Runtime Analysis of the $(1 + (\lambda, \lambda))$ Genetic Algorithm on Random Satisfiable 3-CNF Formulas

Maxim Buzdalov ITMO University 49 Kronverkskiy ave. Saint-Petersburg, Russia 197101

ABSTRACT

The $(1 + (\lambda, \lambda))$ genetic algorithm, first proposed at GECCO 2013, showed a surprisingly good performance on some optimization problems. The theoretical analysis so far was restricted to the ONEMAX test function, where this GA profited from the perfect fitness-distance correlation. In this work, we conduct a rigorous runtime analysis of this GA on random 3-SAT instances in the planted solution model having at least logarithmic average degree, which are known to have a weaker fitness distance correlation.

We prove that this GA with fixed not too large population size again obtains runtimes better than $\Theta(n \log n)$, which is a lower bound for most evolutionary algorithms on pseudo-Boolean problems with unique optimum. However, the self-adjusting version of the GA risks reaching population sizes at which the intermediate selection of the GA, due to the weaker fitness-distance correlation, is not able to distinguish a profitable offspring from others. We show that this problem can be overcome by equipping the self-adjusting GA with an upper limit for the population size. Apart from sparse instances, this limit can be chosen in a way that the asymptotic performance does not worsen compared to the idealistic ONEMAX case. Overall, this work shows that the $(1 + (\lambda, \lambda))$ GA can provably have a good performance on combinatorial search and optimization problems also in the presence of a weaker fitness-distance correlation.

CCS CONCEPTS

•Theory of computation →Theory of randomized search heuristics;

KEYWORDS

Genetic algorithm, theory, satisfiability, random instances, fitnessdistance correlation, runtime analysis, Azuma inequality.

ACM Reference format:

Maxim Buzdalov and Benjamin Doerr. 2017. Runtime Analysis of the $(1 + (\lambda, \lambda))$ Genetic Algorithm on Random Satisfiable 3-CNF Formulas. In *Proceedings of GECCO '17, Berlin, Germany, July 15-19, 2017*, 8 pages. DOI: http://dx.doi.org/10.1145/3071178.3071297

GECCO '17, Berlin, Germany

@ 2017 Copyright held by the owner/author (s). Publication rights licensed to ACM. 978-1-4503-4920-8/17/07... \$15.00

DOI: http://dx.doi.org/10.1145/3071178.3071297

Benjamin Doerr École Polytechnique Laboratoire d'Informatique (LIX) Palaiseau, France

1 INTRODUCTION

The $(1 + (\lambda, \lambda))$ genetic algorithm (GA) was introduced by Doerr, Doerr, and Ebel in 2013 [5]. It builds on the simple idea to first generate from a single parent individual several offspring via standard-bit mutation with higher-than-usual mutation rate, to select the best of these, and to perform a biased crossover with the parent to reduce the destructive effects of the high mutation rate. This use of crossover with the parent as repair mechanism is novel in evolutionary discrete optimization.

The so far moderate number of results on this GA show that it has some remarkable properties. A mathematical runtime analysis on the ONEMAX test function class [2-5] shows that the right parameter setting leads to an optimization time of slightly better than $O(n\sqrt{\log n})$. This is remarkable as all previous runtime analyses of evolutionary algorithms on the ONEMAX test function class showed that these algorithms needed at least $\Omega(n \log n)$ fitness evaluations. The result is remarkable also in that it is the first time that crossover was rigorously shown to give an asymptotic runtime improvement for a simple test function. A third noteworthy property of this GA is that a simple self-adjusting choice of the offspring population size λ inspired by the 1/5-th rule from continuous optimization could further improve the runtime to $\Theta(n)$. Again, this is the first time that a 1/5-th rule type dynamic parameter choice could be proven useful in discrete evolutionary optimization. These mathematical analyses are complemented by an experimental investigation on the ONEMAX test function, on linear functions with random weights, and on royal road functions. Further, Goldman and Punch [7] in the analysis of their parameter-less population pyramid algorithm also used the $(1 + (\lambda, \lambda))$ GA as comparison, and it performed well on random MAX-SAT instances.

It is clear that the intermediate selection of the best mutation offspring becomes most effective if there is a strong fitness-distance correlation. The ONEMAX function, by definition, has a perfect fitness-distance correlation. So the question has to be asked to what degree the positive results for ONEMAX remain true in optimization problems with a weaker fitness distance correlation. To study this question, we analyze the performance of the $(1 + (\lambda, \lambda))$ GA on random satisfiable 3-SAT instances in the planted solution model. For these, the fitness-distance correlation can be scaled via the instance density. In works [6, 11], the analysis of the performance of the (1 + 1) EA on these instances has shown that when the clause-variable ratio is $m/n = \Omega(n)$, then the following strong fitness-distance correlation holds apart from an exponentially small failure probability: For any two search points x, y which are least half as good as a random search point and such that they differ in one bit, the fitness and the distance are perfectly correlated in that the one closer to the optimum has a fitness larger by $\Theta(m/n)$.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

From this, one could easily derive an $O(n \log n)$ optimization time of the (1 + 1) EA. However, when the clause-variable ratio is only logarithmic, this strong fitness-distance correlation is far from being satisfied. Therefore, only the much weaker condition could be shown that each pair (x, y) as above shows a $\Theta(m/n)$ fitness advantage of the closer search point with probability $1 - n^3$. This was enough to show that with high probability (over the joint probability space of instance and algorithm) the (1 + 1) EA finds the optimum in $O(n \log n)$ iterations.

Our results: We conduct a rigorous runtime analysis of the $(1 + (\lambda, \lambda))$ GA (with mutation rate $p = \lambda/n$ and crossover bias $c = 1/\lambda$ linked to the population size λ as recommended in [2]), on the same type of random 3-SAT instance (see Section 2 for the details) as regarded in [6, 11]. We observe that the weaker fitnessdistance correlation of low-density instance indeed poses a problem for the $(1 + (\lambda, \lambda))$ GA when the population size (and thus the mutation rate) is high. In this case, the mutation offspring are distant enough so that the weak fitness-distance correlation prevents the $(1 + (\lambda, \lambda))$ GA to detect an individual closer to the optimum than the typical offspring. For the $(1 + (\lambda, \lambda))$ GA with static value of λ , our experiments and informal considerations suggest that in this case the $(1 + (\lambda, \lambda))$ GA, for large enough *n*, reverts to the behavior similar to the (1 + 1) EA, however, the constant in $O(n \log n)$ is proportional to $1/\lambda$, just like it happens for ONEMAX. Things are worse when the self-adaptive version of the $(1 + (\lambda, \lambda))$ GA is used. In this case, the low probability to find an improving solution lets the value of λ further increase as governed by the 1/5-th rule parameter adaptation. Consequently, the probability to select a profitable individual out of the mutation offspring population further decreases and this negative effect becomes stronger. The result is that λ quickly reaches the allowed maximum of *n* and the performance approaches the one of the (1 + 1) EA with the cost of one iteration *n* times higher than usual.

On the positive side, we make precise that these negative effect can be overcome in most cases by an appropriate choice of λ . We show that when λ is asymptotically smaller that the fourth root of the density, then the probability to find an improving solution is asymptotically the same as for the optimization of ONEMAX. Consequently, when the density is $\omega(\log^2 n)$, then we can still use $\lambda = \sqrt{\log n}$ and obtain an expected optimization time (number of fitness evaluations) of $O(n\sqrt{\log n})$. Note that in this work, we do not try to achieve the later improvement of this runtime guarantee [4] by a factor $\Theta(\sqrt{\log \log \log n}/\log \log n)$ though we are optimistic that such a guarantee can be shown with mildly more effort. For a logarithmic instance density, the smallest regarded here and in [6, 11], for any $\varepsilon > 0$ with a choice of $\lambda = \log^{0.25-\varepsilon} n$ we still obtain a runtime of $O(n \log^{0.75+\varepsilon} n)$ and beat the $O(n \log n)$ performance the (1 + 1) EA has on these instances.

For the self-adjusting version, where a typical run of the $(1 + (\lambda, \lambda))$ GA on ONEMAX uses λ -values of up to \sqrt{n} , we show that adding an upper limit up to which the value of λ can at most grow, overcomes this difficulties sketched above. The runtime increase incurred by such a limit is again manageable. If the upper limit is $\overline{\lambda}$ and it depends on the instance density as $\overline{\lambda} = o((m/n)^4)$, then the runtime of the self-adjusting $(1+(\lambda, \lambda))$ GA is $O(n \cdot \max\{1, \log n/\overline{\lambda}\})$.

Hence already for densities asymptotically larger than $\log^4 n$, we obtain the linear runtime that is valid for the ideal ONEMAX fitness landscape for densities at least logarithmic.

Techniques employed: Our main result that $\lambda = o((m/n)^{1/4})$ suffices for the well-functioning of the $(1 + (\lambda, \lambda))$ GA on our random 3-SAT instances of density m/n is based on a different fitnessdistance correlation result than those used in [6, 11]. Whereas the latter require that with good probability all neighbors of a given solution (not excessively far from the optimum) have a fitness advantage or disadvantage of order m/n (depending on whether they are closer to the optimum or further away), we require this condition only for a certain fraction of the neighbors. This relaxation will not be a problem since it only reduces the probability that the $(1 + (\lambda, \lambda))$ GA finds a certain improvement by a constant factor (being an elitist algorithm, we do not need to care about fitness losses). On the positive side, this relaxation (i) allows us to extend the fitness-distance correlation requirement to all vertices in the λ -neighborhood instead of only all direct neighbors and (ii) gives us that this correlation property with probability $1 - \exp(-\Omega(n))$ holds for all vertices (not excessively far from the optimum). Consequently, the performance results we show hold for all but an exponentially small fraction of the input instances. This fitnessdistance correlation result also implies that the result of [6] for logarithmic densities holds in the same strong version as the one for linear densities, namely that on all but an exponentially small fraction of the input instances the expected runtime is $O(n \log n)$.

To prove our fitness-distance correlation result, we use McDiarmid's bounded differences version [9] of the Azuma martigale concentration inequality in a novel way. To reduce the maximum influence of the independent basic random variables on the discrete quantity of interest, we replace this quantity by a larger continuous function in a way that the influence of each basic random variable is significantly reduced. We are not aware of this type of argument being used before, either in evolutionary computation or randomized algorithms in general.

Most of the proofs in this paper are omitted for space reasons. The full version of this work is available at http://arxiv.org/abs/1704. 04366. The source code for experiments is available at GitHub¹.

2 PRELIMINARIES

In this section, we define the notation, algorithms, and problems we regard in this work. Our notation is standard. We write [a..b] to denote the set of all integers in the real interval [a; b]. We write [r] to denote the integer closest to the real number r, rounding up in case of ties.

2.1 The *k*-CNF SAT Problem

Consider a set *V* of *n* Boolean variables $V = \{x_1, x_2, ..., x_n\}$. A *clause C* over *V* is the logical disjunction of exactly *k* literals, $C = l_1 \lor l_2 ... \lor l_k$, and each literal l_i is either a variable x_j or its negation $\neg x_j$. A *k*-*CNF formula F* is a logical conjunction of exactly *m* clauses $F = C_1 \land C_2 \land ... \land C_m$. A *k*-CNF formula *F* is *satisfiable* if and only if there is an assignment of truth values to the variables such that every clause contains at least one true literal, i.e., the whole expression *F* evaluates to true. We shall only regard the case k = 3,

¹https://github.com/mbuzdalov/papers/tree/master/2017-onell-3cnf

Runtime Analysis of the $(1 + (\lambda, \lambda))$ Genetic Algorithm on Random 3-CNF Formulas

GECCO '17, July 15-19, 2017, Berlin, Germany

but we have no doubts that the main claims are true for any constant $k \ge 3$.

We consider random 3-CNF formulas consisting of *m* clauses of length k = 3 over the *n* variables in *V*. We take the usual assumption that each clause consists of distinct variables. This assumption is very natural since any clause of length 3 that contains repeating variables can be immediately reduced to an equivalent clause of length 2 or, alternatively, to a tautology. However, we explicitly allow repeated clauses in *F*.

Let $\Omega_{n,m}$ be the finite set of all 3-CNF formulas over *n* variables and *m* clauses. We work in the so-called *planted solution* model. Hence there is a target assignment x^* and *F* is a formula chosen uniformly at random among all formulas in $\Omega_{n,m}$ which are satisfied by x^* . We refer to [6, 11] for a justification of this model and a discussion how it relates to other random satisfiability problems.

We shall, without loss of generality, assume that the planted solution is $x^* = (1, ..., 1)$. This is justified by the fact that we only regard unbiased algorithms, that is, algorithms that treat bit positions and the bit values 0 and 1 symmetrically. Hence these algorithms cannot profit from "knowing" the optimal solution. A random formula in this model can be constructed by *m* times (with replacement) choosing a random clause satisfied by x^* , that is, a random clause among all clauses containing at least one positive literal. Note that such a random formula may have other satisfying assignment than x^* . Nevertheless, we denote the structural distance, which is the Hamming distance here, of a solution *x* to the planted solution x^* by $d(x) = |\{i : x_i = 0\}|$. When talking about fitness-distance correlation and related concepts, we shall always refer to this distance.

2.2 3-CNF and Evolutionary Algorithms

An assignment of true/false values to a set of *n* Boolean variables can be represented by a bit string $x \in \{0, 1\}^n$ such that $x_i = 1$ if and only if the *i*-th variable is having the value true. For a length-*m* formula *F* on *n* variables, we define the fitness function $f = f_F$: $\{0, 1\}^n \rightarrow [0..m]$ via $f(x) = |\{C \in F \mid C \text{ is satisfied by } x\}|$, the number of clauses satisfied by the assignment represented by *x*. If *F* is satisfiable, the task of finding a satisfying assignment reduces to the task of maximizing *f*.

In [6], it is proven that when $m/n > c \ln n$ for sufficiently large constant *c*, the runtime of the (1+1) EA is $O(n \log n)$ with probability polynomially close to one. One of the key concepts of the proof is the *fitness-distance correlation*. In the case of logarithmic density, this concept can be formulated as follows:

LEMMA 2.1 (LEMMA 4 FROM [6]). Let $0 < \varepsilon < \frac{1}{2}$. Assume that $m/n > c \ln n$ for sufficiently large constant c. Then there exist two constants c_1 and c_2 such that, for any two solutions x_1 and x_2 such that

- they are different in exactly one bit;
- this bit is set to 1 in x₂;
- the structural distance d(x₁) from x₁ to the planted solution is at most (1/2 + ε)n;

we have $c_1m/n \le f(x_2) - f(x_1) \le c_2m/n$ with probability at least $1 - n^{-3}$.

Algorithm '	1 · (1	$+ (\lambda)$	λ)) GA	with	fixed	integer	nonulation	size)
mgommin .	I • (1	(//,	n_{1} O_{1}	. WILLII	IIACU	nneger	population	SILC /

	()) off while lined lineger population once it
1: $x \leftarrow \text{UniformRan}$	$DOM(\{0,1\}^n)$ > Initialization
2: for $t \leftarrow 1, 2, 3, \ldots$	do ⊳ Optimization
3: $p \leftarrow \lambda/n$	▹ Mutation probability
4: $c \leftarrow 1/\lambda$	▷ Crossover probability
5: $\ell \sim \mathcal{B}(n,p)$	 Mutation strength
6: for $i \in [1\lambda]$ d	o ▷ Phase 1: Mutation
7: $x^{(i)} \leftarrow Mu^{(i)}$	$rate(x, \ell)$
8: end for	
9: $x' \leftarrow \text{Uniform}$	$RANDOM(\{x^{(j)} \mid f(x^{(j)}) = \max\{f(x^{(i)})\}\})$
10: for $i \in [1\lambda]$ d	o ▷ Phase 2: Crossover
11: $y^{(i)} \leftarrow CRC$	PSSOVER(x, x', c)
12: end for	
13: $y \leftarrow \text{Uniform}$	Random($\{y^{(j)} \mid f(y^{(j)}) = \max\{f(y^{(i)})\}\}$)
14: if $f(y) \ge f(x)$	then > Selection
15: $x \leftarrow y$	
16: end if	
17: end for	

2.3 The $(1 + (\lambda, \lambda))$ Genetic Algorithm

The $(1 + (\lambda, \lambda))$ Genetic Algorithm, or the $(1 + (\lambda, \lambda))$ GA for short, was proposed by Doerr, Doerr and Ebel in [5]. Its main working principles are (i) to use mutation with a higher-than-usual mutation rate to speed up exploration and (ii) crossover with the parent to diminish the destructive effects of this mutation. Two versions of the algorithm were proposed, one with static parameters and one with a self-adjusting parameter choice.

The fixed-parameter version is outlined in Algorithm 1. It uses the following two variation operators.

- *ℓ*-bit mutation: The unary mutation operator MUTATE(*x*, *ℓ*) creates from *x* ∈ {0,1}ⁿ a new bit string *y* by flipping exactly *ℓ* bits chosen randomly without replacement.
- biased uniform crossover: The binary crossover operator CROSSOVER(x, x', c) with *crossover bias* $c \in [0, 1]$ constructs a new bit string y from two given bit strings x and x' by choosing for each $i \in [1..n]$ the second argument's value $(y_i = x'_i)$ with probability c and setting $y_i = x_i$ otherwise.

The $(1 + (\lambda, \lambda))$ GA has three parameters, the mutation rate *p*, the crossover bias *c*, and the offspring population size λ . After randomly initializing the one-element parent population $\{x\}$, in each iteration the following steps are performed:

- In the *mutation phase*, λ offspring are sampled from the parent *x* by applying λ times independently the mutation operator MUTATE(*x*, ℓ), where the step size ℓ is chosen at random from the binomial distribution $\mathcal{B}(n, p)$. Consequently, each offspring has the distribution of standard bit mutation with mutation rate *p*, but all offspring have the same Hamming distance from the parent.
- In an *intermediate selection* step, the mutation offspring with maximal fitness, called *mutation winner* and denoted by x', is determined (breaking ties randomly).
- In the crossover phase, λ offspring are created from x and x' using via the biased uniform crossover CROSSOVER(x, x', c).

• *Elitist selection.* The best of the crossover offspring (breaking ties randomly and ignoring individuals equal to *x*) replaces *x* if its fitness is at least as large as the fitness of *x*.

Throughout this paper we use the mutation rate $p = \lambda/n$ and the crossover bias $c = 1/\lambda$ as recommended and justified in [5, Sections 2 and 3] and [2, Section 6].

For this $(1+(\lambda, \lambda))$ GA, a first runtime analysis [5] was conducted on the ONEMAX test function

ONEMAX :
$$\{0,1\}^n \to \mathbb{Z}; x \mapsto \sum_{i=1}^n x_i$$
.

It was shown that for arbitrary λ , possibly being a function of n, the expected optimization time (number of fitness evaluations until the optimum is evaluated for the first time) is $O(\max\{n\lambda, \frac{n \log n}{\lambda}\})$. This expression is minimized for $\lambda = \Theta(\sqrt{\log n})$, giving an upper bound of $O(n\sqrt{\log n})$. The analysis, on which we will build on in this work, uses the fitness level method [12]. Roughly speaking, the arguments are that in an iteration starting with an individual x with fitness distances d = d(x) = n - ONEMAx(x) (i) with probability $\Omega(\min\{1, d\lambda^2/n\})$ the mutation winner is less than ℓ fitness levels worse than the parent, and that (ii) in this case with constant probability the crossover winner is better than the parent. Consequently, the expected number of iterations needed to gain an improvement from x is $O(\max\{1, n/d\lambda^2\})$. Summing over all d and noting that one iteration uses 2λ fitness evaluations gives the claim.

This result is interesting in that this is the first time that crossover was proven to bring an asymptotic speed-up for a simple fitness landscape like ONEMAX. Previously, a constant improvement was shown to be possible for ONEMAX using crossover [10]. Note that all mutation-based algorithms that treat bit positions and bit values symmetrically need at least $\Omega(n \log n)$ fitness evaluations, as this is the unary unbiased black-box complexity of ONEMAX [8].

The $(1 + (\lambda, \lambda))$ GA, which in principle nothing more than a (1 + 1) EA with a complicated mutation operator, in experiments showed a performance superior to the one of the classic (1 + 1) EA on ONEMAX (showing also that the constants hidden in the asymptotic notation are small), on linear functions and on royal road functions [5] as well as on maximum satisfiability instances [7].

Subsequently, the runtime analysis on ONE-MAX was improved [4] and the tight bound of $\Theta(\max\{n \log n/\lambda, n\lambda \log \log \lambda/\log \lambda\})$ was shown for all values of $\lambda \leq n$. This is minimized to $\Theta(n\sqrt{\log n \log \log \log n}/\log \log n)$ when setting $\lambda = \Theta(\sqrt{\log n \log \log \log n})$.

Already in [5], it was observed that a dynamic choice of the parameter λ can reduce the runtime to linear. For this, a fitness dependent choice of $\lambda = \lceil \frac{n}{n - \text{ONEMAX}(x)} \rceil$ suffices. This seems to be the first time that a super-constant speed-up was provably obtained by a dynamic parameter choice (see [1] for a result showing that also the $(1+\lambda)$ EA can profit from a dynamic choice offspring population size when optimizing ONEMAX). Since a fitness-dependent choice as above is unlikely to be guessed by an algorithm user, also a self-adjusting variant setting λ success-based according to a 1/5-th rule was proposed in [5]. That this indeed closely tracks the optimal value of λ and gives a runtime of O(n) was later shown in [3]. The self-adjusting version of the $(1 + (\lambda, \lambda))$ GA is described in Algorithm 2.

Alg	gorithm 2 : $(1 + (\lambda, \lambda))$ GA with	self-adjusting $\lambda \leq \overline{\lambda}$
1:	$F \leftarrow \text{constant} \in (1; 2)$	⊳ Update strength
2:	$U \leftarrow 5$	▶ The 5 from the "1/5-th rule"
3:	$x \leftarrow \text{UniformRandom}(\{0,1\}^n)$	
4:	for $t \leftarrow 1, 2, 3, \dots$ do	
5:	$p \leftarrow \lambda/n, c \leftarrow 1/\lambda, \lambda' \leftarrow [\lambda]$	$, \ell \sim \mathcal{B}(n, p)$
6:	for $i \in [1\lambda']$ do	▹ Phase 1: Mutation
7:	$x^{(i)} \leftarrow \text{Mutate}(x, \ell)$	
8:	end for	
9:	$x' \leftarrow \text{UniformRandom}(\{x^{(i)}\})$	$f^{(j)} \mid f(x^{(j)}) = \max\{f(x^{(i)})\}\}$
10:	for $i \in [1\lambda']$ do	▶ Phase 2: Crossover
11:	$y^{(i)} \leftarrow \text{Crossover}(x, x')$, c)
12:	end for	
13:	$y \leftarrow \text{UniformRandom}(\{y^{(j)}\})$	$f(y^{(j)}) = \max\{f(y^{(i)})\}\}$
14:	if $f(y) > f(x)$ then	 Selection and Adaptation
15:	$x \leftarrow y, \lambda \leftarrow \max\{\lambda/F, 1\}$	}
16:	else if $f(y) = f(x)$ then	
17:	$x \leftarrow y, \lambda \leftarrow \min\{\lambda F^{1/(U)}\}$	$^{(-1)},\overline{\lambda}\}$
18:	else	
19:	$\lambda \leftarrow \min\{\lambda F^{1/(U-1)}, \overline{\lambda}\}$	
20:	end if	
21:	end for	

The main idea of the self-adjusting $(1 + (\lambda, \lambda))$ GA is as follows. If an iteration leads to an increase of the fitness, indicating that progress is easy, then the value of λ is reduced by a constant factor F > 1. If an iteration did not produce a fitness improvement, then λ is increased by a factor of $F^{1/4}$. Consequently, after a series of iterations with an average success rate of 1/5, the algorithm ends up with the initial value of λ . Needless to say, λ can never drop below 1 and never rise above *n* (when using the recommended mutation rate λ/n). Since we will later regard a self-adaptive version with different upper limit, we formulated Algorithm 2 already with the additional parameter $\overline{\lambda}$ as upper limit. Hence the self-adaptive $(1 + (\lambda, \lambda))$ GA as proposed in [5] uses $\overline{\lambda} = n$. We also note that whenever λ should be interpreted as an integer (e.g. when the population size of the current iteration needs to be determined), the value $\lambda' = [\lambda]$ rounded to the closest integer is taken instead.

2.4 Concentration Inequalities

In this section we describe the concentration inequalities which we use in this paper.

Suppose X_1, \ldots, X_n are independent random variables taking values in [-1, 1], however, such that for some $\kappa > 0$ we have $\Pr[X_i = 0] \ge 1 - \kappa$ for all $i \in [1..n]$. For this situation, the following corollary from the Bernstein's inequality, which seems to be rarely used in the theory of evolutionary computation, holds.

$$\Pr[X \le \mu - \delta] \le \max\{\exp(-\delta^2/4\kappa n), \exp(-\frac{3}{8}\delta)\}$$
(1)

Suppose X_1, \ldots, X_n are arbitrary independent random variables, and assume the function f defined over the product of the domains of the X_i satisfies

$$\sup_{x_1,x_2,\ldots,x_n,\hat{x}_i} |f(x_1,\ldots,x_n) - f(x_1,\ldots,\hat{x}_i,\ldots,x_n)| \le c_i$$

Runtime Analysis of the $(1 + (\lambda, \lambda))$ Genetic Algorithm on Random 3-CNF Formulas

for all $1 \le i \le n$. Then, for all $\delta > 0$, the following *McDiarmid's* inequality [9] holds.

$$\Pr[f(X_1, \dots, X_n) - E[f(X_1, \dots, X_n)] \ge +\delta] \le e^{-\frac{2\delta^2}{\sum_{i=1}^n c_i^2}}.$$
 (2)

ALMOST LIKE ONEMAX: CONDITIONS FOR 3 THE $(1 + (\lambda, \lambda))$ GA TO BEHAVE WELL ON **RANDOM 3-CNF FORMULAS**

In this section, we formulate and prove conditions which are enough for the $(1 + (\lambda, \lambda))$ GA to have local fitness-distance correlation.

In the Doerr, Neumann and Sutton paper [6], two results were proven. When the clause density m/n is $\Omega(n)$, then a strong fitness distance correlation is exhibited. Apart from an exponentially small failure probability, the random instance is such that any possible one-bit flip leads to a fitness gain or loss (depending on whether the distance decreases or not) of $\Theta(m/n)$. In particular, the fitness function is such that there are no misleading one-bit flips. On such a function, the classic "multiplicative drift with the fitness" proof for ONEMAX can be imitated and easily yields an $O(n \log n)$ optimization time (in expectation and with high probability) [6, Theorem 2].

For smaller clause densities $m/n = \Omega(\log n)$, with probability 1-o(1), taken over both the random 3-CNF formula and the random decisions of the optimization process, also an $O(n \log n)$ optimization time is observed [6, Theorem 3]. The difference to the previous setting is that now there may be misleading one-bit flips, however, they are rare enough that a typical run does not encounter them.

In both these theorems, the (1 + 1) EA with mutation rate 1/nis regarded. Consequently, with constant probability exactly one bit is flipped. This together with the fitness-distance correlations exhibited is exploited to compute the drift. In contrast to this, the $(1 + (\lambda, \lambda))$ GA uses larger mutation rates, meaning that the typical distance between parent and offspring is larger. In addition, it does not need to estimate whether parent or offspring is closer to the optimum, but it needs to select among several offspring the one with slightly smaller distance.

For example, in the case $\lambda = 10$ consider a certain point in time when the parent individual has already a decent fitness. Then most offspring in the mutation phase will have a genotypic distance that is worse than that of the parent by a margin of 10. An offspring which has flipped a single bit which is missing in the parent will have a distance of 8. To be successful, this 20% advantage has to be visible for the $(1+(\lambda, \lambda))$ GA via a sufficiently strong fitness-distance correlation.

The common sense suggests that the bigger the λ is, the more problems the algorithm should have in detecting a "good" offspring. On the other hand, the bigger the clause density m/n is, the simpler it should be for the algorithm to handle bigger values of λ , as the problem becomes more similar to ONEMAX.

First, we investigate how the average fitness of a search point at a distance d from the optimum looks like, and how a difference of two such values behaves asymptotically depending on the difference between distances.

LEMMA 3.1. The probability for a random 3-CNF clause C, consisting of distinct variables and satisfied by the planted assignment x^* ,

to be also satisfied by an assignment x with the Hamming distance d = d(x) from x^* , is $P(n, d) = \frac{6 \cdot \binom{n}{3} + \binom{n-d}{3}}{7 \cdot \binom{n}{3}}$.

LEMMA 3.2. If $n-d = \Theta(n)$ and $\ell = o(n)$, then $P(n, d) - P(n, d+\ell) =$ $\Theta(\ell/n).$

COROLLARY 3.3. Consider a random 3-CNF formula on n variables and m clauses with the planted assignment x^* . The expected fitness $f_{avg}(d)$ of any search point x with the Hamming distance d = d(x)from x^* is $f_{\text{avg}}(d) = m \cdot \frac{6 \cdot \binom{n}{3} + \binom{n-d}{3}}{7 \cdot \binom{n}{3}}$.

For
$$n - d = \Theta(n)$$
 and $\ell = o(n)$, we have

$$f_{\text{avg}}(d) - f_{\text{avg}}(d-\ell) = \Theta(m\ell/n)$$

As we see, the average fitness values demonstrate a good fitnessdistance correlation. However, the actual fitness values of concrete search points may deviate quite far apart from the average values when the concrete 3-CNF formula is fixed. To show a good performance of the $(1 + (\lambda, \lambda))$ GA, we now show a stronger fitnessdistance correlation for offspring from the same parent. For a number $\overline{\lambda}$ to be made precise later, we define the following.

Definition 3.4. Consider a search point x with a distance d > 0from the planted assignment x^* . Let $\ell \in [1, \overline{\lambda}]$.

The set $X_{\ell}^{-} := X_{\ell}^{-}(x)$ of ℓ -bad offsping is the set of all search points produced from flipping in *x* exactly ℓ bits which coincide in x and x^* .

The set $X_{\ell}^+ := X_{\ell}^+(x)$ of ℓ -good offspring is the set of all search points produced from flipping in x exactly ℓ bits of which at least one is different in x and x^* .

The point *x* is *well-behaved* if for all $\ell \in [1, \overline{\lambda}]$

- (i) there are at most |X⁻_ℓ|/λ̄ elements x⁻ ∈ X⁻_ℓ such that f(x) f(x⁻) ≤ f_{avg}(d) f_{avg}(d + ℓ 1), and
 (ii) there are at most |X⁺_ℓ|/2 elements x⁺ ∈ X⁺_ℓ such that
- $f(x) f(x^+) \ge f_{\text{avg}}(d) f_{\text{avg}}(d + \ell 1).$

The motivation for this definition is that whenever the current search point x is well-behaved, the selection inside the mutation phase of the $(1 + (\lambda, \lambda))$ GA is able to distinguish with good probability an offspring with one of the missing bits guessed right from the offspring with no missing bit flipped.

LEMMA 3.5. Assume x is the current best solution of the (1 + (λ, λ) GA, f(x) < m, x is well-behaved and the current value of λ satisfies $\lambda \leq \overline{\lambda}$. Then, whenever at least one good offspring appears at the mutation phase, the $(1 + (\lambda, \lambda))$ GA will choose one of them with at least constant probability as mutation winner.

Now we show that for all sufficiently small values of $\overline{\lambda}$, all interesting search points x are well-behaved with an overwhelming probability. In a sense, this is the core of the main results of this paper. The proof uses a novel way to make McDiarmids inequality more powerful by artificially reducing the influence a change of the value of a random variable.

THEOREM 3.6. If $\overline{\lambda}$ is an integer satisfying $\overline{\lambda}$ = $o(\min\{n, (m/n)^{1/4}\})$, then with probability $e^{-\omega(n)}$ (taken in the probability space of random 3-CNF formulas) all x with $n - d(x) = \Theta(n)$ are well-behaved.

PROOF. There are 2^n possible search points x in total (including the search points which are too far from the optimum), and for every search point $2\overline{\lambda} = O(n)$ statements about sets $X_{\ell}^-(x)$ and $X_{\ell}^+(x)$ need to be proven. So by a union bound, it is enough to prove that for each $\ell \in [1..\overline{\lambda}]$ and each x with probability $e^{-\omega(n)}$ the sets $X_{\ell}^-(x)$ and $X_{\ell}^+(x)$ are as in the definition of well-behaved.

For the remainder of this proof, we fix an *x* and a value of ℓ . Let d := d(x). We prove only the statement on $X^{-}(x)$. The statement on $X^{+}(x)$ can be proven in the same way and then even yields a sharper bound as required in Definition 3.4.

For convenience, we use the short-hand notation $X^- = X_{\ell}^-(x)$. For an $x^- \in X^-$, the fitness loss $f(x) - f(x^-)$ is a random variable which is determines by the random formula *F*. We denote by $f^- = f_{avg}(d+\ell)$ the average value of $f(x^-)$ and by $f^= = f_{avg}(d+\ell-1)$ the threshold value from Definition 3.4. We use an indicator random variable $V(x^-)$ for the event that x^- violates the fitness constraint in Definition 3.4, that is,

$$V(x^{-}) := \begin{cases} 1, & f(x) - f(x^{-}) \le f_{\text{avg}}(d) - f^{=}, \\ 0, & \text{otherwise.} \end{cases}$$

The statement we want to prove is equivalent to

$$\Pr\left[\sum_{x^- \in X^-} V(x^-) > \frac{|X^-|}{\overline{\lambda}}\right] < e^{-\omega(n)}.$$
(3)

This expression calls for one of the Chernoff bounds to be applied. However, the resulting bounds are not strong enough for our aim to be achieved. For this reason, we introduce a new random variable, $\tilde{V}(x^-)$, which is less sensitive to the change of a single clause in the random instance. Let $f^{\equiv} = (f^- + f^=)/2$. Then

$$\tilde{V}(x^{-}) := \begin{cases} 0, & f(x) - f(x^{-}) \ge f_{\text{avg}}(d) - f^{\equiv}; \\ 1, & f(x) - f(x^{-}) \le f_{\text{avg}}(d) - f^{=}; \\ \frac{f(x) - f(x^{-}) - (f_{\text{avg}}(d) - f^{=})}{f^{=} - f^{\equiv}}, & \text{otherwise.} \end{cases}$$

As $\tilde{V}(x^-) \ge V(x^-)$ with probability one, it is enough to show, instead of (3), that

$$\Pr\left[\sum_{x^- \in X^-} \tilde{V}(x^-) > \frac{|X^-|}{\bar{\lambda}}\right] < e^{-\omega(n)}.$$
 (4)

For brevity, we write $\tilde{V}^{\Sigma} := \sum_{x^- \in X^-} \tilde{V}(x^-)$. We interpret the random variable \tilde{V}^{Σ} as a function of the independent uniformly distributed random variables C_1, \ldots, C_m which define the random formula *F*. Recall that, by Corollary 3.3, $f^= - f^- = \Theta(m/n)$, and so is $f^= - f^{\Xi}$. Thus replacing a single C_i in *F* by C'_i introduces a change of at most $2/(f^= - f^{\Xi}) = \Theta(n/m)$ in the value $\tilde{V}(x^-)$, however, only for those x^- for which *x* and x^- differ in one of the at most 6 variables contained in $C_i \cup C'_i$ (for all other x^- , the value of $\tilde{V}(x^-)$ does not change). The number of the former kind of x^- is at most $\binom{n-d}{\ell} - \binom{n-d-6}{\ell} = \sum_{i=0}^{5} \binom{n-d-(i+1)}{\ell} \leq \binom{n-d-1}{\ell-1} = \binom{n-d}{\ell} \frac{\ell}{n-d} = |X^-| \cdot \Theta(\ell/n)$. The maximum change of \tilde{V}^{Σ} inflicted by changing one clause hence is $|X^-| \cdot O(\ell/m)$.

We use it to apply McDiarmid's inequality, see (2), and obtain

$$\Pr\left[\tilde{V}^{\Sigma} > |X^{-}|/\overline{\lambda}\right] = \Pr\left[\tilde{V}^{\Sigma} > E\left[\tilde{V}^{\Sigma}\right] + \left(|X^{-}|/\overline{\lambda} - E\left[\tilde{V}^{\Sigma}\right]\right)\right]$$

Maxim Buzdalov and Benjamin Doerr

$$\leq \exp\left(-\frac{2\left(\frac{|X^{-}|}{\overline{\lambda}} - E\left[\tilde{V}^{\Sigma}\right]\right)^{2}}{m\left(|X^{-}|O\left(\frac{\ell}{m}\right)\right)^{2}}\right) = \exp\left(-\Omega\left(\frac{m\left(1 - \overline{\lambda}\frac{E[\tilde{V}^{\Sigma}]}{|X^{-}|}\right)^{2}}{\ell^{2}\overline{\lambda^{2}}}\right)\right).$$
(5)

To complete this bound, we need to show that the term $(1 - \overline{\lambda} E[\tilde{V}^{\Sigma}]/|X^{-}|)$ is at least some positive constant. To do this with the least effort, we introduce the random variable $\overline{V}(x^{-})$ defined by

$$\overline{V}(x^{-}) = \begin{cases} 1, & f(x) - f(x^{-}) \le f_{\text{avg}}(d) - f^{\equiv}, \\ 0, & \text{otherwise} \end{cases}$$

for all $x^- \in X^-$, and observe that it dominates $\tilde{V}(x^-)$. Consequently, we have $E[\tilde{V}(x^-)] \leq E[\overline{V}(x^-)] = \Pr[f(x) - f(x^-) \leq f_{avg}(d) - f^{\equiv}]$. By symmetry, these probabilities are identical for all $x^- \in X^-$. As $E[f(x^-)] = f^-$, $f^{\equiv} - f^- = \Theta(m/n)$, and $E[f(x) - f^-] = \Theta(\ell m/n)$, the bound from (1) gives

$$\begin{aligned} \Pr[f(x) - f(x^{-}) &\leq f_{\text{avg}}(d) - f^{\equiv}] \\ &= \Pr[f(x) - f(x^{-}) \leq E[f(x)] - E[f(x^{-})] + f^{-} - f^{\equiv}] \\ &= \Pr[f(x) - f(x^{-}) \leq E[f(x) - f(x^{-})] - \Theta(m/n)] \\ &\leq \max\{\exp(-\Omega(m^2/n^2)/(12\ell m/n)), \exp(-(3/8)\Omega(m/n))\} \\ &= \max\{\exp(-\Omega(m/\ell n)), \exp(-\Omega(m/n))\} = \exp(-\Omega(m/\ell n)). \end{aligned}$$

Note that $f(x) - f(x^{-})$ is the sum of *m* independent random variables describing the influence of each of the *m* random clauses on this expression. With probability at least $1 - 3\ell/n$, a random clause contains none of the λ variables *x* and x^{-} differ in. In this case, the clause contributes equally to f(x) and $f(x^{-})$, hence zero to the difference.

To finish this part, we shall note that we can estimate

$$\begin{split} \overline{\lambda} \frac{E[\overline{V}^{\Sigma}]}{|X^{-}|} &= \frac{\lambda}{|X^{-}|} \sum_{x^{-} \in X^{-}} E[\widetilde{V}(x^{-})] \leq \frac{\lambda}{|X^{-}|} \sum_{x^{-} \in X^{-}} E[\overline{V}(x^{-})] \\ &= \overline{\lambda} E[\overline{V}(x^{-})] = \overline{\lambda} \exp(-\Omega(m/\ell n)) \leq \overline{\lambda} \exp(-\Omega(m/\overline{\lambda}n)). \end{split}$$

Using the assumption $\overline{\lambda} = o((m/n)^{1/4})$, we estimate

$$\overline{\lambda} \cdot \exp(-\Omega(m/\overline{\lambda}n)) = o\left(\left(\frac{m}{n}\right)^{1/4}\right) \cdot \exp\left(-\omega\left(\left(\frac{m}{n}\right)^{3/4}\right)\right) = o(1),$$

which means that $1 - \overline{\lambda} E[\tilde{V}^{\Sigma}]/|X^{-}| = 1 - o(1)$. Consequently, from (5) we obtain

$$\Pr\left[\tilde{V}^{\Sigma} > |X^{-}|/\overline{\lambda}\right] \le \exp\left(-\Omega\left(\frac{m}{\ell^{2}\overline{\lambda^{2}}}\right)\right) \le e^{-\omega(n)},$$

where the second estimate follows from $\ell \leq \overline{\lambda}$ and $\overline{\lambda} = o((m/n)^4)$. This finishes the proof.

Now we are ready to state the main result of this section.

THEOREM 3.7. Consider solving random 3-CNF formulas with planted solutions on n variables and m clauses by the $(1 + (\lambda, \lambda))$ GA, where $m/n > c \log n$ for large enough constant c. If there exists some integer $\overline{\lambda} = o(\min\{n, (m/n)^{1/4}\})$ such that, during the entire run of the $(1 + (\lambda, \lambda))$ GA, $\lambda \le \overline{\lambda}$, and the algorithm starts from a random assignment, then with probability 1 - o(1), taken both over the random 3-CNF formulas and the algorithm decisions, the algorithm will demonstrate, in every point, the same progress, divided by at most constant, as the same algorithm optimising ONEMAX on n variables.

4 RUNNING TIME OF THE ADAPTIVE (1 + (λ, λ)) GA WITH CONSTRAINED λ

In the previous section, we were able to show that the $(1+(\lambda, \lambda))$ GA can cope with a weaker fitness-distance correlation when the offspring population size λ is not too large. Since the self-adjusting version of this GA can increase λ to relatively large values, see [3], we now propose a variant that avoids this. For an integer $\overline{\lambda}$, the $\overline{\lambda}$ -constrained self-adjusting $(1 + (\lambda, \lambda))$ GA is identical to the self-adjusting GA proposed in [5] except that after the adjustment a λ -value larger than $\overline{\lambda}$ is reduced to $\overline{\lambda}$. Consequently, this GA never uses an offspring population size larger than $\overline{\lambda}$. Our main result is that even with a logarithmic upper bound $\overline{\lambda}$, the asymptotic performance on ONEMAX is not worsened. This allows to use the insights of the previous section and prove that the constrained self-adjusting GA performs well also on our random 3-SAT instances.

THEOREM 4.1. Let $\overline{\lambda} = \overline{\lambda}(n)$ be an integer. The expected running time of the $\overline{\lambda}$ -constrained self-adjusting $(1 + (\lambda, \lambda))$ GA on ONEMAX is $O(n \cdot \max\{1, (\log n)/\overline{\lambda}\})$. If $\overline{\lambda} = o(\min\{n, (m/n)^{1/4}\})$, then this result extends to the random 3-SAT instances regarded in this work. In particular, by taking $\overline{\lambda} = \lceil \min\{\sqrt{n}, (m/n)^{1/4-\varepsilon}\} \rceil$ for any sufficiently small constant ε , this runtime becomes $O(n \cdot \max\{1, (\log n)/(m/n)^{1/4-\varepsilon})\})$ and thus linear for $m/n \ge (\log n)^{4+20\varepsilon}$.

We briefly sketch why this result is true. In [5], it was shown that for the optimization of ONEMAX an asymptotically optimal value of λ , given the current Hamming distance d(x) = n - f(x) from the optimum, is $\lambda^* = \sqrt{n/d(x)}$. The self-adjusting GA, as seen in experiments [5], tends to keep λ close to this optimum value, and this effect has been proven in [3].

When d(x) is larger than $n/(2\overline{\lambda})^2$, then this optimal value λ^* is below $\overline{\lambda}$. Hence, very roughly speaking, the $\overline{\lambda}$ -constrained and the unconstrained variants of the self-adjusting GA behave identical (in fact, the constrained version rather profits from the constraint on λ , which prevents uselessly large λ -values). Consequently, the expected time to reduce the *d*-value from its initial value to a value of at most $n/(2\overline{\lambda})^2$ is at most O(n) as for the unconstrained version.

When $d(x) \leq n/(2\overline{\lambda})^2$, then the current value of λ will stay close to $\overline{\lambda}$. Consequently, again very roughly speaking, a run of the constrained self-adjusting GA is similar to one with fixed value $\lambda = \overline{\lambda}$ for the offspring population size. For the latter, the expected time (always counting fitness evaluations) to improve a current *d*-value of d(x) is $O(\max{\lambda, n/d(x)\lambda})$. Hence this second part of the optimization process takes an expected time of at most $O((n \log n)/\overline{\lambda})$.

5 EXPERIMENTS

Although theoretical results give us insights on how powerful our adaptation schemes are, they may be not enough to tell practitioners whether they can be used to solve practical problems. To complement our theoretical research, we conducted a series of experiments, which show practical efficiency of these schemes.

5.1 Results for OneMax

The first experiment was dedicated to evaluation of performance for fixed-size $(1 + (\lambda, \lambda))$ GA, as well as unlimited and constrained adaptation schemes, on the ONEMAX problem. Similar experiments



Figure 1: Plots of median runtimes on OneMax

had been also performed in [5]. In this paper, we evaluated much larger problem sizes, which helped revealing subtle differences in behaviour of various algorithms. We also evaluated more fixed population sizes systematically, and we also included the version of the $(1 + (\lambda, \lambda))$ GA with constrained adaptation $(\lambda \le 2 \log(n + 1))$.

We considered problem sizes to be powers of two, from 2^4 to 2^{24} . The fixed population sizes were chosen to be $\lambda \in [2..20]$. For the adaptive $(1 + (\lambda, \lambda))$ GA schemes, the unlimited (denoted as $\lambda \le n$) and the logarithmically constrained (denoted as $\lambda \le 2 \ln n$) ones were included in the comparison. We also ran the (1 + 1) EA. For every algorithm and problem size, 100 runs were executed, and the median number of function evaluations is reported.

The results are presented in Fig. 1. The abscissa axis, logarithmically scaled, represents the problem size, while the ordinate axis, linearly scaled, represents the number of function evaluations, divided by the problem size. Thus, $\Theta(n \log n)$ algorithms give a straight inclined line, and $\Theta(n)$ algorithms give a horizontal line.

It can be seen that the fixed-size versions of the $(1 + (\lambda, \lambda))$ GA demonstrate a slightly superlinear behaviour at small sizes, and then switches at certain threshold to a strictly $\Theta(n \log n)$ behaviour. This behaviour follows from the $\Theta(\min(n \log n/\lambda, n\lambda \log \log \lambda/\log \lambda))$ bound proven in [4]. The lower envelope of the fixed-size plots exactly corresponds to the optimal fixed population size as a function of the problem size. It can be seen that it behaves like a convex upwards function. This ideally corresponds to the $\Theta(n\sqrt{\log n} \log \log \log n/\log \log n)$ running time estimation proven for the optimal fixed population size in [4].

Both adaptive versions demonstrate easy-to-see linear runtime, with a quotient of approximately 11 for the unlimited adaptation, and of approximately 12.5 for the logarithmically constrained adaptation. Thus, constraining the population size by a logarithmic bound not only preserves the asymptotics of the runtime, but also changes the absolute performance only very little.

5.2 Results for Random 3-CNF Formulas

The second experiment was dedicated to solving random 3-CNF formulas with planted solutions. The same set of algorithms was considered. We took the problem sizes again in the form of 2^t , where $t \in [7..20]$. The rise of the lower bound is due to the fact that, at sizes 2^6 and below, difficult formulas became too often to appear, and studying the behaviour of the $(1 + (\lambda, \lambda))$ GA on such formulas is not a scope of this research. The number of clauses *m* was chosen to be $m(n) = \lfloor 4n \log n \rfloor$.



Figure 2: Plots of median runtimes on random 3-CNF formulas with planted solutions



Figure 3: Example runs with unconstrained λ , $n = 2^{16}$

The results are presented in Fig. 2. The trends seen here are quite similar to the ones for ONEMAX (see Fig. 1). One can see that the constrained adaptive $(1 + (\lambda, \lambda))$ GA, even with the logarithmic adaptation constraint (which has not yet been proven to be helpful) still performs better than all fixed-size variations. However, its runtime is probably not linear for this problem. The $(1 + (\lambda, \lambda))$ GA with unlimited adaptation, on the other hand, performs worse than the (1 + 1) EA. The reasons is that the population sizes $\lambda = \Theta(\sqrt{n})$, as well as mutation probabilities, are too high for Theorem 3.7 to hold, thus the adaptation logic of the $(1 + (\lambda, \lambda))$ GA breaks down.

We performed an additional series of runs, where for each iteration of each run we recorded the distance to the optimum d(x) of the currently best individual x, as well as the value of λ . The size was chosen to be $n = 2^{16}$, and the number of clauses m was, again, $m(n) = \lfloor 4n \log n \rfloor$. We made five runs, for which plots for dependencies of λ on d = d(x) are presented on Fig. 3. For convenience, the values of $\sqrt{n/d}$ were taken instead of d(x) to use with the abscissa axis, and both axes are logarithmic. In these coordinates, the optimal $\lambda(d) = \sqrt{n/d}$ values form a straight line (drawn in black).

Fig. 3 demonstrates that for distances which satisfy $\sqrt{n/d} \ge \log n$, the plots stay around the optimal values for λ . For smaller distances, adaptation starts to diverge, and values of λ tend to be greater than necessary. In the main experiments, the medians for the maximum λ values were as follows: 2109.99 for $n = 2^{13}$, 4747.47 for $n = 2^{14}$, 10681.82 for $n = 2^{15}$ and 27945.01 for $n = 2^{16}$. These values suggest that maximum λ tends to be $\Theta(n)$.

Maxim Buzdalov and Benjamin Doerr

6 CONCLUSION

The runtime analysis of the $(1 + (\lambda, \lambda))$ -GA conducted in this work shows that the GA can cope with a weaker fitness-distance correlation than the perfect one of the ONEMAX test function, the only example mathematically analyzed before. However, a weaker fitness-distance correlation requires that the population size λ is not taken too large, as otherwise the strong mutation rate of λ/n creates offspring that are too far from each other for the GA to find the closest one to the optimum in the intermediate selection step.

Our recommendation on how to use the $(1 + (\lambda, \lambda))$ -GA therefore is to first try a moderate size constant λ , say $\lambda = 5$ or $\lambda = 10$. If this leads to an improved performance, than larger values of λ can be tried. For the self-adjusting version of the GA, we generally recommend using an upper limit for the value which λ can take. For first experiments, this value should be taken around the best static value for λ and then slowly increased.

We remark that the main part of the body of the GA, namely the generation of y from x, can be used as a mutation operator also in other algorithms. We have no experience with this approach so far, but are optimistic that it can give good results as well.

Acknowledgments

This research was started when Maxim Buzdalov stayed at École Polytechnique in May 2016 supported by a grant from the French Consulate in Russia (bourse Metchnikov). Maxim Buzdalov was also financially supported by the Government of Russian Federation, Grant 074-U01. This research also benefited from the support of the "FMJH Program Gaspard Monge in optimization and operations research", and from the support to this program from EDF.

REFERENCES

- Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. 2015. Black-box Complexity of Parallel Search with Distributed Populations. In Proceedings of Foundations of Genetic Algorithms. ACM, 3–15.
- [2] Benjamin Doerr. 2016. Optimal Parameter Settings for the (1 + (λ, λ)) Genetic Algorithm. In Proceedings of Genetic and Evolutionary Computation Conference. ACM, 1107–1114. Full version available at http://arxiv.org/abs/1604.01088.
- [3] Benjamin Doerr and Carola Doerr. 2015. Optimal Parameter Choices Through Self-Adjustment: Applying the 1/5-th Rule in Discrete Settings. In Proceedings of Genetic and Evolutionary Computation Conference. 1335–1342.
- [4] Benjamin Doerr and Carola Doerr. 2015. A Tight Runtime Analysis of the (1 + (λ, λ)) Genetic Algorithm on OneMax. In Proceedings of Genetic and Evolutionary Computation Conference. 1423–1430.
- [5] Benjamin Doerr, Carola Doerr, and Franziska Ebel. 2015. From black-box complexity to designing new genetic algorithms. *Theoretical Computer Science* 567 (2015), 87–104.
- [6] Benjamin Doerr, Frank Neumann, and Andrew M. Sutton. 2015. Improved Runtime Bounds for the (1+1) EA on Random 3-CNF Formulas Based on Fitness-Distance Correlation. In Proceedings of Genetic and Evolutionary Computation Conference. 1415–1422.
- [7] B. W. Goldman and W. F. Punch. 2014. Parameter-less Population Pyramid. In Proceedings of Genetic and Evolutionary Computation Conference. 785–792.
- [8] Per Kristian Lehre and Carsten Witt. 2012. Black-Box Search by Unbiased Variation. Algorithmica 64, 4 (2012), 623–642.
- [9] Colin McDiarmid. 1989. On the Method of Bounded Differences. Surveys in Combinatorics 141 (1989), 148–188.
- [10] Dirk Sudholt. 2012. Crossover speeds up building-block assembly. In Proceedings of Genetic and Evolutionary Computation Conference. 689–696.
- [11] Andrew M. Sutton and Frank Neumann. 2014. Runtime Analysis of Evolutionary Algorithms on Randomly Constructed High-Density Satisfiable 3-CNF Formulas. In *Proceedings of Parallel Problem Solving from Nature*. Lecture Notes in Computer Science, Vol. 8672. Springer, 942–951.
- [12] Ingo Wegener. 2002. Methods for the analysis of evolutionary algorithms on pseudo-Boolean functions. In *Evolutionary Optimization*, S. Sarker, X. Yao, and M. Mohammadian (Eds.). Kluwer, Dordrecht, 349–369.