The (1 +) Evolutionary Algorithm with Self-Adjusting Mutation Rate

Doerr, Benjamin; Gießen, Christian; Witt, Carsten; Yang, Jing

Published in:
Algorithmica

Link to article, DOI:
[10.1007/s00453-018-0502-x](https://doi.org/10.1007/s00453-018-0502-x)

Publication date:
2018

Document Version
Early version, also known as pre-print

[Link back to DTU Orbit](#)

Citation (APA):
Doerr, B., Gießen, C., Witt, C., & Yang, J. (2018). The (1 +) Evolutionary Algorithm with Self-Adjusting Mutation Rate. *Algorithmica*, 81(2), 593–631. <https://doi.org/10.1007/s00453-018-0502-x>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

The $(1+\lambda)$ Evolutionary Algorithm with Self-Adjusting Mutation Rate*

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

Christian Gießen
DTU Compute
Technical University of Denmark
Kgs. Lyngby, Denmark

Carsten Witt
DTU Compute
Technical University of Denmark
Kgs. Lyngby, Denmark

Jing Yang
Laboratoire d'Informatique (LIX)
École Polytechnique
Palaiseau, France

May 28, 2018

Abstract

We propose a new way to self-adjust the mutation rate in population-based evolutionary algorithms in discrete search spaces. Roughly speaking, it consists of creating half the offspring with a mutation rate that is twice the current mutation rate and the other half with half the current rate. The mutation rate is then updated to the rate used in that subpopulation which contains the best offspring.

We analyze how the $(1+\lambda)$ evolutionary algorithm with this self-adjusting mutation rate optimizes the OneMax test function. We prove that this dynamic version of the $(1+\lambda)$ EA finds the optimum in an expected optimization time (number of fitness evaluations) of $O(n\lambda/\log \lambda + n \log n)$. This time is asymptotically smaller than the optimization time of the classic $(1+\lambda)$ EA. Previous work shows that this performance is best-possible among all λ -parallel mutation-based unbiased black-box algorithms.

This result shows that the new way of adjusting the mutation rate can find optimal dynamic parameter values on the fly. Since our adjustment mechanism is simpler than the ones previously used for adjusting the mutation rate and does not have parameters itself, we are optimistic that it will find other applications.

1 Introduction

Evolutionary algorithms (EAs) have shown a remarkable performance in a broad range of applications. However, it has often been observed that this performance depends crucially on the use of the right parameter settings. Parameter optimization and parameter

*An extended abstract of this report appeared in the proceedings of the 2017 Genetic and Evolutionary Computation Conference (GECCO 2017) [22].

control are therefore key topics in EA research. Since these have very different characteristics in discrete and continuous search spaces, we discuss in this work only evolutionary algorithms for discrete search spaces.

Theoretical research has contributed to our understanding of these algorithms with mathematically founded runtime analyses, many of which show how the runtime of an EA is determined by its parameters. The majority of these works investigate *static parameter settings*, i. e., the parameters are fixed before the start of the algorithm and are not changed during its execution. More recently, a number of results were shown which prove an advantage of *dynamic parameter settings*, that is, the parameters of the algorithm are changed during its execution. Many of these rely on making the parameters functionally dependent on the current state of the search process, e.g., on the fitness of the current-best individual. While this provably can lead to better performances, it leaves the algorithm designer with an even greater parameter setting task, namely inventing a suitable functional dependence instead of fixing numerical values for the parameters. This problem has been solved by theoretical means for a small number of easy benchmark problems, but it is highly unclear how to find such functional relations in the general case.

A more designer-friendly way to work with dynamic parameters is to modify the parameters based on simple rules taking into account the recent performance. A number of recent results shows that such *on the fly* or *self-adjusting* parameter settings can give an equally good performance as the optimal fitness-dependent parameter setting, however, with much less input from the algorithm designer. For example, good results have been obtained by increasing or decreasing a parameter depending on whether the current iteration improved the best-so-far solution or not, e.g., in a way resembling the 1/5-th rule from continuous optimization.

Such success-based self-adjusting parameter settings can work well when there is a simple monotonic relation between success and parameter value, e.g., when one speculates that increasing the size of the population in an EA helps when no progress was made. For parameters like the mutation rate, it is not clear what a success-based rule can look like, since a low success rate can either stem from a too small mutation rate (regenerating the parent with high probability) or a destructive too high mutation rate. In [19], a relatively complicated learning mechanism was presented that tries to learn the right mutation strength by computing a time-discounted average of the past performance stemming from different parameter values. This learning mechanism needed a careful trade-off between exploiting the currently most profitably mutation strength and experimenting with other parameter values and a careful choice of the parameter controlling by how much older experience is taken less into account than more recent observations.

1.1 A New Self-Adjusting Mechanism for Population-Based EAs

In this work, we propose an alternative way to adjust the mutation rate on the fly for algorithms using larger offspring populations. It aims at overcoming some of the difficulties of the learning mechanism just described. The simple idea is to create half

the offspring with twice the current mutation rate and the other half using half the current rate. The mutation rate is then modified to the rate which was used to create the best of these offspring (choosing the winning offspring randomly among all best in case of ambiguity). We do not allow the mutation rate to leave the interval $[2/n, 1/4]$, so that the rates used in the subpopulations are always in the interval $[1/n, 1/2]$.

We add one modification to the very basic idea described in the first paragraph of this section. Instead of always modifying the mutation rate to the rate of the best offspring, we shall take this winner's rate only with probability a half and else modify the mutation rate to a random one of the two possible values (twice and half the current rate). Our motivation for this modification is that we feel that the additional random noise will not prevent the algorithm from adjusting the mutation rate into a direction that is more profitable. However, the increased amount of randomness may allow the algorithm to leave a possible basin of attraction of a locally optimal mutation rate. Observe that with probability $\Theta(1/n^2)$, a sequence of $\log_2 n$ random modifications all in the same direction appears. Hence with this inverse-polynomial rate, the algorithm can jump from any mutation rate to any other (with the restriction that only a discrete set of mutation rates can appear). We note that the existence of random modifications is also exploited in our runtime analysis, which will show that the new self-adjusting mechanism selects mutation rates good enough to lead to the asymptotically optimal runtime among all dynamic choices of the mutation rate for the $(1+\lambda)$ EA.

In this first work proposing this mechanism, we shall not spend much effort fine-tuning it, but rather show in a proof-of-concept manner that it can find very good mutation rates. In a real application, it is likely that better results are obtained by working with three subpopulations, namely an additional one using (that is, exploiting) the current mutation rate. Also, it seems natural that more modest adjustments of the mutation rate, that is, multiplying and dividing the rate by a number F that is smaller than the value $F = 2$ used by our mechanism, is profitable. We conduct some elementary experiments supporting this intuition in Section 9.

1.2 Runtime Analysis for the Self-Adjusting $(1+\lambda)$ EA on OneMax

To prove that the self-adjusting mechanism just presented can indeed find good dynamic mutation rates, we conduct a rigorous runtime analysis for our self-adjusting $(1+\lambda)$ EA on the classic test function

$$\text{ONEMAX} : \{0, 1\}^n \rightarrow \mathbb{R}; (x_1, \dots, x_n) \mapsto \sum_{i=1}^n x_i.$$

The runtime of the $(1+\lambda)$ EA with fixed mutation rates on ONEMAX is well understood [13, 27]. In particular, Gießen and Witt [27] show that the expected runtime (number of generations) is $(1 \pm o(1)) \left(\frac{1}{2} \cdot \frac{n \ln \ln \lambda}{\ln \lambda} + \frac{e^r}{r} \cdot \frac{n \ln n}{\lambda} \right)$ when a mutation rate of r/n , r a constant, is used. Thus for λ not too large, the mutation rate determines the leading constant of the runtime, and a rate of $1/n$ gives the asymptotically best runtime.

As a consequence of their work on parallel black-box complexities, Badkobeh, Lehre, and Sudholt [2] showed that the $(1+\lambda)$ EA with a suitable fitness-dependent mutation

rate finds the optimum of ONEMAX in an asymptotically better runtime of $O(\frac{n}{\log \lambda} + \frac{n \log n}{\lambda})$, where the improvement is by a factor of $\Theta(\log \log \lambda)$.¹ This runtime is best-possible among all λ -parallel unary unbiased black-box optimization algorithms. In particular, no other dynamic choice of the mutation rate in the $(1+\lambda)$ EA can achieve an asymptotically better runtime. The way how the mutation rate depends on the fitness in the above result, however, is not trivial. When the parent individual has fitness distance d , then mutation rate employed is $p = \max\{\frac{\ln \lambda}{n \ln(en/d)}, \frac{1}{n}\}$.

Our main technical result is that the $(1+\lambda)$ EA adjusting the mutation rate according to the mechanism described above has the same (optimal) asymptotic runtime. Consequently, the self-adjusting mechanism is able to find on the fly a mutation rate that is sufficiently close to the one proposed in [2] to achieve asymptotically the same expected runtime.

Theorem 1. *Let $\lambda \geq 45$ and $\lambda = n^{O(1)}$. Let T denote the number of generations of the $(1+\lambda)$ EA with self-adjusting mutation rate on ONEMAX. Then,*

$$E(T) = \Theta\left(\frac{n}{\log \lambda} + \frac{n \log n}{\lambda}\right).$$

This corresponds to an expected number of function evaluations of $\Theta(\frac{\lambda n}{\log \lambda} + n \log n)$.

To the best of our knowledge, this is the first time that a simple mutation-based EA achieves a super-constant speed-up via a self-adjusting choice of the mutation rate.

As an interesting side remark, our proofs reveal that a quite non-standard but fixed mutation rate of $r = \ln(\lambda)/2$ also achieves the $\Theta(\log \log \lambda)$ improvement as it implies the bound of $\Theta(n/\log \lambda)$ generations if λ is not too small. Hence, the constant choice $r = O(1)$ as studied in [27] does not yield the asymptotically optimal number of generations unless λ is so small that the $n \log n$ -term dominates.

Lemma 2. *Let $\lambda \geq 45$ and $\lambda = n^{O(1)}$, Let T denote the number of generations of the $(1+\lambda)$ EA with fixed mutation rate $r = \ln(\lambda)/2$. Then,*

$$E(T) = O\left(\frac{n}{\log \lambda} + \frac{n \log n}{\sqrt{\lambda}}\right).$$

This corresponds to an expected number of function evaluations of $O(\frac{\lambda n}{\log \lambda} + \sqrt{\lambda} n \log n)$.

The paper is structured as follows: In Section 2 we give an overview over previous analyses of the $(1+\lambda)$ EA and of self-adjusting parameter control mechanism in EAs from a theoretical perspective. In Section 3 we give the algorithm and the mutation scheme. For convenience, we also state some key theorems that we will frequently use in the rest of the paper. Afterwards, Section 4 gives a high-level overview of our main

¹As usual in the analysis of algorithms, we write $\log$ when there is no need to specify the base of the logarithm, e.g., in asymptotic terms like $\Theta(\log n)$. To avoid the possible risk of an ambiguity for small arguments, as usual, we set $\log x = \max\{1, \log x\}$. All other logarithms are used in their classic meaning, that is, for all $x > 0$, we denote by $\ln x$ the natural logarithm of x and by $\log_2 x$ its binary logarithm.

proof strategy. The following technical sections deal with the runtime analysis of the expected time spent by the $(1+\lambda)$ EA on ONEMAX in each of three regions of the fitness distance d . We label these regions the *far region*, *middle region* and *near region*, each of which will be dealt with in a separate section. The proof of the main theorem and of Lemma 2 is then given in Section 8. Finally, we conclude in Section 10.

2 Related Work

Since this is a theoretically oriented work on how a dynamic parameter choice speeds up the runtime of the $(1+\lambda)$ EA on the test function ONEMAX, let us briefly review what is known about the theory of this EA and dynamic parameter choices in general.

2.1 The $(1+\lambda)$ EA

The first to conduct a rigorous runtime analysis of the $(1+\lambda)$ EA were Jansen, De Jong, and Wegener [30]. They proved, among other results, that when optimizing ONEMAX a linear speed-up exists up to a population size of $\Theta(\log(n) \log \log(n) / \log \log \log(n))$, that is, for $\lambda = O(\log(n) \log \log(n) / \log \log \log(n))$, finding the optimal solution takes an expected number of $\Theta(n \log(n) / \lambda)$ generations, whereas for larger λ at least $\omega(n \log(n) / \lambda)$ generations are necessary. This picture was completed in [13] with a proof that the expected number of generations taken to find the optimum is $\Theta(\frac{n \log n}{\lambda} + \frac{n \log \log \lambda}{\log \lambda})$. The implicit constants were determined in [27], giving the bound of $(1 \pm o(1))(\frac{1}{2} \frac{n \ln \ln \lambda}{\ln \lambda} + \frac{e^r}{r} \frac{n \ln n}{\lambda})$, for any constant r , as mentioned in the introduction.

Aside from the optimization behavior on ONEMAX, not too much is known for the $(1+\lambda)$ EA, or is at least not made explicit (it is easy to see that waiting times for an improvement which are larger than λ reduce by a factor of $\Theta(\lambda)$ compared to one-individual offspring populations). Results made explicit are the $\Theta(n^2/\lambda + n)$ expected runtime (number of generations) on LEADINGONES [30], the worst-case $\Theta(n + n \log(n) / \lambda)$ expected runtime on linear functions [13], and the $O(m^2(\log n + \log w_{\max}) / \lambda)$ runtime estimate for minimum spanning trees valid for $\lambda \leq m^2/n$ [39], where n denotes the number of vertices of the input graph, m the number of edges, and $w_{\max}$ the maximum of the integral edge weights.

2.2 Dynamic Parameter Choices

While it is clear the EAs with parameters changing during the run of the algorithm (dynamic parameter settings) can be more powerful than those only using static parameter settings, only recently considerable advantages of dynamic choices could be demonstrated by mathematical means (for discrete optimization problems; in continuous optimization, step size adaptation is obviously necessary to approach arbitrarily closely a target point). To describe the different ways to dynamically control parameters, we use in the following the language proposed in Eiben, Hinterding, and Michalewicz [26] and its extension from [12].

2.2.1 Deterministic Parameter Control

In this language, *deterministic parameter control* means that the dynamic choice of a parameter does not depend on the fitness landscape. The first to rigorously analyze a deterministic parameter control scheme are Jansen and Wegener [29]. They regard the performance of the (1+1) EA which uses in iteration t the mutation rate k/n , where $k \in \{1, 2, \dots, 2^{\lceil \log_2 n \rceil - 2}\}$ is chosen such that $\log_2(k) \equiv t - 1 \pmod{\lceil \log_2 n \rceil - 1}$. In other words, they cyclically use the mutation rates $1/n, 2/n, 4/n, \dots, K/n$, where K is the largest power of two that is less than $n/2$. Jansen and Wegener demonstrate that there exists an example function where this dynamic EA significantly outperforms the (1+1) EA with any static mutation rate. However, they also observe that for many classic problems, this EA is slower by a factor of $\Theta(\log n)$.

In [40], a rank-based mutation rate was analyzed for the $(\mu+1)$ EA. A previous experimental study [6] suggested that this is a profitable approach, but the mathematical runtime analysis in [40] rather indicates the opposite. While there are artificial examples where a huge runtime gain could be shown and also the worst-case runtime of the $(\mu+1)$ EA reduces from essentially n^n to $O(3^n)$, a rigorous analysis on the ONEMAX function rather suggests that the high rate of offspring generated with a mutation rate much higher than $1/n$ brings a significant risk of slowing down the optimization process.

For two non-standard settings in evolutionary computation, deterministic parameter control mechanisms also gave interesting results. For problems where the solution length is not known [5], more precisely, where the number or the set of bits relevant for the solution quality is unknown, again random mutation rates gave good results [16, 21]. Here however, not a power-law scheme was used, but rather one based on very slowly decreasing summable sequences. For problems where the discrete variables take many values, e.g., the search space is $\{0, \dots, r-1\}^n$ for some large r , the question is how to change the value of an individual variable. The results in [17] suggest that a harmonic mutation strength, that is, changing a variable value by $\pm i$ with i chosen randomly with probability proportional to $1/i$, can be beneficial. This distribution was analyzed earlier in [8] for the one-dimensional case, where it was also shown to give the asymptotically best performance on a ONEMAX type problem.

For randomized search heuristics outside evolutionary computation, Wegener [43] showed that simulated annealing (using a time-dependent temperature) can beat the Metropolis algorithm (using a static temperature).

2.2.2 Adaptive Parameter Control

A parameter control scheme is called *adaptive* if it used some kind of feedback from the optimization process. This can be functionally dependent (e.g., the mutation rate depends on the fitness of the parent) or success-based (e.g., a 1/5th rule).

The first to conduct a runtime analysis for an adaptive parameter control mechanism (and show a small advantage over static choices) were Böttcher, Doerr, and Neumann [3]. They proposed to use the *fitness-dependent* mutation rate of $1/(\text{LEADINGONES}(x) + 1)$ for the optimization of the LEADINGONES test function. They proved that with this

choice, the runtime of the (1+1) EA improves to roughly $0.68n^2$ compared to a time of $0.86n^2$ stemming from the classic mutation rate $1/n$ or a runtime of $0.77n^2$ stemming from the asymptotically optimal static rate of approximately $1.59/n$.

For the $(1 + (\lambda, \lambda))$ GA, a fitness-dependent offspring population size of order $\lambda = \Theta(\sqrt{n/d(x)})$ was suggested in [15], where $d(x)$ is the fitness-distance of the parent individual to the optimum. This choice improves the optimization time (number of fitness evaluations until the optimum is found) on ONEMAX from $\Theta(n\sqrt{\log(n)\log\log\log(n)/\log\log(n)})$ stemming from the optimal static parameter choice [10] to $O(n)$. Since in this adaptive algorithm the mutation rate p is functionally dependent on the offspring population size, namely via $p = \lambda/n$, the dynamic choice of λ is equivalent to a fitness-dependent mutation rate of $1/\sqrt{nd(x)}$.

In the aforementioned work by Badkobeh et al. [2], a fitness-dependent mutation rate of $\max\{\frac{\ln \lambda}{n \ln(en/d(x))}, \frac{1}{n}\}$ was shown to improve the classic runtime of $O(\frac{n \log \log \lambda}{\log \lambda} + \frac{n \log n}{\lambda})$ to $O(\frac{n}{\log \lambda} + \frac{n \log n}{\lambda})$. In [19], the (1+1) EA using a k -bit flip mutation operator together with a fitness-dependent choice of k was shown to give a performance on ONEMAX that is very close to the theoretical optimum (among all unary unbiased black-box algorithms), however, this differs only by lower order terms from the performance of the simple randomized local search heuristic (RLS). For nature-inspired algorithms other than evolutionary ones, Zarges [44, 45] proved that fitness-dependent mutation rates can be beneficial in artificial immune systems.

2.2.3 Self-adjusting and Self-adaptive Parameter Control

While all these results show an advantage of adaptive parameter settings, it remains questionable if an algorithm user would be able to find such a functional dependence of the parameter on the fitness. This difficulty can be overcome via *self-adjusting* parameter choices, where the parameter is modified according to a simple rule often based on the success or progress of previous iterations, or via *self-adaptation*, where the parameter is encoded in the genome and thus subject to variation and selection. The understanding of self-adaptation is still very limited. The only theoretical work on this topic [7], however, is promising and shows examples where self-adaptation can lead to significant speed-ups for non-elitist evolutionary algorithms.²

In contrast to this, the last years have produced a profound understanding of self-adjusting parameter choices. The first to perform a mathematical analysis were Lässig and Sudholt [34], who considered the $(1+\lambda)$ EA and a simple parallel island model together with two self-adjusting mechanisms for population size or island number, including halving or doubling it depending on whether the current iteration led to an improvement or not. These mechanisms were proven to give significant improvements of the “parallel” runtime (number of generations) on various test functions without increasing significantly the “sequential” runtime (number of fitness evaluations).

²Note added in proof: Very recently, it was shown in [25] that the $(1, \lambda)$ EA with self-adaptive mutation rate has a runtime asymptotically equivalent to the one of the algorithm regarded in this work.

In [12] it was shown that the fitness-dependent choice of λ for the $(1 + (\lambda, \lambda))$ GA described above can also be found in a self-adjusting way. To this aim, another success-based mechanism was proposed, which imitates the 1/5-th rule from evolution strategies. With some modifications, this mechanism also works on random satisfiability problems [4]. For the problem of optimizing an r -valued ONEMAX function, a self-adjustment of the step size inspired by the 1/5-th rule was found to find the asymptotically best possible runtime in [18].

These results indicate that success-based dynamics work well for adjusting parameters when a monotonic relation like “if progress is difficult, then increase the population size” holds. For adjusting a parameter like the mutation rate, it is less obvious how to do this. For example, in the search space $\{0, 1\}^n$ both a too large mutation rate (creating a stronger drift towards a Hamming distance of $n/2$ from the optimum) and a too small mutation rate (giving a too small radius of exploration) can be detrimental. For this reason, to obtain a self-adjusting version of their result on the optimal number k to optimize ONEMAX via k -bit flips [19], in [20] a learning mechanism was proposed that from the medium-term past estimates the efficiency of different parameter values. As shown there, this does find the optimal mutation strength sufficiently well to obtain essentially the runtime stemming from the fitness-dependent mutation strength exhibited before.

In the light of these works, our result from the methodological perspective shows that some of the difficulties of the learning mechanism of [20], e.g., the whole book-keeping being part of it and also the setting of the parameters regulating how to discount information over time, can be overcome by the mechanism proposed in this work. In a sense, the use of larger populations enables us to adjust the mutation rate solely on information learned in the current iteration. However, we do also use the idea of [20] to intentionally use parameter settings which appear to be slightly off the current optimum to gain additional insight.

2.2.4 Selection Hyper-Heuristics

We note that so-called *selection hyper-heuristics* may lead to processes resembling dynamic parameter choices. Selection hyper-heuristics are methods that select, during the run of the algorithm, which one out of several pre-specified simpler algorithmic building blocks to use. When the different pre-specified choices are essentially identical apart from an internal parameter, then this selection hyper-heuristic approach could also be interpreted as a dynamic choice of the internal parameter. For example, when only the two mutation operators are available that flip exactly one or exactly two bits, then a selection hyper-heuristic choosing between them could also be interpreted as the *randomized local search* heuristic using a dynamic choice of the number of bits it flips. We briefly review the main runtime analysis works of this flavor.

The first to conduct a rigorous runtime analysis for selection hyper-heuristics were Lehre and Özcan [35]. They show that the $(1+1)$ EA using the *mixed strategy* of choosing the mutation operator randomly between the 1-bit flip operator (with probability p) and the 2-bit flip operator (with probability $1 - p$) optimizes the ONEMAX function in time $\min\{\frac{1}{p}n(1 + \ln n), \frac{1}{1-p}n^2(1 - \frac{1}{n})\}$. It appears to us that, most likely, this result is

not absolutely correct, since, e.g., in the case $p = 0$ the expected optimization time is clearly infinite (if the random initial search point has an odd Hamming distance from the optimum, then the optimum cannot be reached only via 2-bit flips). So most likely, the upper bound should be increased by an additive term of $\frac{1}{p}n$.

Lehre and Özcan [35] further construct an example function `GAPATH`, which has the property that it can only be optimized when both the 1-bit and the 2-bit flip mutation operator occur with positive probability and use this to discuss various static and dynamic ways to randomly decide between the two operators. A similar example for a multi-objective optimization setting was given by Qian, Tang, and Zhou [41].

Note that for $p = \frac{1}{2}$, the randomized selection heuristic is the classic *randomized search heuristic* using 1-bit and 2-bit flips, which was regarded, among others, in [28, 39].

In [1], Alanazi and Lehre extend the previous work to several classic selection hyper-heuristics like *simple random* (take a random low-level heuristic in each iteration), *random gradient* (take a random low-level heuristic and repeat using it as long as a true fitness improvement is obtained), *greedy* (in each iteration, use all low-level heuristic in parallel and continue with best result obtained), *permutation* (generate initially a random cyclic order of the low-level heuristics and then use them in this order). Again for the choice between 1-bit and 2-bit flips, they prove upper and lower bounds for the expected optimization time on the `LEADINGONES` benchmark function. While the results are relatively tight (the corresponding upper and lower bounds deviate by at most a factor of 6), the intervals of possible asymptotic runtimes intersect, so this first runtime analysis of this type does not yet give a conclusive picture on how the different selection heuristics compare.

Given that the probabilities to find a true improvement are very low in this discrete optimization problem, one would expect that all these four heuristics have the same runtime (apart from lower order terms), and this is indeed the first set of results by Lissovoi, Oliveto, and Warwicker [37], who show that the expected runtime in all cases is $(1 + o(1))\frac{1}{2}\ln(3)n^2$. They build on this strong result by proposing to use a slower adaptation. For the random gradient method, they propose to use a random low-level heuristic for up to τ iterations. If an improvement is found, immediately another phase with this operator is started. If a phase of τ iterations does not see a fitness improvement, then a new phase is started with a random operator.

For this *generalized random gradient* heuristic using a phase length of $\tau = cn$ for a constant c , they show (still for the `LEADINGONES` problem and the 1-bit and 2-bit mutation operator) an expected runtime of $(1 + o(1))g(c)n^2$, where $g(c)$ is a constant depending on c only that tends to $\frac{\ln(2)+1}{4}$ when c is tending to infinity. Consequently, this new hyper-heuristic outperforms the previously investigated ones and approaches arbitrarily well the best-possible performance that can be obtained from the two mutation operators, which is, as also shown in [37], $(1 + o(1))\frac{\ln(2)+1}{4}n^2$.

The *generalized random gradient* heuristic was further extended in [24]. There an operator was defined as successful (which leads to another phase using this operator) if it leads to σ improvements in a phase of at most τ iterations. Hence in this language, the previous *generalized random gradient* heuristic uses $\sigma = 1$. By using a larger value of σ ,

the algorithm is able to take more robust decisions on what is a success. This was used in [24] to determine the phase length τ in a self-adjusting manner. While the previous work [37] does not state this explicitly, the choice of τ is crucial. A τ -value of smaller asymptotic order than $\Theta(n)$ leads to typically no improvement within a phase and thus reverts the algorithms to the simple random heuristic. A τ -value of more than $cn \ln(n)$, where c is a suitable constant, results in that both operators are successful in most parts of the search space. Consequently, the algorithm sticks to the first choice for a large majority of the optimization process and thus does not profit from the availability of both operators.

Since the choice of τ is that critical, a mechanism successfully adjusting it to the right value is desirable. In [24] it is shown that by choosing $\sigma \in \Omega(\log^4 n) \cap o(\sqrt{n/\log n})$ —note that this is a quite wide range—the value of τ can be easily adjusted on the fly via a multiplicative update rule resembling vaguely the 1/5-th success rule. This gives again the asymptotically optimal runtime of $(1 + o(1)) \frac{\ln(2)+1}{4} n^2$.

3 Preliminaries

We shall now formally define the algorithm analyzed and present some fundamental tools for the analysis.

3.1 Algorithm

We consider the $(1+\lambda)$ EA with self-adjusting mutation rate for the minimization of pseudo-boolean functions $f : \{0, 1\}^n \rightarrow \mathbb{R}$, defined as Algorithm 1.

The general idea of the mutation scheme is to adjust the mutation strength according to its success in the population. We perform mutation by applying standard bit mutation with two different mutation probabilities $r/(2n)$ and $2r/n$ and we call r the *mutation rate*. More precisely, for an even number $\lambda \geq 2$ the algorithm creates $\lambda/2$ offspring with mutation rate $r/2$ and with $2r$ each.

The mutation rate is adjusted after each selection. With probability a half, the new rate is taken as the mutation rate that the best individual (i.e. the one with the lowest fitness, ties broken uniformly at random) was created with (*success-based adjustment*). With the other 50% probability, the mutation rate is adjusted to a random value in $\{r/2, 2r\}$ (*random adjustment*). Note that the mutation rate is adjusted in each iteration, that is, also when all offspring are worse than the parent and thus the parent is kept for the next iteration.

If an adjustment of the rate results in a new rate r outside the interval $[2, n/4]$, we replace this rate with the corresponding boundary value. Note that in the case of $r < 2$, a subpopulation with rate less than 1 would be generated, which means flipping less than one bit in expectation. At a rate $r > n/4$, a subpopulation with rate larger than $n/2$ would be created, which again is not a very useful choice.

We formulate the algorithm to start with an initial mutation rate r^{init} . The only assumption on r^{init} is to be greater than or equal to 2. The $(1+\lambda)$ EA with this self-adjusting choice of the mutation rate is given as pseudocode in Algorithm 1.

Algorithm 1 $(1+\lambda)$ EA with two-rate standard bit mutation

```

Select  $x$  uniformly at random from  $\{0,1\}^n$  and set  $r \leftarrow r^{\text{init}}$ .
for  $t \leftarrow 1, 2, \dots$  do
    for  $i \leftarrow 1, \dots, \lambda$  do
        Create  $x_i$  by flipping each bit in a copy of  $x$  independently with probability
         $r_t/(2n)$  if  $i \leq \lambda/2$  and with probability  $2r_t/n$  otherwise.
         $x^* \leftarrow \arg \min_{x_i} f(x_i)$  (breaking ties randomly).
        if  $f(x^*) \leq f(x)$  then
             $x \leftarrow x^*$ .
    Perform one of the following two actions with prob.  $1/2$ :
    • Replace  $r_t$  with the mutation rate that  $x^*$  has been created with.
    • Replace  $r_t$  with either  $r_t/2$  or  $2r_t$ , each with probability  $1/2$ .
    Replace  $r_t$  with  $\min\{\max\{2, r_t\}, n/4\}$ .

```

Let us explain the motivation for the random adjustments of the rate. Without such random adjustments, the rate can only be changed into some direction if a winning offspring is generated with this rate. For simple functions like ONEMAX, this is most likely sufficient. However, when the fitness of the best of $\lambda/2$ offspring, viewed as a function of the rate, is not unimodal, then several adjustments into a direction at first not yielding good offspring might be needed to reach good values of the rate. Here, our random adjustments enable the algorithm to cross such a valley of unfavorable rate values. We note that such ideas are not uncommon in evolutionary computation, with standard-bit mutation being the most prominent example (allowing to perform several local-search steps in one iteration to cross fitness valleys).

A different way to implement a mechanism allowing larger changes of the rate to cross unfavorable regions would have been to not only generate offspring with rates $r/2$ and $2r$, but to allow larger deviations from the current rate with some small probability. One idea could be choosing for each offspring independently the rate $r2^{-i}$ with probability $2^{-|i|-1}$ for all $i \in \mathbb{Z}$, $i \neq 0$. This should give similar results, but to us the process appears more chaotic (e.g., with not the same number of individuals produced with rates $r/2$ and $2r$).

The *runtime*, also called the *optimization time*, of the $(1+\lambda)$ EA is the smallest t such that an individual of minimum f -value has been found. Note that t corresponds to a number of iterations (also called generations), where each generation creates λ offspring. Since each of these offspring has to be evaluated, the number of function evaluations, which is a classical cost measure, is by a factor of λ larger than the runtime as defined here. However, assuming a massively parallel architecture that allows for parallel evaluation of the offspring, counting the number of generations seems also a valid cost

measure. In particular, a speed-up on the function $\text{ONEMAX}(x_1, \dots, x_n) := x_1 + \dots + x_n$ by increasing λ can only be observed in terms of the number of generations. Note that for reasons of symmetry, it makes no difference whether ONEMAX is minimized (as in the present paper) or maximized (as in several previous research papers).

Throughout the paper, all asymptotic notation will be with respect to the problem size n .

3.2 Drift Theorems

Our results are obtained by drift analysis, which is also used in previous analyses of the $(1+\lambda)$ EA without self-adaptation on ONEMAX and other linear functions [13, 27].

The first theorems stating upper bounds on the hitting time using variable drift go back to [31, 38]. We take a formulation from [36] but simplify it to Markov processes for notational convenience.

Theorem 3 (Variable Drift, Upper Bound). *Let $(X_t)_{t \geq 0}$, be random variables describing a Markov process over a finite state space $S \subseteq \{0\} \cup [x_{\min}, x_{\max}]$, where $x_{\min} > 0$. Let T be the random variable that denotes the earliest point in time $t \geq 0$ such that $X_t = 0$. If there exists a monotone increasing function $h(x) : [x_{\min}, x_{\max}] \rightarrow \mathbb{R}^+$, where $1/h(x)$ is integrable on $[x_{\min}, x_{\max}]$, such that for all $x \in S$ with $\Pr(X_t = x) > 0$ we have*

$$E(X_t - X_{t+1} \mid X_t = x) \geq h(x)$$

then for all $x' \in S$ with $\Pr(X_0 = x') > 0$

$$E(T \mid X_0 = x') \leq \frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{x'} \frac{1}{h(x)} dx.$$

The variable drift theorem is often applied in the special case of *additive drift* in discrete spaces: assuming $E(X_t - X_{t+1} \mid X_t = x; X_t > 0) \geq \varepsilon$ for some constant ε , one obtains $E(T \mid X_0 = x') \leq x'/\varepsilon$.

Since we will make frequent use of it in the following sections as well, we will also give the version of the *Multiplicative Drift Theorem* for upper bounds, due to [14]. Again, this is implied by the previous variable drift theorem.

Theorem 4 (Multiplicative Drift [14]). *Let $(X_t)_{t \geq 0}$ be random variables describing a Markov process over a finite state space $S \subseteq \mathbb{R}_0^+$ and let $x_{\min} := \min\{x \in S \mid x > 0\}$. Let T be the random variable that denotes the earliest point in time $t \geq 0$ such that $X_t = 0$. If there exist $\delta > 0$ such that for all $x \in S$ with $\Pr(X_t = x) > 0$ we have*

$$E(X_t - X_{t+1} \mid X_t = x) \geq \delta x ,$$

then for all $x' \in S$ with $\Pr(X_0 = x') > 0$,

$$E(T \mid X_0 = x') \leq \frac{1 + \ln\left(\frac{x'}{x_{\min}}\right)}{\delta} .$$

3.3 Chernoff Bounds

For reasons of self-containedness, we state two well-known multiplicative Chernoff bounds and a lesser known additive Chernoff bound that is also known in the literature as Bernstein's inequality. These bounds can be found, e.g., in [9].

Theorem 5 (Bernstein's inequality, Chernoff bounds). *Let $X_1, \dots, X_n$ be independent random variables and $X = \sum_{i=1}^n X_i$.*

- (a) *Let b be such that $E(X_i) - b \leq X_i \leq E(X_i) + b$ for all $i = 1, \dots, n$. Let $\sigma^2 = \sum_{i=1}^n \text{Var}(X_i) = \text{Var}[X]$. Then for all $\Delta \geq 0$,*

$$\Pr(X \geq E(X) + \Delta) \leq \exp\left(-\frac{\Delta^2}{2(\sigma^2 + \frac{1}{3}b\Delta)}\right).$$

- (b) *Assume that for all $i = 1, \dots, n$, the random variable X_i takes values in $[0, 1]$ only. Then*

- $\Pr(X \leq (1 - \delta)E(X)) \leq \exp(-\delta^2 E(X)/2)$ for all $0 \leq \delta \leq 1$,
- $\Pr(X \geq (1 + \delta)E(X)) \leq \exp(-\delta^2 E(X)/(2 + \delta))$ for all $\delta > 0$.

3.4 Occupation Probabilities

As mentioned above, we will be analyzing two depending stochastic processes: the random decrease of fitness and the random change of the mutation rate. Often, we will prove by drift analysis that the rate is drifting towards values that yield an almost-optimal fitness decrease. However, once the rate has drifted towards such values, we would also like the rates to stay in the vicinity of these values in subsequent steps. To this end, we apply the following theorem from [33]. Note that in the paper a slightly more general version including a self-loop probability is stated, which we do not need here.

Theorem 6 (Theorem 7 in [33]). *Let a Markov process $(X_t)_{t \geq 0}$ on $\mathbb{R}_0^+$, where $|X_t - X_{t+1}| \leq c$, with additive drift of at least d towards 0 be given (i.e., $E(X_t - X_{t+1} \mid X_t; X_t > 0) \geq d$), starting at 0 (i.e. $X_0 = 0$). Then we have, for all $t \in \mathbb{N}$ and $b \in \mathbb{R}_0^+$,*

$$\Pr(X_t \geq b) \leq 2e^{\frac{2d}{3c}(1-b/c)}.$$

We can readily apply this theorem in the following lemma that will be used throughout the paper to bound the rate r_t .

Lemma 7. *If there is a point $a \geq 4$ such that $\Pr(r_{t+1} < r_t \mid r_t > a) \geq 1/2 + \varepsilon$ for some constant $\varepsilon > 0$, then for all $t' \geq \min\{t \mid r_t \leq a\}$ and all $b \geq 2$ it holds $\Pr(r_{t'} \geq a \cdot 2^{b+1}) \leq 2e^{-2b\varepsilon/3}$.*

Proof. Apply Theorem 6 to the process $X_t := \max\{0, \lfloor \log_2(r_t/a) \rfloor\}$. Note that this process is on $\mathbb{N}_0$, moves by an absolute value of at most 1 and has drift $E(X_t - X_{t+1} \mid X_t; X_t > 0) \geq 2\varepsilon$. We use $c := 1$ and $d := 2\varepsilon$ in the theorem and estimate $1 - b \leq -b/2$. $\square$

3.5 Useful Inequalities

The following well-known estimates will be used regularly in this work.

Lemma 8. (a) For all $x > 0$, we have $\ln x \leq x/e$.

(b) For all $x \in \mathbb{R}$, we have $1 + x \leq e^x$.

(c) For all $0 \leq x \leq 2/3$, we have $1 - x \geq e^{-x-x^2} \geq e^{-2x}$.

Proof. Since $(\ln x - x/e)' = 1/x - 1/e, x > 0$, then $(e, 0)$ is the maximum point of $\ln x - x/e, x > 0$. Thus part (a) holds.

We notice that $(1 + x - e^x)' = 1 - e^x$, then $(0, 0)$ is the maximal point of $1 + x - e^x$. Therefore part (b) holds.

For the proof of part (c), the second inequality can be easily obtained, since e^x monotonically increases and $-x - x^2 \geq -2x$ when $0 \leq x \leq 2/3$. To prove the first inequality, let $f(x) = 1 - x - e^{-x-x^2}$ for $0 \leq x \leq 2/3$. Then $f'(x) = (1 + 2x)e^{-x-x^2} - 1$ and $f''(x) = (1 - 4x - 4x^2)e^{-x-x^2}$. We notice that $f''(x)$ has only one zero point between 0 and $2/3$. Let x_0 denote the zero point. Thus $f''(x)$ is positive before x_0 , and then becomes negative after x_0 . Using the fact $f'(0) = 0$ and $f'(2/3) < 0$, we can easily see that $f'(x)$ first increases from 0 to $f'(x_0) > 0$, and then decreases below 0. This means the minimum of $f(x)$ is attained at 0 or $2/3$. Due to the fact that $f(0) = 0$ and $f(2/3) > 0$, the first inequality in part (c) holds. $\square$

4 Proof Overview

Since the following runtime analysis of our self-adjusting $(1+\lambda)$ EA on the ONEMAX function is slightly technical, let us outline the main proof ideas here in an informal manner. Let always x denote the parent individual of the current iterations and $k := f(x)$ its fitness distance from the optimum (recall that we are minimizing the ONEMAX function, hence the fitness distance equals the objective function value which in turn is the Hamming distance from the optimum $x^* = (0, \dots, 0) \in \{0, 1\}^n$ and thus the number of ones in x).

The outer proof argument is variable drift, that is, for each fitness value k we estimate the expected fitness gain (“progress”) in an iteration starting with a parent individual x with $f(x) = k$ and then we use the variable drift theorem (Theorem 3) to translate this information on the progress into an expected runtime (number of generations until the optimum is found).

To obtain sufficiently strong lower bounds on the expected progress, we need to argue that the rate self-adjustment sufficiently often sets the current rate to a sufficiently good value. This is the main technical difficulty as it needs a very precise analysis of the quality of the offspring in both subpopulations. Since this requires different arguments depending on the current fitness distance k , we partition the process into three regimes.

In the *far region* covering the fitness distance values $k \geq n/\ln(\lambda)$, we need a rate r that is at least almost logarithmic in λ to ensure that the average progress is high enough.

Note that to gain the $\Theta(n)$ fitness levels from the initial value of approximately $n/2$ to $n/\ln \lambda$ in at most a time of $O(n/\log \lambda)$, we need an average progress of $\Omega(\ln \lambda)$ per iteration. We shall not be able to obtain this progress throughout this region, but we will obtain an expected progress of

$$\Omega\left(\frac{\log \lambda}{\log \frac{en}{k}}\right)$$

in an iteration with initial fitness distance k . This will be sufficient to reach a fitness distance of $n/\log \lambda$ in the desired $O(n/\log \lambda)$ iterations.

As said, the main difficulty is arguing that the self-adjusting mechanism keeps the rate sufficiently often in the range we need, which is roughly

$$\left[\frac{1}{2 \ln \frac{en}{k}} \ln \lambda, \frac{4n^2}{(n-2k)^2} \ln \lambda \right]$$

for all $k \in [n/\ln \lambda, n/2 - 1]$. Note that these range boundaries, in particular the upper one, depend strongly on k . Hence for arguing that our rate is sufficiently often in this range, we need to consider both the rate changes from the self-adjustments and the changing k -value. In this region we profit from the fact that we work with relatively high rates, which lead to strong concentration behaviors. This allows to argue that, for example, exceeding the upper boundary already by small constant factors is highly unlikely.

The *middle region* covering the fitness distances $k \in [n/\lambda, n/\ln \lambda]$ is small enough so that we do not require the algorithm to find a near-optimal rate very often. In fact, it suffices that the rate is below $\frac{1}{2} \ln \lambda$ with constant probability. This ensures an expected progress of at least $\min\{\frac{1}{8}, \sqrt{\lambda k}/32n\}$, which is sufficient to traverse also this region in time $O(n/\log \lambda)$. Consequently, in this region we only need to argue that the rate does not become too large, whereas there is no lower bound on r which we require. Further, the upper bound of $\frac{1}{2} \ln \lambda$ is large enough so that strong concentration arguments can be exploited, which show that deviations above the upper bound are highly unlikely. Since the upper bound does not change over time, we can now conveniently use an occupation probability argument to show that the rate with constant probability is only a small constant factor over the desired range, which is enough since the random fluctuations of the rate with constant probability reduce the rate further to the desired range.

In the *near region* covering the remaining fitness distance values $k \leq n/\lambda$, the parent individual is already so close to the optimum that any rate higher than the minimal rate $r = 2$ is sub-optimal. Hence in this region, we know precisely the optimum value for the rate. Nevertheless, it is not very easy to show that the subpopulation using the smaller rate $r/2$ has a higher chance of containing the new parent individual. Due to the small rates predominant in this region and the fact that progress generally is difficult (due to the proximity to the optimum), often both subpopulations will contain best offspring (which are in fact copies of the parent). Hence the typical reason for the winning offspring stemming from the smaller-rate subpopulation is not anymore that

the other subpopulation contains only worse individuals, but only that the smaller-rate subpopulation contains more best offspring.

By quantifying this effect we establish a drift of the rate down to its minimum value 2 in the near region, and another occupation probability argument is used to show that the rate with sufficiently high probability will take this value in subsequent sets. Nevertheless, the concentration of the rate around this target value is weaker than in the other regions, and additional care has to be taken to handle iterations in the near region in which the unlikely, but still possible event occurs that the rate exceeds $\ln \lambda$. This value is a critical threshold in the near region since rates larger than $\ln \lambda$ will typically make all offspring worse than the parent. As an additional obstacle, there is a small interval of rates above the threshold in which we cannot show a drift of the rate to its minimum. To show that this small interval nevertheless is left towards its lower end with probability $\Omega(1)$, the occupation probability argument is combined with a potential function whose shape exploits the random adjustments that the $(1+\lambda)$ EA performs with 50% probability.

5 Far Region

In this first of three technical sections, we analyze the optimization behavior of our self-adjusting $(1+\lambda)$ EA in the regime where the fitness distance k is at least $n/\ln \lambda$. Since we are relatively far from the optimum, it is relatively easy to make progress. On the other hand, this regime spans the largest number of fitness levels (namely $\Theta(n)$), so we need to exhibit a sufficient progress in each iteration. Also, this is the regime where the optimal mutation rate varies most drastically. Since it is not important for the following analysis, we remark without proof that the optimal rate is n for $k \geq n/2 + \omega(\sqrt{n \log \lambda})$, $(1 + o(1))n/2$ for $k = n/2 \pm o(\sqrt{n \log \lambda})$, and then quickly drops to $r = \Theta(\log \lambda)$ for $k \leq n/2 - \varepsilon n$. Despite these difficulties, our $(1+\lambda)$ EA manages to find sufficiently good mutation rates to be able to reach a fitness distance of $k = n/\ln \lambda$ in an expected number of $O(n/\log \lambda)$ iterations.

Lemma 9. *Let n be sufficiently large and $0 < k < n/2$. We define $c_1(k) = (2 \ln(en/k))^{-1}$ and $c_2(k) = 4n^2/(n - 2k)^2$.*

- (a) *If $n/\ln \lambda \leq k$ and $r \leq c_1(k) \ln \lambda$, then the probability that a best offspring has been created with rate $2r$ is at least 0.64.*
- (b) *Let $\lambda \geq 100$. If $c_2(k) \ln \lambda \leq r \leq n/4$, then the probability that all best offspring have been created with rate $r/2$ is at least 0.51.*
- (c) *If $r \geq 2(1 + \gamma)c_2(k) \ln \lambda$, then the probability that all best offspring are worse than the parent is at least $1 - \lambda^{-\gamma}$.*

Proof. To prove part (a), let $q(k, i, r)$ and $Q(k, i, r)$ be the probabilities that standard bit mutation with mutation rate $p = r/n$ creates from a parent with fitness distance k

an offspring with fitness distance exactly $k - i$ and at most $k - i$. Then

$$q(k, i, r) = \sum_{j=0}^{k-i} \binom{k}{i+j} \binom{n-k}{j} \left(\frac{r}{n}\right)^{i+2j} \left(1 - \frac{r}{n}\right)^{n-i-2j}$$

and $Q(k, i, r) = \sum_{j=i}^k q(k, j, r)$. We first show that the probability of not achieving $i \geq 2r$ is less than 0.2. This is because for large enough n and $r \geq 2$ we have $(1 - o(1))\binom{k}{2r} \geq (1 - o(1))4!(k/(2r))^{2r} > 4(k/(2r))^{2r}$ and for $r \leq c_1(k) \ln \lambda = o(\sqrt{n})$ by Lemma 8, part (c) we have $(1 - 2r/n)^n \geq e^{-2r-4r^2/n} = (1 - o(1))e^{-2r}$, thus

$$\begin{aligned} Q(k, 2r, 2r) &\geq q(k, 2r, 2r) \geq \binom{k}{2r} \left(\frac{2r}{n}\right)^{2r} \left(1 - \frac{2r}{n}\right)^n \\ &\geq 4 \left(\frac{k}{2r}\right)^{2r} \left(\frac{2r}{n}\right)^{2r} e^{-2r} \geq 4 \left(\frac{k}{en}\right)^{2r} \geq 4 \left(\frac{k}{en}\right)^{2c_1(k) \ln \lambda} = \frac{4}{\lambda}. \end{aligned}$$

Considering $\lambda/2$ offspring using rate $2r$, the probability that none of them achieves a fitness improvement of at least $2r$ is less than $(1 - 4/\lambda)^{\lambda/2} < \exp(-4/2) < 0.2$. We next argue that an offspring having a progress of $2r$ or more with good probability comes from the $2r$ -subpopulation. We notice that

$$\frac{q(k, i, 2r)}{q(k, i, r/2)} \geq 4^i \left(\frac{1 - 2r/n}{1 - r/(2n)}\right)^n \geq (1 - o(1))4^i e^{-1.5r}.$$

Thus for $r \geq 2$ and $i \geq 2r$ we have

$$\frac{q(k, i, 2r)}{q(k, i, r/2)} \geq (1 - o(1)) \left(\frac{4}{e}\right)^{2r} > 4.$$

Therefore if the best progress among all λ offspring is i and $i \geq 2r$, the conditional probability that an offspring having fitness distance $k - i$ is generated with rate $2r$ is at least $q(k, i, 2r)/(q(k, i, 2r) + q(k, i, r/2)) \geq 4/5$. In total, the probability that a best offspring has been created with rate $2r$ is at least

$$(1 - 0.2) \cdot 4/6 = 0.64.$$

For the proof of part (b), let $\lambda \geq 100$ and $c_2(k) \ln \lambda \leq r \leq n/4$. Our idea is to show a probability of $o(1/\lambda)$ for the event that an offspring with rate $2r$ is not worse than the expected fitness of an offspring with rate $r/2$. Let $X(k, r)$ denote the random decrease of the fitness distance when applying standard bit mutation with probability $p = r/n$ to an individual with k ones. Then

$$\begin{aligned} E(X(k, r)) &= kp - (n - k)p = \frac{(2k - n)r}{n}, \\ \text{Var}(X(k, r)) &= np(1 - p) = r \left(1 - \frac{r}{n}\right). \end{aligned}$$

According to Bernstein's inequality (Theorem 5 (a)), for any $\Delta > 0$ we have

$$\Pr(X \geq E(X) + \Delta) \leq \exp\left(\frac{-\Delta^2}{2\text{Var}(X) + 2\Delta/3}\right) \leq \exp\left(\frac{-\Delta^2}{2\text{Var}(X) + 2\Delta}\right)$$

We apply this bound with $X = X(k, 2r)$ and $\Delta = E(X(k, r/2)) - E(X(k, 2r)) = (n - 2k)1.5r/n > 0$. Then,

$$\begin{aligned} \Pr(X(k, 2r) \geq E(X(k, r/2))) &\leq \exp\left(\frac{-\Delta^2}{2\text{Var}(X(k, 2r)) + 2\Delta}\right) \\ &= \exp\left(\frac{-9(n-2k)^2 r}{4n(7n-8r-6k)}\right) \leq \exp\left(-\frac{9(n-2k)^2 c_2(k) \ln \lambda}{28n^2}\right) = \frac{1}{\lambda^{9/7}}. \end{aligned}$$

We notice that we have $7n - 8r - 6k > 7n - 4n - 3n = 0$ in the second inequality. From a union bound we see that with probability less than $\lambda^{-9/7}(\lambda/2) < 100^{-2/7}/2 < 0.14$, the best offspring created with rate $2r$ is at least as good as the expected fitness of an individual created with rate $r/2$.

We estimate the probability that the best offspring using rate $r/2$ has a fitness distance at most $E(X(k, r/2))$. Let y be an offspring obtained from x via standard bit mutation with mutation rate $r/(2n)$. Let X^+ and X^- be the number of one-bits flipped and zero-bits flipped, respectively.

$$\begin{aligned} X^+ &:= |\{i \in \{1, \dots, n\} \mid x_i = 1 \wedge y_i = 0\}|, \\ X^- &:= |\{i \in \{1, \dots, n\} \mid x_i = 0 \wedge y_i = 1\}|. \end{aligned}$$

Both X^+ and X^- are binomially distributed and $X(k, r/2) = X^+ - X^-$. We aim at a lower bound for $\Pr(X(k, r/2) \geq E(X(k, r/2)))$. We have $\Pr(X^+ \geq E(X^+) - 1) \geq 1/2$, since the median of X^+ is between $\lfloor E(X^+) \rfloor$ and $\lceil E(X^+) \rceil$ by [32]. It remains to bound $\Pr(X^- \leq E(X^-) - 1)$. We notice that $E(X^-) = (n - k)r/(2n) > \ln \lambda > 1$ and $E(X^-) \leq (n - k)/4$. Applying Theorem 12 in [11] to binomial random variable $\tilde{X}^- := (n - k) - X^-$, we obtain $\Pr(X^- \leq E(X^-) - 1) = \Pr(\tilde{X}^- \geq E(\tilde{X}^-) + 1) \geq 0.037$. Therefore

$$\begin{aligned} \Pr(X(k, r/2) \geq E(X(k, r/2))) &\geq \Pr(X^- \leq E(X^-) - 1) \Pr(X^+ \geq E(X^+) - 1) \\ &\geq 0.037 \cdot 0.5 = 0.0185. \end{aligned}$$

For $\lambda/2$ offspring using rate $r/2$, the probability that the best one has a fitness distance at most $E(X(k, r/2))$ is more than $1 - (1 - 0.0185)^{\lambda/2} \geq 1 - 0.9815^{50} \geq 0.6$. Therefore with probability at least $0.6 \cdot (1 - 0.14) > 0.51$, all best offspring are from the $r/2$ -subpopulation. This proves the second statement of the lemma.

For proof of part (c), let $r \geq 2(1 + \gamma)c_2(k) \ln \lambda$. An offspring created with mutation rate r/n is at least as good as its parent if and only if $X(k, r) \geq 0$. By using Bernstein's

inequality (Theorem 5 (a)) with $X = X(k, r)$ and $\Delta = -E(X(k, r))$, we have

$$\begin{aligned}\Pr(X(k, r) \geq 0) &\leq \exp\left(-\frac{E(X(k, r))^2}{2\text{Var}(X(k, r)) + 2E(X(k, r))}\right) \\ &= \exp\left(-\frac{(n-2k)^2 r}{2n(2n-2k-r)}\right) \leq \exp\left(-\frac{(n-2k)^2 r}{4n^2}\right).\end{aligned}$$

Since $r \geq 2(1+\gamma)c_2(k)\ln\lambda$, the corresponding probabilities for rate $r/2$ and $2r$ are at most $1/\lambda^{1+\gamma}$ and $1/\lambda^{4+4\gamma}$, respectively. By a union bound, with probability at most $1/\lambda^\gamma$, the best offspring is at least as good as its parent. This proves part (c). $\square$

The lemma above shows that the rate r is subject to a constant drift towards the interval $[c_1(k)\ln n, c_2(k)\ln n]$. Unfortunately, we cannot show that we obtain a sufficient fitness progress for all r -values in this range. However, we can do so for a range smaller only by constant factors. This is what we do now (for large values of k) and in Lemma 11 (for smaller values of k). This case distinction is motivated by the fact that $c_2(k)$ becomes very large when k approaches $n/2$. Having a good fitness drift only for such a smaller range of r -values is not a problem since the random movements of r let us enter the smaller range with constant probability. This is what we will exploit in Theorem 12 and its proof.

Let $2 \leq r \leq n/4$ be the current rate and let $\tilde{r} \in \{r/2, 2r\}$. Let $\tilde{\Delta}(\lambda/2, k, \tilde{r})$ denote the fitness gain of the best of $\lambda/2$ offspring generated with rate $\tilde{r}$ from a parent x with fitness distance $k := \text{ONEMAX}(x)$ and the parent itself. More precisely, let $x^{(i)}, i \in \{1, \dots, \lambda/2\}$, be independent offspring generated from x by flipping each bit independently with probability $\tilde{r}/n$. Then the random variable $\tilde{\Delta}(\lambda/2, k, \tilde{r})$ is defined by $\max\{0, k - \min\{\text{ONEMAX}(x^{(i)}) \mid i \in \{1, \dots, \lambda/2\}\}\}$. We use $\Delta := \max\{\tilde{\Delta}(\lambda/2, k, r/2), \tilde{\Delta}(\lambda/2, k, 2r)\}$ to denote the fitness gain in a iteration which uses x as parent and r as mutation rate.

We next show that a region contained in $[c_1(k)\ln\lambda, c_2(k)\ln\lambda]$ provides at least a logarithmic (in λ) drift on fitness.

Lemma 10. *Let n be sufficiently large, $2n/5 \leq k < n/2$ and $\lambda \geq e^5 > 148$. If $2 \leq \ln(\lambda) \leq r \leq \min\{n^2 \ln(\lambda)/(25(n-2k)^2), n/4\}$, then $E(\Delta \mid k) \geq 10^{-3} \ln \lambda$.*

Proof. We first notice that $\ln(\lambda) \geq 5$ and $n^2/(25(n-2k)^2) \geq 1/(25 \cdot 0.2^2) = 1$ for all $2n/5 \leq k < n/2$. We aim to prove $\tilde{\Delta}(\lambda/2, k, \tilde{r}) \geq 10^{-3} \ln \lambda$ with $\tilde{r} = r/2$. Let X^+ and X^- be the number of one-bits flipped and zero-bits flipped, respectively, in an offspring using rate $p = \tilde{r}/n$. X^+ and X^- follow binomial distributions $\text{Bin}(k, p)$ and $\text{Bin}(n-k, p)$, respectively. Let $u := E(X^+) = kp$. Then $u \leq k/8$ and $u \geq (k/n) \ln(\lambda)/2 \geq 0.4 \cdot 0.5 \cdot \ln \lambda = 0.2 \ln \lambda \geq 1$. Furthermore, let

$$\begin{aligned}B(x) &:= \Pr(X^+ = x) = \binom{k}{x} p^x (1-p)^{k-x} \text{ for all } x \in \{0, 1, \dots, k\}, \\ F(y) &:= \Pr(X^+ \geq y) = \sum_{i=\lceil y \rceil}^k B(i) \text{ for all } y \in [0, k].\end{aligned}$$

Applying Theorem 10 in [11], we obtain $F(u) = \Pr(X^+ \geq E(X^+)) \geq 1/4$. Similarly $\Pr(X^- \leq E(X^-)) \geq 1/4$. We prove that for $\delta := E(X^-) - E(X^+) + 0.05 \ln \lambda = (n - 2k)p + 0.05 \ln \lambda$, we have $F(u + \delta) \geq \lambda^{-0.8}/9$, and thus $\Pr(X^+ - X^- \geq 0.05 \ln \lambda) \geq \lambda^{-0.8}/36$. Since for any $x \in \mathbb{Z}_{\geq u}$ we have

$$\frac{B(x+1)}{B(x)} = \frac{k-x}{x+1} \cdot \frac{p}{1-p} \leq \frac{u-up}{u-up+1-p} < 1,$$

we obtain $B(\lceil u \rceil) > B(\lceil u \rceil + 1) > \dots > B(k)$, and thus $F(u + \delta) \geq \lceil \delta \rceil B(\lceil u + 2\delta \rceil)$ as well as $\lceil \delta \rceil B(\lceil u \rceil) \geq F(u) - F(u + \delta)$. We see that

$$\begin{aligned} \frac{B(\lceil u + 2\delta \rceil)}{B(\lceil u \rceil)} &= \frac{(k - \lceil u \rceil) \cdots (k - \lceil u + 2\delta \rceil + 1)}{(\lceil u \rceil + 1) \cdots (\lceil u + 2\delta \rceil)} \cdot \frac{p^{2\delta}}{(1-p)^{2\delta}} \\ &\geq \left(\frac{k - u - 2\delta}{k(1-p)} \right)^{2\delta} \frac{u^{2\delta}}{\lceil u + 1 \rceil \cdots \lceil u + 2\delta \rceil} \geq \left(1 - \frac{2\delta}{3u} \right)^{2\delta} \frac{u^{2\delta}}{\lceil u + 1 \rceil \cdots \lceil u + 2\delta \rceil}, \end{aligned}$$

where we used $k(1-p) = k - u \geq 8u - u = 7u$. Using a sharp version of Stirling's approximation due to Robbins [42], we compute

$$\begin{aligned} \lceil u + 2\delta \rceil! &\leq \sqrt{2\pi \lceil u + 2\delta \rceil} \left(\frac{\lceil u + 2\delta \rceil}{e} \right)^{\lceil u + 2\delta \rceil} \exp \left(\frac{1}{12(\lceil u + 2\delta \rceil)} \right), \\ \lceil u \rceil! &\geq \sqrt{2\pi \lceil u \rceil} \left(\frac{\lceil u \rceil}{e} \right)^{\lceil u \rceil} \exp \left(\frac{1}{12(\lceil u \rceil + 1)} \right), \\ \frac{1}{\lceil u + 1 \rceil \cdots \lceil u + 2\delta \rceil} &= \frac{\lceil u \rceil!}{\lceil u + 2\delta \rceil!} \geq \sqrt{\frac{\lceil u \rceil}{\lceil u + 2\delta \rceil}} \frac{\lceil u \rceil^{\lceil u \rceil} e^{2\delta}}{\lceil u + 2\delta \rceil^{\lceil u + 2\delta \rceil}} \geq \frac{\sqrt{\frac{u}{u+2\delta+1}} u^u e^{2\delta}}{(u + 2\delta + 1)^{u+2\delta+1}}. \end{aligned}$$

We notice that $u \geq (k/n) \ln(\lambda)/2 \geq 0.4 \cdot 0.5 \cdot \ln \lambda = 0.2 \ln \lambda$ and

$$\delta = (n - 2k)p + 0.05 \ln \lambda \leq 0.2np + 0.05 \ln \lambda \leq 0.5u + 0.05 \ln \lambda < u.$$

Thus $2\delta/(3u) \leq 2/3$ and $2\delta/(u + 2\delta) \leq 2\delta/(\delta + 2\delta) = 2/3$. By Lemma 8 (c) we have $\ln(1 - x) \geq -x - x^2 \geq -2x$ for $0 \leq x \leq 2/3$. Hence

$$\begin{aligned} \frac{B(\lceil u + 2\delta \rceil)}{B(\lceil u \rceil)} &\geq \sqrt{\frac{u}{u + 2\delta + 1}} \left(1 - \frac{2\delta}{7u} \right)^{2\delta} \frac{u^{u+2\delta} e^{2\delta}}{(u + 2\delta + 1)^{u+2\delta+1}} \\ &\geq \sqrt{\frac{1}{4}} \exp \left(2\delta \ln \left(1 - \frac{2\delta}{7u} \right) + (u + 2\delta + 1) \ln \left(1 - \frac{2\delta}{u + 2\delta + 1} \right) + 2\delta \right) \\ &\geq \frac{1}{2} \exp \left(-\frac{2(2\delta)^2}{7u} - \frac{(2\delta)^2}{u + 2\delta + 1} \right) \geq \frac{1}{2} \exp \left(-\frac{8\delta^2}{7u} - \frac{4\delta^2}{u} \right) \\ &= \frac{1}{2} \exp \left(\frac{-6\delta^2}{u} \right) \geq \frac{1}{2} \lambda^{-0.8}, \end{aligned}$$

where in the last inequality, using $\ln(\lambda)/2 \leq \tilde{r} \leq n^2 \ln(\lambda)/(50(n-2k)^2)$, we estimate

$$\begin{aligned} \frac{6\delta^2}{u} &= \frac{6((n-2k)\tilde{r}/n + 0.05 \ln \lambda)^2}{k\tilde{r}/n} \\ &= \frac{6n}{k} \left(\frac{(n-2k)^2 \tilde{r}}{n^2} + \frac{2(n-2k)0.05 \ln \lambda}{n} + \frac{(0.05 \ln \lambda)^2}{\tilde{r}} \right) \\ &\leq \frac{6}{0.4} \left(\frac{\ln \lambda}{50} + 0.2 \cdot 0.1 \ln \lambda + \frac{0.05^2 \ln \lambda}{0.5} \right) < 0.8 \ln \lambda. \end{aligned}$$

Recalling $F(u + \delta) \geq \lceil \delta \rceil B(\lceil u + 2\delta \rceil)$ and $\lceil \delta \rceil B(\lceil u \rceil) \geq F(u) - F(u + \delta)$, we compute

$$F(u + \delta) \geq \lceil \delta \rceil B(\lceil u + 2\delta \rceil) \geq \lceil \delta \rceil (0.5\lambda^{-0.8})B(\lceil u + 2\delta \rceil) \geq (0.5\lambda^{-0.8})(F(u) - F(u + \delta)).$$

Since $F(u) \geq 0.25$ and $\lambda > 148$, we obtain

$$F(u + \delta) \geq \frac{0.5\lambda^{-0.8}F(u)}{1 + 0.5\lambda^{-0.8}} \geq \frac{0.5 \cdot 0.25 \cdot \lambda^{-0.8}}{1 + 0.5 \cdot 148^{-0.8}} > \frac{\lambda^{-0.8}}{9}.$$

Using $\Pr(X^+ - X^- \geq 0.05 \ln \lambda) \geq \lambda^{-0.8}/36$, we bound

$$\begin{aligned} \Pr(\tilde{\Delta}(\lambda/2, k, r/2) \geq 0.05 \ln \lambda \mid k) &\geq 1 - (1 - \lambda^{-0.8}/36)^{\lambda/2} \\ &\geq 1 - (1 - 148^{-0.8}/36)^{148/2} > 0.02. \end{aligned}$$

Finally $E(\Delta \mid k) \geq 0.02 \cdot 0.05 \ln \lambda = 10^{-3} \ln \lambda$. $\square$

We now extend the lemma to the whole region of $n/\ln \lambda \leq k < n/2$. If $k < 2n/5$ the situation becomes easier because $4 \leq c_2(k) < 100$ and every r in the smaller range $[c_1(k) \ln \lambda, \ln(\lambda)/2]$ provides at least an expected fitness improvement that is logarithmic in λ . Together with the previous lemma, we obtain the following statement for the drift in the whole region $n/\ln \lambda \leq k < n/2$.

Lemma 11. *Let n be sufficiently large, $n/\ln \lambda \leq k < n/2$ and $\lambda \geq e^5$. If $c_1(k) \ln \lambda \leq r \leq c_2(k) \ln(\lambda)/100$ with $c_1(k), c_2(k)$ defined as in Lemma 9, then*

$$E(\Delta \mid k) \geq 10^{-3} \ln(\lambda)/\ln(en/k).$$

Proof. If $r > \ln(\lambda)$, then $c_2(2n/5) = 100$ implies $k > 2n/5$ and the claim follows from Lemma 10. Hence let us assume $r \leq \ln(\lambda)$ in the remainder and compute $\tilde{\Delta}(\lambda/2, k, \tilde{r})$ with $\tilde{r} = r/2$.

We consider the probability $Q(k, i, \tilde{r})$ of creating from a parent with distance k an offspring with fitness distance at most $k-i$ via standard bit mutation with mutation rate $p = \tilde{r}/n$. Let $i := \max\{1, \lfloor c_1(k) \ln(\lambda)/2 \rfloor\} \leq \max\{1, r/2\} = \tilde{r}$. If $\lfloor c_1(k) \ln(\lambda)/2 \rfloor < 1$, using $(1-p)^n \geq e^{-pn-p^2n} \geq (1-o(1))e^{-\tilde{r}}$ by Lemma 8(c), we compute

$$Q(k, 1, \tilde{r}) \geq kp(1-p)^n \geq \frac{k\tilde{r}}{n}(1-p)^n \geq \frac{(1-p)^n}{\ln \lambda} \geq \frac{(1-o(1))e^{-\ln(\lambda)/2}}{\ln \lambda} \geq \frac{1}{\lambda}.$$

Otherwise if $\lfloor c_1(k) \ln(\lambda)/2 \rfloor > 1$, using $\binom{k}{i} \geq (k/i)^i$, we obtain

$$\begin{aligned} Q(k, i, \tilde{r}) &\geq \binom{k}{i} p^i (1-p)^n \geq \left(\frac{k}{i} \cdot \frac{\tilde{r}}{n}\right)^i e^{-\tilde{r}-p^2 n} \geq \left(\frac{k}{n}\right)^i e^{-\tilde{r}-p^2 n} \\ &\geq \left(\frac{k}{en}\right)^i e^{-\tilde{r}} \geq \left(\frac{k}{en}\right)^{\frac{\ln \lambda}{4 \ln(en/k)}} e^{-\tilde{r}} = e^{-\ln(\lambda)/4 - \tilde{r}} \geq e^{-\ln \lambda} = \frac{1}{\lambda}. \end{aligned}$$

Hence, $\Pr(\tilde{\Delta}(\lambda/2, k, \tilde{r}) \geq i \mid k) \geq 1 - (1 - 1/\lambda)^{\lambda/2} \geq 1 - e^{-1/2} > 0.3$. We notice that $i \geq c_1(k) \ln(\lambda)/4$, since $i = 1 \geq c_1(k) \ln(\lambda)/4$ when $c_1(k) \ln \lambda < 2$ and $i = \lfloor c_1(k) \ln(\lambda)/2 \rfloor \geq c_1(k) \ln(\lambda)/4$ when $c_1(k) \ln \lambda \geq 2$. Consequently

$$E(\Delta \mid k) > \Pr(\tilde{\Delta}(\lambda/2, k, \tilde{r}) \geq i \mid k) \cdot i \geq 0.3 \cdot c_1(k) \ln(\lambda)/4 = (0.3/8) \ln(\lambda)/\ln(en/k).$$

□

If we only consider generations that use a rate within the right region, we can bound the expected runtime to reach $k \leq n/\ln \lambda$ by $O(n/\log \lambda)$ since the drift on the fitness is of order $\log \lambda$. The following theorem shows that the additional time spent with adjusting the rate towards the right region does not change this bound on the expected runtime.

Theorem 12. *The $(1+\lambda)$ EA with self-adjusting mutation rate reaches a ONEMAX-value of $k \leq n/\ln \lambda$ within an expected number of $O(n/\log \lambda)$ iterations. This bound is valid regardless of the initial mutation rate.*

Proof. Let us denote by $k(x)$ the fitness distance of a search point $x \in \{0, 1\}^n$.

We first argue that it takes an expected number of at most $O(\sqrt{n})$ iterations to reach a fitness distance of less than $n/2$. To this end, consider a parent x with fitness distance $k(x) \geq n/2$. Let $2 \leq r \leq n/4$ be the current rate and let $\tilde{r} \in \{r/2, 2r\}$. Let y be an offspring obtained from x via standard bit mutation with mutation rate $\tilde{r}/n$. Let

$$\begin{aligned} X^+ &:= |\{i \in \{1, \dots, n\} \mid x_i = 1 \wedge y_i = 0\}|, \\ X^- &:= |\{i \in \{1, \dots, n\} \mid x_i = 0 \wedge y_i = 1\}|. \end{aligned}$$

Then the fitness improvement is $k(x) - k(y) = X^+ - X^-$ and both X^+ and X^- follow binomial distributions. From elementary properties of the binomial distribution, see, e.g., [11], we have $\Pr(X^+ \geq E(X^+) + 1) = \Omega(1)$ and $\Pr(X^- \leq E(X^-)) = \Omega(1)$, and these are independent events. Hence with constant probability, we have

$$\begin{aligned} k(y) &= k(x) - X^+ + X^- \leq k(x) - E(X^+) - 1 + E(X^-) \\ &\leq k(x) - 1 - (2k(x) - n) \frac{\tilde{r}}{n} \leq k(x) - 1. \end{aligned}$$

Clearly, for the best offspring z out of the λ offspring generated in this iteration, we have $k(z) \leq k(y)$. Consequently, in each iteration starting with a parent x with $k(x) \geq n/2$, with constant probability we gain at least one fitness level. By the additive drift theorem (see Theorem 3 and the subsequent discussion), from the initial random

search point x_0 it takes an expected $O(\max\{0, k(x_0) - n/2 + 1\})$ iterations to reach a parent with fitness distance less than $n/2$. Since $X := k(x_0) \sim \text{Bin}(n, 1/2)$, we have

$$\begin{aligned} E(\max\{0, X - n/2\}) &= \frac{1}{2}E(|X - E(X)|) \leq \frac{1}{2}\sqrt{E((X - E(X))^2)} \\ &= \frac{1}{2}\sqrt{\text{Var}(X)} = \frac{1}{4}\sqrt{n}, \end{aligned}$$

where we first exploited the symmetry of X and then the well-known estimate $E(Y)^2 \leq E(Y^2)$ valid for all random variables Y , in particular, for $Y = |X - E(X)|$. By the law of total expectation, the expected time to reach a search point with k -value below $n/2$ is $O(\sqrt{n})$.

Without loss of generality we can now assume $k < n/2$ for the initial state. Our intuition is that once we begin to use a rate $r \in [(c_1(k)/2) \ln \lambda, c_2(k) \ln \lambda]$ at some distance level k , we will have a considerable drift on the ONEMAX-value and the strong drift on the rate keeps r within or close to this interval. After we make progress and k decreases to a new level, the corresponding c_1 and c_2 decrease, and the algorithm may take some time to readjust r into new bounds.

We consider the stochastic process $X_t = \lfloor \log_2(r_t) \rfloor$ and the current ONEMAX-value Y_t . According to Lemma 9 (a) and (b), there exists $\varepsilon = \Omega(1)$ such that

$$\begin{aligned} \Pr(X_{t+1} - X_t \mid X_t; X_t < \log_2(c_1(K_t) \ln \lambda) - 1) &\geq \varepsilon, \\ \Pr(X_t - X_{t+1} \mid X_t; X_t > \log_2(c_2(K_t) \ln \lambda)) &\geq \varepsilon. \end{aligned}$$

Let $k_0 > k_1 > \dots > k_N$ be all the different ONEMAX-values taken by Y_t until for the first time $Y_t \leq n/\ln \lambda$. By the additive drift theorem, it takes at most $O(\log n)$ iterations to have $r_t \in [(c_1(k_0)/2) \ln \lambda, c_2(k_0) \ln \lambda]$, regardless of how we set the initial rate. Since $c_1(Y_t)$ is non-increasing, $r_t \geq c_1(Y_t) \ln(\lambda)/2$ implies $r_t \geq c_1(Y_{t+1}) \ln(\lambda)/2$. Thus, no readjustments are necessary to satisfy the lower bound $r_t \geq c_1(Y_t) \ln(\lambda)/2$ once we have $r_t \geq c_1(k_0) \ln(\lambda)/2$. We now regard the upper bound condition $r_t \leq c_2(Y_t) \ln(\lambda)$. Let (k_{s_i}) be the subsequence consisting of all k_{s_i} such that $\{t \mid r_t \leq c_2(Y_t) \ln(\lambda), Y_t = k_{s_i}\} \neq \emptyset$. If k_i is in the subsequence, once $r_t \leq c_2(Y_t) \ln(\lambda)$ is achieved, the runtime until the first time $Y_t = k_{i+1}$ can be computed using the occupation lemma and the variable drift theorem. After that, it takes at most $O(\log(c_2(k_i)/c_2(k_{i+1})))$ more iterations to readjust r_t from $c_2(k_i) \ln(\lambda)$ to $c_2(k_{i+1}) \ln(\lambda)$ by the additive drift theorem. Similar arguments hold for k_{i+1} if it is also in the subsequence. Otherwise let $k_j < k_{i+1}$ denote the next ONEMAX-values after k_i in the subsequence. By definition, the fitness distance decreases from k_{i+1} to k_j before the rate being adjusted into the corresponding range. This means that improvements in distance are made during readjustment. Analogous to above, the expected (readjusting) time to decrease the rate from $c_2(k_{i+1}) \ln(\lambda)$ to $c_2(k_j) \ln(\lambda)$ is $O(\log(c_2(k_{i+1})/c_2(k_j)))$. Therefore, the expected number of total readjusting time is at most

$$\sum_{i=1}^{s_{i+1} \leq N} O\left(\log\left(\frac{c_2(k_{s_i})}{c_2(k_{s_{i+1}})}\right)\right) = O\left(\log\left(\frac{c_2(k_0)}{c_2(k_N)}\right)\right) = O(\log n).$$

When computing the expected number of non-adjusting iterations until $Y_t = k_N$, we choose a constant $b \in \mathbb{N}_{\geq 2}$ large enough such that $2e^{-2b\varepsilon/3} \leq 1/2 - \delta/2$ holds for some positive constant $\delta > 0$. When $r_t \in [(c_1(k)/2) \ln \lambda, c_2(k) \ln \lambda]$ and $Y_t = k$ for some t , then by Lemma 7, we obtain

$$\Pr(r_{t'} \geq 2^{b+1} c_2(k) \ln \lambda) \leq 2e^{-2b\varepsilon/3} \text{ and } \Pr(r_{t'} \leq 2^{-b-2} c_1(k) \ln \lambda) \leq 2e^{-2b\varepsilon/3} \text{ for } t' \geq t.$$

Hence, we obtain that $r_{t'} \in [2^{-b-2} c_1(k) \ln \lambda, 2^{b+1} c_2(k) \ln \lambda]$ happens with probability at least δ . We see that there are at most $\lceil \log_2(100) \rceil$ steps between being in the range $[(c_1(k)/2) \ln \lambda, c_2(k) \ln \lambda]$ and being in the smaller range $[c_1(k) \ln \lambda, c_2(k) \ln \lambda / 100]$ which is described in Lemma 11. If $r_t \in [2^{-b-2} c_1(k) \ln \lambda, 2^{b+1} c_2(k) \ln \lambda]$, it takes at most a constant number of iterations α in expectation to reach $[c_1(k) \ln \lambda, c_2(k) \ln \lambda / 100]$ because our mutation scheme employs a 50% chance to perform a random step of the mutation rate. Based on Lemma 11, the fitness drift at distance k of all rates within the narrow region is at least $10^{-3} \ln(\lambda) / \ln(en/k)$. This contributes to an average drift of at least

$$10^{-3} \cdot \frac{\ln(\lambda)}{\ln(en/k)} \cdot \frac{\delta}{1 + \alpha} = \Omega\left(\frac{\log(\lambda)}{\log(n/k)}\right)$$

for all random rates at distance k . Using the variable drift theorem (Theorem 3), we estimate the runtime as

$$\begin{aligned} & O\left(\frac{\log(\log \lambda)}{\log \lambda} + \int_{n/\log \lambda}^{n/2} \frac{\log(n/k) dk}{\log \lambda}\right) = O\left(\frac{1}{\log \lambda} \int_{n/\log \lambda}^{n/2} (\log(n) - \log(k)) dk\right) \\ &= O\left(\frac{1}{\log \lambda} \left(\log(n)k - k \log(k) + \log(k)\right) \Big|_{n/\log \lambda}^{n/2}\right) \\ &= O\left(\frac{1}{\log \lambda} \left(\log(n) \left(\frac{n}{2} - \frac{n}{\log \lambda}\right) - \frac{n}{2} \log\left(\frac{n}{2}\right) + \frac{n}{\log \lambda} \log\left(\frac{n}{\log \lambda}\right) + \log\left(\frac{n/2}{n/\log \lambda}\right)\right)\right) \\ &= O\left(\frac{1}{\log \lambda} \left(\frac{n \log 2}{2} - \frac{n \log(\log \lambda)}{\log \lambda} + \log(\log \lambda)\right)\right) = O\left(\frac{n}{\log \lambda}\right). \end{aligned}$$

We notice that the expected runtime for adjusting r_t is $O(\log n) = O(n/\log \lambda)$. Therefore, the total runtime is $O(n/\log \lambda)$ in expectation. $\square$

6 Middle Region

In this section we estimate the expected number of generations until the number of one-bits has decreased from $k \leq n/\ln \lambda$ to $k \leq n/\lambda$. We first claim that the right region for r is $1 \leq r \leq \ln(\lambda)/2$. Hence, the $(1+\lambda)$ EA is not very sensitive to the choice of r here. Intuitively, this is due to the fact that a total fitness improvement of only $O(n/\log \lambda)$ suffices to cross the middle region, whereas an improvement of $\Omega(n)$ is needed for the far region.

We estimate the drift of the fitness in Lemma 13 and apply that result afterwards to estimate the number of generations to cross the region.

Lemma 13. *Let $n/\lambda \leq k \leq n/\ln \lambda$, $\lambda \geq 26$ and $2 \leq r \leq \ln \lambda$. Then*

$$E(\Delta \mid k) \geq \min \left\{ \frac{1}{8}, \frac{\sqrt{\lambda}k}{32n} \right\}.$$

Proof. We aim at computing $\tilde{\Delta}(\lambda/2, k, \tilde{r})$ with $\tilde{r} = r/2$. The probability that no zero-bit flips in a single offspring stemming from rate $\tilde{r}$ is $(1 - \tilde{r}/n)^{n-k} \geq e^{-\tilde{r}} \geq 1/\sqrt{\lambda}$. We regard the number Z of offspring of rate $\tilde{r}$ that have no flipped zeros. The expectation $E(Z)$ is at least $(1/\sqrt{\lambda}) \cdot (\lambda/2) = \sqrt{\lambda}/2$. Applying Chernoff bounds (Theorem 5), we observe that Z exceeds $\lambda_0 := \sqrt{\lambda}/4$ with probability at least $1 - \exp(-\sqrt{\lambda}/16) > 1/4$ since $\lambda \geq 26$. Assuming this to happen, we look at the first λ_0 offspring without flipped zeros. For $i \in \{1, \dots, \lambda_0\}$ let X_i be the number of flipped ones in the i -th offspring. Then the X_i are i.i.d. with $X_i \sim \text{Bin}(k, \tilde{r}/n)$. Let $X^* := \max\{X_i\}$. The expectation of X^* is analyzed in Gießen and Witt [27, Lemma 4, part 2]; more precisely the following statement is shown for constant $\tilde{r}$:

$$\text{If } \frac{\lambda_0 k \tilde{r}}{n} \geq \alpha \text{ then } E(X^*) \geq \frac{\alpha}{1 + \alpha}.$$

Since $\tilde{r}$ is not necessarily assumed constant here, we generalize the result by closely following the proof in [27]. Note that $X^* \geq 1$ if at least one of the offspring does not flip a zero-bit. Hence,

$$\begin{aligned} E(X^*) &\geq 1 - \left(\left(1 - \frac{\tilde{r}}{n}\right)^k \right)^{\lambda_0} \geq 1 - \left(\left(1 - \frac{\tilde{r}}{n}\right)^{\alpha n / (\lambda_0 \tilde{r})} \right)^{\lambda_0} \\ &\geq 1 - \left(e^{-\alpha / \lambda_0} \right)^{\lambda_0} = 1 - e^{-\alpha} \geq 1 - \frac{1}{1 + \alpha} = \frac{\alpha}{1 + \alpha}, \end{aligned}$$

where the second inequality used the assumption $\lambda_0 k \tilde{r}/n \geq \alpha$, the third one $e^x \geq 1 + x$ for $x \in \mathbb{R}$. and the fifth one the equivalent to the previous inequality $e^{-x} \leq 1/(1 + x)$. Hence, Lemma 4, part 2 in [27] holds also for arbitrary $\tilde{r}$.

Setting $\alpha := \lambda_0 k \tilde{r}/n$, we now distinguish between two cases. If $\alpha \geq 1$, we obtain $E(X^*) \geq \alpha/(1 + \alpha) \geq 1/2$; otherwise we have $1 + \alpha < 2$, hence $E(X^*) \geq \alpha/2 = \lambda_0 k \tilde{r}/(2n)$. Thus,

$$E(X^*) \geq \min\{1/2, \sqrt{\lambda}k/(8n)\}.$$

Hence, using the law of total probability, we obtain the lower bound on the drift for the middle region. $\square$

We now use our result on the drift to estimate the time spent in this region. We notice that $c_2(k) = 4 + o(1)$ when $k = o(n)$. This means we will often have $r_t \in [2, \ln(\lambda)]$ which provides the drift we need.

Theorem 14. *Let $\lambda \geq 26$ and $\lambda = n^{O(1)}$. Assume $k \leq n/\ln \lambda$ for the current ONEMAX-value of the self-adjusting $(1+\lambda)$ EA. Then the expected number of generations until $k \leq n/\lambda$ is $O(n/\log \lambda)$.*

Proof. For $k \leq n/\ln \lambda$ the upper bound from Lemma 9 is $c_2(k) = 4/(1 - 2/\ln \lambda) < 11$. According to the lemma we have for $X_t := \lceil \log_2 r_t \rceil$ that $E(X_t - X_{t+1} \mid X_t; X_t > \log_2(c_2(k) \ln \lambda)) \geq 2\varepsilon = \Omega(1)$. The additive drift theorem yields that in $O(\log n)$ time we have $r_t \leq c_2(k) \ln \lambda$. We choose b large enough such that $2e^{-2b\varepsilon/3} \leq 1 - \delta$ holds for some positive constant $\delta > 0$ and note that b is constant. Applying Lemma 7 we obtain $\Pr(r_t \geq 2^b c_2(k) \ln \lambda) \leq 2e^{-2b\varepsilon/3}$ and $r_t \leq 2^b c_2(k) \ln \lambda$ happens with probability at least δ . Once $r_t < 2^b c_2(k) \ln \lambda < 2^b(11 \ln \lambda)$ it takes at most a constant number of iterations α in expectation to draw r_t to $\ln(\lambda)$ or less. According to Lemma 13 this ensures a drift of at least $(1/4) \min\{1/2, \sqrt{\lambda}k/(8n)\} \geq (1/32) \min\{1, \sqrt{\lambda}k/n\}$, which implies an average drift of at least $c \min\{1, \sqrt{\lambda}k/n\}$ over all random rates at distance k , where $c > 0$ is a constant. Considering $\min\{1, \sqrt{\lambda}k/n\}$, the minimum is taken on the first argument if $k > n/\sqrt{\lambda}$, and on the second if $k < n/\sqrt{\lambda}$.

We are interested in the expected time to reduce the ONEMAX-value to at most n/λ . To ease the application of drift analysis, we artificially modify the process and make it create the optimum when the state (ONEMAX-value) is strictly less than n/λ . Clearly, the first hitting time of state at most n/λ does not change by this modification. Applying the variable drift theorem (Theorem 3) with $x_{\min} = n/\lambda$, $X_0 = k \leq n/\ln \lambda$ and $h(x) = c \min\{1, \sqrt{\lambda}k/n\}$, the expected number of generations to reach state at most n/λ is bounded from above by

$$\frac{n/\lambda}{c\sqrt{\lambda}(n/\lambda)/n} + \int_{n/\lambda}^{n/\sqrt{\lambda}} \frac{n}{c\sqrt{\lambda}x} dx + \int_{n/\sqrt{\lambda}}^{n/\ln \lambda} \frac{1}{c} dx = O\left(\frac{n}{\sqrt{\lambda}}\right) + O\left(\frac{n \log \lambda}{\sqrt{\lambda}}\right) + O\left(\frac{n}{\log \lambda}\right),$$

which is $O(n/\log \lambda)$. The overall expected number of generations spent is $O(\log n + n/\log \lambda) = O(n/\log \lambda)$ since $\lambda = n^{O(1)}$ by assumption. $\square$

7 Near Region

In the near region, we have $k \leq n/\lambda$. Hence, the fitness is so low that we can expect only a constant number of offspring to flip at least one of the remaining one-bits. This assumes constant rate. However, higher rates are detrimental since they are more likely to destroy the zero-bits of the few individuals flipping one-bits. Hence, we expect the rate to drift towards constant values, as shown in the following lemma.

Lemma 15. *Let $k \leq n/\lambda$, $\lambda \geq 45$ and $4 \leq r_t \leq \ln(\lambda)/4$. Then the probability that $r_{t+1} = r_t/2$ is at least 0.5099.*

Proof. To prove the claim we exploit the fact that only few one-bits are flipped in both subpopulations. Using $r := r_t$, we shall argue as follows. With sufficiently high (constant) probability, (i) the $2r$ -subpopulation contains no individual strictly better than the parent, that is, with fitness less than k , and (ii) all $2r$ -offspring with fitness

at least k are identical to the parent. Conditional on this, either the $r/2$ -population contains individuals with fitness less than k and the winning individual surely stems from this subpopulation, or the $r/2$ -population contains no better offspring. In the latter case, we argue that there are many more individuals with fitness exactly k in the $r/2$ -population than in the $2r$ -population, which gives a sufficiently high probability for taking the winning individual from this side (as it is chosen uniformly at random from all offspring with fitness k).

Let $N_{r/2}$ and N_{2r} be the number of offspring that did not flip any zero-bits using rate $r/2$ and $2r$, respectively. Then $E(N_{r/2}) = (\lambda/2)(1 - r/(2n))^{n-k} \geq (1 - o(1))\lambda e^{-r/2}/2$, since $k \geq 1$ and

$$\left(1 - \frac{r}{2n}\right)^{n-1} \geq e^{-\frac{r}{2}} \left(1 - \frac{r}{2n}\right)^{\frac{r}{2}-1} \geq e^{-\frac{r}{2}} \left(1 - \frac{c \ln n}{8n}\right)^{\frac{c \ln n}{8}-1} = (1 - o(1))e^{-\frac{r}{2}},$$

where we used that $r \leq \ln \lambda/4$ and $\lambda = n^{O(1)}$, i. e. $\ln \lambda \leq c \ln n$ for some constant c . Using $k \leq n/\lambda$ we get $E(N_{2r}) = (\lambda/2)(1 - 2r/n)^{n-k} \leq (\lambda/2)e^{-2r(1-1/\lambda)}$. In fact, we can discriminate $N_{r/2}$ and N_{2r} by using Theorem 5 in the following way: we have

$$\begin{aligned} \Pr\left(N_{r/2} \leq \frac{\lambda}{4}e^{-\frac{r}{2}}\right) &\leq \exp\left(-(1 - o(1))^3 \frac{\lambda}{16}e^{-\frac{r}{2}}\right) \\ &\leq \exp\left(-(1 - o(1))\frac{1}{16}\lambda^{7/8}\right) \\ &< 0.175, \end{aligned}$$

for sufficiently large n , since $r \leq (\ln \lambda)/4$ and $\lambda \geq 45$. Similarly, we obtain

$$\Pr\left(N_{2r} \geq \lambda e^{-2r(1-\frac{1}{\lambda})}\right) \leq \exp\left(-\frac{1}{6}\lambda^{1-\frac{1}{2}(1-\frac{1}{\lambda})}\right) < 0.312.$$

Note that $e^{-2r(1-1/\lambda)} < e^{-r/2}/4$ holds for all $r \geq 4$ and $\lambda \geq 45$. Since the offspring are generated independently, the events $N_{r/2} > \lambda e^{-r/2}/4$ and $N_{2r} < \lambda e^{-2r(1-1/\lambda)}$ happen together with probability at least $(1-0.175) \cdot (1-0.312) \geq 0.567 =: 1 - p_{\text{err}_1}$. Conditioning on this and by using a union bound the probability p_{err_2} that at least one of the N_{2r} offspring that do not flip any zero-bits flips at least one one-bit can be upper bounded by

$$p_{\text{err}_2} := N_{2r} \cdot \frac{2kr}{n} \leq 2re^{-2r(1-\frac{1}{\lambda})} \leq 0.004$$

using $k/n \leq 1/\lambda$ in the first and $r \geq 4$ and $\lambda \geq 45$ in the last inequality. By using a union bound, we find the probability p_{err_3} that at least one $2r$ -offspring flips exactly one one-bit and exactly one zero-bit to be at most

$$\begin{aligned} p_{\text{err}_3} &:= \frac{\lambda}{2} \frac{2rk}{n} \left(1 - \frac{2r}{n}\right)^{k-1} \frac{2r(n-k)}{n} \left(1 - \frac{2r}{n}\right)^{n-k-1} \\ &\leq 2r^2 \left(1 - \frac{2r}{n}\right)^{n-2} \leq (1 + o(1))2e^{2(\ln(r)-r)} < (2 + o(1))e^{-\frac{6}{5}r} < 0.017, \end{aligned}$$

for sufficiently large n , using $k/n \leq 1/\lambda$ for the first inequality. The third inequality is due to $\ln x - x \leq -(1 - e^{-1})x < -(3/5)x$ for all $x > 0$ (see Lemma 8.a) and the last inequality stems from $r \geq 4$. The second inequality follows from

$$\left(1 - \frac{2r}{n}\right)^{n-2} \leq e^{-\frac{2r}{n}(n-2)} = e^{-2r(1-\frac{2}{n})} = (1 + o(1))e^{-2r},$$

using again $r \leq \ln \lambda \leq c \ln n$ for some constant c . Let M_r be the number of such offspring. Any other fitness-decreasing flip-combinations of zeroes and ones in the $2r$ -subpopulation require an offspring to flip at least two one-bits. The probability that such an offspring is created is at most

$$p_{\text{err}_4} := \frac{\lambda}{2} \binom{k}{2} \left(\frac{2r}{n}\right)^2 \leq \frac{\ln^2 \lambda}{16\lambda} < 0.021,$$

using $k \leq n/\lambda$ and $r \leq \ln \lambda/4$ and the fact that $(\ln^2 x)/x$ is decreasing for $x \geq e^2$ and $\lambda \geq 45 > e^2$.

The events $N_{r/2} > \lambda e^{-r/2}/4$, $N_{2r} < \lambda e^{-2r(1-\frac{1}{\lambda})}$, $M_r = 0$, and the event that no fitness-decreasing offspring is created in the $2r$ -subpopulation are sufficient to ensure that the best individual is either surely from the $r/2$ -population or chosen uniformly at random from the $N_{r/2} + N_{2r}$ offspring. Conditioning on these events, we have that the probability that the best offspring is chosen from the $r/2$ -population is at least

$$\frac{N_{r/2}}{N_{r/2} + N_{2r}} \geq \frac{1}{1 + 4e^{-r(2(1-\frac{1}{\lambda})-\frac{1}{2})}} > 0.988.$$

Hence, using a union bound for the error probabilities, the unconditional probability is at least

$$\frac{(1 - p_{\text{err}_1} - p_{\text{err}_2} - p_{\text{err}_3} - p_{\text{err}_4}) \cdot 0.988 + \frac{1}{2}}{2} > 0.5099.$$

□

We note that the restriction $r_t \geq 4$ in the lemma above is not strictly necessary. Also for smaller r_t , the probability that the winning individual is chosen from the $r_t/2$ -population is by an additive constant larger than $1/2$. Showing this, however, would need additional proof arguments as for smaller r_t , the event that both subpopulations contain individuals with fitness $k - 1$ becomes more likely. We avoid this additional technicality by only arguing for $r_t \geq 4$, which is enough since any constant r_t is sufficient for the fitness drift we need (since we do not aim at making the leading constant precise).

In the following proof of the analysis of the near region, we use the above lemma (with quite some additional arguments) to argue that the r -value quickly reaches 4 or less and from then on regularly returns to this region. This allows to argue that in the near region we have a speed-up of a factor of $\Theta(\lambda)$ compared to the (1+1) EA, since every offspring only has a probability of $O(1/\lambda)$ of making progress (see also [13, 27]).

Theorem 16. Assume $k \leq n/\lambda$ for the current ONEMAX-value of the self-adjusting $(1+\lambda)$ EA. Then the expected number of generations until the optimum is reached is $O(n \log(n)/\lambda + \log n)$.

Proof. The aim is to estimate the ONEMAX-drift at the points in time (generations) t where $r_t = O(1)$. To bound the expected number of generations until the mutation rate has entered this region, we basically consider the stochastic process $Z_t := \max\{0, \lfloor \log_2(r_t/a) \rfloor\}$, where $a := 4$, which is the lower bound on r_t from Lemma 15. However, as we do not have proved a drift of Z_t towards smaller values in the region $L := (\ln \lambda)/4 \leq r_t \leq 16 \ln \lambda =: U$ (where 16 is an upper bound on $c_2(k)$ from Lemma 11), we use the potential function

$$X_t(Z_t) := \begin{cases} Z_t & \text{if } a \leq r_t \leq L \\ \log_2(L/a) + \sum_{i=1}^{\lfloor \log_2(r_t/L) \rfloor} 4^{-i} & \text{if } L < r_t < U \\ \log_2(L/a) + \sum_{i=1}^{\log_2(U/L)+1} 4^{-i} + 4^{-\log_2(U/L)-1} \lfloor \log_2(r_t/a) \rfloor & \text{otherwise.} \end{cases}$$

assuming that L and U have been rounded down and up to the closest power of 2, respectively. We note that $X_t \geq \log_2(r_t/a) - 1 = \log_2(r_t/(2a))$ if $r_t \leq L$ and $X_t \geq (L/U)^2 \log_2(r_t/(2a))$ in all three cases.

The potential function has a slope of 1 for $a \leq r_t \leq L$. Lemma 15 gives us the drift $E(X_t - X_{t+1} \mid X_t; a \leq r_t \leq L) \geq (0.5099 - 0.4901)/(2a) = \Omega(1)$. The function satisfies $X_t(Z_t) - X_t(Z_t - 1) \geq 4(X_t(Z_{t+1}) - X_t(Z_t))$ if $L < r_t < U$, which corresponds to the region where the probability of decreasing r_t by a factor of 1/2 has only be bounded from below by 1/4 due to the random steps. Still, $E(X_t - X_{t+1} \mid X_t; L < r_t < U) \geq (1/4)4^{-\log_2(U/L)-1} - (3/4)4^{-\log_2(U/L)} = \Omega(1)$ in this region due to the concavity of the potential function. Finally, $E(X_t - X_{t+1} \mid X_t; r_t \geq U) = 4^{-\log_2(U/L)-1}(1/(2a))\Omega(1) = \Omega(1)$ by Lemma 11. Hence, altogether $E(X_t - X_{t+1} \mid X_t; r_t \geq a) \geq \kappa$ for some constant $\kappa > 0$. As $X_0 = O(\log n)$, additive drift analysis yields an expected number of $O(\log n)$ generations until for the first time $X_t = 0$ holds, corresponding to $r_t \leq a$. We denote this hitting time by T .

We now consider an arbitrary point of time $t \geq T$. The aim is to show a drift on the ONEMAX-value, depending on the current ONEMAX-value Y_t , which satisfies $Y_t \leq n/\lambda$ with probability 1. To this end, we will use Lemma 7. We choose b large enough such that $2e^{-2b \cdot \kappa/4} \leq 1 - \delta$ holds for some positive constant $\delta > 0$ and note that b is constant. We consider two cases. If $X_t \leq b$, which happens with probability at least δ according to the lemma, then the bound $X_t \geq (L/U)^2(\log_2(r_t/(2a)))$ implies $r_t \leq 2^{(U/L)^2} 2a2^b$. Hence, we have $r_t = O(1)$ in this case and obtain a probability of at least

$$1 - \left(1 - \binom{Y_t}{1} \left(\frac{2r_t}{n}\right) \left(1 - \frac{2r_t}{n}\right)^{n-1}\right)^\lambda \geq 1 - \left(1 - \Theta\left(\frac{Y_t}{n}\right)\right)^\lambda = \Omega(\lambda Y_t/n)$$

to improve the ONEMAX-value by 1, using that $Y_t = O(n/\lambda)$ and pessimistically assuming a rate of $2r_t$ in all offspring. If $X_t > b$, we bound the improvement from below by 0. Using the law of total probability, we obtain

$$E(Y_t - Y_{t+1} \mid Y_t; Y_t \leq n/\lambda) = \delta \Omega(\lambda Y_t/n) = \Omega(\lambda Y_t/n).$$

Now a multiplicative drift analysis with respect to the stochastic process on the Y_t , more precisely Theorem 4 using $\delta = \Theta(\lambda/n)$ and minimum state 1, gives an expected number of $O((n/\lambda) \log Y_0) = O(n \log(n)/\lambda)$ generations until the optimum is found. Together with the expected number $O(\log n)$ until the r -value becomes at most a , this proves the theorem. $\square$

8 Putting Everything Together

In this section, we put together the analyses of the different regimes to prove our main result.

Proof of Theorem 1. The lower bound actually holds for all unbiased parallel black-box algorithms, as shown in [2].

We add up the bounds on the expected number of generations spent in the three regimes, more precisely we add up the bounds from Theorem 12, Theorem 14 and Theorem 16, which gives us $O(n/\log \lambda + n \log(n)/\lambda + \log n)$ generations. Due to our assumption $\lambda = n^{O(1)}$ the bound is dominated by $O(n/\log \lambda + n \log(n)/\lambda)$ as suggested. $\square$

Proof of Lemma 2. We basically revisit the regions of different ONEMAX-values analyzed in this paper and bound the time spent in these regions under the assumption $r = \ln(\lambda)/2$. In the far region, Lemmas 10 and 11, applied with this value of r , imply a fitness drift of $\Omega(\log(\lambda)/\log(en/k))$ per generation, so the expected number of generations spent in the far region is $O(n/\log \lambda)$ as computed by variable drift analysis in the proof of Theorem 12.

The middle region is shortened at the lower end. For $k \geq n/\sqrt{\lambda}$, Lemma 13 gives a fitness drift of $\Omega(1)$, implying by additive drift analysis $O(n/\log \lambda)$ generations to reduce the fitness to at most $n/\sqrt{\lambda}$.

In the near region, which now starts at $n/\sqrt{\lambda}$, we have to argue slightly differently. Note that every offspring has a probability of at least $(1-r)^n \geq e^{-\ln(\lambda)/2 + O(1)} = \Omega(\lambda^{-1/2})$ of not flipping a zero-bit. Hence, we expect $\Omega(\sqrt{\lambda})$ such offspring. We pessimistically assume that the other individuals do not yield a fitness improvement; conceptually, this reduces the population size to $\Omega(\sqrt{\lambda})$ offspring, all of which are guaranteed not to flip a zero-bit. Adapting the arguments from the proof of Theorem 16, the probability that at least one of these individuals flips at least a one-bit is at least

$$1 - \left(1 - \binom{Y_t}{1} \left(\frac{r_t}{n}\right)\right)^{\Omega(\sqrt{\lambda})} \geq 1 - \left(1 - \Theta\left(\frac{Y_t}{n}\right)\right)^{\Omega(\sqrt{\lambda})} = \Omega(\sqrt{\lambda} Y_t / n),$$

which is a lower bound on the fitness drift. Using the multiplicative drift analysis, the expected number of generations in the near region is $O(n \log(n)/\sqrt{\lambda})$. Putting the times for the regions together, we obtain the lemma. $\square$

9 Experiments

Since our analysis is asymptotic in nature we performed some elementary experiments in order to see whether besides the asymptotic runtime improvement (showing an improvement for an unspecified large problem size n) we also see an improvement for realistic problem sizes. For this purpose we implemented the $(1+\lambda)$ EA in C using the GNU Scientific Library (GSL) for the generation of pseudo-random numbers. In this section our performance measure is the runtime, represented by the number of generations until the optimum is found for the first time.

The first plot in Figure 1 displays the average runtime over 10000 runs of the self-adjusting $(1+\lambda)$ EA on ONEMAX for $n = 5000$ as given in Algorithm 1 over $\lambda = 100, 200, \dots, 1000$; the second plot shows the corresponding interquartile ranges to support that the results are statistically significant. We set the initial mutation rate to 2, i.e., the minimum mutation rate the algorithm can attain. Moreover, the plot displays the average runtimes of the classic $(1+\lambda)$ EA using static mutation probabilities of $1/n$ and of $(\ln \lambda)/(2n)$, the latter of which is asymptotically optimal for large λ according to Lemma 2.

The average runtimes of both algorithms profit from higher offspring population sizes λ leading to lower average runtimes as λ increases. Interestingly, the two static settings of the classic $(1+\lambda)$ EA outperform the self-adjusting $(1+\lambda)$ EA for small values of λ up to $\lambda = 400$. For higher offspring population sizes the self-adjusting $(1+\lambda)$ EA outperforms the classic ones, indicating that the theoretical performance gain of $\ln \ln \lambda$ can in fact be relevant in practice. Furthermore, we implemented the self-adjusting $(1+\lambda)$ EA without the random steps, that is, when the rate is always adjusted according to how the best offspring are distributed over the two subpopulations. The experiments show that this variant of the self-adjusting $(1+\lambda)$ EA performs generally slightly better on ONEMAX. Since the ONEMAX fitness landscape is structurally very simple, this result is not totally surprising. It seems very natural that the fitness of the best of $\lambda/2$ individuals, viewed as a function in the rate, is a unimodal function. In this case, the advantage of random steps to be able to leave local optima of this function is not needed. On the other hand, of course, this observation suggests to try to prove our performance bound rigorously also for the case without random rate adjustments. We currently do not see how to do this. Lastly, we implemented the $(1+\lambda)$ EA using the fitness-depending mutation rate $p = \max\{\frac{\ln \lambda}{n \ln(en/d)}, \frac{1}{n}\}$ as presented in [2]. The experiments suggest that this scheme outperforms all other variants considered.

Additionally we implemented another variant of the self-adjusting $(1+\lambda)$ EA using three equally-sized subpopulations i.e. the additional one is using (that is, exploiting) the current mutation rate. We compared this variant with the self-adjusting $(1+\lambda)$ EA, both with and without using random steps. The results are shown in Figure 2, again with respect to average runtimes and interquartile ranges. The experiments suggest that the variant using three subpopulations outperforms the self-adjusting $(1+\lambda)$ EA slightly for small population sizes. For very high population sizes, using just two subpopulations seems to be a better choice.

Figure 1: Static and Self-adjusting $(1+\lambda)$ EA average runtime comparison ($n = 5000$) on ONEMAX and the corresponding interquartile ranges.

Figure 2: Average runtime of the self-adjusting $(1+\lambda)$ EA with two and three subpopulations each with and without random steps on ONEMAX ($n = 5000$) with corresponding interquartile ranges.

Figure 3: Average runtime of the self-adjusting $(1+\lambda)$ EA with different mutation rate update factors F on ONEMAX ($n = 5000$) with corresponding interquartile ranges

To gain some understanding on how the parameters influence the runtime, we implemented the self-adjusting $(1+\lambda)$ EA using different mutation rate update factors, that is, we consider the self-adjusting $(1+\lambda)$ EA as given in Algorithm 1 where the mutation rate r_t is increased or decreased by some factor F (instead of the choice $F = 2$ made in Algorithm 1). Note that we do not change the rule that we use the rates $r/2$ and $2r$ to create the subpopulations. Furthermore, after initialization, the algorithm starts with rate F and the rate is capped below by F and above by $1/(2F)$ during the run, accordingly.

The results are shown in Figure 3. The plot displays the average runtime over 10000 runs of the self-adjusting $(1+\lambda)$ EA on ONEMAX for $n = 5000$ over $\lambda = 100, 200, \dots, 1000$ using the update factors $F = 2.0, 1.5, 1.01$. The plot suggests that lower values of F yield a better performance. This result is not immediately obvious. Clearly, a large factor F implies that the rate changes a lot from generation to generation (namely by a factor of F). These changes prevent the algorithm from using a very good rate for several iterations in a row. On the other hand, a small value for F implies that it takes longer to adjust the rate to value that is far from the current one. We also performed experiments for $F > 2$ and observed even worse runtimes than for $F = 2.0$. The figure does not display the outcomes of these additional experiments since this would impair the readability of the diagram.

Finally, to illustrate the nontrivial development of the rate during a run of the algorithm we plotted the rate of three single runs of the self-adjusting $(1+\lambda)$ EA using different factors F over the fitness in Figure 4. Since the algorithm initialized with the rate F , the rate increases after initialization and decreases again with decreasing fitness-distance to the optimum. The plot suggests that for higher values of F the rate is more unsteady due to the greater impact of the rate adjustments while smaller rate updates yield a more stable development of the rate. Interestingly, for all three values of F , the rates seem to correspond to the same rate after the initial increasing phase from F . Note that this illustration does not indicate the actual runtime. In fact, the specific runtimes are 19766 for $F = 1.01$, 19085 for $F = 1.05$ and 19857 for $F = 1.1$. A similar, more pronounced behaviour can be seen for $F = 2.0$; we chose these particular values of F for illustrative purposes since for $F = 2.0$ the variance in the rate can be visually confusing for the reasons given above.

While we would draw from this experiment the conclusion that a smaller choice of F is preferable in a practical application of our algorithm, the influence of the parameter on the runtime is not very large. So it might not be worth optimizing it and rather view Algorithm 1 as a parameter-less algorithm.

10 Conclusions

We proposed and analyzed a new simple self-adjusting mutation scheme for the $(1+\lambda)$ EA. It consists of creating half the offspring with a slightly larger and the rest with a slightly smaller mutation rate. Based on the success of the subpopulations, the mutation rate is adjusted. This simple scheme overcomes difficulties of previous self-adjusting choices,

Figure 4: Development of the rate over the fitness of three example runs of the self-adjusting $(1+\lambda)$ EA on ONEMAX ($n = 100000$, $\lambda = 1000$), using different factors F

e.g., the careful choice of the exploration-exploitation balance and the forgetting rate in the learning scheme of [20].

We proved rigorously that this self-adjusting $(1+\lambda)$ EA optimizes the ONEMAX test function in an expected number of $O(n\lambda/\log \lambda + n \log n)$ fitness evaluations. This matches the runtime shown in [2] for a careful fitness-dependent choice of the mutation rate, which was also shown to be asymptotically optimal among all λ -parallel black-box optimization algorithms. Hence our runtime result indicates that the self-adjusting mechanism developed in this work is able to find very good mutation rates. To the best of our knowledge, this is the first time that a self-adjusting choice of the mutation rate speeds up a mutation-based algorithm on the ONEMAX test function by more than a constant factor.

The main technical challenge in this work is to analyze the quality of the best offspring. In contrast to most previous runtime analyses, where only the asymptotic order of the fitness gain was relevant, we needed a much higher degree of precision as we needed to make statements about in which sub-population the best offspring is, or, in case of multiple best offspring, how they are distributed over the two subpopulations. Note that the quality of the best offspring is not as strongly concentrated around its expectation as, e.g., the average quality.

As a side-result of our analyses, we have observed that using a fixed rate of $r = \ln(\lambda)/2$ gives the bound $O(n/\log \lambda + n \log(n)/\lambda^{1/2})$, which is also asymptotically optimal unless λ is small. However, this setting is far off the usual constant choice of r . It is the first time that a significantly larger mutation rate was shown to be useful in a simple mutation-based algorithm for a simple fitness landscape. Previously, it was only observed that larger mutation rates can be helpful to leave local optima [23].

From this work, a number of open problems arise. A technical challenge is to prove that our algorithm also without the random rate adjustments performs well. This requires an even more precise analysis of the qualities of the offspring in the two subpopulations, for which we currently do not have the methods. From the view-point of understanding the mutation rate for population-based algorithms, two interesting questions are (i) to what extent our observation that larger mutation rates are beneficial for the $(1+\lambda)$ EA on ONEMAX generalizes to other algorithms and problems, and (ii) for which other problems our self-adjusting choice of the mutation rate gives an improvement over the classic choice of $1/n$ or other static choices.

Acknowledgments This work was supported by a public grant as part of the Investissement d’avenir project, reference ANR-11-LABX-0056-LMH, LabEx LMH, and by a grant by the Danish Council for Independent Research (DFF-FNU 4002-00542).

References

- [1] Fawaz Alanazi and Per Kristian Lehre. Runtime analysis of selection hyper-heuristics with classical learning mechanisms. In *Proc. of CEC ’14*, pages 2515–

2523. IEEE, 2014.

- [2] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. Unbiased black-box complexity of parallel search. In *Proc. of PPSN '14*, pages 892–901. Springer, 2014.
- [3] Süntje Böttcher, Benjamin Doerr, and Frank Neumann. Optimal fixed and adaptive mutation rates for the LeadingOnes problem. In *Proc. of PPSN '10*, pages 1–10. Springer, 2010.
- [4] Maxim Buzdalov and Benjamin Doerr. Runtime analysis of the $(1 + (\lambda, \lambda))$ genetic algorithm on random satisfiable 3-CNF formulas. In *Proc. of GECCO '17*, pages 1343–1350. ACM, 2017.
- [5] Stephan Cathabard, Per Kristian Lehre, and Xin Yao. Non-uniform mutation rates for problems with unknown solution lengths. In *Proc. of FOGA '11*, pages 173–180. ACM, 2011.
- [6] Jorge Cervantes and Christopher R. Stephens. Rank based variation operators for genetic algorithms. In *Proc. of GECCO '08*, pages 905–912. ACM, 2008.
- [7] Duc-Cuong Dang and Per Kristian Lehre. Self-adaptation of mutation rates in non-elitist populations. In *Proc. of PPSN '16*, pages 803–813. Springer, 2016.
- [8] Martin Dietzfelbinger, Jonathan E. Rowe, Ingo Wegener, and Philipp Woelfel. Tight bounds for blind search on the integers and the reals. *Combinatorics, Probability & Computing*, 19:711–728, 2010.
- [9] Benjamin Doerr. Analyzing randomized search heuristics: tools from probability theory. In Anne Auger and Benjamin Doerr, editors, *Theory of Randomized Search Heuristics*, pages 1–20. World Scientific Publishing, 2011.
- [10] Benjamin Doerr. Optimal parameter settings for the $(1 + (\lambda, \lambda))$ genetic algorithm. In *Proc. of GECCO '16*, pages 1107–1114. ACM, 2016.
- [11] Benjamin Doerr. An elementary analysis of the probability that a binomial random variable exceeds its expectation. *Probability and Statistics Letters*, 2018. To appear. Preprint *arXiv:1712.00519*.
- [12] Benjamin Doerr and Carola Doerr. Optimal parameter choices through self-adjustment: applying the 1/5-th rule in discrete settings. In *Proc. of GECCO '15*, pages 1335–1342. ACM, 2015.
- [13] Benjamin Doerr and Marvin Künnemann. Optimizing linear functions with the $(1+\lambda)$ evolutionary algorithm – different asymptotic runtimes for different instances. *Theoretical Computer Science*, 561:3–23, 2015.
- [14] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64:673–697, 2012.

- [15] Benjamin Doerr, Carola Doerr, and Franziska Ebel. From black-box complexity to designing new genetic algorithms. *Theoretical Computer Science*, 567:87–104, 2015.
- [16] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Solving problems with unknown solution length at (almost) no extra cost. In *Proc. of GECCO '15*, pages 831–838. ACM, 2015.
- [17] Benjamin Doerr, Carola Doerr, and Timo Kötzing. The right mutation strength for multi-valued decision variables. In *Proc. of GECCO '16*, pages 1115–1122. ACM, 2016.
- [18] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Provably optimal self-adjusting step sizes for multi-valued decision variables. In *Proc. of PPSN '16*, pages 782–791. Springer, 2016.
- [19] Benjamin Doerr, Carola Doerr, and Jing Yang. Optimal parameter choices via precise black-box analysis. In *Proc. of GECCO '16*, pages 1123–1130. ACM, 2016.
- [20] Benjamin Doerr, Carola Doerr, and Jing Yang. k -bit mutation with self-adjusting k outperforms standard bit mutation. In *Proc. of PPSN '16*, pages 824–834. Springer, 2016.
- [21] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Unknown solution length problems with no asymptotically optimal run time. In *Proc. of GECCO '17*, pages 1367–1374. ACM, 2017.
- [22] Benjamin Doerr, Christian Gießen, Carsten Witt, and Jing Yang. The $(1+\lambda)$ evolutionary algorithm with self-adjusting mutation rate. In *Proc. of GECCO '17*, pages 1351–1358. ACM, 2017.
- [23] Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. Fast genetic algorithms. In *Proc. of GECCO '17*, pages 777–784. ACM, 2017.
- [24] Benjamin Doerr, Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. On the runtime analysis of selection hyper-heuristics with adaptive learning periods. In *Proc. of GECCO '18*. ACM, 2018. To appear.
- [25] Benjamin Doerr, Carsten Witt, and Jing Yang. Runtime analysis for self-adaptive mutation rates. In *Proc. GECCO '18*. ACM, 2018. To appear.
- [26] Agoston Endre Eiben, Robert Hinterding, and Zbigniew Michalewicz. Parameter control in evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 3:124–141, 1999.
- [27] Christian Gießen and Carsten Witt. The interplay of population size and mutation probability in the $(1+\lambda)$ EA on OneMax. *Algorithmica*, 78:587–609, 2017.
- [28] Oliver Giel and Ingo Wegener. Evolutionary algorithms and the maximum matching problem. In *Proc. of STACS '03*, pages 415–426. Springer, 2003.

- [29] Thomas Jansen and Ingo Wegener. On the analysis of a dynamic evolutionary algorithm. *Journal of Discrete Algorithms*, 4:181–199, 2006.
- [30] Thomas Jansen, Kenneth A. De Jong, and Ingo Wegener. On the choice of the offspring population size in evolutionary algorithms. *Evolutionary Computation*, 13:413–440, 2005.
- [31] Daniel Johannsen. *Random combinatorial structures and randomized search heuristics*. PhD thesis, Saarland University, 2010.
- [32] Rob Kaas and Jan M. Buhrman. Mean, median and mode in binomial distributions. *Statistica Neerlandica*, 34:13–18, 1980.
- [33] Timo Kötzing, Andrei Lissovoi, and Carsten Witt. (1+1) EA on generalized dynamic OneMax. In *Proc. of FOGA '15*, pages 40–51. ACM, 2015.
- [34] Jörg Lässig and Dirk Sudholt. Adaptive population models for offspring populations and parallel evolutionary algorithms. In *Proc. of FOGA '11*, pages 181–192. ACM, 2011.
- [35] Per Kristian Lehre and Ender Özcan. A runtime analysis of simple hyper-heuristics: To mix or not to mix operators. In *Proc. of FOGA '13*, pages 97–104. ACM, 2013.
- [36] Per Kristian Lehre and Carsten Witt. Concentrated hitting times of randomized search heuristics with variable drift. In *Proc. of ISAAC '14*, pages 686–697. Springer, 2014.
- [37] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. On the runtime analysis of generalised selection hyper-heuristics for pseudo-Boolean optimisation. In *Proc. of GECCO '17*, pages 849–856. ACM, 2017.
- [38] Boris Mitavskiy, Jonathan E. Rowe, and Chris Cannings. Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *International Journal of Intelligent Computing and Cybernetics*, 2:243–284, 2009.
- [39] Frank Neumann and Ingo Wegener. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Theoretical Computer Science*, 378:32–40, 2007.
- [40] Pietro Simone Oliveto, Per Kristian Lehre, and Frank Neumann. Theoretical analysis of rank-based mutation - combining exploration and exploitation. In *Proc. of CEC '09*, pages 1455–1462. IEEE, 2009.
- [41] Chao Qian, Ke Tang, and Zhi-Hua Zhou. Selection hyper-heuristics can provably be helpful in evolutionary multi-objective optimization. In *Proc. of PPSN '16*, pages 835–846. Springer, 2016.

- [42] Herbert Robbins. A remark on Stirling's formula. *The American Mathematical Monthly*, 62:26–29, 1955.
- [43] Ingo Wegener. Simulated annealing beats Metropolis in combinatorial optimization. In *Proc. of ICALP '05*, pages 589–601. Springer, 2005.
- [44] Christine Zarges. Rigorous runtime analysis of inversely fitness proportional mutation rates. In *Proc. of PPSN '08*, pages 112–122. Springer, 2008.
- [45] Christine Zarges. On the utility of the population size for inversely fitness proportional mutation rates. In *Proc. of FOGA '09*, pages 39–46. ACM, 2009.