the 8th International Conference on Hybrid Intelligent Systems, pp. 678–683 (2008)
61. Pereira, F., Marques, J.: A study on diversity for cluster geometry optimization. Evol. Intell. **2**, 121–140 (2009)
62. Pereira, F., Marques, J., Leitao, T., Tavares, J.: Analysis of locality in hybrid evolutionary cluster optimization. In: Evolutionary Computation, 2006. CEC 2006. IEEE Congress on, pp. 2285–2292 (2006)
63. Pereira, F.B., Marques, J., Leitao, T., Tavares, J.: Designing efficient evolutionary algorithms for cluster optimization: a study on locality. In: Advances in Metaheuristics for Hard Optimization. Natural Computing Series, pp. 223–250. Springer, Berlin (2008)
64. Pullan, W.: Genetic operators for the atomic cluster problem. Comput. Phys. Commun. **107**(1–3), 137–148 (1997)
65. Pullan, W.: An unbiased population-based search for the geometry optimization of Lennard-Jones clusters: $2 \leq n \leq 372$. J. Comput. Chem. **26**(9), 899–906 (2005)
66. Rata, I., Shvartsburg, A.A., Horoi, M., Frauenheim, T., Siu, K.W.M., Jackson, K.A.: Single-parent evolution algorithm and the optimization of Si clusters. Phys. Rev. Lett. **85**, 546–549 (2000)
67. Roberts, C., Johnston, R.L., Wilson, N.T.: A genetic algorithm for the structural optimization of morse clusters. Theor. Chem. Acc. Theory Comput. Model. **104**, 123–130 (2000)

68. Sastry, K., Xiao, G.: Cluster optimization using extended compact genetic algorithm. Urbana **51**, 61,801 (1989)
69. Sekaj, I.: Robust parallel genetic algorithms with re-initialisation. In: Yao, X., Burke, E.K., Lozano, J.A., Smith, J., Guervós, J.J.M., Bullinaria, J.A., Rowe, J.E., Tiňo, P., Kabán, A., Schwefel, H.P. (eds.) Parallel Problem Solving from Nature—PPSN VIII, vol. 3242, pp. 411–419. Springer, Berlin (2004)
70. Shao, X., Cheng, L., Cai, W.: A dynamic lattice searching method for fast optimization of Lennard-Jones clusters. J. Comput. Chem. **25**(14), 1693–1698 (2004)
71. Skolicki, Z.: An analysis of island models in evolutionary computation. In: Proceedings of the 2005 Workshops on Genetic and Evolutionary Computation, GECCO '05, pp. 386–389. ACM (2005)
72. Skolicki, Z., De Jong, K.: The influence of migration sizes and intervals on island models. In: Proceedings of the 2005 Conference on Genetic and Evolutionary Computation, GECCO '05, pp. 1295–1302. ACM (2005)
73. Smirnov, B.M., Strizhev, A.Y., Berry, R.S.: Structures of large Morse clusters. J. Chem. Phys. **110**, 7412–7420 (1999)
74. Smith, J.: On replacement strategies in steady state evolutionary algorithms. Evol. Comput. **15**, 29–59 (2007)
75. Stillinger, F.H.: Exponential multiplicity of inherent structures. Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Top. **59**, 48–51 (1999)
76. Taillard, É.D., Waelti, P., Zuber, J.: Few statistical tests for proportions comparison. Eur. J. Oper. Res. **185**(3), 1336–1350 (2008)
77. Tanese, R.: Parallel genetic algorithms for a hypercube. In: Proceedings of the Second International Conference on Genetic Algorithms and Their Application, pp. 177–183. L. Erlbaum Associates Inc. (1987)
78. Tanese, R.: Distributed genetic algorithm. In: Proceedings of the Third International Conference on Genetic Algorithms, pp. 434–439 (1989)
79. Tanese, R.: Distributed genetic algorithms for function optimization. Tech. rep. (1989)
80. Tsai, C.J., Jordan, K.D.: Use of the histogram and jump-walking methods for overcoming slow barrier crossing behavior in Monte Carlo simulations: Applications to the phase transitions in the $(Ar)_{13}$ and $(H_2O)_8$ clusters. J. Chem. Phys. **99**, 6957–6970 (1993)
81. Whitley, D., Rana, S., Heckendorn, R.: The island model genetic algorithm: on separability, population size and convergence. J. Comput. Inf. Technol. **7**, 33–48 (1999)
82. Whitley, D., Starkweather, T.: Genitor II: a distributed genetic algorithm. J. Exp. Theor. Artif. Intell. **2**(3), 189–214 (1990)
83. Whitley, L.D.: The genitor algorithm and selection pressure: why rank-based allocation of reproductive trials is best. In: Proceedings of the 3rd International Conference on Genetic Algorithms, pp. 116–123 (1989)
84. Xiao, Y., Williams, D.E.: Genetic algorithm: a new approach to the prediction of the structure of molecular clusters. Chem. Phys. Lett. **215**, 17–24 (1993)
85. Zeiri, Y.: Prediction of the lowest energy structure of clusters using a genetic algorithm. Phys. Rev. E **51**(4), R2769 (1995)
86. Zhao, J., Xie, R.H.: Genetic algorithms for the geometry optimization of atomic and molecular clusters. J. Comput. Theor. Nanosci. **1**(2), 117–131 (2004)