

Empirical Analysis of Diversity-Preserving Mechanisms on Example Landscapes for Multimodal Optimisation

Edgar Covantes $\operatorname{Osuna}^{(\boxtimes)}$ and Dirk Sudholt

University of Sheffield, Sheffield S1 4DP, UK {ecovantes1,d.sudholt}@sheffield.ac.uk

Abstract. Many diversity-preserving mechanisms have been developed to reduce the risk of premature convergence in evolutionary algorithms and it is not clear which mechanism is best. Most multimodal optimisation problems studied empirically are restricted to real-parameter problems and are not accessible to theoretical analysis, while theoreticians analyse the simple bimodal function TwoMAX. This paper looks to narrow the gap between both approaches. We perform an extensive empirical study involving 9 common diversity mechanisms on Jansen-Zarges multimodal function classes (Jansen and Zarges, PPSN 2016) that allow to control important problem features while still being amenable to theoretical analysis. This allows us to study functions with various degrees of multimodality and to explain the results in the light of previous theoretical works. We show which mechanisms are able to find and maintain a large number of distant optima, escape from local optima, and which fail to locate even a single peak.

Keywords: Diversity-preserving mechanisms Evolutionary algorithms \cdot Multimodal optimisation \cdot Empirical study Theory

1 Introduction

Many optimisation problems are multimodal, and finding global optima or high-quality local optima can become a challenge for any optimisation algorithm [15,17]. Evolutionary algorithms (EAs) are well suited to dealing with multimodal problems due to their use of a population. A diverse population can explore several hills in the fitness landscape simultaneously and offer several good solutions to the user, a feature desirable for decision making, in multi-objective

© Springer Nature Switzerland AG 2018

The authors would like to thank the Consejo Nacional de Ciencia y Tecnología - CONACYT (the Mexican National Council for Science and Technology) for the financial support under the grant no. 409151 and registration no. 264342.

A. Auger et al. (Eds.): PPSN 2018, LNCS 11102, pp. 207–219, 2018. https://doi.org/10.1007/978-3-319-99259-4_17

optimisation and in dynamic optimisation. However, a major difficulty when applying EAs is that the population may converge to a sub-optimal individual before the fitness landscape is explored properly.

Many diversity-preserving mechanisms have been developed to reduce the risk of such premature convergence, including fitness sharing, clearing, avoiding duplicates, fitness diversity, crowding methods, restricted tournament selection, and many others [4,7,16,18]. These mechanisms seek to enable EAs to visit many different regions of the search space and generate solutions that differ from those seen before [10]. Given such a variety of mechanisms to choose from, it is often not clear which mechanism is the best choice for a particular problem.

Previous empirical analyses have considered real-parameter multimodal optimisation problems [5] like the 4 one-dimensional, five-peaked, sinusoidal, multimodal functions called M_{1-4} defined in [11, Sect. 5.3]. The single variable x is restricted to the real-value range [0, 1] encoded using binary representation and decoded by interpreting the bit string as unsigned binary integer and dividing it by $2^n - 1$, where n is the length of the bit string. Other studies used Gray codes [15]. The drawback of real-valued encodings is that it is not obvious how phenotypic features such as local optima appear in genotype space; for example what Hamming distance local optima have and how likely it is that mutation jumps from one basin of attraction to another. This makes the analysis of the population dynamics a very challenging task for theoretical analysis.

Previous theoretical studies that considered multimodal problems [1,2,6,13] compared the expected running time of different diversity mechanisms to find both global optima on the bimodal function TwoMAX $(x) := \{\sum_{i=1}^{n} x_i, n - \sum_{i=1}^{n} x_i\}$ that has a straightforward mapping between genotypes (bit strings) and phenotypes (number of 1-bits). TWOMAX consists of two different symmetric slopes (or branches) ZEROMAX and ONEMAX with 0^n and 1^n as global optima, respectively, and the goal is to evolve a population that contains both optima¹. This is challenging as the two optima have the maximum possible Hamming distance. Studying TWOMAX led to insights into the capabilities and weaknesses of various diversity mechanisms (see Sect. 2 and Sudholt's survey [19]), however a question left open is how diversity mechanisms deal with many local optima.

Jansen and Zarges [9] addressed the need for more general classes of functions for multimodal optimisation for both empirical and theoretical analysis by defining multimodal landscapes with straightforward binary encodings (see Sect. 3). We seek to narrow the gap between theory and practice by performing an empirical study on the Jansen-Zarges multimodal function classes, complementing existing rigorous theory for TWOMAX [1,2,6,13] with empirical results of more complex functions with multiple different peaks, slopes and heights. The main goal is to provide insights into the working principles of these mechanisms by testing their ability to find and maintain many local optima in the population as well as their ability to escape from local optima with different basins of

¹ In [6] an additional fitness value for 1^n was added to distinguish between a local optimum 0^n and a unique global optimum 1^n . The discussion of previous work from [6] is adapted to a TWOMAX with two optima [1,2,13] (see Table 1 and [19]).

\mathbf{A}	gorit	hm 1	. (L	$\iota + 1$) EA	
			- \ <i>r</i>		/	-

- 1: Initialise P with μ individuals chosen uniformly at random.
- 2: while stopping criterion not met do
- 3: Choose $x \in P$ uniformly at random.
- 4: Create y by flipping each bit in x independently with probability 1/n.
- 5: Choose $z \in P$ uniformly at random from all individuals with worst fitness in P.
- 6: **if** $f(y) \ge f(z)$ **then** $P = P \setminus \{z\} \cup \{y\}.$

attraction. We use previous theoretical results to inform the choice of algorithm parameters and to discuss in how far our empirical results agree or disagree with theoretical results obtained for TwoMAX.

2 Diversity Mechanisms and Previous Results for TwoMax

Following previous theoretical work, we consider diversity mechanisms embedded in a bare-bones EA: a $(\mu+1)$ EA with uniform parent selection in a population of size μ , using standard mutation (and no crossover). The offspring y replaces a worst individual z from the population if $f(y) \ge f(z)$ (see Algorithm 1).

Table 1 summarises previous work for the $(\mu+1)$ EA with diversity mechanisms on TWOMAX (details for each $(\mu+1)$ EA variant can be found in the respective publications and in [19]). Some mechanisms succeed in finding both optima on TWOMAX efficiently, that is, in (expected) time $O(\mu n \log n)$. Others have a very low success probability. Friedrich, Oliveto, Sudholt, and Witt [6] showed that the plain $(\mu+1)$ EA (PL, Algorithm 1) is not able to maintain individuals on both branches for a long time; the whole population is likely to converge to one of the two peaks [6, Theorem 1].

Introducing other simple mechanisms to the $(\mu+1)$ EA like avoiding genotype duplicates (NGD), where (after initialisation) identical copies of individuals are prevented from entering the population [6, Algorithm 2], and avoiding fitness duplicates (NFD), rejecting individuals with the same fitness [6, Algorithm 3], are not able to prevent the extinction of one branch, ending with the population converging to one optimum with high probability [6, Theorem 2 and 3, resp.].

In the $(\mu+1)$ EA with probabilistic crowding (PC) [2, Algorithm 1], an offspring competes with its parent and the survivor is chosen with a probability proportional to their fitness [12]. Covantes Osuna and Sudholt showed that this mechanism is unable to evolve solutions of significantly higher fitness than that obtained during initialisation (or, equivalently, through random search), even when given exponential time [2, Theorem 2.2]. The reason is that fitnessproportional selection between parent and offspring results in an almost uniform choice as both have very similar fitness, hence fitness-proportional selection degrades to uniform selection for replacement. In *deterministic crowding* (*DC*), the offspring competes against its parent and replaces it if the offspring is at least as good [11]. For the (μ +1) EA with deterministic crowding [6, Algorithm 4],

Table 1. Overview of runtime analyses for the $(\mu+1)$ EA with diversity mechanisms on TwoMAX, adapted from [2]. The success probability is the probability of finding both optima within (expected) time $O(\mu n \log n)$. Conditions include restrictions on the population size μ , the sharing/clearing radius σ , the niche capacity κ , window size w, and $\mu' := \min(\mu, \log n)$.

	Diversity mechanism	Success prob.	Conditions
PL	Plain $(\mu+1)$ EA [6]	<i>o</i> (1)	$\mu = o(n/\log n)$
NGD	No Genotype Duplicates [6]	<i>o</i> (1)	$\mu = o(\sqrt{n})$
NFD	No Fitness Duplicates [6]	o(1)	$\mu = \operatorname{poly}(n)$
PC	Probabilistic Crowding [2]	$2^{-\Omega(n)}$	all μ
DC	Deterministic Crowding [6]	$1 - 2^{-\mu + 1}$	all μ
RTS	Restricted Tournament Selection [2]	$\geq 1 - 2^{-\mu'+3}$	$w \geq 2.5 \mu \ln n$
PFS	Population-based Fitness Sharing [*] ($\sigma = n/2$) [6]	1	$\mu \ge 2$
FS	Individual-based Fitness Sharing [*] ($\sigma = n/2$) [13]	1	$\mu \geq 3$
\mathbf{CL}	Clearing $(\sigma = n/2)$ [1]	1	$\mu \geq \kappa n^2$

*Fitness sharing uses phenotypic sharing based on the number of ones

this mechanism with a sufficiently large population is able to reach both optima with high probability in expected time $O(\mu n \log n)$ [6, Theorem 4].

In restricted tournament selection (RTS), for every offspring created, RTS selects uniformly at random (u. a. r.) w (window size) members from the population with replacement. Each offspring competes with the closest element from this set and the offspring replaces it if its fitness is at least as good [8]. For the $(\mu+1)$ EA with RTS [2, Algorithm 2], the mechanism succeeds in finding both optima of TwoMAX in the same way as deterministic crowding, provided that w is chosen large enough [2, Theorem 3.1]. However, if w is too small, then it cannot prevent one branch taking over the other, leading to exponential running times with high probability [2, Theorem 3.4].

Fitness sharing derates the real fitness of an individual by an amount that represents the similarity to other population members. A population-based fitness sharing (PFS) approach [6, Algorithm 5], constructing the best possible new population amongst parents and offspring is able to find both optima in expected time $O(\mu n \log n)$ for any population size $\mu \ge 2$ [6, Theorem 5]. A drawback of this approach is that examining all possible new populations is computationally expensive. The conventional fitness sharing (FS), where selection is based on individuals, was studied by Oliveto, Sudholt, and Zarges [13, Algorithm 1]. Population size $\mu = 2$ is not sufficient to find both optima in polynomial time; the success probability is only $1/2 - \Omega(1)$ [13, Theorem 1]. However, with $\mu \ge 3$, the $(\mu+1)$ EA finds both optima in expected time $O(\mu n \log n)$ [13, Theorem 3]. In all the above results, fitness sharing used a phenotypic distance: the distance between two search points x and y is the absolute difference in their number of ones. This choice is tailored to TWOMAX and is not applicable in our scenario. Hence our experiments must rely on fitness sharing with genotypic distances (Hamming distance), for which no runtime analyses are available.

In *clearing (CL)*, individuals are sorted in decreasing fitness and are processed in this order. Each individual is compared against other individuals according to its fitness with distance $< \sigma$ (*clearing radius*) which determines if both individuals belong to the same subpopulation (niche) or not. Then, the procedure iterates through all remaining individuals (i.e., those with lower or equal fitness) that haven't been cleared yet, until κ (*niche capacity*) best individuals (also called winners) have been found, and all remaining individuals from the same niche are cleared to the minimum fitness value possible [14]. Finally, the individuals with best fitness are selected (set of winners) and individuals coming from the new generation are preferred [1, Algorithm 1 and 2]. Clearing, with a clearing radius of $\sigma = n/2$, niche capacity $\kappa = 1$, and $\mu \ge \kappa n^2$ is able to find both optimal in expected time $O(\mu n \log n)$ [1, Theorem 5.6].

3 Jansen-Zarges Multimodal Function Classes

Jansen and Zarges [9] introduced several problem classes spanned by k peaks $p_1, p_2, \ldots, p_k \in \{0, 1\}^n$ for an arbitrary number $k \in \mathbb{N}$ of peaks. Each peak *i* has a position $p_i \in \{0,1\}^n$, a slope $a_i \in \mathbb{R}^+$, and an offset $b_i \in \mathbb{R}_0^+$. The fitness value of a search point depends on peaks in its vicinity as defined as follows.

Definition 1 (Definition 3 in [9]). Let $k \in \mathbb{N}$ and k peaks $(p_1, a_1, b_1), (p_2, a_2, b_1)$ $(b_2), \ldots, (p_k, a_k, b_k)$ be given, then

- $\begin{array}{l} \ \mathrm{JZ}_1(x) := a_{\mathrm{cp}(x)} \cdot \mathrm{G}(x, p_{\mathrm{cp}(x)}) + b_{\mathrm{cp}(x)}, \ called \ nearest \ peak \ function, \\ \ \mathrm{JZ}_2(x) := \max_{i \in \{1, 2, \dots, k\}} a_i \cdot \mathrm{G}(x, p_i) + b_i, \ called \ weighted \ nearest \ peak \ function, \end{array}$

where $cp(x) := argmin_{i \in \{1,2,\dots,k\}} H(x,p_i)$ is defined by the closest peak to a search point, and $G(x, p_i) := n - H(x, p_i)$ indicates the proximity of x to p_i .

For the nearest peak function, $JZ_1(x)$, the fitness of a search point x is determined by the proximity to the closest peak i = cp(x) along with its slope a_i and its offset b_i . In cases where multiple *i* minimise $H(x, p_i)$, *i* should additionally maximise $a_i \cdot \mathbf{G}(x, p_i) + b_i$.

The weighted nearest peak function, $JZ_2(x)$, takes the height of peaks into account. The peak i yielding the largest value $a_i \cdot G(x, p_i) + b_i$ determines the function value. The bigger the height of a peak, the bigger its influence on the search space in comparison to smaller peaks. Note that, in case of equal slopes $a_1 = \cdots = a_k$ and equal heights $b_1 = \cdots = b_k$, both functions JZ₁ and JZ₂ using parameters $a_1, \ldots, a_k, b_1, \ldots, b_k$ are identical as for JZ₂ the maximum over all terms $a_i \cdot G(x, p_i) + b_i$ for all $1 \le i \le k$ is attained for the closest peak i = cp(x).

Theorem 2. For JZ_1 and JZ_2 using the same parameters $a_1 = \cdots = a_k$ and $b_1 = \cdots = b_k$ we have $JZ_1 = JZ_2$.

In the case of two peaks p_1 and p_2 , if these peaks are complementary, that is, $p_2 = \overline{p_1}$, then JZ₁ and JZ₂ generalise the TWOMAX function, with TWOMAX being the special case of $p_1 = 0^n$, $p_2 = 1^n$, $a_1 = a_2 = 1$ and $b_1 = b_2 = 0$ [9]. This setting was studied for the $(\mu+1)$ EA with clearing in [3].

We consider peaks being placed independently and u. a. r., as this strategy is simple, fair, and it scales towards an arbitrary number of peaks. The slopes are chosen equal to 1 for all peaks for the sake of simplicity. Even though the peaks are placed randomly, if the peaks have moderately similar heights, the resulting fitness landscape has a clear structure: with high probability all peaks are local optima, and all search points within a Hamming ball of radius $\Omega(n)$ belong to a peak's basin of attraction. This holds for both functions JZ_1 and JZ_2 as they have equal fitness values within the mentioned Hamming balls (but may have different values on other search points).

Theorem 3. Assume k peaks p_1, \ldots, p_k chosen independently and u. a. r. from $\{0,1\}^n$. If $a_1 = \cdots = a_k = 1$ and $\max_{1 \le i \le k} b_i - \min_{1 \le i \le k} b_i \le cn$ for a constant c < 1/2 then with probability $1 - k^2 e^{-\Omega(n)}$ for radius $r := (1/2 - c)/3 \cdot n$ we have:

- 1. all k peaks p_1, \ldots, p_k are local optima in both f_1 and f_2 ,
- 2. for all $1 \leq i \leq k$, all search points in $\mathcal{B}_i := \{x \mid H(x, p_i) \leq r\}$, a Hamming ball of radius r around p_i , belong to the basin of attraction of p_i with respect to both JZ_1 and JZ_2 , that is, there is a Hamming path from x to p_i on which the values of JZ_1 and JZ_2 are strictly increasing, and
- 3. for all search points $x \in \bigcup_{i=1}^k \mathcal{B}_i$, $JZ_1(x) = JZ_2(x)$.

Proof. Assume without loss of generality $\min_{1 \le i \le k} b_i = 0$ (as adding a fixed value does not affect the problem structure), hence $b_i \le cn$ for all $1 \le i \le k$.

By Chernoff bounds, the probability that two different peaks will have Hamming distance at most n/2 - r is $e^{-\Omega(n)}$. By the union bound, the probability that this holds for any pair of peaks is at most $k^2 \cdot e^{-\Omega(n)}$. We assume in the following that every two peaks have a Hamming distance larger than n/2 - r.

Now consider a search point $x \in \mathcal{B}_i$, that is, $\mathrm{H}(x, p_i) \leq r$. Since $r \leq n/6$ we have $n/2 \geq 3r$, and thus for all $j \neq i$ we have $\mathrm{H}(x, p_j) \geq \mathrm{H}(p_i, p_j) - \mathrm{H}(x, p_i) > n/2 - r - r \geq r \geq \mathrm{H}(x, p_i)$. So p_i is a unique closest peak, $\mathrm{cp}(x) = i$. By definition of JZ₁, the second statement follows for JZ₁ as on every shortest Hamming path from x to p_i , subsequently decreasing the Hamming distance to p_i increases the fitness by $a_i = 1$. Since $r \geq 1$ if n is large enough, p_i is a local optimum for JZ₁.

It only remains to show the third statement as then the first two statements also apply to JZ₂. To prove that $JZ_2(x) = JZ_1(x)$ for $x \in \mathcal{B}_i$, we need to show that the maximum over terms $a_j \cdot G(x, p_j) + b_j = n - H(x, p_j) + b_j$ from the definition of JZ₂ is attained for j = i. We have $n - H(x, p_i) + b_i \ge n - r$ as $H(x, p_i) \le r$ and $b_i \ge 0$. For $j \ne i$ we have $n - H(x, p_j) + b_j < n/2 + 2r + cn$ as $b_j \le cn$ and $H(x, p_j) \ge H(p_i, p_j) - H(x, p_i) > n/2 - r - r = n/2 - 2r$. Noting that n/2 + 2r + cn = n/2 + 3r + cn - r = n/2 + (n/2 - cn) + cn - r = n - r establishes $n - H(x, p_j) + b_j \le n - H(x, p_i) + b_i = JZ_1(x)$.

4 Experimental Analysis

For the experimental analysis we test each of the algorithms from Table 1 (referred to by the acronyms defined in the first column) on Jansen-Zarges multimodal function classes. We consider a problem size n = 100, genotypic distance for all algorithms that require a dissimilarity measure and stop runs after $10\mu n \ln n$ generations. This time limit is motivated by [2, Lemma 3.3] stating that, loosely speaking, $2e\mu n \ln n \approx 5.44\mu n \ln n$ generations are sufficient to perform hill climbing on two peaks with high probability.

The experimental framework is divided in 3 experimental set-ups. In Sect. 4.1 we assess the ability of each mechanism to find many peaks with equal height, and in Sect. 4.2, we assess the ability of each mechanism to maintain the population diversity when considering peaks with different heights to yield global and local optima. For both sections, the number of peaks was increased exponentially as $k = \{2, 4, 8, \ldots, 64\}$. For each k, we generated 100 different instances choosing k peaks u. a. r. from $\{0, 1\}^n$. In each experiment, all algorithms are tested on the same set of 100 instances to ensure a fair comparison. The challenge for each mechanism is to find and maintain as many peaks as possible before reaching the $10\mu n \ln n$ generations; we record the fraction of the peaks found. The population size is chosen large enough ($\mu = 100$) to be able to accommodate all peaks.

The analysis in Sect. 4.3 is inspired by [3, Sect. 7.3] and focusses on landscapes with two peaks. In this section we take a closer look at the ability of the diversity mechanisms to deal with different basins of attraction, including a wider range of two-peaked landscapes than the ones likely to be generated by placing peaks u. a. r.. The goal is to observe which mechanisms are able to escape from local optima by tunnelling through the fitness valley that separates two peaks. We choose $\mu = 32$ as in [3, Sect. 7.3] and also consider the same two initialisations: the standard uniform random initialisation and biased initialisation where the whole population is initialised with copies of one peak (0ⁿ for TwoMAX). Biased initialisation is used in order to observe how the mechanisms are able to escape from a local optimum and how fast it is compared to a random initialisation.

Based on the theoretical analysis in [3] we define the window size $w = 2.5\mu \ln n$ for RTS. We know from [6,13] that both FS approaches with phenotypic sharing and $\sigma = n/2$ are always efficient on TwoMAX but no theory for genotypic sharing is available. Preliminary experiments for genotypic sharing and $\sigma = n/2$ on TwoMAX yielded poor results; however with $\sigma = n$ (which implies that all individuals always share fitness) both peaks were found in most runs. This makes sense on other landscapes as well as if σ is set smaller than the radius or basin of attraction around a local optimum, then FS is unable to push individuals away from said local optimum. Thus it seems best to err on the side of choosing σ too large rather than too small.

For CL the situation is different. If σ is chosen too large, such that there are several optima within a distance of σ , then global optima may be cleared, making it impossible to maintain many optima in the population. So for CL it seems best to err on the side of choosing σ too small rather than too large. We choose $\sigma = n/3$ for Sects. 4.1 and 4.2 as with high probability every two

different peaks will have a Hamming distance larger than n/3 (cf. Theorem 3). For Sect. 4.3, we use the recommendation $\sigma = \min\{H(p_1, p_2), n/2\}$ from [3].

4.1 Finding Peaks of Equal Height

We consider the JZ₁ function with equal slopes $a_1 = \cdots = a_k = 1$ and offsets $b_1 = \cdots = b_k = 0$. We know from Theorem 2 that with equal parameters JZ₁ = JZ₂. In Fig. 1 (blue/left box plots), we show the fraction of peaks obtained in each of the 100 instances and its variance for each choice of k.

As can be seen, the PL, NGD and NFD perform poorly; these have already been proven to perform poorly on TWOMAX [6]. PC as predicted in [2] is not able to find even one peak. FS performs best for an intermediate number of peaks, $k \in \{4, 8, 16\}$, but still far worse than the best mechanisms. This is in contrast to theoretical results [6, 13] where FS in both variants was shown to be very effective on TwoMAX. These differences may be down to the differences between TwoMAX and JZ_1 with random peaks and/or they may be caused by the differences between phenotypic and genotypic sharing. Interestingly, PFS performs far worse than the conventional FS. This is surprising as PFS uses a significant amount of computation time to search for the best possible population (in terms of shared fitness) it can create out of all parents and offspring, hence we would have expected it to perform better than FS. A possible explanation for the poor performance of FS is that even when the population is able to locate basins of attraction of several peaks, we found several individuals scattered around each peak, apparently repelling each other and preventing each other from reaching the peak.

Finally, DC, RTS and CL perform surprisingly well: they find all optima most of the time for $k \leq 16$, and find most optima for k = 32. Only for a large number of k = 64 peaks, performance deteriorated to around 80% of peaks found. This deterioration is not surprising as the population size was fixed to $\mu = 100$. RTS with $w = 2.5\mu \ln n$ seems to behave similarly to DC as predicted in [2].

4.2 Finding Peaks with Different Height

For this case we make use of the JZ₂ function with $a_1 = \cdots = a_k = 1$ and $b_1 \cdots b_k$ chosen independently and u. a. r. from $[0, 1, \ldots, n/3]$. This range is motivated by Theorem 3, as here two peaks differ in their heights by at most n/3, choosing the leading constant c := 1/3 as the simplest constant smaller than 1/2. Theorem 3 then yields that all search points within Hamming balls of radius n/18 centred at a peak are located in the peak's basin of attraction. Figure 1 (red/centre box plots) shows the fraction of peaks obtained in each of the 100 instances and its variance for each $k = \{2, 4, 8, \ldots, 64\}$ peaks. To gauge the quality of the peaks found, we also plot the normalised best fitness found (green/right box plots), formally f_i^*/opt_i where f_i^* is the fitness of the best peak found on instance i and opt_i is the optimal value of instance i.

In this setting PL, NGD and NFD manage to find the global optimum in up to 80% of instances. This suggests that on this function class it is fairly easy to find a global optimum. However, they rarely find more than one peak, hence they seem to suffer from premature convergence. PC continues to show the worst performance of all mechanisms. PFS and FS find fewer peaks on JZ_2 compared to JZ_1 . This makes sense since the former setting is more difficult than the latter; both mechanisms seem to suffer from the issues mentioned in Sect. 4.1.

Finally, DC, RTS and CL also find fewer peaks due to the difficulty of this setting, but still show the best performance of all mechanisms analysed in this paper and they manage to find the global optimum in all instances. For $k \leq 8$ is not possible to always find all peaks any more, but they still manage to find at least 50% of the peaks. Then, for $k \geq 16$ the performance deteriorates in such a way that it is not possible to reach any more 50% of the peaks but still the mechanisms manage to find some of the peaks. The general cause of the drop in the performance seems to be that all mechanisms struggle to escape from the optimum found, also that low-quality optimums are being dropped when better peaks have been found.

4.3 Escaping from Local Optima

Theorem 3 and its proof suggest that when peaks are chosen u. a. r., they will have a Hamming distance close to n/2. We would like to investigate how the diversity mechanisms behave if peaks have different Hamming distances. Following [2], we focus on two peaks and vary their Hamming distance between 1 and n by choosing $p_1 = 0^n$ and $p_2 \in \{0^{n-1}1, 0^{n-2}1^2, \ldots, 1^n\}$, along with $a_1 = a_2 = 1$ and $b_1 = b_2 = 0$. As argued in [2], this captures the performance across all possible JZ₁ functions with two complementary peaks and the given slopes and offsets. In particular, it includes many functions that only have an exponentially small probability to be generated when choosing peaks independently and u. a. r. With biased initialisation, the algorithms have to find the other optimum by tunnelling through the fitness valley that separates these two peaks. This is a much harder task compared to hill climbing on various hills, where the aim is for the population to maintain a good spread over the search space.

We use the set-up and empirical data for CL from [3] and report the average number of generations of 100 runs, with two stopping criteria: both optima have been found or $t = 10 \mu n \ln n$ generations were reached.

From Fig. 2a all mechanisms are effective when the Hamming distance is so small that the peaks are very close together such that the second peak can be found by a mutation of the first peak found (except for PC, that is not able to reach a single peak). But as the distance increases, the time for some mechanisms increases rapidly; they are inefficient on all non-trivial settings. DC and RTS seem to be agnostic of Hamming distances as they show a very stable and equal performance across the whole range of Hamming distances. This make sense as DC climbs up both peaks with equal probability (cf. the analysis on TwoMAX [6]) and RTS behaves similarly to DC. CL is very effective and only mildly worse than DC and RTS. We see that for FS with genotypic sharing is only effective if the peaks have a Hamming distance that is very close to n or trivially small. For intermediate values, FS fails badly.



Fig. 1. Experimental results for all $(\mu+1)$ EA variants from Table 1 among 100 instances generated u. a. r. for each number of peaks $k = \{2, 4, 8, \ldots, 64\}$, $\mu = 100$ and n = 100, stopping runs after $10\mu n \ln n$ generations. Blue/left: fraction of peaks found on JZ₁ with peaks of equal height. Red/centre: fraction of peaks found on JZ₂ with peaks with different heights, $b_1 \ldots b_k$ chosen u. a. r. from $\{0, 1, \ldots, n/3\}$. Green/right: normalised best fitness found on JZ₂ experiments. Squares indicate median values. (Color figure online)



Fig. 2. The average number of generations among 100 runs for finding both peaks $p_1 = 0^n$ and $p_2 = \{0^{n-1}1, 0^{n-2}1^2, \ldots, 1^n\}$ on the fitness landscape defined by JZ_2 with $a_1 = a_2 = 1$ and $b_1 = b_2 = 0$ or $t = 10\mu n \ln n$ generations were reached, for all $(\mu+1)$ EA variants mentioned in Table 1, using n = 100 and $\mu = 32$. Results for both random and biased initialisation are shown.

With biased initialisation (Fig. 2b), CL is the only mechanism able to escape from local optima with different basins of attraction. As shown theoretically in [3], this is because cleared individuals are able to explore the fitness landscape by performing a random walk. We know from [6,13] that both FS approaches with phenotypic sharing and sharing radius $\sigma = n/2$ are able to escape from local optima as well, if the two peaks are complementary. With genotypic sharing both FS approaches perform very poorly and seem unable to escape from local optima. Also the other mechanisms fail as they are unable to accept worse search points.

5 Conclusions

We have performed an extensive empirical study involving 9 common diversity mechanisms on Jansen-Zarges multimodal function classes, covering various degrees of multimodality from 2 to 64 peaks and peaks having equal or different heights, reflected in their basins of attraction. Our results show that the plain $(\mu+1)$ EA, the simple mechanisms: avoiding genotype and fitness duplicates cannot maintain subpopulations on several peaks; once a peak has been found it seems impossible to escape from such a peak. Probabilistic crowding shows a terrible performance as it is unable to locate even a single peak. These findings are in line with theoretical results on TwoMAX [2,6].

Previous theoretical results have shown that both fitness sharing approaches are always efficient on TwoMAX if phenotypic distances are being used and parameters are set appropriately [6, 13]. This includes the ability to climb down a peak and to tunnel through fitness valleys to reach other niches. Unfortunately this is not the case for fitness sharing with genotypic distance. Only when the peaks have a Hamming distance that is trivially small or very close to n they seem to be effective; for any other intermediate case they show a poor performance.

Deterministic crowding, restricted tournament selection and clearing perform well for peaks with the same slope and height, much better than all other diversity mechanisms. Only for large numbers of peaks (k = 64) and different heights the performance starts to deteriorate. Finally, only clearing has shown the ability to escape from local optima since all other mechanisms seem unable to accept worse search points or unable to tunnel through fitness valleys.

References

- 1. Covantes Osuna, E., Sudholt, D.: Analysis of the clearing diversity-preserving mechanism. In: Proceedings of FOGA 2017, pp. 55–63. ACM (2017)
- Covantes Osuna, E., Sudholt, D.: Runtime analysis of probabilistic crowding and restricted tournament selection for bimodal optimisation. In: Proceedings of GECCO 2018. (2018 to appear). http://arxiv.org/abs/1803.09766
- Covantes Osuna, E., Sudholt, D.: On the runtime analysis of the clearing diversitypreserving mechanism. Evol. Comput. http://arxiv.org/abs/1803.09715
- Črepinšek, M., Liu, S.H., Mernik, M.: Exploration and exploitation in evolutionary algorithms: a survey. ACM Comput. Surv. 45(3), 35:1–35:33 (2013)
- Das, S., Maity, S., Qu, B.Y., Suganthan, P.: Real-parameter evolutionary multimodal optimization – a survey of the state-of-the-art. Swarm Evol. Comput. 1(2), 71–88 (2011)
- Friedrich, T., Oliveto, P.S., Sudholt, D., Witt, C.: Analysis of diversity-preserving mechanisms for global exploration. Evol. Comput. 17(4), 455–476 (2009)
- Glibovets, N.N., Gulayeva, N.M.: A review of niching genetic algorithms for multimodal function optimization. Cybern. Syst. Anal. 49(6), 815–820 (2013)
- Harik, G.R.: Finding multimodal solutions using restricted tournament selection. In: Proceedings of the 6th ICGA, pp. 24–31. Morgan Kaufmann Publishers Inc. (1995)
- Jansen, T., Zarges, C.: Example landscapes to support analysis of multimodal optimisation. In: Handl, J., Hart, E., Lewis, P.R., López-Ibáñez, M., Ochoa, G., Paechter, B. (eds.) PPSN 2016. LNCS, vol. 9921, pp. 792–802. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-45823-6_74
- Lozano, M., García-Martínez, C.: Hybrid metaheuristics with evolutionary algorithms specializing in intensification and diversification: overview and progress report. Comput. Oper. Res. 37(3), 481–497 (2010)
- 11. Mahfoud, S.W.: Niching methods for genetic algorithms. Ph.D. thesis, University of Illinois at Urbana-Champaign (1995)
- Mengsheol, O., Goldberg, D.: Probabilistic crowding: deterministic crowding with probabilistic replacement. In: Proceedings of GECCO 1999, pp. 409–416 (1999)
- Oliveto, P.S., Sudholt, D., Zarges, C.: On the runtime analysis of fitness sharing mechanisms. In: Bartz-Beielstein, T., Branke, J., Filipič, B., Smith, J. (eds.) PPSN 2014. LNCS, vol. 8672, pp. 932–941. Springer, Cham (2014). https://doi. org/10.1007/978-3-319-10762-2_92
- Pétrowski, A.: A clearing procedure as a niching method for genetic algorithms. In: Proceedings of ICEC 1996, pp. 798–803 (1996)
- Sareni, B., Krahenbuhl, L.: Fitness sharing and niching methods revisited. IEEE Trans. Evol. Comput. 2(3), 97–106 (1998)

- Shir, O.M.: Niching in evolutionary algorithms. In: Rozenberg, G., Bäck, T., Kok, J.N. (eds.) Handbook of Natural Computing, pp. 1035–1069. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-540-92910-9_32
- Singh, G., Deb, K.: Comparison of multi-modal optimization algorithms based on evolutionary algorithms. In: Proceedings of GECCO 2006, pp. 1305–1312. ACM (2006)
- Squillero, G., Tonda, A.: Divergence of character and premature convergence: a survey of methodologies for promoting diversity in evolutionary optimization. Inf. Sci. 329, 782–799 (2016)
- 19. Sudholt, D.: The benefits of population diversity in evolutionary algorithms: a survey of rigorous runtime analyses (2018). http://arxiv.org/abs/1801.10087