The Effects of Locality on the Dynamics of Decoder-Based Evolutionary Search

Jens Gottlieb
Department of Computer Science
Technical University of Clausthal
Julius-Albert-Str. 4, 38678 Clausthal–Zellerfeld
Phone: ++ 49 5323 727155 (Germany)
gottlieb@informatik.tu-clausthal.de

Günther R. Raidl
Institute of Computer Graphics
Vienna University of Technology
Favoritenstraße 9–11/1861, 1040 Vienna
Phone: ++ 43 1 58801-18616 (Austria)
raidl@apm.tuwien.ac.at

Abstract

The success of decoder-based evolutionary algorithms (EAs) strongly depends on the achieved locality of operators and decoders. Most approaches to investigate locality properties are static and consider only parts of the complex interactions within an EA, and sometimes, these techniques give misleading results. We suggest an explicit analysis of the dynamic behavior, emphasizing the effects of locality on evolutionary search. The impact of our methodology is twofold since it confirms previous statically obtained results and allows to gain reliable additional insight about the actual dynamics. The approach is successfully applied to four EAs for the multidimensional knapsack problem, and it can easily be adapted to other problems and EAs.

1 Introduction

Locality is an important prerequisite to prevent evolutionary search resembling pure random search. Usually, locality is implicitly considered throughout a serious design process of evolutionary algorithms (EAs) in a rather intuitive fashion. Only few approaches are known to characterize the achieved locality properties in more detail. Several static investigations rely on fitness landscapes, e.g. operator correlation (Manderick et al. (1991)) and fitness distance correlation (Jones and Forrest (1995)). However, these approaches may be misleading since they do not consider the complete dynamics of evolutionary search (Altenberg (1997)).

In particular decoder-based EAs are hard to analyze in a theoretical way due to the complex interactions of variation operators and the decoding procedure. Therefore, we suggested in (Gottlieb and Raidl (1999)) a methodology based on random sampling and distance measures in genotype and phenotype space to characterize locality properties without performing actual EA runs. For several EAs, this approach is able to successfully predict bad performance due to weak locality and hence represents a useful method of a priori analyzing static features of decoder-based EAs, leading to hypotheses concerning the actual dynamic behavior.

Nevertheless, this approach does not consider aspects like the used selection method, replacement strategy, and population size. Locality properties of the variation operators highly depend on the distribution of parental solutions in the population. Our previous approach regarded different population diversities by artificially generating random solutions of certain distances, but obviously, this method is only a rough approximation of the conditions in a real population.

Due to the strong relation between locality and the search dynamics, there is a need for tracing locality-related properties continuously during real runs. We propose statistical measures to allow an empirical investigation of the population dynamics. This paper complements our previous study on locality by (i) verifying and supporting previous hypotheses based on a priori measurements and (ii) gaining more reliable insight about the actual effects of locality on the search dynamics and hence the overall success of an EA.

Empirical results for selected decoder-based EAs for the *multidimensional knapsack problem* (MKP) will be presented. The MKP is stated as

$$\text{maximize} \quad \sum_{j \in J} p_j x_j \tag{1}$$

$$\text{subject to} \quad \sum_{j \in J} r_{ij} x_j \le c_i, \quad i \in I \tag{2}$$

$$x_j \in \{0,1\}, \qquad j \in J \tag{3}$$

with $I = \{1, \dots, m\}$ and $J = \{1, \dots, n\}$ denoting sets of resources and items, respectively. Each resource i is limited by its capacity $c_i > 0$, and each item j yields a profit $p_j > 0$ and requires a certain amount $r_{ij} \geq 0$ of each resource i. The goal is to find a subset of items with maximum total profit that does not exceed the resource capacities. As the MKP has a wide range of applications (Martello and Toth (1990)) and is NP-complete (Garey and Johnson (1979)), several heuristics – and in particular EAs – were proposed, see (Chu and Beasley (1998), Gottlieb (1999)) for comprehensive surveys. Recently, Leguizamón and Michalewicz (1999) presented an ant system for this problem.

The decoder-based EAs we consider here are described in Sect. 2. Section 3 proposes general statistical measures for tracing locality-related properties during an EA run. Empirical results are presented and discussed in Sect. 4, and conclusions are given in Sect. 5.

2 Decoder-Based EAs for the MKP

The best EAs for the MKP we are aware of employ direct encoding, heuristic repair, and local optimization methods (Chu and Beasley (1998), Gottlieb (1999), Raidl (1998)). However, several decoder-based EAs are also capable of obtaining high-quality solutions (Hinterding (1999)). Generally, the latter approaches are based on an individual's duality of genotype and phenotype, where an arbitrary genotypic search space G is mapped into the phenotypic search space $P = \{0, 1\}^n$ by some decoder; while the EA explores G in an explicit fashion, P is explored implicitly.

The considered decoder-based EAs employ a common general setup, namely a population of size 100, parent selection via binary tournaments, producing one offspring per generation by always performing crossover and mutation, steady-state replacement (deleting the worst individual), phenotypic duplicate elimination, and an evaluation limit of 1 000 000 non-duplicate solutions. Phenotypic duplicate elimination means that an offspring is rejected if its phenotype is already represented by some genotype in the current population (Raidl and Gottlieb (1999)). Obviously, this mechanism enforces a minimum population diversity. Overviews of four selected EAs are presented in the following, together with results concerning final solution qualities and locality properties obtained from our previous static studies (Gottlieb and Raidl (1999)).

2.1 Permutation Based EA

The *permutation based EA* (PBEA) has been proposed by Hinterding (1994) for the unidimensional knapsack problem and has also been applied to the MKP (Gottlieb (2000), Raidl (1998), Thiel and Voss (1994)). Permutations $\pi : J \to J$ of the items form the genotypic search space and are decoded as follows. Starting with the feasible solution $x = (0, \ldots, 0)$, all variables x_j are traversed in the order determined by π, increasing each variable from 0 to 1 if this does not violate any resource constraint. We employ standard permutation operators, namely uniform order based crossover and swap mutation, which randomly exchanges two different positions. This operator setup was suggested by Hinterding (1994) and confirmed to be effective by Gottlieb (2000).

2.2 Ordinal Representation Based EA

The *ordinal representation based EA* (OREA) has originally been examined in the context of the traveling salesperson problem (TSP) (Grefenstette et al. (1985)), but is easily adapted to the MKP. Solution candidates for the MKP are represented by vectors v with $v_a \in \{1, \ldots, n - a + 1\}$ for $a \in J = \{1, \ldots, n\}$. The decoder initially generates a list containing all items in some predefined order and starts with the feasible solution $x = (0, \ldots, 0)$. The vector v is traversed from the first to its last position, interpreting each entry v_a as a position in the current list. Such position identifies the next item j which is removed from the list and then checked for inclusion in the current MKP solution; the corresponding variable x_j is increased if the resource capacities are not exceeded. This representation allows the use of classical one-point crossover since the decoder ensures to generate feasible solutions only. We employ a simple mutation operator which randomly chooses a position a and then draws v_a uniformly from $\{1, \ldots, n - a + 1\}$. OREA fails to achieve a high degree of locality since some change in a single position of v modifies the meaning of all following genes and, therefore, often leads to a huge phenotypic change (Gottlieb and Raidl (1999)).

2.3 Surrogate Relaxation Based EA

The *surrogate relaxation based EA* (SREA) was suggested by Raidl (1999). Solution candidates are represented by vectors of real-valued weights, which are used to temporarily modify the profits p_j in the objective function (1) yielding a similar but slightly different MKP instance. This biased problem is solved by a surrogate duality based heuristic that has originally been proposed by Pirkul (1987). The heuristic starts with the feasible solution $x = (0, \ldots, 0)$ and traverses all items according to decreasing profit/pseudo-resource consumption ratio. Variables x_j are increased if the resource constraints remain satisfied. Pseudo-resource consumptions are determined via reasonable surrogate multipliers obtained from the result of the linear programming (LP) relaxed MKP. As the resource constraints (2) are not affected by the real-valued weights, the decoded solution is feasible with respect to the original constraints. Raidl (1999) proposed to determine the surrogate multipliers only once for the original problem in a preprocessing step to decrease the computational effort. SREA uses uniform crossover and a mutation operator which modifies three randomly chosen weights by resetting them to new random values. The results obtained for SREA are the best among all decoder-based EAs for the MKP we are aware of.

2.4 Lagrangian Relaxation Based EA

The *Lagrangian relaxation based EA* (LREA) was also proposed by Raidl (1999) and employs the same representation and variation operators as SREA. However, LREA employs a different heuristic to generate a solution for the biased problem, namely the Lagrangian relaxation based procedure introduced by Magazine and Oguz (1984). As exact Lagrange multipliers are difficult to obtain, some reasonable (but usually suboptimal) multipliers are calculated by a simple heuristic. Each obtained solution is then locally improved by traversing the variables according to decreasing profit and increasing them if feasibility can be maintained.

2.5 Comparison of the EAs

We compared the considered decoder-based EAs on selected problems of Chu's test suite of MKP benchmarks introduced in (Chu and Beasley (1998)) and available from the OR-Library[1]. Ten runs were performed for the first problem instances of sizes $m \in \{5, 10, 30\}$, $n \in \{100, 250, 500\}$ and tightness ratios $\alpha \in \{0.25, 0.5, 0.75\}$ (which means that $c_i = \alpha \sum_{j \in J} r_{ij}$ for all $i \in I$). The solution quality is measured by the relative *gap* of the objective value to the optimal value of the LP-relaxed problem, i.e. $1 - max^{\mathrm{EA}} / opt^{\mathrm{LP}}$ with max^{EA} and opt^{LP} denoting the best objective value found by the EA and the optimal value of the LP relaxation of MKP, respectively. The *duplicate ratio* represents the ratio of rejected duplicates among all generated solutions.

Table 1 presents obtained average results. In particular SREA yielded most of the time the best results. The solution qualities achieved by OREA are significantly worse than those of the other EAs. As discussed in the empirical studies in Sect. 4, a major reason for the difference of an order of magnitude between OREA and the other EAs is the weak locality of OREA, which does not allow a meaningful exploration of the search space, see also (Gottlieb and Raidl (1999)). Furthermore, it will be shown that OREA also suffers from a lack of efficiency since many duplicates are produced due to missing innovation capabilities of the variation operators. The remaining sections of this work explicitly focus on the search dynamics to analyze the effects of locality and related concepts such as innovation in greater detail.

[1] `http://mscmga.ms.ic.ac.uk/info.html`

Table 1: Obtained average gaps and duplicate ratios for the EAs on Chu's benchmark suite

	PBEA	OREA	SREA	LREA
gap [%]	0.74	2.65	0.58	0.62
duplicate ratio [%]	5.74	36.24	6.65	3.27

3 Statistical Measures

In the following, several statistical measures are proposed that describe locality-related properties of binary crossover and mutation. These measures are supposed to be continuously traced over the generations of an evolutionary search in order to gather information about the search dynamics. As will be shown in Sect. 4, important strengths and weaknesses of specific encodings and evolutionary operators can be revealed and moreover, typical behaviors of decoder-based evolutionary search can be explained with these data.

For the purpose of quantifying the similarity of two different solutions, a problem-dependent distance measure is needed. For the MKP, the definition of the *phenotypic distance metric*

$$d(x, y) := \sum_{j \in J} |x_j - y_j| \quad \text{for } x, y \in P$$

is straightforward. The Hamming distance counts the number of variables with different values – i.e. phenotypic properties – in the two solutions. This definition of $d(x, y)$ satisfies the metric conditions, namely identity, symmetry, and the triangular inequality.

For other combinatorial optimization problems, a phenotypic distance usually needs to be defined in a different, meaningful way, which might not always be as obvious as for the MKP. E.g. in case of the TSP, the total number of different edges might be an appropriate measure since edges can be seen as the most important phenotypic properties of TSP solutions (Grefenstette et al. (1985)).

In the following, we propose several measures dealing with binary crossover. Let $x^{p_1}, x^{p_2} \in P$ be the selected parent solutions that undergo crossover to generate an offspring $x^c \in P$.

3.1 Parent Distance PD^t

The behavior and locality properties of crossover are in general strongly influenced by the similarity of the two selected parents. We therefore define the *parent distance* as

$$PD^t := d(x^{p_1}, x^{p_2})$$

and regard it as random variable which depends on the EA's population at generation t (especially its diversity) and the used selection technique.

In the special case $x^{p_1} = x^{p_2}$, i.e. the same solution is selected twice, crossover is usually not able to create a new, meaningful solution different to its parents or degenerates to some kind of mutation. We denote the probability of this unwanted case as $P(PD^t = 0)$, and high values thereof obviously indicate premature convergence or selection pressure that is too high. For avoiding a bias of other measures by this ineffective case, we consider in the following the meaningful case $x^{p_1} \neq x^{p_2}$, i.e. $PD^t > 0$, only. The expected value $E(PD^t | PD^t > 0)$, which again depends on the population at generation t, is then a measure for the degree of population diversity from the crossover viewpoint. In a typical evolutionary search, $E(PD^t | PD^t > 0)$ is high at the beginning of a run and decreases over time.

3.2 Crossover Innovation CI^t

For $PD^t > 0$ we define the *crossover innovation*

$$CI^t := \min(d(x^c, x^{p_1}), d(x^c, x^{p_2}))$$

as the phenotypic distance of the offspring x^c to its closer parent. CI^t is viewed as random variable depending on the selected parents – therefore strongly on PD^t – and the crossover operator. Obviously, CI^t is 0 if either $x^c = x^{p_1}$ or $x^c = x^{p_2}$. Letting $P(CI^t = 0)$ be the likelihood for crossover generating an offspring that is phenotypicly identical to one of its parents, we expect $P(CI^t = 0)$ to be small when $E(PD^t | PD^t > 0)$ is high. Obviously, high values for $P(CI^t = 0)$ degrade performance. A high $P(CI^t = 0)$ for an at least moderate $E(PD^t | PD^t > 0)$ indicates that crossover either does not mix the two parental genotypes well enough or there is a high degree of redundancy in the genotype space G. There are two possible reasons for such high redundancy: Firstly, $|G|$ might be significantly larger than $|P|$. Often such a representation redundancy decreases performance, but sometimes it may also be beneficial and lead to better final results (Ronald (1997)). Secondly, the decoder might contain local improvement techniques or heuristics that always or mostly map genotypes to preferred phenotypes in a restricted subset $P' \subset P$. We call this effect *heuristic bias*. In this case, solutions $x \in P \setminus P'$ cannot be represented or have substantially smaller probabilities to be generated. While such a restriction of P might sometimes be advantageous, it must be ensured that promising areas and particularly the global optima are covered (Ronald (1997)). The four EAs of Sect. 2 work with such heuristic bias since they restrict the search space to the boundary of the feasible region.

Considering only the case of crossover actually producing new, distinct solutions, i.e. $CI^t > 0$, the expectation $E(CI^t | CI^t > 0)$ and corresponding standard deviation $\sigma(CI^t | CI^t > 0)$ are indicators for locality during crossover: In case of strong locality, $E(CI^t | CI^t > 0)$ should be relatively large for large PD^t and become increasingly smaller for smaller PD^t. In particular when PD^t is small, large values for both, $E(CI^t | CI^t > 0)$ and $\sigma(CI^t | CI^t > 0)$, imply weak locality.

3.3 Crossover Loss CL^t

In addition to the ability to generate new solutions with adequate distances to the parents, another important aspect of crossover is that an offspring mainly consists of phenotypic properties inherited from its parents; only few new properties should be introduced. Only under this condition, meaningful building blocks can emerge as described by the building-block hypothesis (Holland (1975)). To consider this aspect, we define for $PD^t > 0$ the *crossover loss* CL^t in general as the number of phenotypic properties of the offspring x^c that are newly introduced and not inherited from either of the parents x^{p_1} or x^{p_2}. In the case of MKP,

$$CL^t \quad := \quad \sum_{j \in J} \delta(x_j^c, x_j^{p_1}, x_j^{p_2})$$

$$\text{with} \quad \delta(x_j^c, x_j^{p_1}, x_j^{p_2}) = \begin{cases} 0 & \text{if } x_j^c = x_j^{p_1} \text{ or } x_j^c = x_j^{p_2} \\ 1 & \text{otherwise.} \end{cases}$$

Using the proposed phenotypic distance metric, we can rewrite the crossover loss alternatively as

$$CL^t := \frac{1}{2}(d(x^c, x^{p_1}) + d(x^c, x^{p_2}) - d(x^{p_1}, x^{p_2}))\,.$$

Note that $CI^t = 0$ implies $CL^t = 0$. To prevent a bias by that case in which crossover is not able to produce a new, distinct solution, we actually consider the expected value $E(CL^t | CI^t > 0)$ only. Large values immediately indicate weak locality.

3.4 Mutation Innovation MI^t

In order to analyze the effects of the mutation operator, x^m is assumed to be the solution resulting from mutating solution x. We define the *mutation innovation* as the phenotypic distance between x and x^m,

$$MI^t := d(x, x^m) \ .$$

This random variable describes how much phenotypic "innovation" is introduced by the mutation and immediately reflects several important aspects concerning locality of mutation. Similarly to the crossover innovation, we consider the measures $P(MI^t = 0)$, $E(MI^t | MI^t > 0)$, and $\sigma(MI^t | MI^t > 0)$. Large values of $P(MI^t = 0)$ indicate that either mutation often does not change any genotypic properties or that the mapping $G \to P$ induces a high degree of redundancy (see Sect. 3.2). Large values for the expectation $E(MI^t | MI^t > 0)$ or the standard deviation $\sigma(MI^t | MI^t > 0)$ indicate weak locality for mutation.

3.5 Duplicate Probability P^t_{dup}

We further consider the duplicate probability P^t_{dup} that a solution newly generated by the evolutionary operators is phenotypicly identical to any other solution already contained in the population. P^t_{dup} depends mainly on $P(CI^t = 0)$, $P(MI^t = 0)$, and the crossover and mutation probabilities. Clearly, a high duplicate probability immediately implies weak efficiency of the EA. Note that the average duplicate probability throughout the evolutionary search equals the duplicate ratio which was introduced in Sect. 2.5.

4 Empirical Analysis

The statistical measures introduced in Sect. 3 strongly depend on the current population characteristics and are therefore considered as functions of the generation number t. We apply an efficient empirical approach to obtain estimations for the considered measures $P(PD^t = 0)$, $E(PD^t | PD^t > 0)$, $P(CI^t = 0)$, $E(CI^t | CI^t > 0)$, $\sigma(CI^t | CI^t > 0)$, $E(CL^t | CI^t > 0)$, $P(MI^t = 0)$, $E(MI^t | MI^t > 0)$, $\sigma(CL^t | CI^t > 0)$, and P^t_{dup} during an actual EA run. At each generation sample values for the basic random variables PD^t, CI^t, CL^t and MI^t are determined when applying crossover and mutation, respectively. The complete dynamics of the measures are approximated by dividing a run into consecutive intervals of generations and determining estimations of the measures independently within each interval. Since the population dynamics usually change faster in early phases of a run, we increase the size of the generation intervals over time. In the empirical analysis of the EAs we consider here, we start with intervals of size 10 and multiply the size by the factor 10 after the generations 100, 1 000, 10 000, and 100 000. Note that the obtained approximations are more confident for higher generation numbers due to these different interval sizes. In order to increase the overall approximation confidence, we use data collected from 10 independent runs instead of just one single run.

Figure 1 shows resulting plots for an MKP instance with $m = 10$, $n = 250$, and $\alpha = 0.5$, namely the first problem of Chu's test suite with these parameters. Although different absolute values have been obtained as results for other instances, this specific problem is representative in the sense that the same basic tendencies have been observed for all other instances, too.

Due to the duplicate elimination strategy and tournament selection being used in all four EAs, the probability $P(PD^t = 0)$ of selecting two identical parents for crossover is constant (≈1.3%). It depends only on the population size and group size of tournament selection and is therefore not shown in the figure.

The dynamics of $E(PD^t | PD^t > 0)$ are more interesting since they are good indicators for the diversity in the population. PBEA and OREA start with the largest values and therefore have a significantly higher diversity in their early populations. An obvious reason for the lower diversity of LREA and in particular SREA is their stronger heuristic bias inside the decoder; this heuristic bias focuses the search on high-quality regions of P already from the beginning. After about 10 000 generations, all four curves meet at a lower bound (≈ 12). Obviously, the duplicate elimination avoids smaller values and the total loss of diversity. After reaching a minimal value at about generation 20 000 in the case of PBEA, SREA, and LREA, $E(PD^t | PD^t > 0)$ and hence the population diversity increase slightly but consequently again. A reason for this behavior seems to be that the population has already converged in highly fit regions of the search space in this phase of a run. The neighborhoods of identified local optima have already been searched, and the best solutions of these regions are contained in the population. New solutions will only remain in the population for a longer time if they are at least as good as the other solutions in the population. This implies that such solutions usually do not lie in those regions that have already been searched intensively. Thus, these solutions have a higher distance from the current population, leading to an increase in population diversity. We call this effect *post-convergence diversity increase.* This phenomenon is an immediate

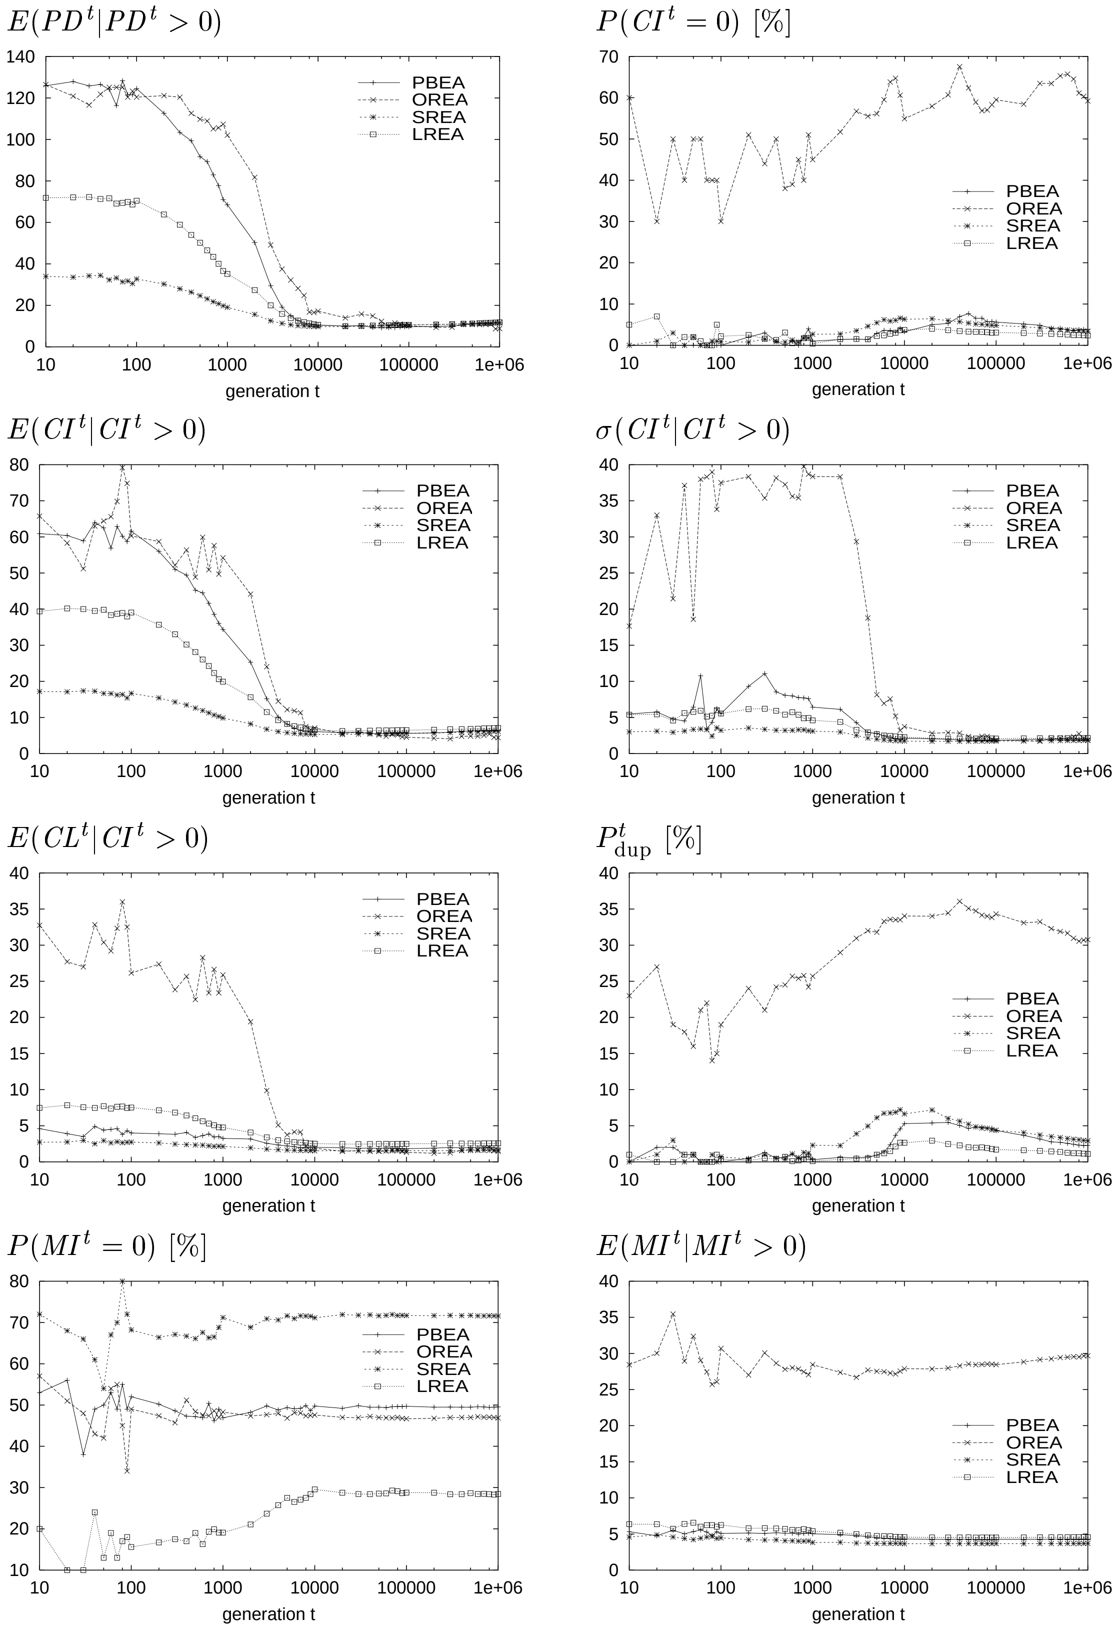

Figure 1: Empirical results for PBEA, OREA, SREA, and LREA applied to an MKP instance with 250 items

consequence of phenotypic duplicate elimination.

The duplicate probability P^t_{dup} reflects population diversity, too. While P^t_{dup} is relatively low as long as the diversity is high, more duplicates are produced when the population has converged and diversity is low. Because of the post-convergence diversity increase, P^t_{dup} decreases again in late stages of the run.

Looking at the probability $P(CI^t = 0)$ for crossover generating a solution which is identical to one of its parents, it is striking that OREA always yields large values from about 40% up to 65%, while the other EAs yield small values below 10% during the whole run. The reason is that OREA uses one-point crossover, which might frequently exchange genes having no effect on the decoded phenotype, because the phenotypic properties are mainly determined by the first genes. Note that the high $P(CI^t = 0)$ of OREA also implies a high duplicate ratio P^t_{dup}.

For each EA, $P(CI^t = 0)$ is relatively small at the beginning and increases due to the decreasing population diversity until the population has converged. During the post-convergence diversity increase, $P(CI^t = 0)$ slightly decreases again.

The dynamics of $E(CI^t|CI^t > 0)$, the distance of a crossover-offspring to its nearer parent, are always strongly correlated to the expected parent distance: $E(CI^t|CI^t > 0) \approx E(PD^t|PD^t > 0)/2$. Contrary to what might have been expected, $E(CI^t|CI^t > 0)$ gives no indication of poor locality of any EA here. But the partly very large standard deviation $\sigma(CI^t|CI^t > 0)$ of OREA reveals missing robustness, i.e. beside offsprings with small distances to one parent, offsprings with large distances to both parents are also generated frequently, which implies weak locality for the crossover of OREA. In comparison to OREA, standard deviations of the other EAs are always relatively low, and therefore we can expect stronger locality.

The plots for the crossover loss $E(CL^t|CI^t > 0)$ visualize more locality properties of crossover. PBEA, SREA, and LREA always yield relatively low values, i.e. offsprings consist mainly of properties inherited from the parents. For OREA, $E(CL^t|CI^t > 0)$ is very high, in particular at the beginning, which indicates weak locality. When the population has converged, i.e. PD^t is low, all EAs exhibit small values for $E(CL^t|CI^t > 0)$. Since SREA nearly always yields the smallest values for $E(CI^t|CI^t > 0)$, $\sigma(CI^t|CI^t > 0)$, and $E(CL^t|CI^t > 0)$, we claim that this best performing EA provides also the strongest crossover-locality.

Regarding mutation, we can observe for all EAs except LREA nearly constant probabilities $P(MI^t = 0)$ of the case that mutation leaves the phenotype unmodified. These probabilities are surprisingly high (SREA: $\approx 70\%$, PBEA and OREA: $\approx 50\%$, LREA: $\approx 15\%$ to 30%) and can be explained by the high degree of decoding redundancy because of heuristic bias (see Sect. 3.2). In the case of LREA, $P(MI^t = 0)$ is initially small and increases during the run; hence it is more difficult for mutation to produce distinct solutions from phenotypes of a converged population than from those appearing in early generations.

Regarding $E(MI^t|MI^t > 0)$ and $\sigma(MI^t|MI^t > 0)$ (which is not depicted here), we observe constantly large values for OREA indicating poor locality. PBEA, SREA, and LREA always exhibit nearly the same small $E(MI^t|MI^t > 0)$ below 7 and $\sigma(MI^t|MI^t > 0)$ below 3, and thus, these EAs provide stronger locality.

5 Conclusions

We proposed a new technique for analyzing the dynamics of decoder-based evolutionary search with particular emphasis on the effects of locality. In contrast to previous approaches for characterizing locality properties, the suggested statistical measures allow to investigate all the dynamic interactions between the variation operators, namely binary crossover and mutation, and the population with its selection and replacement strategies on phenotypic level. Thus, encodings and operators of weak locality can be identified more reliably, provided that a suitable problem-dependent phenotypic distance measure is defined. The methodology of tracing several EA runs and considering samples of subsequent generations provides a simple yet effective way to obtain good approximations for the measures.

Empirical results were presented for four EAs applied to the multidimensional knapsack problem. These results confirm several hypotheses about dynamic behaviors of these EAs raised in our previous study that is based on random sampling (Gottlieb and Raidl (1999)). Most essentially, locality proved once again to be a crucial requirement for any effective evolutionary search. Moreover, the explicit analysis of the dynamics revealed new aspects of the considered EAs. Of particular interest is the phenomenon which we called the post-convergence diversity increase: The population diversity decreases relatively fast due to the heuristic bias, which also introduces a high redundancy in the mapping $G \to P$. Because of the phenotypic duplicate elimination strategy, the diversity is lower-bounded and increases slightly again during the remaining generations. The proposed statistical measures clearly indicate reasons for the poor performance of OREA:

Besides weak locality, OREA tends to produce an offspring phenotypicly identical to one of its parents or an offspring that does not share many similarities with its parents. Thus, OREA cannot perform a meaningful search. The other three EAs were confirmed to achieve the desired level of locality that enables them to perform a meaningful search. SREA, which performed best, also provides the strongest locality.

In general, this work complements our previous static analysis concerning locality. Both studies together provide a very useful methodology to analyze decoder-based EAs. We expect our approach to be helpful in the design of decoder-based EAs for other problem domains, too. Recent (but yet unpublished) results obtained by the first author for the fixed charge transportation problem confirm this expectation.

References

L. Altenberg (1997): *Fitness Distance Correlation Analysis: An Instructive Counterexample*, in Proc. of the 7th Int. Conf. on Genetic Algorithms, pp. 57 – 64

P. C. Chu and J. E. Beasley (1998): *A Genetic Algorithm for the Multidimensional Knapsack Problem*, Journal of Heuristics 4, pp. 63 – 86

M. D. Garey and D. S. Johnson (1979): *Computers and Intractability: A Guide to the Theory of NP-Completeness*, Freeman, San Francisco

J. Gottlieb (1999): *On the Effectivity of Evolutionary Algorithms for Multidimensional Knapsack Problems*, in Proc. of Artificial Evolution 1999, to appear as Springer LNCS

J. Gottlieb (2000): *Permutation-Based Evolutionary Algorithms for Multidimensional Knapsack Problems*, to appear in Proc. of 2000 ACM Symposium on Applied Computing, Como, Italy

J. Gottlieb and G. R. Raidl (1999): *Characterizing Locality in Decoder-Based EAs for the Multidimensional Knapsack Problem*, in Proc. of Artificial Evolution 1999, to appear as Springer LNCS

J. J. Grefenstette, R. Gopal, B. Rosmaita, and D. Van Gucht (1985): *Genetic Algorithms for the Traveling Salesman Problem*, in Proc. of the 1st Int. Conf. on Genetic Algorithms, pp. 160 – 168

J. H. Holland (1975): *Adaption in Natural and Artificial Systems*, University of Michigan Press, Ann Arbor

R. Hinterding (1994): *Mapping, Order-Independent Genes and the Knapsack Problem*, in Proc. of the 1st IEEE Int. Conf. on Evolutionary Computation, pp. 13 – 17

R. Hinterding (1999): *Representation, Constraint Satisfaction and the Knapsack Problem*, in Proc. of the 1999 IEEE Congress on Evolutionary Computation, pp. 1286 – 1292

T. Jones and S. Forrest (1995): *Fitness Distance Correlation as a Measure of Problem Difficulty for Genetic Algorithms*, in Proc. of the 6th Int. Conf. on Genetic Algorithms, pp. 184 – 192

M. J. Magazine and O. Oguz (1984): *A Heuristic Algorithm for the Multidimensional Zero–One Knapsack Problem*, European Journal of Operational Research 16, pp. 319 – 326

B. Manderick, M. de Weger, and P. Spiessens (1991): *The Genetic Algorithm and the Structure of the Fitness Landscape*, in Proc. of the 4th Int. Conf. on Genetic Algorithms, pp. 143 – 150

S. Martello and P. Toth (1990): *Knapsack Problems: Algorithms and Computer Implementations*, J. Wiley & Sons

G. Leguizamón and Z. Michalewicz (1999): *A New Version of Ant System for Subset Problems*, in Proc. of the 1999 IEEE Congress on Evolutionary Computation, pp. 1459 – 1464

H. Pirkul (1987): *A Heuristic Solution Procedure for the Multiconstrained Zero-One Knapsack Problem*, Naval Research Logistics 34, pp. 161 – 172

G. R. Raidl (1998): *An Improved Genetic Algorithm for the Multiconstrained 0–1 Knapsack Problem*, in Proc. of the 1998 IEEE Int. Conf. on Evolutionary Computation, pp. 207 – 211

G. R. Raidl (1999): *Weight-Codings in a Genetic Algorithm for the Multiconstraint Knapsack Problem*, in Proc. of the 1999 IEEE Congress on Evolutionary Computation, pp. 596 – 603

G. R. Raidl and J. Gottlieb (1999): *On the Importance of Phenotypic Duplicate Elimination in Decoder-Based Evolutionary Algorithms*, in Proc. of the 1999 Genetic and Evolutionary Computation Conference, Late breaking papers, pp. 204 – 211

S. Ronald (1997): *Robust Encodings in Genetic Algorithms*, in D. Dasgupta and Z. Michalewicz (eds.), Evolutionary Algorithms in Engineering Applications, pp. 29 – 44, Springer

J. Thiel and S. Voss (1994): *Some Experiences on Solving Multiconstraint Zero-One Knapsack Problems with Genetic Algorithms*, INFOR 32, pp. 226 – 242