



Blending Dynamic Programming with Monte Carlo Simulation for Bounding the Running Time of Evolutionary Algorithms

Kirill Antonov, Maxim Buzdalov, Arina Buzdalova, Carola Doerr

► To cite this version:

Kirill Antonov, Maxim Buzdalov, Arina Buzdalova, Carola Doerr. Blending Dynamic Programming with Monte Carlo Simulation for Bounding the Running Time of Evolutionary Algorithms. IEEE Congress on Evolutionary Computation (CEC'21), Jun 2021, Krakow, Poland. hal-03233843

HAL Id: hal-03233843

<https://hal.sorbonne-universite.fr/hal-03233843>

Submitted on 25 May 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Blending Dynamic Programming with Monte Carlo Simulation for Bounding the Running Time of Evolutionary Algorithms

Kirill Antonov
ITMO University,
Saint Petersburg, Russia

Maxim Buzdalov
ITMO University,
Saint Petersburg, Russia

Arina Buzdalova
ITMO University,
Saint Petersburg, Russia

Carola Doerr
Sorbonne Université,
CNRS, LIP6, Paris, France

Abstract—With the goal to provide absolute lower bounds for the best possible running times that can be achieved by $(1 + \lambda)$ -type search heuristics on common benchmark problems, we recently suggested a dynamic programming approach that computes optimal expected running times and the regret values inferred when deviating from the optimal parameter choice.

Our previous work is restricted to problems for which transition probabilities between different states can be expressed by relatively simple mathematical expressions. With the goal to cover broader sets of problems, we suggest in this work an extension of the dynamic programming approach to settings in which the transition probabilities cannot necessarily be computed exactly, but in which they can be approximated numerically, up to arbitrary precision, by Monte Carlo sampling.

We apply our hybrid Monte Carlo dynamic programming approach to a concatenated jump function and demonstrate how the obtained bounds can be used to gain a deeper understanding into parameter control schemes.

I. INTRODUCTION

Running time analysis of evolutionary algorithms and other search heuristics contributes to our understanding of black-box optimization not only by providing insights into the basic working principles that drive algorithms’ performance, but also by providing *lower bounds* for the performance of broad classes of black-box approaches [1], [2]. Typically expressed as black-box complexity bounds, these lower bounds can be seen as a baseline against which we can compare state-of-the-art solvers, with the goal to quantify the potential of further algorithm development. By comparing black-box complexity bounds for different classes of algorithms, we obtain insight into the impact of certain algorithm features, such as their degree of parallelism [3], [4], their memory [5], [11], their selection principles [7], properties of their sampling strategies [8]–[10], etc. More refined lower bounds can be obtained by specifying a family of algorithms, the available choices for configuring it, and by studying the dependency of the running time on the choices of these configurable components. This has been classically studied in the asymptotic sense [12]–[15], however, some recent works provide tight bounds also for concrete problem dimensions [16], [17].

Most of these works assume a *static* setting of the parameters. In practice, however, it is well known that *dynamic parameter settings* can greatly improve the efficiency of evolutionary algorithms [18]–[20]. Precise lower bounds for algorithms

with dynamic parameter settings are rare, see [21] for a survey of rigorous results for algorithms with dynamic parameter choices and [22]–[24] for a few exceptions that provide precise lower bounds for algorithms with dynamic parameter settings.

Proving precise lower bounds for algorithms with dynamic parameter choices is challenging for several reasons: (i) lower bounds require to make statements about *all* possible settings (as opposed to upper bounds, in which one concrete parameter setting is analyzed), (ii) the (often very complex and trajectory-based) dependency of the parameters on the state of the algorithm needs to be taken into account, and (iii) all bounds need to be very tight to obtain an overall bound that is precise enough to give meaningful results for concrete problem dimensions (and not just for the asymptotic behavior). In lack of precise methods to derive optimal algorithm behavior, several author teams have resorted to compute actions that maximize the step-wise progress, in hope that this “drift-maximization” is close to the true optimal parameter setting policy. This step-wise progress maximization was either done *exactly* [25] or via *Monte Carlo simulations* [26]. It is known, however, that even for the ONEMAX problem with perfect fitness-distance correlation, the drift-maximization strategy is only close to being optimal [23], but not strictly optimal [27], [28]. A different approach to compute precise lower bounds for algorithms with dynamic parameter settings was therefore taken in [27]. For the problem of minimizing the expected running time of $(1+1)$ -type evolutionary algorithms on ONEMAX, the authors first derive an exact method to compute the optimal dynamic parameter choices and then evaluate this formula for concrete problem dimensions using a dynamic programming approach. The work was greatly extended in [28], by considering different $(1+\lambda)$ -type heuristics and by not only computing the optimal parameter choices at each state, but also the cost of deviating from it.

Our contribution: We demonstrate in this work how the approach taken in [27], [28] can be extended to settings in which we do not necessarily have a closed-form expression from which we can derive exact parameter settings. Our overall approach is still based on dynamic programming, but we replace the exact formulae that model the probabilities to transition between different states by Monte Carlo simulations. That is, we take a similar approach as in [26], but we

do not simply maximize the expected drift, but use these approximated probabilities to derive strategies that minimize the expected running time – the quantity that most running time analyses focus on.

We apply our approach to the RUGGEDNESS function suggested in [29]. RUGGEDNESS extends the classical jump function [30] by introducing several small jumps that the algorithms need to make in order to find the optimal solution, see Section IV-A for a formal definition. We have chosen this example because in [31] we observed that classical parameter control schemes essentially fail on RUGGEDNESS, in the sense that they show worse performance than the $(1 + \lambda)$ evolutionary algorithm (EA) with static mutation rate $1/n$ (with n denoting the dimension of the problem). Interestingly, a better performance is obtained when mutation rates are capped at $1/n$ instead of $1/n^2$; see Figure 1 for an illustration. That is, despite having less flexibility, the algorithms perform better. High-performing parameter control mechanisms should not suffer from more flexibility, and we therefore see this example as an interesting use-case to identify deficiencies of state-of-the-art parameter control techniques, and to develop appropriate remedies to mitigate these.

We suggested in [28] heatmaps to illustrate the behavior of parameter control schemes. While indeed being very useful to bound the regret (i.e., the cost of deviating from the optimal parameter setting) of a single iteration, these heatmaps turned out to be suboptimal when analyzing the whole search trajectory. We therefore suggest in this work to complement them with plots that illustrate the regret per iteration. With these plots, one can now make a clear distinction between quick cheap “hiccups” and long-term harmful deviations.

Implications for Ruggedness: The application of our method to RUGGEDNESS shows that the $(1 + \lambda)$ EA with static mutation rate $1/n$ is around 50-75% worse than the $(1 + \lambda)$ EA variant using optimal dynamic mutation rates. We also illustrate the usefulness of our regret plots by analyzing why the $(1 + \lambda)$ EA(A, b) with generalized 1/5-th success rule, shown efficient for other problems in [6], fails on this problem.

Impact and limitations of our work: While the precise values of the lower bound and of the individual regrets for RUGGEDNESS may be of interest to a rather specialized community only, we believe that our general approach of hybridizing Monte Carlo simulation with dynamic programming should be of interest to everyone willing to derive precise lower bounds against which sampling-based optimization algorithms can be compared.

At present, our work is restricted to settings in which states are never visited twice. In practice, this limits our work to greedy (“elitist”) algorithms. Extensions to algorithms and/or problems in which states may be visited more than once forms one of the most challenging directions for future work, on which we comment in more detail in Section V.

Structure of the paper: The $(1 + \lambda)$ EA framework and its main variants studied in [31] are introduced in Section II. In Section III we describe our hybrid dynamic programming approach using Monte Carlo simulation. In Section IV we

Algorithm 1: A family of $(1 + \lambda)$ -type algorithms

Data: n : problem size; $f : \{0, 1\}^n \rightarrow \mathbb{R}$: function to maximize; λ : population size; $\mathcal{D}(p)$: a family of parameterized distributions over $[0..n]$

- 1 Sample parent $x \in \{0, 1\}^n$ uniformly at random;
- 2 **for** $t \leftarrow 1, 2, \dots$ **do**
- 3 **for** $i \in [1..\lambda]$ **do**
- 4 Choose a distribution parameter p_i^t ;
- 5 Sample $k_i \sim \mathcal{D}(p_i^t)$, the number of bits to flip;
- 6 Create y_i by flipping k_i bits in x chosen uniformly at random without replacement;
- 7 Select $x \leftarrow \arg \max_{z \in \{x, y_1, \dots, y_\lambda\}} f(z)$;

present the results of this approach applied to derive the dynamic parameter settings which minimize the expected running time of the $(1 + \lambda)$ EA on RUGGEDNESS. We comment on further extensions of our approach in Section V. Section VI concludes the paper.

II. $(1 + \lambda)$ EVOLUTIONARY ALGORITHMS

We present in Algorithm 1 a fairly general family of $(1 + \lambda)$ EAs, for the context of maximizing a “fitness” function $f : \{0, 1\}^n \rightarrow \mathbb{R}$ defined on bit strings. Algorithms covered by this framework are initialized by selecting their first solution candidate uniformly at random. In each iteration, λ “offspring” are sampled, by modifying the best-so-far solution x . The modification is done by unary unbiased variation operators [8], which –according to a characterization provided in [23]– are defined by fixing a distribution \mathcal{D} over the set of possible mutation strengths $k \in [0..n] := \{0, 1, 2, \dots, n\}$. The operator first samples a mutation strength k from this distribution and then flips k uniformly selected, pairwise different bits in x .¹ The best-so-far solution “survives” and determines the center of variation in the following iteration. Ties can be broken arbitrarily; in practice and in the remainder of this work we assume uniform selection among the best offspring (if at least one of them is at least as good as the parent).

We do not specify in Algorithm 1 a stopping criterion, because we are – in line with the majority of running time analysis papers – interested in the *expected optimization time*, i.e., the average number of function evaluation that an algorithm performs until it evaluates for the first time a solution x^* with $f(x^*) \geq f(y)$ for all $y \in \{0, 1\}^n$.

In [31], the following $(1 + \lambda)$ EA variants were studied:

(1) Shift $(1 + \lambda)$ EA: The standard $(1 + \lambda)$ EA with mutation rate p is Algorithm 1 with $\mathcal{D}(p_i^t) = \mathcal{B}(n, p)$ for all $t \in \mathbb{N}$, $i \in [1..\lambda]$. It was argued in [32], [33] that for $(1 + \lambda)$ schemes this operator should not be used in practice, since it assigns probability mass to flipping 0 bits, which cannot advance the search.

¹Examples: The two most commonly used unary unbiased variation operators are 1-bit-flips and standard bit mutation (SBM). 1-bit-flips, as used within Randomized Local Search (RLS), correspond to the 1-point distribution $\mathbb{P}[k = 1] = 1$, whereas SBM (used by many evolutionary algorithms) is exactly the operator characterized by the binomial distribution $\mathcal{B}(n, p)$.

Two strategies were suggested to mitigate this unfavorable behavior: “shifting” all probability mass from $k = 0$ to $k = 1$ and “resampling” k until $k > 0$. We use the shift strategy in all algorithms considered in this work, i.e., we use the mutation operator characterized by the distribution $\mathcal{B}_{0 \rightarrow 1}(n, p)$, which has $\mathbb{P}[\ell = 1 \mid \ell \sim \mathcal{B}_{0 \rightarrow 1}(n, p)] = \mathbb{P}[\ell \leq 1 \mid \ell \sim \mathcal{B}(n, p)]$ and $\mathbb{P}[\ell = k \mid \ell \sim \mathcal{B}_{0 \rightarrow 1}(n, p)] = \mathbb{P}[\ell = k \mid \ell \sim \mathcal{B}(n, p)]$ for $k \geq 2$. Where p is not specified, we tacitly assume $p = 1/n$.

(2) $(1 + \lambda)$ **EA**(A, b): This algorithm, analyzed in [6], uses a multiplicative parameter update scheme inspired by the 1/5-th success rule [34]. Our version uses shift mutation (as introduced above) instead of resampling mutation, but the update of p remains the same as in [6] with $A = 2$ and $b = 1/2$. That is, p is initialized as $1/n$. At the end of each iteration p is updated to $2p$ if $\max\{f(y_1), \dots, f(y_\lambda)\} \geq f(x)$ and to $p/2$ otherwise.

(3) $(1 + \lambda)$ **EA** _{$r/2, 2r$} : This algorithm was suggested in [35]. It creates $y_1, \dots, y_{\frac{\lambda}{2}}$ by the shift mutation with mutation rate $p/2$ and the other offspring, $y_{\frac{\lambda}{2}+1}, \dots, y_\lambda$ by the shift mutation with mutation rate $2p$. Let $f_h = \max\{f(y_1), \dots, f(y_{\frac{\lambda}{2}})\}$ and $f_d = \max\{f(y_{\frac{\lambda}{2}+1}), \dots, f(y_\lambda)\}$. If $f_h > f_d$, p is updated to $p/2$ with probability $3/4$ and to $2p$ otherwise. If $f_d > f_h$, p is updated to $2p$ with probability $3/4$ and to $p/2$ otherwise. If $f_d = f_h$, p is updated to either $p/2$ or $2p$ equiprobably.

(4) $(1 + \lambda)$ **HQEA**: This algorithm was suggested in [31]. It hybridizes reinforcement learning with the multiplicative update used within the $(1 + \lambda)$ EA(A, b). We cannot present details, for reasons of space, and refer the interested reader to [31] for details. For this work, it is interesting to note that the algorithm also uses shift mutation, with a mutation rate p that is updated after each iteration. All λ offspring are sampled from the same distribution.

Capping of the mutation rate p : An advantage of the shift mutation operator is that it converges to the 1-bit-flip operator when $p \rightarrow 0$. It can therefore be used to interpolate between global and local search. However, we typically want to maintain *some* probability of escaping local optima. In practice, p is therefore often capped to remain within an interval $[p_{\min}, p_{\max}]$. In our work, when considering parameter control methods, we fix $p_{\max} = 1/2$. As discussed in the introduction, p_{\min} can have a decisive influence on the efficiency of the algorithms, as can be seen in the example on the right side of Figure 1. We study two values, $p_{\min} = 1/n$ and $p_{\min} = 1/n^2$.

Scope of our work: We will focus in the following on computing a lower bound for the family of algorithms that follow Algorithm 1, but which use identical distributions to sample the λ offspring (i.e., we require that $\mathcal{D}(p_i^t) = \mathcal{D}(p_j^t)$ for all $i, j \in [\lambda]$). We also restrict the algorithms to those in which $\mathcal{D}(p^t)$ (using the previous convention, we will from now on drop the index i) may depend on n and on $f(x)$, but not on x ,² nor on t , nor on any other information.

²This assumption is irrelevant for our use-case, but the structure of x can have an important impact in general, as is easily seen for the LEADINGONES function, for which $f(0 \dots 0) = f(01 \dots 1) = 0$, but the optimal mutation rate for $(0 \dots 0)$ is one, whereas the optimal one for $(01 \dots 1)$ is p_{\min} .

III. ALGORITHM DESCRIPTION

In this section we describe our algorithm that computes good choices of distribution parameters p for each possible fitness value of the parent.

A. High-Level Description

Our algorithm is based on dynamic programming. We iterate over the possible fitness values, starting from the second best and continuing towards the smallest value. For each fitness value we aim at finding the best possible parameter value that minimizes the remaining expected running time, as well as the remaining time itself, assuming that the $(1 + \lambda)$ EA will subsequently choose the optimal parameter values for all better fitness values. As a side effect we also compute the remaining expected running times for a number of parameter values, which will be later useful to evaluate regrets associated with particular parameter control schemes.

While doing it, we assume, similarly to [27], [28], that for all higher fitness values the best possible expected running times are already computed. However, since we aim at dealing with various fitness functions, we use the Monte Carlo approach to approximate transition probabilities instead.

B. Requirements for the Fitness Function

Before diving into details, we discuss the limitation of this algorithm first. Our main limitation follows from our assumption that, apart from the problem size, the optimal choice of the distribution parameter p depends on the fitness value exclusively. Hence, if two individuals have the same fitness but different structure, it may result in different transition probabilities to higher fitness values. Failing to account for that may result in both overestimation and underestimation of running times. For this reason, we formulate our requirement to the fitness function as follows.

Requirement 1. For any two individuals x_1 and x_2 , such that $f(x_1) = f(x_2)$:

- either both x_1 and x_2 shall be optima;
- or for each fitness value $f' > f(x_1)$, assuming $Y = \{y \mid f(y) = f'\}$ is a set of bit strings with the fitness f' , there shall exist a bijective mapping $\pi : Y \rightarrow Y$ such that such that for each $y \in Y$ the following transition probabilities shall be equal: $\mathbb{P}[x_1 \rightarrow y] = \mathbb{P}[x_2 \rightarrow \pi(y)]$.

When applied inductively, this requirement informally means that the fitness value unambiguously determines the further stochastic behavior of the evolutionary algorithm.

Note that some popular benchmark fitness functions, such as ONEMAX, as well as the function RUGGEDNESS used in this study, satisfy this requirement. Some other functions satisfy it only partially: for instance, with certain definitions of the JUMP function, the individuals that form the small fitness valley may have the same fitness but be structurally different. However, in this particular case the probability that the parent of the $(1 + \lambda)$ EA becomes such an individual during the optimization run is overwhelmingly small, so we can still

Algorithm 2: High-level outline of the hybrid algo.

```
1  $f_{\min}, f_{\max} \leftarrow$  minimum and maximum fitness values;
2 Initialize optimal times:  $T_{f_{\max}}^* \leftarrow 0$ ;
3 for  $f \leftarrow f_{\max} - 1, \dots, f_{\min}$  do
4   for  $p \in \{p_1^{(f)}, p_2^{(f)}, \dots, p_{m_f}^{(f)}\}$  do
5     Compute approximate probabilities  $(\tilde{p}_i)_{i=0,1,\dots}$ 
      of increasing fitness by  $i$  with mutation rate  $p$ 
      using the Monte Carlo approach;
6      $T_{f,p} \leftarrow \frac{1}{1 - \tilde{p}_0} \left( 1 + \sum_{i>f} T_i \cdot \tilde{p}_{(i-f)} \right)$ ;
7   Store optimal time:  $T_f^* \leftarrow \min_p(T_{f,p})$ ;
8   Store optimal rate:  $P_f^{\text{opt}} \leftarrow \arg \min_p(T_{f,p})$ ;
9 return  $\{P^{\text{opt}}, T^*, T\}$ 
```

apply our algorithm and get the results that are imprecise only up to a factor whose difference from one is negligible.

C. Detailed Description

The high-level pseudocode of the proposed algorithm is given in Algorithm 2.

We maintain T_f^* , our approximation of the optimal expected remaining time starting at fitness f until the optimum is hit. This value for the maximum fitness, f_{\max} that corresponds to the optimum, is obviously zero: $T_{f_{\max}}^* = 0$. We compute the remaining values starting from $f_{\max} - 1$ and stepping down until the minimum fitness value f_{\min} is processed. Since the algorithms from the $(1 + \lambda)$ EA family are elitist, T_f^* depends only on $T_{f'}^*$ with $f' > f$, but not on other entries.

To evaluate a particular T_f^* , we evaluate the expected remaining times $T_{f,p}$ starting at fitness f until the optimum is hit, provided that (i) while the parent's fitness is f , the $(1 + \lambda)$ EA uses the mutation rate p , and (ii) when the parent's fitness is updated, the $(1 + \lambda)$ EA uses the previously computed optimal mutation rates. Since we aim at an easy-to-compute approximation scheme that does not depend on the structure of the fitness function, we employ the following simplifications.

- Instead of testing all possible p , which is computationally infeasible, we use a finite set of mutation rates $\{p_1^{(f)}, \dots, p_{m_f}^{(f)}\}$ that may depend on f . To obtain a good approximation, this set shall be dense enough around the assumed optimal mutation rate, but in the same time it should have a small enough size so that the computation finishes in affordable times. We detail our choices of these sets later in Section III-D.
- Instead of analytically computing transition probabilities of the $(1 + \lambda)$ EA from the current fitness f to each larger fitness f' (which may be very hard, error-prone and computationally demanding for certain problems) we use the Monte Carlo approach, which amounts to simulation of a part of the run of the $(1 + \lambda)$ EA.

The Monte Carlo evaluation of $T_{f,p}$ is performed as follows.

- 1) Choose a large number of modeled iterations, N_I .

- 2) Choose a large number N_T of *successful* iterations (the iterations that resulted in a strict fitness increase) to wait for: if the transitions are easy enough, N_T successes already make a good estimation of the majority of likely transitions, so we do not execute all N_I iterations and hence save computational resources.
- 3) Find an individual with fitness f and set it as the parent individual for the next iterations of the $(1 + \lambda)$ EA.
- 4) Perform N_I iterations of the $(1 + \lambda)$ EA (or fewer if N_T is hit earlier), however, at the end of each iteration do *not* update the parent regardless of the outcome. Instead, for the j -th iteration we record the fitness f'_j of the individual that would otherwise become the next parent.

Note that the correctness of not updating the parent even in the case of neutral move is motivated by Requirement 1: any individual with the fitness f induces the same behavior, so execution of a neutral move may be safely omitted.

Let $N_A \leq N_I$ be the actual number of simulated iterations. The probability estimations \tilde{p}_i of increasing the fitness by i are computed as $\tilde{p}_i = \frac{1}{N_A} \cdot |\{j \mid 1 \leq j \leq N_A; f'_j = f + i\}|$. Note that, as the $(1 + \lambda)$ EA is elitist, \tilde{p}_0 counts also the occasions where the best offspring was worse than the parent.

Finally, $T_{f,p}$ is computed based on \tilde{p}_i and $T_{f'}^*$ for $f' > f$ using the equation in line 6 of Algorithm 2, which is a solution of the recurrent relation $T_{f,p} = 1 + \tilde{p}_0 \cdot T_{f,p} + \sum_{i>0} \tilde{p}_i \cdot T_{f+i}^*$.

D. Choice of Mutation Rate Sets

The choice of the mutation rate sets $\{p_1^{(f)}, \dots, p_{m_f}^{(f)}\}$ has to be a trade-off between the accuracy of the resulting values and the computation time. However, the freedom to choose different sets for different fitness values partially reduces the effects of this trade-off. In particular, one may conduct preliminary experiments to see which probabilities for which fitness values are most promising, and increase the coverage near these probabilities in more involved experiments. Technically, this opens a possibility of a self-adaptive scheme, where first some predefined grid is used for each fitness value, and then it is refined in the most promising regions to get better results. In this paper, however, we used a more conservative scheme detailed below.

In our experiments, we used a union of two grids: the multiplicative grid of the form $p_i = p_{\text{base}} \cdot \alpha^i$, which we use for large fitness values and which spans relatively small probabilities, and the additive grid of the form $p_i = p_{\text{base}} + i \cdot p_{\text{step}}$ to cover the whole ranges of fitness values and probabilities. As the particular parameter values need to depend on the problem size, we give these values, as well as the Monte Carlo simulation parameters N_I and N_T , in the next section.

IV. EXAMPLE APPLICATION OF OUR APPROACH

In this section we outline several kinds of insights that can be derived from the results computed by the proposed algorithm. Some of them, namely lower bounds for parameter control methods, plots of optimal parameter values, and parameter efficiency heatmaps, have been previously proposed in [27], [28], and the regret plots are new to this paper.

A. The Ruggedness Function

We have chosen as use-case the RUGGEDNESS function introduced in [29] (function F6 there). It is basically a function of concatenated jumps. That is, assuming that the unique optimum is located in some bit string z , with fitness value n , all points at Hamming distance one from z have fitness $n - 2$, while those at distance two have fitness $n - 1$, those at distance three have fitness $n - 4$, those at distance four have fitness $n - 3$, and so on. Formally, letting $\text{OM}_z(x) := \{i \in [1..n] \mid z_i = x_i\}$, RUGGEDNESS assigns $r(x) = n$ if $\text{OM}_z(x) = n$, $r(x) = \text{OM}_z(x) + 1$ if $\text{OM}_z(x) \equiv n \pmod 2$, and $r(x) = \text{OM}_z(x) - 1$ otherwise.

As mentioned above, all the mutation rate control methods considered in [31] perform worse on RUGGEDNESS than the shift $(1 + \lambda)$ EA with static mutation rate $p = 1/n$, see Fig. 1. This raises two important questions, which we will answer in the remainder of this section: (i) How far are the algorithms benchmarked in [31] from the best possible performance? (ii) What is the impact of the individual parameter choices that are made during the optimization process?

Our experiments consistently use problem size $n = 100$. This choice is motivated by three main factors: it is large enough to see absolute differences in algorithms' behavior, yet small enough to finish the Monte Carlo simulations, and it allows a straightforward comparison with the previous works [28], [31].

To compute the necessary data for RUGGEDNESS, we use the following parameters for our Monte Carlo approach: $N_I = 10^6$, $N_T = 5 \cdot 10^4$. We use the same set of probabilities for each fitness, which is constructed using a multiplicative grid with the following parameters: $p_{\text{base}} = 10^{-4}$ to match $p_{\min} = 1/n^2 = 10^{-4}$, and $\alpha = 10^{1/25}$, so that 100 probabilities less than 1 are employed, whose logarithms are evenly distributed.

B. Lower Runtime Bounds for Parameter Control Methods

The near-optimal runtime values T_f^* for each fitness f obtained from the proposed algorithm may be directly used to construct the lower bounds on the expected running times possible for parameter control methods. For this, we compute for each fitness value f the probability p_f^{init} of hitting it with the first created individual. The value $T = \sum_f T_f^* p_f^{\text{init}}$, which is the expected runtime of the $(1 + \lambda)$ EA assuming near-optimal parameter choices, is then used as lower bound.

Fig. 1 presents the comparison of these lower bounds, measured in iterations, with the performance of the standard $(1 + \lambda)$ EA, as well as of the $(1 + \lambda)$ EA employing a few existing parameter control methods, on RUGGEDNESS and on ONEMAX for comparison, all using the shift mutation. For ONEMAX we use the exact method for computing the lower bounds from [28]. For the parameter control methods we additionally employ two different lower bounds on the mutation rate, $p_{\min} = 1/n$ and $p_{\min} = 1/n^2$.

Now we compare the insights achieved for RUGGEDNESS with the ones achieved before. In [31] one was able to see only that some parameter control methods are far away from $(1 + \lambda)$ EA when using $p_{\min} = 1/n^2$, and the situation becomes better when $p_{\min} = 1/n$ is used. More precisely, with the less generous lower bound on the mutation rate, all the considered methods are able to perform just as $(1 + \lambda)$ EA. But is it possible to come up with a more efficient parameter control method that would outperform $(1 + \lambda)$ EA? With the just obtained lower bounds on expected runtimes, represented with green lines in Fig. 1, one can conclude that even if one makes the optimal mutation rate choices for each fitness value, only a relatively small constant-factor improvement would be possible. This gives an important bound which was not previously known.

One of the interesting observations on RUGGEDNESS, which can be derived from Fig. 1, is that the $(1 + \lambda)$ EA(A, b) has a much worse performance than both the $(1 + \lambda)$ EA $_{r/2, 2r}$ and

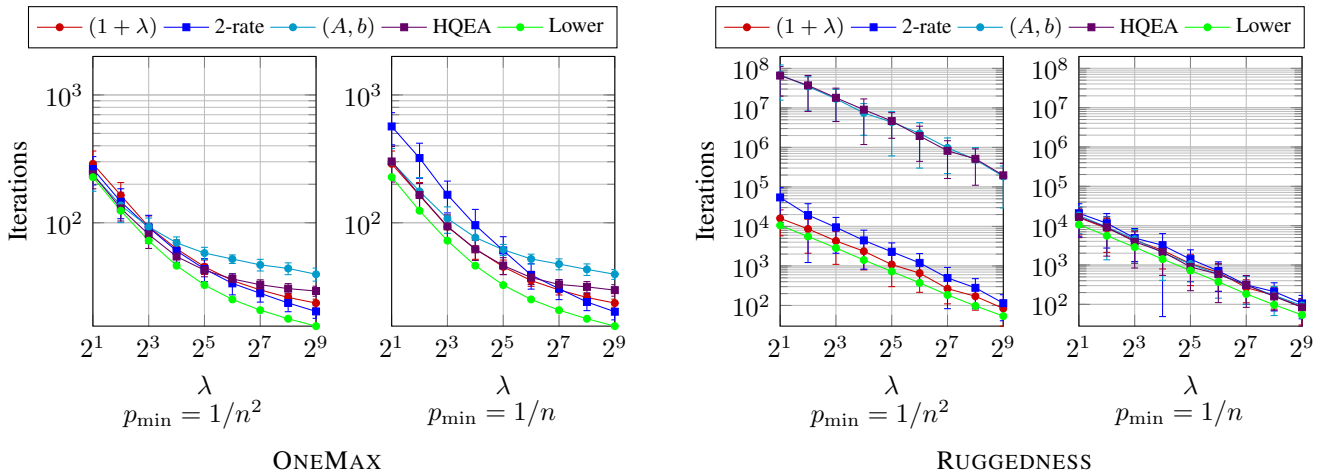


Fig. 1. Lower bounds for the expected running times of the $(1 + \lambda)$ EA using binomial distribution with varying p for mutation, compared with the performance of various parameter control methods averaged over 100 independent runs. Problem size is $n = 100$ for all four plots. Legend entries are: $(1 + \lambda)$: the $(1 + \lambda)$ EA with fixed $p = 1/n$, 2-rate: the $(1 + \lambda)$ EA $_{r/2, 2r}$, (A, b) : the $(1 + \lambda)$ EA(A, b), HQEA: the $(1 + \lambda)$ HQEA.




Fig. 2. Best mutation rates for different values of λ and fitness values. Problem size is $n = 100$ for both plots.

the $(1+\lambda)$ EA. However, the reason for such a behavior cannot be derived from Fig. 1. We return to this question further in this section by looking more closely at how these methods choose mutation rates and how they relate to the best choices.

C. Plots of Optimal Parameter Values

The observation of running times alone cannot provide deeper insights about the problem structure, including the optimal mutation rates for different fitness values. Such a knowledge can shed some light on which issues a parameter control method may face when the $(1+\lambda)$ EA solves the problem in question. The plots of near-optimal mutation rates, which the proposed algorithm provides, may serve this purpose. What is more, one can infer the influence of the population size λ on the optimal parameters from such plots.

Fig. 2 presents the plots of near-optimal mutation rates as a function of the fitness value for ONEMAX, computed as in [28], and RUGGEDNESS. One can see that for ONEMAX these plots demonstrate a steady tendency of decreasing while getting closer to the optimum, which matches the general expectations. At the same time, most the plots for RUGGEDNESS, while having a similar general trend, show regular oscillations with a period of 2. Indeed, a small mutation rate is better for the $(1+\lambda)$ EA to improve from the fitness value that has a different parity than n (by flipping a single bit), but at least two bits need to be flipped to improve from the fitness value with the same parity as n , which requires a larger mutation rate that tends to $2/n$ as one gets closer to the optimum.

Note that such optimal parameters may be difficult to be tracked precisely by most parameter control methods that assume a slow change of the optimal mutation rate.

D. Parameter Efficiency Heatmaps

Since Algorithm 2 provides expected runtimes $T_{f,p}$ for each fitness value and mutation rate, assuming that for higher fitness values the optimal mutation rates are chosen, we can use this

data to estimate the relative efficiency of mutation rate choices for each fitness value. In Fig. 3 we present this information as a heatmap for $n = 100$ and $\lambda = 512$. Each cell of a heatmap corresponds to a pair of f and p , where the following p are chosen: $p = 10^{-4+i/25}$ for $0 \leq i \leq 100$.

Colors of the cells represent the relative efficiency of the corresponding p among all mutation rates for the corresponding f . They are derived from the value $C_{f,p} = \exp(\alpha_f \cdot (T_f^* - T_{f,p}))$, where $\alpha_f \leq 1$ is chosen in such a way that at least a half of $C_{f,p}$ are at least 0.5. This is done to display enough information about at least 50% of possible choices, while not artificially emphasizing differences if they are negligible.

Using Fig. 3, one can make the following observations about dynamic mutation rates for the $(1+\lambda)$ EA on RUGGEDNESS:

- Overall, near-optimal mutation rates start with a rather high value at the beginning of the optimization and show a trend to gradually decrease towards the optimum.
- This trend is non-monotone. By comparison with ONEMAX, one can see that for even fitness values RUGGEDNESS shows behavior similar to ONEMAX. However, for odd fitness values the best mutation rates are higher.

These observations agree well with the findings from the previous subsection and Fig. 2. However, one can also see that in most of fitness values, except for a few last ones, there is a wide enough range of nearly equally good parameters for each fitness value, which shows that most parameter control methods have, in theory, a fair chance of maintaining good mutation rates until the last moments.

Fig. 3 also shows traces of the actual mutation rates chosen by the $(1+\lambda)$ EA_{r/2,2r} and $(1+\lambda)$ EA(A, b) algorithms on RUGGEDNESS, both with $p_{\min} = 1/n^2$, for the corresponding fitness values. The following observations can be made:

- Both algorithms quickly, in less than 10 iterations, reach near-optimal mutation rates at fitness of roughly 70.
- For fitness of roughly 80, both algorithms choose rea-




Fig. 3. Parameter efficiency heatmaps for $(1 + \lambda)$ EA, $n = 100$, $\lambda = 512$. Heatmap for RUGGEDNESS also contains traces of mutation rate choices in runs of the $(1 + \lambda)$ EA $_{r/2,2r}$ (black) and the $(1 + \lambda)$ EA (A, b) (red), both with $p_{\min} = 1/n^2$. Color denotes the relative efficiency, the larger the better.

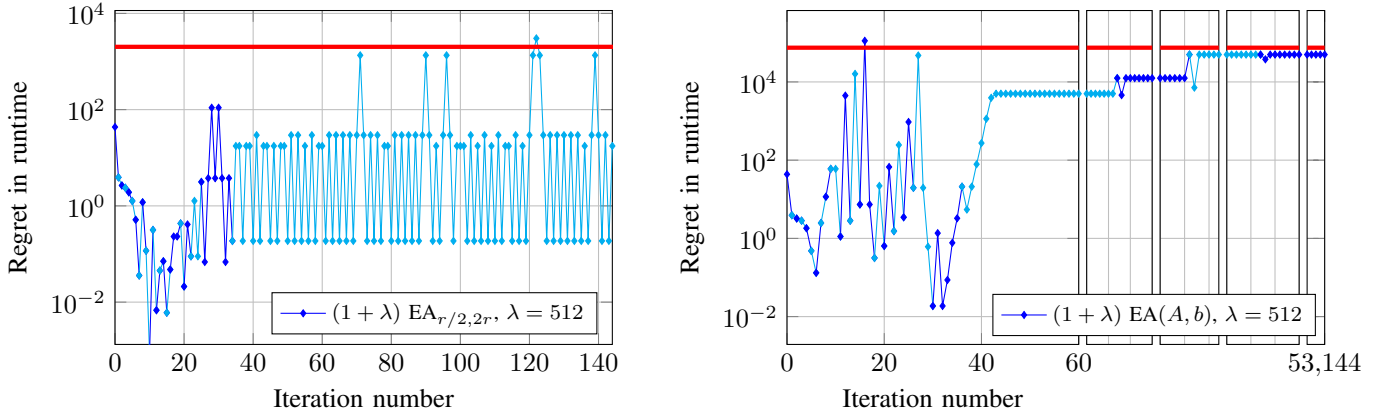


Fig. 4. Difference between the expected runtimes for the best mutation rate choice and the one actually chosen by the parameter control algorithms in example runs. Left plot shows a complete run of the $(1 + \lambda)$ EA $_{r/2,2r}$. Right plot shows selected iterations of the $(1 + \lambda)$ EA (A, b) , the omitted iterations have the same fitness and regret. Each red line denotes an artificial visual boundary between finite expected runtimes (below) and infinite ones (above).

sonably good values often enough, which ensures some progress, with the $(1 + \lambda)$ EA $_{r/2,2r}$ doing slightly better.

- For last few fitness values, both algorithms visit odd fitness values exclusively, but the $(1 + \lambda)$ EA $_{r/2,2r}$ still chooses good mutation rates, whereas the $(1 + \lambda)$ EA (A, b) gets stuck with the too small rates.

The latter observation clearly points at the behavior that causes the poor performance of the $(1 + \lambda)$ EA (A, b) seen at Fig. 1. It appears that the (A, b) adaptation rule forces a switch to local search as progress slows down, which converges the mutation rate to $p_{\min} = 1/n^2$. With such rate it makes the chance of finding two-bit improvements $\Theta(n^2)$ smaller, which explains the large performance gap observed in Fig. 1.

E. Regret Plots

Overlaying parameter choices of particular parameter control methods over the parameter efficiency heatmaps may indicate the presence of problems in a method, and how large

the deviation is. However, it cannot show how much of the performance the method loses from acting suboptimally.

A possible solution is to show how much the method loses to the optimal choice in each iteration. More precisely, for each iteration we take the parameter value p chosen by the parameter control method in question that corresponds to the current parent's fitness value, and plot the value $|T_{f,p} - T_f^*|$ for that iteration. Such *regret plots* are presented in Fig. 4. We change the plot color whenever the fitness value changes, so that one can see which iterations share the same fitness value.

We can observe in the example runs in Fig. 4 that both methods initially quickly reduce their regret. However, the $(1 + \lambda)$ EA $_{r/2,2r}$ keeps it at a low level, whereas the $(1 + \lambda)$ EA (A, b) shows a more complicated behavior. First it features two peaks of local regret maxima, which match two occasions of too high mutation rates seen in Fig. 3. It then quickly moves through a good region and falls down to small mutation rates, where it spends most of its time with very large regrets, whose sum dominates the total running time.

Note that, occasionally, a parameter control method may enter a region where $T_{f,p}$ reaches too high values, or where it is even infinite (which happens twice in Fig. 4). If the method is quick enough to return to better mutation rates, this does not harm the performance, as is the case for the $(1 + \lambda)$ EA_{r/2,2r}.

V. GENERALIZATION OF THE HYBRID APPROACH

As discussed in the introduction, our hybrid approach is at the moment restricted to settings in which states are not visited more than once. This can be resolved, at the price of computational complexity, by construction of Markov chains on all states with equal fitness and solving the resulting system of equations. However, this may also require a revision of used definitions, such as the one for regret: if the mutation rate is fixed to some value for one of the states, globally *suboptimal* choices for other states may actually make the runtime *smaller*.

Our approach is, in fact, not limited to $(1 + \lambda)$ type algorithms: as it simulates an iteration of the algorithm as a parameterized black-box, a wide range of algorithms may be modeled, which is much harder with exact approaches.

The number of states may become too large to have them considered explicitly for certain problems, as well as for population-based algorithms. While we see no obvious solution for this problem yet, state-merging approaches might be a solution for some of the cases.

VI. CONCLUSIONS

We have extended in this work the dynamic programming approach for computing optimal state-dependent, dynamic mutation rates suggested in [27], [28] to settings in which the transition probabilities cannot necessarily be computed by closed-form expressions, but where they need to be approximated by Monte Carlo simulations. We have applied this approach to derive optimal parameter choices for the $(1 + \lambda)$ EA on RUGGEDNESS. We have also introduced *regret plots*, which not only show deficiencies in parameter control methods, but indicate their impact on the running time.

We plan to extend our work to more complicated setting detailed in Section V, as well as to use the results presented in this work to improve some of the parameter control methods.

Acknowledgments. The reported study was funded by RFBR and CNRS, project number 20-51-15009, and by the Paris Ile-de-France region.

REFERENCES

- [1] B. Doerr and F. Neumann, Eds., *Theory of Evolutionary Computation—Recent Developments in Discrete Optimization*. Springer, 2020.
- [2] C. Doerr, “Complexity theory for black-box optimization heuristics,” in *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*. Springer, 2020.
- [3] G. Badkobeh, P. K. Lehre, and D. Sudholt, “Unbiased black-box complexity of parallel search,” in *PPSN*. Springer, 2014.
- [4] P. K. Lehre and D. Sudholt, “Parallel black-box complexity with tail bounds,” *IEEE Trans. Evol. Comp.*, vol. 24, no. 6, 2020.
- [5] S. Droste, T. Jansen, and I. Wegener, “Upper and lower bounds for randomized search heuristics in black-box optimization,” *Theory of Computing Systems*, vol. 39, 2006.
- [6] C. Doerr and M. Wagner, “Simple on-the-fly parameter selection mechanisms for two classical discrete black-box optimization benchmark problems,” in *GECCO*. ACM, 2018.
- [7] C. Doerr and J. Lengler, “Introducing elitist black-box models: When does elitist behavior weaken the performance of evolutionary algorithms?” *Evol. Comp.*, vol. 25, 2017.
- [8] P. K. Lehre and C. Witt, “Black-box search by unbiased variation,” *Algorithmica*, vol. 64, no. 4, 2012.
- [9] B. Doerr, T. Kötzing, J. Lengler, and C. Winzen, “Black-box complexities of combinatorial problems,” *Theor. Comput. Sci.*, vol. 471, 2013.
- [10] J. Rowe and M. Vose, “Unbiased black box search algorithms,” in *GECCO*. ACM, 2011.
- [11] B. Doerr and C. Winzen, “Playing Mastermind with constant-size memory,” *Theory of Computing Systems*, vol. 55, 2014.
- [12] C. Witt, “Tight bounds on the optimization time of a randomized search heuristic on linear functions,” *Combinatorics, Probability & Computing*, vol. 22, 2013.
- [13] D. Sudholt, “A new method for lower bounds on the running time of evolutionary algorithms,” *IEEE Trans. Evol. Comp.*, vol. 17, 2013.
- [14] C. Gießen and C. Witt, “The interplay of population size and mutation probability in the $(1 + \lambda)$ EA on OneMax,” *Algorithmica*, vol. 78, no. 2, 2017.
- [15] P. S. Oliveto, D. Sudholt, and C. Witt, “A tight lower bound on the expected runtime of standard steady state genetic algorithms,” in *GECCO*. ACM, 2020.
- [16] F. Chicano, A. M. Sutton, L. D. Whitley, and E. Alba, “Fitness probability distribution of bit-flip mutation,” *Evolutionary Computation*, vol. 23, no. 2, 2015.
- [17] C. Gießen and C. Witt, “Optimal mutation rates for the $(1 + \lambda)$ EA on OneMax through asymptotically tight drift analysis,” *Algorithmica*, vol. 80, 2018.
- [18] G. Karafotias, M. Hoogendoorn, and A. Eiben, “Parameter control in evolutionary algorithms: Trends and challenges,” *IEEE Trans. Evol. Comp.*, vol. 19, 2015.
- [19] A. Aleti and I. Moser, “A systematic literature review of adaptive parameter control methods for evolutionary algorithms,” *ACM Computing Surveys*, vol. 49, 2016.
- [20] A. E. Eiben, R. Hinterding, and Z. Michalewicz, “Parameter control in evolutionary algorithms,” *IEEE Trans. Evol. Comp.*, vol. 3, 1999.
- [21] B. Doerr and C. Doerr, “Theory of parameter control mechanisms for discrete black-box optimization: Provable performance gains through dynamic parameter choices,” in *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization*. Springer, 2020.
- [22] S. Böttcher, B. Doerr, and F. Neumann, “Optimal fixed and adaptive mutation rates for the LeadingOnes problem,” in *PPSN*. Springer, 2010.
- [23] B. Doerr, C. Doerr, and J. Yang, “Optimal parameter choices via precise black-box analysis,” *Theor. Comput. Sci.*, vol. 801, 2020.
- [24] A. Lissovoi, P. S. Oliveto, and J. A. Warwicker, “Simple hyper-heuristics control the neighbourhood size of randomised local search optimally for LeadingOnes,” *Evol. Comp.*, vol. 28, 2020.
- [25] T. Bäck, “The interaction of mutation rate, selection, and self-adaptation within a genetic algorithm,” in *PPSN*. Elsevier, 1992.
- [26] Á. Fialho, L. D. Costa, M. Schoenauer, and M. Sebag, “Extreme value based adaptive operator selection,” in *PPSN*. Springer, 2008.
- [27] N. Buskalic and C. Doerr, “Maximizing drift is not optimal for solving OneMax,” *Evolutionary Computation*, 2021, to appear. DOI: 10.1162/evco_a_00290.
- [28] M. Buzdalov and C. Doerr, “Optimal mutation rates for the $(1 + \lambda)$ EA on OneMax,” in *PPSN*. Springer, 2020.
- [29] C. Doerr, F. Ye, N. Hoesesh, H. Wang, O. M. Shir, and T. Bäck, “Benchmarking discrete optimization heuristics with IOHprofiler,” *Applied Soft Computing*, vol. 88, 2020.
- [30] S. Droste, T. Jansen, and I. Wegener, “On the analysis of the $(1+1)$ evolutionary algorithm,” *Theor. Comput. Sci.*, vol. 276, 2002.
- [31] A. Buzdalova, C. Doerr, and A. Rodionova, “Hybridizing the 1/5-th success rule with q-learning for controlling the mutation rate of an evolutionary algorithm,” in *PPSN*. Springer, 2020.
- [32] E. Carvalho Pinto and C. Doerr. (2018) Towards a more practice-aware runtime analysis of evolutionary algorithms. [Online]. Available: <https://arxiv.org/abs/1812.00493>
- [33] E. C. Pinto and C. Doerr, “A simple proof for the usefulness of crossover in black-box optimization,” in *PPSN*. Springer, 2018.
- [34] I. Rechenberg, *Evolutionssstrategie*. Stuttgart: Friedrich Fromman Verlag (Günther Holzboog KG), 1973.
- [35] B. Doerr, C. Gießen, C. Witt, and J. Yang, “The $(1 + \lambda)$ evolutionary algorithm with self-adjusting mutation rate,” *Algorithmica*, vol. 81, 2019.