Self-Adjusting Offspring Population Sizes Outperform Fixed Parameters on the Cliff Function*

Mario Alejandro Hevia Fajardo
Department of Computer Science
University of Sheffield, Sheffield, UK

Dirk Sudholt
Chair of Algorithms for Intelligent Systems
University of Passau, Passau, Germany

ABSTRACT

In the discrete domain, self-adjusting parameters of evolutionary algorithms (EAs) has emerged as a fruitful research area with many runtime analyses showing that self-adjusting parameters can outperform the best fixed parameters. Most existing runtime analyses focus on elitist EAs on simple problems, for which moderate performance gains were shown. Here we consider a much more challenging scenario: the multimodal function Cliff, defined as an example where a $(1, \lambda)$ EA is effective, and for which the best known upper runtime bound for standard EAs is $O(n^{25})$.

We prove that a $(1, \lambda)$ EA self-adjusting the offspring population size λ using success-based rules optimises Cliff in $O(n)$ expected generations and $O(n \log n)$ expected evaluations. Along the way, we prove tight upper and lower bounds on the runtime for fixed λ (up to a logarithmic factor) and identify the runtime for the best fixed λ as n^η for $\eta \approx 3.9767$ (up to sub-polynomial factors). Hence, the self-adjusting $(1, \lambda)$ EA outperforms the best fixed parameter by a factor of at least $n^{2.9767}$ (up to sub-polynomial factors).

CCS CONCEPTS

• **Theory of computation** → *Theory of randomized search heuristics.*

KEYWORDS

Parameter control, runtime analysis, non-elitism, drift analysis, multimodal optimisation

ACM Reference Format:
Mario Alejandro Hevia Fajardo and Dirk Sudholt. 2021. Self-Adjusting Offspring Population Sizes Outperform Fixed Parameters on the Cliff Function. In *Foundations of Genetic Algorithms XVI (FOGA '21), September 6–8, 2021, Virtual Event, Austria.* ACM, New York, NY, USA, 15 pages. https://doi.org/10.1145/3450218.3477306

1 INTRODUCTION

Evolutionary algorithms (EAs) as well as other Randomised Search Heuristics are used to solve a wide range of problems in part owing to their ease of implementation and their effectiveness in problems with little a priori knowledge. When applying an EA, a crucial task is to determine suitable parameters for the problem in hand such as mutation rate, crossover probability, population sizes, among others. In fact, it is well understood that the efficiency of the algorithms may depend drastically on their parameters [9], even to the point where small changes to the parameters can increase the runtime from polynomial to exponential [22, 33]. To make matters worse, the optimal parameters may change during the optimisation problem, hence any static parameter choice may be sub-optimal [9].

Parameter control mechanisms aim to solve this problem by using dynamic parameter settings that adjust to the current state of the optimisation identifying and tracking the optimal parameter settings. Doerr and Doerr [9] have classified them into several types; here we focus on *success-based* (also called *self-adjusting*) parameter control mechanisms for their simplicity. In continuous optimisation, parameter control mechanisms have been used as standard for several decades because it is crucial to ensure convergence to the optimum. In contrast, in the discrete domain parameter control had not been as widely used in the past. In recent years it has become more common in part owing to runtime analyses showing that parameter control mechanisms can outperform the best static parameter settings, see the survey by Doerr and Doerr [9].

Here we highlight some examples relevant to this work where parameter control mechansims have been proposed, along with proven performance guarantees. Böttcher et al. [3] considered the test function $\text{LeadingOnes}(x) := \sum_{i=1}^{n} \prod_{j=1}^{i} x_i$ that counts the number of consecutive ones at the start of the bit string x. They showed that fitness-dependent mutation rates can improve the performance of the $(1 + 1)$ EA on LeadingOnes by a constant factor. Badkobeh et al. [1] presented an adaptive strategy for the mutation rate in the $(1+\lambda)$ EA that, for all values of λ, leads to provably optimal performance on OneMax. Lässig and Sudholt [20] presented adaptive schemes for choosing the offspring population size in $(1+\lambda)$ EAs and the number of islands in an island model. Doerr et al. [10] proposed a fitness-dependent offspring population size of $\lambda = \sqrt{n/(n - f(x))}$ for the $(1 + (\lambda, \lambda))$ GA showing that it optimises OneMax in $O(n)$ evaluations which is asymptotically faster than any static parameter choice and proposed a self-adjusting mechanism based on the one-fifth rule that tracked the optimal parameter in experiments. Later, Doerr and Doerr [8] proved that the self-adjusting mechanism in the $(1 + (\lambda, \lambda))$ GA has the same asymptotic runtime on OneMax as the fitness-dependent mechanism. Hevia Fajardo and Sudholt [16] studied modifications to the self-adjusting mechanism in the $(1 + (\lambda, \lambda))$ GA on Jump functions, showing that they can perform nearly as well as the $(1 + 1)$ EA with the optimal mutation rate. Doerr et al. [12] showed that a success-based parameter control mechanism is able to identify and track the optimal mutation rate in the $(1+\lambda)$ EA on OneMax, matching the

*This research has been supported by CONACYT under the grant no. 739621 and registration no. 843375.

FOGA '21, September 6–8, 2021, Virtual Event, Austria

ACM ISBN 978-1-4503-8352-3/21/09…$15.00
https://doi.org/10.1145/3450218.3477306

performance of the best known fitness-dependent parameter [1]. Mambrini and Sudholt [26] adapted the migration interval in island models and showed that adaptation can reduce the communication effort beyond the best possible fixed parameter. Doerr et al. [11] proved that a success-based parameter control mechanism based on the one-fifth rule is able to achieve an asymptotically optimal runtime on LeadingOnes. Lissovoi et al. [25] proposed a Generalised Random Gradient Hyper-Heuristic that uses a learning period of τ steps that can learn to adapt the neighbourhood size of Randomised Local Search optimally during the run on LeadingOnes. They proved that it has the best possible runtime achievable by any algorithm that uses the same low level heuristics. This result required the correct selection of the learning period τ; this was later solved using a self-adjusting mechanism adapting the learning period having an optimal asymptotic expected runtime on LeadingOnes [13], OneMax and Ridge [24]. Rajabi and Witt [30] used a self-adjusting asymmetric mutation that can provide a constant-factor speed-up on OneMax over asymmetric mutations [19] and obtained the same asymptotic performance on the generalised OneMax function. Rajabi and Witt [31] proposed a stagnation detection mechanism that raises the mutation rate when the algorithm is likely to have encountered a local optima. The mechanism can be added to any existing EA; when added to the $(1+1)$ EA, the SD–$(1+1)$ EA has the same asymptotic runtime on Jump as the optimal parameter setting. In a follow up study, Rajabi and Witt [32] added the stagnation detection mechanism to the RLS obtaining a constant factor speed-up from the SD–$(1+1)$ EA.

Most of the above analyses concern elitist EAs; there are very few studies of parameter control mechanisms for EAs using non-elitist selection. The first runtime analysis to show an asymptotic speedup for parameter control mechanisms in non-elitist EAs was presented by Dang and Lehre [6], showing that for a tailored multimodal function a self-adaptive EA is able to adjust the mutation rate, leading to exponential speedup over EAs with static mutation rates. Case and Lehre [4] showed a speedup for a self-adaptive (μ, λ) EA on the LeadingOnes problem with unknown solution length. Similarly Doerr et al. [14] proved that a self-adaptive mechanism for the mutation rate in the (1,λ) EA with a sufficiently large λ has the same asymptotic expected runtime on OneMax as in [1]. Lissovoi et al. [23] proposed a hyper-heuristic that chooses between elitist and non-elitist heuristics that achieves the best known expected runtime for general purpose randomised search heuristics on the problem class Cliff_d. The present authors proved that a self-adjusting mechanism based on the one-fifth rule is able to find and maintain suitable parameter values of λ for the $(1, \lambda)$ EA leading to an asymptotically optimal runtime on OneMax [17].

We argue that providing runtime analyses of parameter control mechanisms for non-elitist EAs is an important direction for research. One reason is that non-elitist algorithms are frequently used in practice and understanding the dynamics of non-elitist EAs is vital to narrow the gap between theory and practice. Furthermore, and perhaps more importantly, many existing theoretical studies concern fairly easy problems on which algorithms with static parameters already run efficiently, and so the performance gains obtained through parameter control are often moderate at best. More research effort should be devoted to considering non-elitist EAs on more challenging problems since the potential for performance improvements is much greater.

1.1 Our contribution

In this work we provide an example of significant performance improvements through parameter control for a multimodal problem. We study the $(1, \lambda)$ EA on the multimodal problem Cliff [18] with a mechanism to self-adjust the choice of the offspring population size λ. The function Cliff (formally defined in Section 2) typically requires EAs to jump down a "cliff" by accepting a huge loss in fitness, and then to climb up a slope towards the global optimum. Elitist EAs are unable to accept this fitness loss and typically need to jump directly to the global optimum (Theorem 8 in [29] gives a tight bound for the (1+1) EA). The $(1, \lambda)$ EA is able to jump down the cliff if and only if *all* offspring are generated at the bottom of the cliff. Hence, the smaller the population size, the higher the probability of jumping down the cliff. However, the $(1, \lambda)$ EA also needs to be able to climb up a OneMax-like slope towards the cliff and towards the global optimum. The offspring population size λ must be large enough to enable hill climbing. Rowe and Sudholt [33] showed that there is a phase transition on OneMax at $\log_{\frac{e}{e-1}} n$. More specifically, $\lambda \geq \log_{\frac{e}{e-1}} n$ is sufficient to optimise OneMax efficiently, in expected $O(n \log(n) + \lambda n)$ evaluations, but all $\lambda \leq (1-\varepsilon) \log_{\frac{e}{e-1}} n$ lead to exponential optimisation times. Every fixed value of λ must strike a delicate balance to enable jumps down the cliff and at the same time being able to hill climb. Jägersküpper and Storch [18] showed a bound of $O(e^{5\lambda})$ for $\lambda \geq 5 \ln n$, which gives an upper bound of $O(n^{25})$ for $\lambda = 5 \ln n$. To our knowledge, this is the best known upper bound for the runtime of the $(1, \lambda)$ EA to date. A lower bound of $\min\{n^{n/4}, e^{\lambda/4}\}/3$ for all λ was shown in [18]. Comparing the term $e^{\lambda/4} \approx 1.284^{\lambda}$ to the upper bound of order $e^{5\lambda} \approx 148.413^{\lambda}$, the exponents (to the base of e) differ by a factor of 20 and the bases to the power of λ differ by a factor larger than 115. This leaves a large polynomial gap between upper and lower bounds for $\lambda = \Theta(\log n)$.

We refine results from [18] and show that the runtime is $\Omega(\xi^{\lambda})$ and $O(\xi^{\lambda} \log n)$ for a base of $\xi \approx 6.196878$, for reasonable values of λ. For the best fixed λ, we show that the expected runtime is $O(n^{\eta} \log n)$ for a constant $\eta \approx 3.976770136$, and that it grows faster than any polynomial of degree less than η.

More importantly, we then show that parameter control is highly beneficial in this scenario. We present a self-adjusting $(1, \lambda)$ EA that self-adjusts λ and prove that it is able to optimise Cliff in $O(n)$ expected generations and $O(n \log n)$ expected evaluations. This is faster than any static parameter choice by a factor of $\Omega(n^{2.9767}/\log n)$ and it is asymptotically the best possible runtime for any unary unbiased black-box algorithm [21].

We remark that this is the first bound of $O(n \log n)$ for a standard evolutionary algorithm on Cliff; previously, $O(n \log n)$ bounds were only achieved by using additional mechanisms such as ageing [5] and hyper-heuristics [23].

Our analysis builds on our previous work [17] that analysed a similar algorithm on the simple OneMax function. The considered algorithm works using a variant of the famous one-fifth success rule: in a generation in which the current fitness is increased (a success), λ is decreased by a factor of F, where F is a parameter. In an unsuccessful generation, λ is increased by a factor of $F^{1/s}$. The parameter s is called the *success rate* and it implies that, if on average one in $s+1$ generations is successful, the current value of λ is maintained (as we have one success and s unsuccessful generations and so $\lambda_{t+s+1} = \lambda_t \cdot (F^{1/s})^s \cdot 1/F = \lambda_t$).

In [17] we showed that with a success rate of $0 < s < \frac{e-1}{e}$ the self-adjusting $(1,\lambda)$ EA optimises OneMax in $O(n)$ expected generations and $O(n \log n)$ expected evaluations. We also showed that larger parameters lead to exponential times. Hence, we use the same restriction $0 < s < \frac{e-1}{e}$ in this work.

To make the self-adjusting $(1,\lambda)$ EA work in multimodal optimisation, we need to tackle an important challenge that requires a redesign of the self-adjusting $(1,\lambda)$ EA in [17]. Success-based parameter control mechanisms can be problematic on multimodal problems because once a local optimum is reached the success of previous generations does not give a good indication of what parameters are needed to escape the local optimum. Strategies have been proposed and analysed to solve this problem, showing a good performance. Some examples from other contexts are: systematically increasing the mutation rate once the neighbourhood has been searched in order to increase the radius of exploration [31, 32] or resetting the parameter once it has reached a certain value, allowing the algorithm to cycle through the possible parameter values [16].

We enhance the self-adjusting $(1,\lambda)$ EA from [17] with a resetting mechanism: whenever λ exceeds a predefined maximum of $\lambda_{\max}$, it is reset to $\lambda = 1$. When such a reset happens at the top of the cliff, there is a good probability of jumping down the cliff.

Note that the resetting mechanism may have unwanted side effects: resets may happen in difficult fitness levels, for instance on Cliff resets may happen when climbing up the slope to the global optimum and successes become rare. Hence we need to choose $\lambda_{\max}$ sufficiently large to mitigate this risk and enhance the analysis of the self-adjusting $(1,\lambda)$ EA on OneMax with additional arguments.

1.2 Outline

The paper is structured as follows. Section 2 gives necessary definitions and bounds transition probabilities.Since the current best known upper and lower bounds for the $(1,\lambda)$ EA with static λ on Cliff from [18] are far from tight, we first provide refined upper and lower bounds that are tight up to a logarithmic factor in Section 3. Our upper and lower bounds are then used to determine the degree η of the polynomial term in the runtime for the best fixed value of λ.

In Section 4 we show that despite the possibility of resetting λ near the optimum, the self-adjusting $(1,\lambda)$ EA is able to optimise Cliff in expected $O(n)$ generations and $O(n \log n)$ expected evaluations. We divide the optimisation in several phases showing that the algorithm is able to hill-climb effectively to both the local and global optimum. When it encounters the local optimum, λ typically increases to its maximum, increasing its selection pressure and behaving like an elitist EA. But then λ is reset to 1, reducing the selection pressure of the algorithm allowing it to escape local optima. The behaviour in the local optima is similar to the behaviour of the meta-heuristic from [23] where the algorithm changes from elitism to non-elitism to jump out of local optimum and later on it behaves roughly as a purely elitist algorithm. Some proofs are omitted due to space limitations.

2 PRELIMINARIES

We present a runtime analysis of the self-adjusting $(1,\lambda)$ EA[1] on the n-dimensional pseudo-Boolean benchmark function Cliff.

2.1 The Cliff Function Class

We write $|x|_1$ to denote the number of one-bits in the bit string x. The Cliff benchmark function was first proposed by Jägersküpper and Storch [18] as an example where non-elitism helps the optimisation process. The function is designed to guide hill-climbing algorithms towards a local optimum (*cliff*) where the global optimum is the only other search point with a higher fitness value but it is at a linear distance from the local optimum. An elitist algorithm then needs to perform a large jump to find the global optimum; a non-elitist algorithm instead can escape the local optimum by performing a fitness-decreasing step that leads to another slope guiding to the global optimum. An instance of Cliff with $n = 90$ is shown in Figure 1. We define the Cliff function as follows:

$$\textsc{Cliff}(x) := \begin{cases} |x|_1 & \text{if } |x|_1 \leq 2n/3, \\ |x|_1 - n/3 + 1/2 & \text{otherwise.} \end{cases}$$

Following Lissovoi et al. [23] our definition of Cliff differs from [18] in that the cliff is located at $2n/3$ one-bits instead of $2n/3 - 1$. We choose this definition because it resembles the definition of the Jump class functions and improves the readability of the runtime analysis. Throughout the analysis we assume n is divisible by 3. Also following [23], we say that all search points x with $|x|_1 \leq 2n/3$ form the *first slope* and all other search points form the *second slope*.

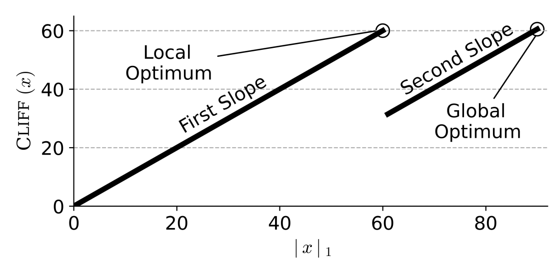

Figure 1: Cliff(x) with $n = 90$

The Cliff function was used as a benchmark in several other works, including studies of the Strong Selection Weak Mutation (SSWM) model of evolution [29], artificial immune systems [5] and hyper-heuristics [23].

2.2 The Self-Adjusting $(1,\lambda)$ EA

The self-adjusting $(1,\lambda)$ EA was first proposed in [17] and studied on OneMax. The algorithm uses the generalised success based rule (one-$(s+1)$-th success rule) to adjust the offspring population size λ. If the fittest offspring y is better than the parent x, the offspring population size is divided by the *update strength* F, and multiplied by $F^{1/s}$ otherwise, with s being the *success rate*.

[1] named $(1, \{F^{1/s}\lambda, \lambda/F\})$ EA in [17]

In this work we consider a variation of the self-adjusting $(1, \lambda)$ EA with a resetting mechanism for λ (see Algorithm 1) where λ is reset to 1 if $\lambda = \lambda_{\max}$ and there is an unsuccessful generation. This strategy has also been successfully applied to the self-adjusting mechanism of the $(1 + (\lambda, \lambda))$ GA in [16]. In addition, the strategy is similar to the stagnation detection from [31, 32] in that if $\lambda_{\max}$ is set appropriately when $\lambda = \lambda_{\max}$ the algorithm is likely to be in a local optimum and the behaviour of the algorithm changes. In this case when λ is large enough the algorithm maintains its fitness with high probability, but when λ is reset to 1 the behaviour changes to a pure random walk allowing the algorithm to escape local optima.

Algorithm 1: Self-adjusting $(1, \{F^{1/s}\lambda, \lambda/F\})$ EA resetting λ.

Initialization: Choose $x \in \{0, 1\}^n$ uniformly at random (u.a.r.) and set $\lambda := 1$;
Optimization: **for** $t \in \{1, 2, \dots\}$ **do**
  Mutation: **for** $i \in \{1, \dots, \lfloor \lambda \rceil\}$ **do**
    $y'_i \in \{0, 1\}^n \leftarrow$ standard bit mutation(x);
  Selection: Choose $y \in \{y'_1, \dots, y'_{\lfloor \lambda \rceil}\}$ with $f(y) = \max\{f(y'_1), \dots, f(y'_{\lfloor \lambda \rceil})\}$ u.a.r.;
  Update: $x \leftarrow y$;
  if $f(y) > f(x)$ **then** $\lambda \leftarrow \max\{\lambda/F, 1\}$;
  if $f(y) \leq f(x) \wedge \lambda = \lambda_{\max}$ **then** $\lambda \leftarrow 1$;
  if $f(y) \leq f(x) \wedge \lambda \neq \lambda_{\max}$ **then** $\lambda \leftarrow \min\{\lambda F^{1/s}, \lambda_{\max}\}$;

Note that we regard λ to be a real value, so that changes by factors of $1/F$ or $F^{1/s}$ happen on a continuous scale. Following Doerr and Doerr [8] and Hevia Fajardo and Sudholt [17], we assume that, whenever an integer value of λ is required, λ is rounded to a nearest integer. For the sake of readability, we often write λ as a real value even when an integer is required.

In our analysis we define $X_0, X_1, \dots$ as the sequence of states of the algorithm, where $X_t = (x_t, \lambda_t)$ describes the current search point x_t and the offspring population size λ_t at generation t. We often omit the subscripts t when the context is obvious.

2.3 Transition Probabilities

We now define and estimate transition probabilities that apply to all $(1, \lambda)$ EA variants (with or without self-adjustment) in the context of OneMax and Cliff.

Definition 2.1. For all $\lambda \in \mathbb{N}$ we state the following definitions from [17] in the context of OneMax:

$$p^{+}_{i,\lambda} = \Pr\left(|x_{t+1}|_1 > i \mid |x_t|_1 = i\right)$$
$$p^{-}_{i,\lambda} = \Pr\left(|x_{t+1}|_1 < i \mid |x_t|_1 = i\right)$$
$$\Delta^{-}_{i,\lambda} = \mathrm{E}\left(i - |x_{t+1}|_1 \mid |x_t|_1 = i \text{ and } |x_{t+1}|_1 < i\right)$$

The following probabilities and expectations are tailored to the Cliff function. This includes probabilities for jumping down the cliff ($p^{\downarrow}_{i,\lambda}$), that is, jumping from the first slope to the second slope, jumping back up the cliff ($p^{\uparrow}_{i,\lambda}$), that is, jumping from the second slope to the first slope, increasing the fitness without jumping back up the cliff ($p^{\rightarrow}_{i,\lambda}$), and decreasing the fitness without jumping down the cliff ($p^{\leftarrow}_{i,\lambda}$).

Definition 2.2. For all $\lambda \in \mathbb{N}$ we define:

$$p^{\downarrow}_{i,\lambda} = \begin{cases} \Pr\left(|x_{t+1}|_1 > 2n/3 \mid 2n/3 \geq |x_t|_1 = i\right) & \text{if } i \leq 2n/3 \\ 0 & \text{otherwise.} \end{cases}$$

$$p^{\uparrow}_{i,\lambda} = \begin{cases} \Pr\left(|x_{t+1}|_1 \leq 2n/3 \mid 2n/3 < |x_t|_1 = i\right) & \text{if } i > 2n/3 \\ 0 & \text{otherwise.} \end{cases}$$

$$p^{\rightarrow}_{i,\lambda} = \begin{cases} \Pr\left(i < |x_{t+1}|_1 \leq 2n/3 \mid |x_t|_1 = i\right) & \text{if } |x|_1 \leq 2n/3, \\ \Pr\left(i < |x_{t+1}|_1 \mid |x_t|_1 = i\right) & \text{otherwise.} \end{cases}$$

$$p^{\leftarrow}_{i,\lambda} = \begin{cases} \Pr\left(|x_{t+1}|_1 < i \mid |x_t|_1 = i\right) & \text{if } i \leq 2n/3, \\ \Pr\left(2n/3 < |x_{t+1}|_1 < i \mid |x_t|_1 = i\right) & \text{otherwise.} \end{cases}$$

$$\Delta^{\leftarrow}_{i,\lambda} = \begin{cases} \mathrm{E}\left(i - |x_{t+1}|_1 \mid |x_t|_1 = i, |x_{t+1}|_1 < i\right) & \text{if } i \leq 2n/3, \\ \mathrm{E}\left(i - |x_{t+1}|_1 \mid |x_t|_1 = i, 2n/3 < |x_{t+1}|_1 < i\right) & \text{otherwise.} \end{cases}$$

Finally, for $i > 2n/3$,

$$\Delta^{\uparrow}_{i,\lambda} = \mathrm{E}\left(i - |x_{t+1}|_1 \mid 2n/3 < |x_t|_1 = i \text{ and } |x_{t+1}|_1 \leq 2n/3\right).$$

The events underlying the probabilities from Definition 2.2 are mutually disjoint. They relate to the probabilities defined for OneMax in [17] as follows:

$$p^{+}_{i,\lambda} = p^{\rightarrow}_{i,\lambda} + p^{\downarrow}_{i,\lambda} \tag{1}$$
$$p^{-}_{i,\lambda} = p^{\leftarrow}_{i,\lambda} + p^{\uparrow}_{i,\lambda} \tag{2}$$

The following lemma collects bounds on the above transition probabilities that will be used throughout the remainder.

Lemma 2.3. *For any $(1, \lambda)$ EA on Cliff, the quantities from Definition 2.2 are bounded as follows:*

For all $i \in \{1, \dots, n\}$,

$$p^{\leftarrow}_{i,\lambda} \leq p^{-}_{i,\lambda} \leq \left(\frac{e-1}{e}\right)^{\lambda} \tag{3}$$
$$\Delta^{\leftarrow}_{i,\lambda} \leq \Delta^{-}_{i,\lambda} \leq \frac{e}{e-1}. \tag{4}$$

For all $i \in \{2n/3, \dots, n\}$, letting $d := i - 2n/3$,

$$p^{\uparrow}_{i,\lambda} \leq \frac{\lambda (i/n)^d}{d!} \tag{5}$$
$$\Delta^{\uparrow}_{i,\lambda} \leq d + 1. \tag{6}$$

Finally,

$$p^{\rightarrow}_{i,\lambda} \geq \begin{cases} 1 - \left(1 - \frac{1}{3e}\right)^{\lambda} & \text{if } i < 2n/3, \\ 1 - \left(1 - \frac{n-i}{en}\right)^{\lambda} - p^{\uparrow}_{i,\lambda} & \text{if } i > 2n/3. \end{cases} \tag{7}$$

Proof. The first inequality in (3), $p^{\leftarrow}_{i,\lambda} \leq p^{-}_{i,\lambda}$, follows from (1). The stated upper bound on $p^{-}_{i,\lambda}$ was shown in [17].

We have $\Delta^{\leftarrow}_{i,\lambda} \leq \Delta^{-}_{i,\lambda}$ since for $i \leq 2n/3$, $\Delta^{\leftarrow}_{i,\lambda} = \Delta^{-}_{i,\lambda}$ by definition of $\Delta^{\leftarrow}_{i,\lambda}$ and otherwise $\Delta^{\leftarrow}_{i,\lambda}$ is capped as only target search points with more than $2n/3$ ones are considered. The second inequality in (4) was shown in [17].

To bound $p^{\uparrow}_{i,\lambda}$, we argue that a necessary requirement for creating an offspring with at most $2n/3$ ones is that d one-bits flip. There are $\binom{i}{d}$ ways for choosing d one-bits and the probability that d specific bits are flipped is $(1/n)^d$. Thus,

$$p^{\uparrow}_{i,1} \leq \binom{i}{d}\left(\frac{1}{n}\right)^d \leq \frac{i^d}{d!}\cdot\left(\frac{1}{n}\right)^d \leq \frac{(i/n)^d}{d!}.$$

Using the union bound over all offspring, we get

$$p^{\uparrow}_{i,\lambda} \leq \frac{\lambda(i/n)^d}{d!}.$$

To bound $\Delta^{\uparrow}_{i,\lambda}$ we bound the number of one-bits flipped by the number of bit-flips during a generation conditional on flipping d bits. Let B denote the random number of flipping bits in a standard bit mutation with mutation probability $1/n$, then using Lemma 1.7.3 from [7] we get $\mathrm{E}(B \mid B \geq d) \leq d + \mathrm{E}(B) = d + 1$. Increasing the offspring population size does not increase $\Delta^{\uparrow}_{i,\lambda}$ because if multiple offspring jump up the cliff, the algorithm will transition to an offspring with a maximum number of ones on the first slope.

To bound $p^{\rightarrow}_{i,\lambda}$ we argue that, for all $i < 2n/3$, if there is an offspring with $i+1$ ones then the number of ones increases. For $|x_t|_1 < 2n/3$ the probability that an offspring flips only one 0-bit is

$$p^{+}_{i,1} \geq \frac{n-i}{n}\left(1-\frac{1}{n}\right)^{n-1} \geq \frac{n-i}{en}. \tag{8}$$

The probability that any of the λ offspring flips only one 0-bit is

$$1-\left(1-p^{+}_{i,1}\right)^{\lambda} \geq 1-\left(1-\frac{n-i}{en}\right)^{\lambda} \geq 1-\left(1-\frac{1}{3e}\right)^{\lambda}.$$

This proves the claimed lower bound on $p^{\rightarrow}_{i,\lambda}$ if $|x_t|_1 < 2n/3$. If $|x_t|_1 > 2n/3$ then there is an offspring with $i+1$ ones with probability

$$1-\left(1-\frac{n-i}{en}\right)^{\lambda}.$$

In this case either the number of ones increases or the algorithm goes back up the cliff. Hence,

$$p^{\rightarrow}_{i,\lambda} \geq 1-\left(1-\frac{n-i}{en}\right)^{\lambda} - p^{\uparrow}_{i,\lambda}. \qquad \square$$

We also give a lower bound for a mutation creating an offspring from the top of the cliff (that is, a parent with $2n/3$ ones) that increases the number of ones by at least $\frac{c\log\log n}{\log\log\log n}$, for an arbitrary constant $c > 0$.

Lemma 2.4. *For every constant $c > 0$ and all $n \geq 2^{2^{2^{3c}}}$ the probability that a standard bit mutation of a search point with $2n/3$ ones yields an offspring with at least $2n/3 + \frac{c\log\log n}{\log\log\log n}$ ones is at least $(\log n)^{-c}$.*

Proof. Let $\kappa := \frac{c\log\log n}{\log\log\log n}$, then a sufficient condition for the sought event is that κ 0-bits flip. Since there are $\binom{n/3}{\kappa}$ ways to choose these flipping bits, the probability is at least

$$\binom{n/3}{\kappa}\left(\frac{1}{n}\right)^{\kappa} \geq \left(\frac{n/3}{\kappa}\right)^{\kappa}\left(\frac{1}{n}\right)^{\kappa} = \left(\frac{1}{3\kappa}\right)^{\kappa} = (3\kappa)^{-\kappa}.$$

Plugging in κ, we get

$$(3\kappa)^{-\kappa} = \left(\frac{3c\log\log n}{\log\log\log n}\right)^{-\frac{c\log\log n}{\log\log\log n}} = 2^{-\frac{c\log\log n}{\log\log\log n}\cdot\log\left(\frac{3c\log\log n}{\log\log\log n}\right)}.$$

By assumption on n, we have $\log\log\log n \geq 3c$, thus $\frac{3c\log\log n}{\log\log\log n} \leq \log\log n$ and we bound the sought probability by

$$2^{-\frac{c\log\log n}{\log\log\log n}\cdot\log\log\log n} = 2^{-c\log\log n} = (\log n)^{-c}. \qquad \square$$

In the remainder, we sometimes tacitly use the following argument. If in an iterative process there is an event that happens independently in each step with probability at most p, the probability that the event happens during a phase of T steps, T a random variable with $\mathrm{E}(T) < \infty$, is at most

$$\sum_t \Pr(T=t)\cdot tp = p\cdot\mathrm{E}(T). \tag{9}$$

3 STATIC PARAMETER SETTINGS

We first consider the performance of the $(1,\lambda)$ EA with a static choice of λ. The main result in this section is the following theorem that gives upper and lower bounds for the expected optimisation time of the $(1,\lambda)$ EA on Cliff.

Theorem 3.1. *The expected optimisation time $E(T)$ of the $(1,\lambda)$ EA with static λ on Cliff is*

$$E(T) = \Omega\left(\xi^{\lambda}\right) \quad \text{if } \lambda = O(n) \text{ and}$$
$$E(T) = O\left(\xi^{\lambda}\cdot\log n\right) \quad \text{if } \log_{\frac{e}{e-1}} n \leq \lambda = O(\log n),$$

where $\xi^{-1} := \frac{1}{e}\sum_{a=0}^{\infty}\sum_{b=a+1}^{\infty}\left(\frac{2}{3}\right)^a\left(\frac{1}{3}\right)^b\frac{1}{a!b!} \approx 0.1613715804$ and thus $\xi \approx 6.196878$.

The lower bound is exponential in λ for all $\lambda = O(n)$. The constant ξ^{-1} roughly represents the probability of increasing the number of ones in a mutation of a parent at the top of the cliff, i. e. a parent with $2n/3$ ones. For $\lambda \leq (1-\varepsilon)\log_{\frac{e}{e-1}} n$, that is, if the offspring population size is too small to allow for hill climbing on OneMax, a much stronger lower bound of $2^{cn^{\varepsilon/2}}$, for some constant $c > 0$, was shown in [33] for all functions with a unique optimum. In the (arguably more interesting) parameter regime $\lambda > (1-\varepsilon)\log_{\frac{e}{e-1}} n$ our result improves upon the only other lower bound we are aware of: [18] showed a lower bound for all λ of $\min\{n^{n/4}, e^{\lambda/4}\}/3$. The term $e^{\lambda/4}$ is roughly 1.284^{λ} and hence considerably lower than our lower bound of 6.196^{λ}. In this parameter regime our lower bound is nearly tight: for the best known values of λ for OneMax [33], the upper bounds only differ from the lower bound by a logarithmic factor.

To prove Theorem 3.1, we first show that, for all search points with at most $3n/4$ ones, improvements are found easily if $\lambda \geq \log_{\frac{e}{e-1}} n$. Note that the considered set of search points includes search points on the first slope as well as search points on the second slope at a linear distance past the top of the cliff. We will see that, once a search point with at least $3n/4$ ones has been reached, the algorithm will not jump back up the cliff, with high probability. The choice of the constant $3/4$ is somewhat arbitrary; we could have chosen any other constant in $(2/3, 1)$.

Lemma 3.2. *Consider the* $(1, \lambda)$ *EA with* $\log_{\frac{e}{e-1}} n \le \lambda = O(\log n)$ *and a current search point* x_t *with* $|x_t|_1 \le 3n/4$. *Then the following statements hold.*

(1) *The probability of creating an offspring with* $|x_t|_1 + 1$ *ones is at least* $1 - n^{-0.2}$.
(2) *For all* x_t *with* $|x_t|_1 \in [0, 3n/4] \setminus [2n/3, 2n/3+\log n]$, *the drift in the number of ones is* $E(|x_{t+1}|_1 - |x_t|_1 \mid x_t) \ge 1 - o(1)$.
(3) *For every* $\kappa \in [0, n/12]$, *the expected time until a search point with exactly* $2n/3$ *ones or at least* $2n/3 + \kappa$ *ones is reached is* $O(n)$.

Proof. The probability that one fixed offspring has more than $|x|_1$ ones is at least $(n/4)/(en) = 1/(4e)$. The probability that there is an offspring that increases the number of ones is at least

$$1 - \left(1 - \frac{1}{4e}\right)^{\log_{\frac{e}{e-1}} n} \ge 1 - \left(\frac{4e}{4e-1}\right)^{-\log_{\frac{4e}{4e-1}}(n)/\log_{\frac{4e}{4e-1}}\left(\frac{e}{e-1}\right)}$$
$$= 1 - n^{-1/\log_{\frac{4e}{4e-1}}\left(\frac{e}{e-1}\right)} \ge 1 - n^{-0.2}.$$

For the second statement, let A^{+1} denote the event from the first statement, that is, creating an offspring with $|x_t|_1 + 1$ ones and let $A^{\uparrow}$ be the event that $|x_t|_1 > 2n/3$ and $|x_{t+1}|_1 \le 2n/3$. By the law of total probability, abbreviating $\Delta := (|x_{t+1}|_1 - |x_t|_1 \mid x_t)$,

$$\begin{aligned} \mathrm{E}(\Delta) = \mathrm{E}\left(\Delta \mid A^{+1}, \overline{A^{\uparrow}}\right) \cdot \Pr\left(A^{+1}, \overline{A^{\uparrow}}\right) \\ + \mathrm{E}\left(\Delta \mid \overline{A^{+1}}, \overline{A^{\uparrow}}\right) \cdot \Pr\left(\overline{A^{+1}}, \overline{A^{\uparrow}}\right) \\ + \mathrm{E}\left(\Delta \mid A^{\uparrow}\right) \cdot \Pr\left(A^{\uparrow}\right). \end{aligned}$$

The first line is at least $1 \cdot \left(1 - \Pr\left(\overline{A^{+1}}\right) - \Pr\left(A^{\uparrow}\right)\right) \ge 1 - n^{-0.2} - p_{i,\lambda}^{\uparrow}$ by the first statement and a union bound. The second line is at least $-(\log(n) + n^{-\omega(1)} \cdot n)n^{-0.2} = -o(1)$, since the probability of flipping at least $\log n$ bits is $n^{-\omega(1)}$, also under conditions $\overline{A^{+1}}, \overline{A^{\uparrow}}$, and using the trivial bound n if this happens nevertheless. The third line is at least $-\Delta_{i,\lambda}^{\uparrow} p_{i,\lambda}^{\uparrow}$. Plugging this together, we get

$$\begin{aligned} \mathrm{E}(\Delta) &\ge 1 - n^{-0.2} - p_{i,\lambda}^{\uparrow} - o(1) - \Delta_{i,\lambda}^{\uparrow} p_{i,\lambda}^{\uparrow} \\ &= 1 - n^{-0.2} - (\Delta_{i,\lambda}^{\uparrow} + 1)p_{i,\lambda}^{\uparrow} - o(1). \end{aligned}$$

For $|x_t|_1 < 2n/3$, $p_{i,\lambda}^{\uparrow} = 0$ and the claim follows. For $|x_t|_1 > 2n/3 + \log n$, by Lemma 2.3 with $d \ge \log n$,

$$(\Delta_{i,\lambda}^{\uparrow} + 1)p_{i,\lambda}^{\uparrow} \le \frac{\lambda(d+2)}{d!} \le \frac{\lambda(\log(n)+2)}{(\log n)!} = n^{-\omega(1)}.$$

This implies the second statement.

For the third statement, we first note that the second statement also holds for the drift on the function OneMax, for all x_t with $|x_t|_1 \le 3n/4$, as the negative terms involving $p_{i,\lambda}^{\uparrow}$ disappear. Let us first consider the case that the current search point has at most $2n/3$ ones. Then, by the additive drift theorem [15], the expected time until a search point with at least $2n/3$ ones is reached is $O(n)$. Note that is it possible (though unlikely) that the top of the cliff is skipped and the algorithm jumps down the cliff from a search point with at most $2n/3 - 1$ ones. By the first statement of this lemma, the conditional probability of this happening, conditional on increasing the number of ones, is $O(n^{-0.2})$. If it happens regardless, we consider the following case of having more than $2n/3$ ones.

If the current search point has more than $2n/3$ ones, we argue that on OneMax, by the same drift arguments as above, the expected time to reach a search point with at least $2n/3 + \kappa$ ones is $O(\kappa) = O(n)$. We only see a difference to OneMax if the algorithm jumps back up the cliff. Then we are left with a search point of at most $2n/3$ ones and we apply the above arguments.

If $T(n)$ denotes the worst-case time with respect to the initial number of ones $i < 2n/3 + \kappa$, we have shown the recurrence: $T(n) \le O(n) + O(n^{-0.2}) \cdot T(n)$. It is easy to see that $T(n) = O(n)$. □

Another important step for proving Theorem 3.1 is estimating the probability of a standard bit mutation of a parent at the top of the cliff increasing the number of ones, $p_{2n/3,1}^{+}$. This is because in order to jump down the cliff, all offspring must increase the number of ones, which has a probability of $(p_{2n/3,1}^{+})^{\lambda}$. To prove the claimed upper and lower bounds in Theorem 3.1 we need precise estimations of $p_{2n/3,1}^{+}$ as it appears in the base of an expression exponential in λ; the commonly used inequalities $\frac{n-2n/3}{en} \le p_{2n/3,1}^{+} \le \frac{n-2n/3}{n}$ (that is, $\frac{1}{3e} \le p_{2n/3,1}^{+} \le \frac{1}{3}$) are too loose. The following lemma gives precise upper and lower bounds on $p_{i,1}^{+}$ for almost all values of i as this generality is achieved quite easily and the lemma may be of independent interest. The proof is omitted due to space restrictions.

Lemma 3.3. *For all* $i \in \{0, \ldots, n-1\}$,

$$p_{i,1}^{+} \le \left(1 - \frac{1}{n}\right)^{n-2} \sum_{a=0}^{\infty} \sum_{b=a+1}^{\infty} \left(\frac{i}{n}\right)^{a} \left(\frac{n-i}{n}\right)^{b} \frac{1}{a!b!}. \tag{10}$$

For all $i \in \{\lceil \log n \rceil, \ldots, n - \lceil \log n \rceil\}$,

$$p_{i,1}^{+} \ge \left(1 - \frac{1}{n}\right)^{n-2} \sum_{a=0}^{\infty} \sum_{b=a+1}^{\infty} \left(\frac{i}{n}\right)^{a} \left(\frac{n-i}{n}\right)^{b} \frac{1}{a!b!} \left(1 - \frac{2\log^2 n}{\min\{i, n-i\}}\right). \tag{11}$$

For the specific value $i = 2n/3$, we obtain the following special case. Along with $\frac{1}{e} \le \left(1 - \frac{1}{n}\right)^{n-2} \le \frac{1}{e} \cdot (1 + O(1/n))$, Equations (10) and (11) in Lemma 2.3 imply the following.

Corollary 3.4.

$$p_{2n/3,1}^{+} = \frac{1}{e} \sum_{a=0}^{\infty} \sum_{b=a+1}^{\infty} \left(\frac{2}{3}\right)^{a} \left(\frac{1}{3}\right)^{b} \frac{1}{a!b!} \pm O\left(\frac{\log^2 n}{n}\right)$$

which is approximately $0.1613715804 \pm O(\log^2(n)/n)$.

Now we are ready to prove Theorem 3.1.

Proof of Theorem 3.1. By Chernoff bounds, with probability $1 - e^{-\Omega(n)}$, the initial search point has at most $2n/3$ ones. In the following, we assume that in the first $\Theta(\xi^{\lambda})$ expected generations, no mutation ever flips at least $n/3$ bits. The probability of flipping at least $n/3$ bits in one mutation is at most $1/((n/3)!) = n^{-\Omega(n)}$ and a union bound over λ offspring and $\Theta(\xi^{\lambda})$ expected generations (cf.(9)) still yields a probability of $n^{-\Omega(n)}$.

Under this assumption, a necessary condition for finding the optimum is that a transition from a search point with at most

$2n/3$ ones to a search point with more than $2n/3$ ones is made (i.e. a jump down the cliff). By Lemma 1 in [35], the probability of this event is maximised if the parent has exactly $2n/3$ ones; then it is $(p^+_{2n/3,1})^\lambda$. By Equation (10) in Lemma 2.3, along with $(1-1/n)^{n-2} \le 1/e \cdot (1-1/n)^{-2} \le 1/e \cdot (1+2/n)$,

$$(p^+_{2n/3,1})^\lambda \le \left(\xi^{-1}\left(1+\frac{2}{n}\right)\right)^\lambda \le \xi^{-\lambda} \cdot e^{2\lambda/n} = O(\xi^{-\lambda}).$$

The expected waiting time for this transition to happen is at least $\Omega(\xi^\lambda)$. This establishes a lower bound of $(1 - e^{-\Omega(n)} - n^{-\Omega(n)}) \cdot \Omega(\xi^\lambda) = \Omega(\xi^\lambda)$.

Now we show the upper bound. We consider the time T_κ until a search point with at least $2n/3+\kappa$ ones is found, for $\kappa := \frac{2\log\log n}{\log\log\log n}$, assuming that the current search point is at the top of the cliff, that is, the current search point has $2n/3$ ones.

The expected time to return to the top of the cliff, or to find a search point with at least $2n/3+\kappa$ ones, is bounded by $O(n)$ by Lemma 3.2.

Let p^- denote the probability of accepting a search point with less than $2n/3$ ones from the top of the cliff and let $p^+ = (p^+_{2n/3,1})^\lambda$ denote the probability of accepting a search point with more than $2n/3$ ones from the top of the cliff.

By Lemma 2.4 with $c := 2$, the probability of creating an offspring at distance at least $\kappa := \frac{2\log\log n}{\log\log\log n}$ from the cliff is at least $1/\log^2 n$. This clearly also holds under the condition of the event underlying p^+, that is, that all offspring have more than $2n/3$ ones. The probability that there is one such offspring is at least

$$1 - \left(1 - \frac{1}{\log^2 n}\right)^\lambda \ge \frac{\lambda/\log^2 n}{1+\lambda/\log^2 n} \ge \frac{\lambda}{2\log^2 n} \ge \frac{1}{4\log n},$$

where the first inequality follows from Lemma 10 in [2] and the last inequality follows from $\lambda \ge \log_{\frac{e}{e-1}} n \ge \frac{1}{2}\log_2 n$.

Together, we have established a recurrence for $\mathrm{E}(T_\kappa)$:

$$\begin{aligned}\mathrm{E}(T_\kappa) \le{}& 1 + p^-(O(n) + \mathrm{E}(T_\kappa)) \\ &+ p^+\left(O(n) + \left(1 - \frac{1}{4\log n}\right)\mathrm{E}(T_\kappa)\right) + (1-p^+-p^-)\mathrm{E}(T_\kappa).\end{aligned}$$

This is equivalent to

$$\frac{p^+}{4\log n} \cdot \mathrm{E}(T_\kappa) \le 1 + p^- O(n) + p^+ O(n). \tag{12}$$

We argue that, at the top of the cliff, the probability of moving to a search point with a different number of ones is $p^- + p^+ = O(1/n)$. This is because if there is a mutation that does not flip any bits, the next search point will be at the top of the cliff again. Hence, in order to move to a search point with a different number of ones, all offspring must flip at least one bit. The probability of this event is at most, using $(1-1/n)^n = (1-1/n)\cdot(1-1/n)^{n-1} \ge 1/e - 1/(en)$,

$$\begin{aligned}p^- + p^+ &\le \left(1 - \left(1-\frac{1}{n}\right)^n\right)^\lambda \\ &\le \left(1 - \frac{1}{e} + \frac{1}{en}\right)^\lambda = \left(1-\frac{1}{e}\right)^\lambda \left(1 + \frac{1}{(e-1)n}\right)^\lambda \\ &= \left(\frac{e-1}{e}\right)^{\log_{\frac{e}{e-1}} n} \left(1 + \frac{1}{(e-1)n}\right)^\lambda \\ &\le \frac{1}{n} \cdot e^{\lambda/((e-1)n)} = O(1/n)\end{aligned}$$

using $\lambda = O(n)$ in the last step. Plugging this in to (12) and multiplying both sides by $4\log(n)/p^+$, we get

$$\mathrm{E}(T_\kappa) \le O\left(\frac{\log n}{p^+}\right).$$

Now assume that a search point with $2n/3+d$ ones has been reached, where $\kappa \le d \le \log n$. By Lemma 3.2 and Lemma 2.3, the probability that in one generation the number of ones is increased (i.e. the algorithm does not jump up the cliff again) is at least

$$1 - n^{-0.2} - \frac{\lambda(3/4)^d}{d!}.$$

By a union bound, the probability that this happens for all $d = \kappa \dots n^{0.1}$ is at least

$$1 - \frac{n^{0.1}}{n^{0.2}} - \lambda \sum_{d=\kappa}^{n^{0.1}} \frac{(3/4)^d}{d!}.$$

The sum is bounded from above, using $d! \ge (d/e)^e$, as

$$\begin{aligned}\sum_{d=\kappa}^{n^{0.1}} \frac{(3/4)^d}{d!} \le \sum_{d=\kappa}^{n^{0.1}} \left(\frac{3e}{4d}\right)^d &\le \sum_{d=\kappa}^{\infty} \left(\frac{3e}{4\kappa}\right)^d \\ &= \left(\frac{3e}{4\kappa}\right)^\kappa \cdot \sum_{d=0}^{\infty}\left(\frac{3e}{4\kappa}\right)^d \\ &= \left(\frac{3e}{4\kappa}\right)^\kappa \cdot \frac{1}{1-\frac{3e}{4\kappa}} \le 2\left(\frac{3e}{4\kappa}\right)^\kappa.\end{aligned}$$

Plugging in κ, we get

$$\begin{aligned}2\left(\frac{3e\log\log\log n}{4\log\log n}\right)^{\frac{2\log\log n}{\log\log\log n}} &\le 2\cdot 2^{\frac{2\log\log n}{\log\log\log n}\cdot\log\left(\frac{3e\log\log\log n}{4\log\log n}\right)} \\ &= 2\cdot 2^{-\frac{2\log\log n}{\log\log\log n}\cdot\log\left(\frac{4\log\log n}{3e\log\log\log n}\right)}.\end{aligned}$$

For large enough n, we have

$$\frac{4\log\log n}{3e\log\log\log n} \ge 2(\log\log n)^{1/2}$$

and then

$$\log\left(\frac{4\log\log n}{3e\log\log\log n}\right) \le 1 + \frac{1}{2}\log\log\log n.$$

Plugging this in, we bound the sum by

$$2\cdot 2^{-\frac{2\log\log n}{\log\log\log n}\cdot\left(1+\frac{1}{2}\log\log\log n\right)} \le 2\cdot 2^{-\log\log n - \frac{2\log\log n}{\log\log\log n}} = o\left(\frac{1}{\log n}\right).$$

Thus, the probability of reaching a search point with at least $2n/3 + n^{0.1}$ ones before going back up the cliff is at least

$$1 - \frac{n^{0.1}}{n^{0.2}} - \lambda\sum_{d=\kappa}^{\log n}\frac{(3/4)^d}{d!} \ge 1 - o(1).$$

From that point on, for any search point with at least $2n/3 + n^{0.1}/2$ ones, we can use arguments from the analysis of the $(1,\lambda)$ EA on OneMax in [33] since the $(1,\lambda)$ EA must flip at least $n^{0.1}/2$ bits to return to the cliff, and this has exponentially small probability. It is not difficult to show using the negative drift theorem [27, 28] and the second statement of Lemma 3.2, that, once we have reached a distance of at least $n^{0.1}$ from the cliff, reducing that distance to at most $n^{0.1}/2$ has an exponentially small probability. Assuming we never go back up the cliff, the remaining expected optimisation time is $O(n \log n)$ [33].

In total, the expected optimisation time is

$$O(1) \cdot \mathrm{E}\left(T_\kappa\right) = O\left(\frac{\log n}{p^+}\right) = O\left(\frac{\log n}{(p^+_{2n/3,1})^\lambda}\right).$$

This implies the claim since $p^+_{2n/3,1} = \xi^{-1} - O((\log^2 n)/n)$ by Corollary 3.4 and

$$\left(p^+_{2n/3,1}\right)^\lambda = \left(\xi^{-1} - \frac{O(\log^2 n)}{n}\right)^\lambda = \xi^{-\lambda}\left(1 - \frac{O(\log^2 n)}{n}\right)^\lambda \geq \xi^{-\lambda}\left(1 - \frac{O(\lambda \log^2 n)}{n}\right) = \Omega(\xi^{-\lambda}). \qquad \square$$

Along with the exponential lower bound from [33] for small λ-values, Theorem 3.1 shows that the expected optimisation time for any λ grows faster than any polynomial with degree less than $\eta \approx 3.976770136$.

Theorem 3.5. *Let $\eta := 1/\log_\xi\left(\frac{e}{e-1}\right) \approx 3.976770136$. The expected optimisation time of the $(1, \lambda)$ EA with static λ is $\omega(n^{\eta-\varepsilon})$ for every constant $\varepsilon > 0$ and every λ and $O(n^\eta \log n)$ for $\lambda = \lceil \log_{\frac{e}{e-1}} n \rceil$.*

Proof. By Theorem 3.1, the expected optimisation time for $\lambda = \lceil \log_{\frac{e}{e-1}} n \rceil$ is $O(\xi^\lambda \log n)$. Using

$$\xi^\lambda \geq \xi^{\log_{\frac{e}{e-1}} n} = \xi^{\log_\xi(n)/\log_\xi\left(\frac{e}{e-1}\right)} = n^{1/\log_\xi\left(\frac{e}{e-1}\right)} = n^\eta$$

this establishes the claimed upper bound.

For the lower bound, we exploit that for every constant $\varepsilon' > 0$, for all $\lambda \leq (1 - \varepsilon') \log_{\frac{e}{e-1}} n$ the expected optimisation time of the $(1, \lambda)$ EA on every function with a unique optimum is at least $2^{cn^{\varepsilon'/2}}$, for some constant $c > 0$, by Theorem 10 in [33]. This is clearly in $\omega(n^\eta)$. For $\lambda = \omega(n)$ the lower bound $\min\{n^{n/4}, e^{\lambda/4}\}/3$ from Theorem 8 in [18] is exponential. It thus suffices to consider $\lambda > (1-\varepsilon') \log_{\frac{e}{e-1}} n$ and $\lambda = O(n)$. The lower bound from Theorem 3.1 then becomes $\Omega(\xi^{(1-\varepsilon') \log_{\frac{e}{e-1}} n}) = \Omega(n^{(1-\varepsilon')\eta})$. Choosing $\varepsilon' := \varepsilon/(2\eta)$, this is $\Omega(n^{\eta-\varepsilon/2}) = \omega(n^{\eta-\varepsilon})$ as claimed. $\square$

4 SELF-ADJUSTING OFFSPRING POPULATIONS ARE EFFICIENT ON CLIFF

In this section we show that the self-adjusting $(1, \lambda)$ EA is faster than the $(1, \lambda)$ EA with static parameter choice by a polynomial factor of $\Theta(n^{2.9767}/\log n)$, achieving the best possible asymptotic runtime for any unary unbiased black-box algorithm of $O(n \log n)$ [21]. The main result of this section is shown in Theorem 4.1.

Theorem 4.1. *Let the update strength $F > 1$ and the success rate $0 < s < \frac{e-1}{e}$ be constants and $enF^{1/s} \leq \lambda_{\max} = \mathrm{poly}(n)$. Then for any initial search point and any initial $\lambda_0 \leq \lambda_{\max}$ the self-adjusting $(1, \lambda)$ EA resetting λ optimises Cliff in $O(n)$ expected generations and $O(\lambda_{\max} \log n)$ expected function evaluations.*

For $\lambda_{\max} = \lceil enF^{1/s} \rceil$ we get $O(n \log n)$ evaluations in expectation.

The proof of our result is divided in four phases: reaching the cliff, jumping down the cliff, climbing away from the cliff and finding the global optimum.

4.1 Reaching the cliff

We note that the algorithm studied here is the same as the algorithm studied in [17] as long as all generations that use $\lambda = \lambda_{\max}$ are successful. Here we show that before reaching the cliff the probability of an unsuccessful generation with $\lambda = \lambda_{\max}$ is sufficiently small to not affect the optimisation. Additionally, the results from [17] on OneMax can be applied when only considering improvements that increase the fitness by 1. Hence, they can be translated to the Cliff function to calculate the time the algorithm takes to reach the cliff, giving the following bounds.

Lemma 4.2 (Adapted from Theorem 3.5 in [17]). *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. For every initial offspring population size $\lambda_0 \leq \lambda_{\max}$ and every initial search point x_0 with $|x_0|_1 = 2n/3 - a$ for $a \geq 1$ the algorithm evaluates a solution x_t with $|x_t|_1 \geq 2n/3$ using in expectation $O(a + \log n)$ generations and $O(a + \log n + \lambda_0)$ evaluations.*

We translate the results from [17] on OneMax to the first slope of Cliff, therefore, before giving a proof for Lemma 4.2 we first show that w.o.p the self-adjusting $(1, \lambda)$ EA does not reset λ in this region and consequently it behaves as the algorithm studied in [17]. By (9) this holds for any random time period of polynomial expected length. The following lemma concerns a larger region of search points with up to $3n/4$ ones as this will be useful later on.

Lemma 4.3. *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. The probability that in a generation t with $|x_t|_1 \leq 3n/4$ and $|x_t|_1 \neq 2n/3$ the self-adjusting $(1, \lambda)$ EA resets λ is at most $e^{-\Omega(n)}$.*

Proof. In order to reset λ a generation using $\lambda = \lambda_{\max}$ must not increase the fitness. By Lemma 2.3 with $\lambda = \lambda_{\max} \geq enF^{1/s}$ the probability of this event is at most

$$1 - p^{\rightarrow}_{\lambda_{\max}} \leq \left(1 - \frac{1}{3e}\right)^{enF^{1/s}} \leq \exp\left(-\frac{nF^{1/s}}{3}\right) = e^{-\Omega(n)}. \qquad \square$$

We now show the relevant definitions and lemmas adapted from [17] including the necessary modifications to their proofs to translate them to Cliff.

Definition 4.4. We define the potential function $g(X_t)$ as in [17]:

$$g(X_t) = |x_t|_1 - \frac{se}{e-1} \log_F \left(\max\left(\frac{enF^{1/s}}{\lambda_t}, 1\right)\right).$$

Using Definition 4.4 we can see that given that $\lambda_{\max} \geq enF^{1/s}$ the drift of the potential does not change as long as λ does not reset to 1. Hence the following lemma still holds.

Lemma 4.5 (Adapted from Lemma 3.4 in [17]). *Consider the self-adjusting $(1, \lambda)$ EA resetting λ as in Theorem 4.1 and assume that the event stated in Lemma 4.3 does not occur. Then for every generation t with $|x_t|_1 < 2n/3$,*

$$E\left(g(X_{t+1}) - g(X_t) \mid X_t\right) \geq \frac{1}{e} - \frac{s}{e-1} > 0$$

for large enough n.

Again, as long as λ does not reset, the following lemma taken from [17] that describes the expected value of λ holds.

Lemma 4.6 (Lemma 3.13 in [17]). *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1 and assume that the event stated in Lemma 4.3 does not occur. If the best-so-far fitness at time t is at most i then*

$$E(\lambda_t \mid \lambda_0) \leq \lfloor \lambda_0/F^t \rfloor + \frac{en}{n-i} \cdot \left(F^{1/s} + \frac{F^{1/s}}{\ln F}\right).$$

For completeness we state the following lemma taken from [17].

Lemma 4.7 (Lemma 3.3 in [17]). *For all generations t, $|x|_1$ and the potential are related as: $|x|_1 - \frac{se}{e-1}\log_F(enF^{1/s}) \leq g(X_t) \leq |x|_1$.*

With the previous lemmas we can now prove Lemma 4.2.

Proof of Lemma 4.2. Following the arguments of the proof of Theorem 3.5 in [17] to bound the number of generations to reach $|x_t|_1 \geq 2n/3$ we use the potential function $g(X_t)$. To fit the perspective of the additive drift theorem [15] we switch to the potential function $\overline{g}(X_t) := \max(2n/3 - g(X_t), 0)$ and stop when $\overline{g}(X_t) = 0$ (which implies that $|x_t|_1$ is least $2n/3$) or $|x_t|_1$ of at least $2n/3$ is reached beforehand. Note that the maximum caps the effect of generations that jump down the cliff. Lemma 4.5 shows that the potential $g(X_t)$ has a positive constant drift whenever $|x_t|_1 < 2n/3$, and given that the drift bound for $g(X_t)$ still holds when only considering fitness improvements by 1 it also holds for $\overline{g}(X_t)$.

The initial value $\overline{g}(X_0)$ is at most $a + \frac{se}{e-1}\log_F\left(enF^{1/s}\right)$ by Lemma 4.7. Using Lemma 4.5 and the additive drift theorem [15], the expected number of generations is

$$E(T_1) \leq \frac{a + \frac{se}{e-1}\log_F\left(enF^{1/s}\right)}{\frac{1}{e} - \frac{s}{e-1}} = O(a + \log n).$$

The expected number of function evaluations during this time is $E\left(\lambda_0 + \lambda_1 + \cdots + \lambda_{T_1-1}\right) = E\left(\sum_{t=0}^{T_1-1} \lambda_t \mid \lambda_0\right)$. We bound all summands by Lemma 4.6, applied with a worst case fitness of $i := 2n/3$. This yields a random variable λ^* with

$$E\left(\lambda^*\right) \leq \frac{en}{n/3} \cdot \left(F^{1/s} + \frac{F^{1/s}}{\ln F}\right) = e/3 \cdot \left(F^{1/s} + \frac{F^{1/s}}{\ln F}\right)$$

and $E\left(\lambda^*\right) \geq E\left(\lambda_t \mid \lambda_0\right) - \lfloor \lambda_0/F^t \rfloor$ for all $t < T_1$. Thus, the expected time can be bounded by T_1 i.i.d. variables λ^* and $\sum_{t=0}^{\infty} \lfloor \lambda_0/F^t \rfloor \leq \frac{F\lambda_0}{F-1} = O(\lambda_0)$. Since T_1 is itself a random variable, we apply Wald's equation [34] to conclude that

$$O(\lambda_0) + E\left(\sum_{t=0}^{T_1-1} \lambda^*\right) = O(\lambda_0) + E(T_1) \cdot E\left(\lambda^*\right) = O(a + \log n + \lambda_0).$$

Finally, if the failure from Lemma 4.3 occurs we restart the analysis with a worst-case value of n for a. Since the failure has an exponentially small probability, this does not affect the claimed expectations. □

4.2 Jumping down the cliff

After reaching the cliff, the algorithm needs to jump down the cliff. This requires a generation in which *all* offspring lie on the second slope. We have seen in Section 3 that this probability is exponentially small in λ_t. The resetting mechanism implies that when λ reaches its maximum value $\lambda_{\max}$ and the following generation is unsuccessful, we reach small values of λ_t and jumps down the cliff become likely.

We also know from Section 3 that we need a sufficiently large jump to prevent the algorithm from jumping straight back up the cliff. This probability decreases with the distance to the cliff (that is, $|x_t|_1 - 2n/3$) and it increases with λ_t as many offspring can amplify the probability of a jump back up the cliff. The following definition captures states from which the probability of jumping up the cliff is sufficiently small.

Definition 4.8. Given some even value $\kappa \in \mathbb{N}$, a state (x_t, λ_t) is called *κ-safe* if $|x_t|_1 \geq 2n/3 + \kappa$ and $\lambda_t \leq 2^{-\kappa/2} \cdot (\kappa/2)!$.

Note that $2^{-\kappa/2} \cdot (\kappa/2)!$ is non-decreasing in κ.

In this subsection we give upper bounds on the expected number of generations and the expected number of function evaluations to reach a κ-safe state for the specific value $\kappa := \frac{\log\log n}{\log\log\log n}$. We also consider search points with at least $3n/4$ ones as safe, regardless of the value of λ_t; reaching such a search point will be the goal of the following phase.

Lemma 4.9. *Consider the self-adjusting $(1, \lambda)$ EA with resetting λ as in Theorem 4.1. Let $\kappa := \frac{\log\log n}{\log\log\log n}$. For every initial $\lambda_0 \leq \lambda_{\max}$ and every initial search point x_0 with $|x_0|_1 \geq 2n/3$ the algorithm reaches a κ-safe state, or a search point with at least $3n/4$ ones, in $O(\log(\lambda_{\max})\log n)$ expected generations and $O(\lambda_{\max}\log n)$ expected evaluations.*

The main idea of the proof of Lemma 4.9 is that, once the algorithm reaches a local optimum with $2n/3$ ones, λ will increase until it resets to 1 and this is repeated until the algorithm leaves its local optimum. Every time $\lambda = 1$ the algorithm will accept any offspring, then we can wait until a lucky mutation step can directly jump to a search point with at least $2n/3 + \frac{\log\log n}{\log\log\log n}$ ones. Finally, we account for the time that the algorithm takes to reset λ to 1, including the time spent outside of local optima in states that are not safe.

One such set of non-safe states is that of states with at least $2n/3 + \kappa$ ones but violating the upper bound for λ_t from Definition 4.8. For these states we cannot exclude that the algorithm will return to the first slope before it reaches a safe state, but we can bound the time the algorithm spends before either event happens and later on account for this time. This is done in the following lemma.

Lemma 4.10. *Consider the self-adjusting $(1, \lambda)$ EA with resetting λ as in Theorem 4.1. Let $\kappa := \frac{\log\log n}{\log\log\log n}$. From any state (x_0, λ_0) with $|x_0|_1 \geq 2n/3 + \kappa$ in expectation the algorithm needs at most $O(\log\lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations to reach a κ-safe state, to return to the first slope or to find a search point x_t with $|x_t|_1 \geq 3n/4$.*

The proof of Lemma 4.10 is omitted due to space limitations. The main proof idea is to show that with a large value of λ, with high probability all generations are successful and λ is quickly reduced to a safe value, unless any of the other two events happens.

With Lemma 4.10 we are able to prove Lemma 4.9.

Proof of Lemma 4.9. We define a cycle as a sequence of generations that begins after λ is reset to 1 and $|x_t|_1 = 2n/3$. Since $|x_0|_1 \geq 2n/3$ and there is no assumption on λ_0 we need to take into a account the time taken to start the first cycle; this will be accounted later on.

By Lemma 2.4 with $c := 1$ in the first generation of every cycle (with $\lambda = 1$) there is a probability of at least $1/\log n$ to create and accept an offspring with at least $2n/3 + \frac{\log\log n}{\log\log\log n}$ ones. The next generation will have $\lambda = F^{1/s}$, which for a sufficiently large n satisfies $\lambda \leq 2^{-\kappa/2} \cdot (\kappa/2)!$. Hence, in expected $\log n$ cycles the sought event will happen. Now it remains to bound the expected number of generations and evaluations in each cycle.

If during a cycle all generations maintain the current fitness value, after j generations the offspring population size is $F^{j/s}$. For $j := \lceil s \log_F \lambda_{\max} \rceil$, we get an offspring population size of

$$F^{\lceil s \log_F \lambda_{\max} \rceil / s} \geq F^{\log_F \lambda_{\max}} = \lambda_{\max}.$$

Therefore, the number of generations needed to reset λ is at most $\lceil s \log_F \lambda_{\max} \rceil + 1$. Using $\lceil F^{j/s} \rceil \leq 2F^{j/s}$, during these generations, the number of evaluations is at most

$$\begin{aligned}\sum_{j=0}^{\lceil s \log_F \lambda_{\max} \rceil} \left\lceil F^{j/s} \right\rceil &\leq 2 \sum_{j=0}^{\lceil s \log_F \lambda_{\max} \rceil} \left(F^{1/s}\right)^j \\ &= 2 \cdot \frac{(F^{1/s})^{\lceil s \log_F (\lambda_{\max}) \rceil + 1} - 1}{F^{1/s} - 1} \\ &\leq 2 \cdot \frac{(F^{1/s})^{s \log_F (\lambda_{\max}) + 2}}{F^{1/s} - 1} \\ &= \frac{2F^{2/s}}{F^{1/s} - 1} \cdot \lambda_{\max} = O(\lambda_{\max}).\end{aligned}$$

We now show that when the algorithm is in a local optimum, with constant probability all following generations will maintain the current fitness value until λ is reset to 1. When $|x_t|_1 = 2n/3$, in order for a generation to maintain the number of one-bits, it is sufficient to create at least one copy of the parent. Hence, the probability of the event is at least

$$\begin{aligned}1 - \left(1 - \left(1 - \frac{1}{n}\right)^n\right)^\lambda &\geq \frac{\lambda \left(1 - \frac{1}{n}\right)^n}{1 + \lambda \left(1 - \frac{1}{n}\right)^n} = \frac{1}{1 + \frac{1}{\lambda (1 - \frac{1}{n})^n}} \\ &\geq \frac{1}{\exp\left(\frac{1}{\lambda (1 - \frac{1}{n})^n}\right)} = \exp\left(-\frac{1}{\lambda \left(1 - \frac{1}{n}\right)^n}\right).\end{aligned}$$

The probability that a cycle is comprised only of generations that maintain the fitness value is at least

$$\begin{aligned}\prod_{j=0}^{\lceil s \log_F \lambda_{\max} \rceil} \exp\left(-\frac{1}{F^j \left(1 - \frac{1}{n}\right)^n}\right) &\geq \prod_{j=0}^{\infty} \exp\left(-\frac{1}{F^j \left(1 - \frac{1}{n}\right)^n}\right) \\ &= \exp\left(-\frac{1}{\left(1 - \frac{1}{n}\right)^n} \sum_{j=0}^{\infty} F^{-j}\right) = \exp\left(-\frac{1}{\left(1 - \frac{1}{n}\right)^n} \cdot \frac{F}{F-1}\right) = \Omega(1).\end{aligned}$$

Therefore, in each cycle the algorithm will directly increase λ to $\lambda_{\max}$ and then reset to 1 with constant probability using $O(\log \lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations.

Since the self-adjusting $(1, \lambda)$ EA is non-elitist, there is still a possibility for the algorithm to either jump down the cliff (but not to the desired distance of κ) or to reduce the number of ones. If the number of one-bits is reduced, by Lemma 3.7 in [17] with probability $1 - O(1/n)$ the number of one-bits will never drop below $2n/3 - O(\log n)$ ones before reaching a point with $2n/3$ ones. Hence, by Lemma 4.2, it will take $O(\log n)$ generations and $O(\log n + \lambda_0) = O(\lambda_{\max})$ evaluations in expectation to return to a local optimum. If the algorithm jumps down the cliff, using the same arguments as in Lemma 4.2 we can see that the algorithm will use $O(\log n)$ generations and $O(\lambda_{\max})$ evaluations to either find a solution with at least $2n/3 + \frac{\log\log n}{\log\log\log n}$ ones or jump back to the first slope. If the algorithm finds a solution with at least $2n/3 + \frac{\log\log n}{\log\log\log n}$ ones but with $\lambda > 2^{-\kappa/2} \cdot (\kappa/2)!$ then by Lemma 4.10 it will take $O(\log \lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations in expectation to either reduce λ to $\lambda \leq 2^{-\kappa/2} \cdot (\kappa/2)!$ or to return to the first slope.

Finally, when jumping back to the first slope, by Lemma 2.3, in expectation the number of ones is reduced by $\Delta^{\uparrow}_{i,\lambda} \leq d + 1$, that is, in expectation the algorithm jumps to a point that has $a \leq 1$ less ones than the local optimum. Let T' be the time to go back to $|x|_1 = 2n/3$ from $|x|_1 = 2n/3 - a$. By the law of total expectation and Lemma 4.2, $\mathrm{E}(T') = \mathrm{E}(\mathrm{E}(T \mid a)) \leq \mathrm{E}(O(a + \log n + \lambda_0)) = O(\mathrm{E}(a)) + O(\log n + \lambda_0)$. Given that $\mathrm{E}(a) \leq 1$ we obtain $\mathrm{E}(T') = O(\log n + \lambda_0) = O(\lambda_{\max})$ evaluations. For the number of generations we use the same arguments and obtain $O(\log n)$ generations. Therefore, if the algorithm moves out of the local optimum it will return to it in $O(\log n + \log \lambda_{\max}) = O(\log \lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations. In expectation this will happen only a constant number of times before λ is reset to 1. This implies that in expectation each cycle is comprised of $O(\log \lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations.

It remains to account for the time before the first cycle. Using the same arguments as before for any $\lambda_0 \leq \lambda_{\max}$ and $|x_0|_1 \geq 2n/3$ the algorithm will spend $O(\log \lambda_{\max})$ generations and $O(\lambda_{\max})$ evaluations to start the first cycle or reach the desired state. Noting that the expected number of cycles is $\log n$ proves the claim. □

4.3 After jumping down the cliff

Now we show that, with probability $\Omega(1)$, we reach a search point with at least $3n/4$ ones when starting from a κ-safe state with $\kappa := \frac{\log\log n}{\log\log\log n}$. As in Section 3, the target of reaching $3n/4$ ones is chosen such that the probability of an improving mutation is always $\Omega(1)$.

Proving this claim is not straightforward for several reasons. It is always possible to have a mutation jumping back up the cliff. This probability decreases with the distance to the cliff (that is, $|x_t|_1 - 2n/3$) and it increases with λ_t as many offspring can amplify the probability of a jump back up the cliff (cf. Lemma 2.3). Fortunately, the notion of κ-safe states implies that we start with a distance of at least κ to the cliff and a small λ_t, so that initially this probability amplification does not pose a huge risk.

But small values of λ are risky for another reason. Since the number of ones is larger than $2n/3$ and hence significantly larger than $n/2$, the expected number of ones in any offspring is smaller than that of its parent. With $\lambda_t \approx 1$ there is a constant negative drift towards decreasing the number of zeros and "slipping down" the second slope. Fortunately, λ_t will increase during unsuccessful generations and we will see that this effect prevents the algorithm from slipping down the second slope.

In [17] we showed that, on OneMax, for the potential function from Definition 4.4 there is a positive drift in the potential throughout the run: on OneMax, Lemma 4.5 holds for all non-optimal search points. We further constructed a so-called "ratchet argument", arguing that significant decreases in the potential are unlikely. In our approach, the potential and the fitness only differ by a term of $\Theta(\log n)$, as stated in Lemma 4.7. Hence we concluded that, with high probability, the best fitness never decreases by a term of $r \log n$, for some constant $r > 0$.

Unfortunately, this ratchet argument is not directly applicable here, since we can only guarantee a distance of $\kappa := \frac{\log\log n}{\log\log\log n} \ll r \log n$ to the cliff. Hence the ratchet argument from [17] has far too much slack.

The proof of the ratchet argument in [17] applies the negative drift theorem [27, 28] to an interval on the potential scale of size $\Theta(\log n)$, in order to obtain failure probabilities that are polynomially small (that is, exponentially small in the interval length).

We refine the ratchet argument here by defining a revised potential function tailored to a fitness range up to $3n/4$ ones, where the fitness and the potential only differ by an additive term of $\Theta(1)$. Then we apply the negative drift theorem [27, 28] to an interval of size $\kappa/2 - O(1) = \frac{\log\log n}{2\log\log\log n} - O(1)$ on the potential scale to show that the number of ones does not drop below $2n/3 + \kappa/2$ in a time that is exponential in the interval length. More specifically, the time period will be determined as γ^κ, for some constant $\gamma > 1$. During this time, the potential has a positive drift and with good probability the algorithm moves sufficiently far away from the cliff, that is, to a distance of $\Theta(\gamma^\kappa)$.

Since $\gamma^{\frac{\log\log n}{\log\log\log n}} = o(\log n)$, we can only guarantee a sub-logarithmic increase in the distance and the failure probability from the negative drift theorem is $\omega(1/\log n)$. Thus, we iterate this argument three times, with exponentially increasing values for κ, until we reach a search point with at least $3n/4$ ones (or we return to the first slope).

Throughout these arguments, we also show that λ_t is bounded from above as $\lambda_t \le 2^{-\kappa/2}\cdot(\kappa/2)!$ as in the definition of κ-safe states. This definition requires a distance of at least κ from the top of the cliff, however we can only guarantee a distance of at least $\kappa/2$. We call such states *weakly κ-safe*.

Definition 4.11. A state (x_t, λ_t) is called *weakly κ-safe* if $|x_t|_1 \ge 2n/3 + \kappa/2$ and $\lambda_t \le 2^{-\kappa/2}\cdot(\kappa/2)!$.

We start by revising the potential function from [17] as follows.

Definition 4.12. Let $\varepsilon := \frac{e-1}{e} - s$. We define the potential function $h(X_t)$ as

$$h(X_t) = |x_t|_1 - \frac{se}{e-1}\log_F\left(\max\left(\frac{8e + \log_{\frac{e}{e-1}}(2/\varepsilon)F^{1/s}}{\lambda_t}, 1\right)\right).$$

Lemma 4.13. *Consider the self-adjusting $(1,\lambda)$ EA as in Theorem 4.1. For all states (x_t, λ_t) with $d := |x_t|_1 - 2n/3 = \omega(1)$, $|x_t|_1 \le 3/4$ and $\lambda_t \le 2^{-d/2}\cdot(d/2)!$,*

$$E\left(h(X_{t+1}) - h(X_t) \mid X_t\right) \ge \frac{1}{2e} - \frac{s}{2(e-1)} > 0$$

for large enough n. This also holds when only considering improvements that increase the fitness by 1.

Proof. The proof follows the proof of Lemma 3.4 in [17], using additional arguments to consider jumps up the cliff and the possibility that λ is reset when $\lambda_t = \lambda_{\max}$. In this proof, we use $p^0_{i,\lambda} := 1 - p^{\rightarrow}_{i,\lambda} - p^{\leftarrow}_{i,\lambda} - p^{\uparrow}_{i,\lambda}$ to denote the probability of not changing the current number of ones. We first assume that $\lambda_{\max} > 2^{-d/2}\cdot(d/2)!$, which implies that resets are impossible under the assumption $\lambda_t \le 2^{-d/2}\cdot(d/2)!$.

We first consider the case $\lambda_t \le 8e + \log_{\frac{e}{e-1}}(2/\varepsilon)$ as then $\lambda_{t+1} \le 8e + \log_{\frac{e}{e-1}}(2/\varepsilon)F^{1/s}$ and $h(X_{t+1}) = |x_{t+1}|_1 - \frac{se}{e-1}(\log_F(8e + \log_{\frac{e}{e-1}}(2/\varepsilon)F^{1/s}) - \log_F(\lambda_{t+1}))$. When the number of ones increases, in expectation they do so by $\Delta^{\rightarrow}_{i,\lambda}$ and since $\lambda_{t+1} = \lambda_t/F$, the penalty term $\frac{se}{e-1}(\log_F(8e + \log_{\frac{e}{e-1}}(2/\varepsilon)F^{1/s}) - \log_F(\lambda_t))$ increases by $\frac{se}{e-1}$ (unless $\lambda_{t+1} = 1$ is reached, in which case the increase might be lower). When the number of ones does not change, the penalty decreases by $\frac{e}{e-1}$. When the number of ones decreases, conditional on $|x_{t+1}|_1 > 2n/3$, the expected decrease is at most $\Delta^{\leftarrow}_{i,\lambda}$ and the penalty decreases by $\frac{e}{e-1}$. Finally, then when the algorithm creates an offspring up the cliff the expected decrease in the number of ones is $\Delta^{\uparrow}_{i,\lambda}$ and the penalty increases by $\frac{se}{e-1}$. Together,

$$\begin{aligned}
&\mathrm{E}\left(h(X_{t+1}) - h(X_t) \mid X_t, \lambda_t \le 8e + \log_{\frac{e}{e-1}}(2/\varepsilon)\right) \\
&\ge p^{\rightarrow}_{i,\lambda}\left(\Delta^{\rightarrow}_{i,\lambda} - \frac{se}{e-1}\right) + p^0_{i,\lambda}\cdot\frac{e}{e-1} + p^{\leftarrow}_{i,\lambda}\left(-\Delta^{\leftarrow}_{i,\lambda} + \frac{e}{e-1}\right) \\
&\qquad + p^{\uparrow}_{i,\lambda}\left(-\Delta^{\uparrow}_{i,\lambda} - \frac{se}{e-1}\right).
\end{aligned}$$

Using $\Delta^{\rightarrow}_{i,\lambda} \ge 1$ (which also holds when only considering fitness increases by 1), $\Delta^{\leftarrow}_{i,\lambda} \le \frac{e}{e-1}$ and $\Delta^{\uparrow}_{i,\lambda} \le d+1$ by Lemma 2.3, this is at least

$$p^{\rightarrow}_{i,\lambda}\left(1 - \frac{se}{e-1}\right) + p^0_{i,\lambda}\cdot\frac{e}{e-1} - p^{\uparrow}_{i,\lambda}\left(d + 1 + \frac{se}{e-1}\right).$$

We bound the second summand from below using $\frac{e}{e-1} > 1 - \frac{se}{e-1}$ (note that the left-hand side is larger than 1 and the right-hand side is less than 1) and obtain a lower bound of

$$\begin{aligned}
&p^{\rightarrow}_{i,\lambda}\left(1 - \frac{se}{e-1}\right) + p^0_{i,\lambda}\left(1 - \frac{se}{e-1}\right) - p^{\uparrow}_{i,\lambda}\left(d + 1 + \frac{se}{e-1}\right) \\
&= (1 - p^{\leftarrow}_{i,\lambda} - p^{\uparrow}_{i,\lambda})\left(1 - \frac{se}{e-1}\right) - p^{\uparrow}_{i,\lambda}\left(d + 1 + \frac{se}{e-1}\right) \\
&= (1 - p^{\leftarrow}_{i,\lambda})\left(1 - \frac{se}{e-1}\right) - p^{\uparrow}_{i,\lambda}\left(d + 2\right).
\end{aligned}$$

Lemma 2.3 shows that $p_{i,\lambda}^{\leftarrow} \leq \frac{e-1}{e}$ for all λ, hence $1 - p_{i,\lambda}^{\leftarrow} \geq \frac{1}{e}$ and $p_{i,\lambda}^{\uparrow} \leq \frac{\lambda(3/4)^d}{d!}$. If $\lambda = O(1)$ and $d = \omega(1)$ then $p_{i,\lambda}^{\uparrow}(d+2) = o(1)$, hence for a sufficiently large n,

$$\begin{aligned}&\mathrm{E}\left(h(X_{t+1}) - h(X_t) \mid X_t, \lambda_t \leq 8e + \log_{\frac{e}{e-1}}(2/\varepsilon)\right)\\&\geq \frac{1}{e} - \frac{s}{e-1} - o(1) \geq \frac{1}{2e} - \frac{s}{2(e-1)}.\end{aligned}$$

For the case $8e + \log_{\frac{e}{e-1}}(2/\varepsilon) < \lambda_t < 2^{-d/2} \cdot (d/2)!$, in an unsuccessful generation the penalty term is capped at its maximum and we pessimistically bound the positive effect on the potential from below by 0. However, the probability of increasing the number of ones is large enough to show a positive drift.

By assumption $\lambda_t \leq 2^{-d/2} \cdot (d/2)!$. Along with (5) from Lemma 2.3,

$$p_{i,\lambda}^{\uparrow} \leq \frac{\lambda(3/4)^d}{d!} \leq \left(\frac{3}{8}\right)^d.$$

We also have $\lambda_t \geq 8e + \log_{\frac{e}{e-1}}(2/\varepsilon) \geq 8e$ since $1/\varepsilon \geq \frac{e}{e-1}$. Then, by Lemma 2.3, $8e + \log_{\frac{e}{e-1}}(2/\varepsilon) < \lambda_t \leq 2^{-d/2} \cdot (d/2)!$ implies the following two statements.

$$p_{i,\lambda}^{\rightarrow} \geq 1 - \left(1 - \frac{1}{4e}\right)^{8e} - \left(\frac{3}{8}\right)^d \geq 1 - \frac{1}{2e} - \left(\frac{3}{8}\right)^d$$

$$\begin{aligned}p_{i,\lambda}^{\leftarrow}\Delta_{i,\lambda}^{\leftarrow} &\leq \left(\frac{e-1}{e}\right)^{8e+\log_{\frac{e}{e-1}}(2/\varepsilon)} \cdot \frac{e}{e-1}\\&= \left(\frac{e-1}{e}\right)^{8e-1+\log_{\frac{e}{e-1}}(2/\varepsilon)} \geq \left(\frac{e-1}{e}\right)^{\log_{\frac{e}{e-1}}(2/\varepsilon)} = \frac{\varepsilon}{2}.\end{aligned}$$

Together,

$$\begin{aligned}&\mathrm{E}\left(h(X_{t+1}) - h(X_t) \mid X_t, 8e + \log_{\frac{e}{e-1}}(2/\varepsilon) < \lambda_t < 2^{-d/2} \cdot (d/2)!\right)\\&\geq\ p_{i,\lambda}^{\rightarrow}\left(1 - \frac{se}{e-1}\right) + p_{i,\lambda}^{\leftarrow}\left(-\Delta_{i,\lambda}^{\leftarrow}\right) - p_{i,\lambda}^{\uparrow}\left(d + 1 + \frac{se}{e-1}\right)\\&\geq\ \left(1 - \frac{1}{2e} - \left(\frac{3}{8}\right)^d\right)\left(1 - \frac{se}{e-1}\right) - \frac{\varepsilon}{2} - \left(\frac{3}{8}\right)^d (d+2)\end{aligned}$$

given that $d = \omega(1)$,

$$\geq\ \left(1 - \frac{1}{e}\right)\left(1 - \frac{se}{e-1}\right) - \frac{\varepsilon}{2} - o(1)$$

using the definition $\varepsilon := \frac{e-1}{e} - s$,

$$\begin{aligned}&=\ \left(1 - \frac{1}{e}\right)\left(1 - \frac{se}{e-1}\right) - \frac{e-1}{2e} + \frac{s}{2} - o(1)\\&=\ \frac{1}{2}\left(\frac{e-1}{e} - s\right) - o(1) =\ \frac{e-1-es}{2e} - o(1)\end{aligned}$$

for sufficiently large n,

$$\geq\ \frac{e-1-es}{2e(e-1)} =\ \frac{1}{2e} - \frac{s}{2(e-1)}.$$

Now, if $\lambda_{\max} \leq 2^{-d/2} \cdot (d/2)!$ then resets may happen if $\lambda_t = \lambda_{\max}$. A reset decreases the potential by at most $n + O(1)$ as this is the range of the potential scale. The probability of a reset is at most $e^{-\Omega(n)}$ by Lemma 4.3. Hence, this only affects the drift by an additive term $-O(n) \cdot e^{-\Omega(n)} = -o(1)$, which can easily be absorbed in the $-o(1)$ terms from the above calculations. □

The following lemma shows that λ typically does not grow beyond the threshold $2^{-\kappa/2} \cdot (\kappa/2)!$ from the definition of weakly κ-safe states.

Lemma 4.14. *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. Then for all $\kappa \geq 324F^{1/s}$ the following holds. If the current state (x_t, λ_t) has $\lambda_t \leq 2^{-\kappa/2} \cdot (\kappa/2)!$ and $2n/3 < |x_t|_1 < 3n/4$, then with probability at least $1 - 2^{-2\kappa}$ we have $\lambda_{t+1} \leq 2^{-\kappa/2} \cdot (\kappa/2)!$.*

Proof. Since $\lambda_t \leq 2^{-\kappa/2} \cdot (\kappa/2)!$, a necessary condition for $\lambda_{t+1} > 2^{-\kappa/2} \cdot (\kappa/2)!$ is that generation t is unsuccessful. Since $|x_t|_1 < 3n/4$, the probability of finding an improvement in any mutation is at least $1/(4e)$ and the probability of an unsuccessful generation is at most

$$\begin{aligned}\left(1 - \frac{1}{4e}\right)^{F^{-1/s}2^{-\kappa/2}(\kappa/2)!} &\leq \left(\frac{4e}{4e-1}\right)^{-F^{-1/s}(2e)^{-\kappa/2}(\kappa/2)^{\kappa/2}}\\&= 2^{-F^{-1/s}(\kappa/(4e))^{\kappa/2}\log\left(\frac{4e}{4e-1}\right)}. \quad (13)\end{aligned}$$

The condition $\kappa \geq 324F^{1/s}$ implies

$$\frac{\kappa}{4e} \geq \left(1 + \frac{4}{\log\left(\frac{4e}{4e-1}\right)}\right)F^{1/s} \geq \left(1 + \frac{4F^{1/s}}{\log\left(\frac{4e}{4e-1}\right)}\right).$$

We bound the absolute value of the exponent in (13) using $(1+y)^x \geq xy$ for $x \in \mathbb{N}_0$, $y \geq 0$, as follows.

$$\begin{aligned}&F^{-1/s}(\kappa/(4e))^{\kappa/2}\log\left(\frac{4e}{4e-1}\right)\\&\geq F^{-1/s}\left(1 + \frac{4F^{1/s}}{\log\left(\frac{4e}{4e-1}\right)}\right)^{\kappa/2}\log\left(\frac{4e}{4e-1}\right)\\&\geq F^{-1/s}\frac{4F^{1/s}}{\log\left(\frac{4e}{4e-1}\right)} \cdot \kappa/2 \cdot \log\left(\frac{4e}{4e-1}\right) = 2\kappa.\end{aligned}$$

Hence the probability of an unsuccessful generation is at most $2^{-2\kappa}$. □

The following lemma now generalises and refines the "ratchet argument" from [17].

Lemma 4.15. *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. Let $T_{3n/4} = \inf\{t \mid |x_t|_1 \geq 3n/4\}$ be the number of generations until a search point with at least $3n/4$ ones is reached.*

There are constants $\gamma := \gamma(s, F) \in (1, 2]$ and $\kappa_0 := \kappa_0(s, F, \gamma) \geq 2$ such that for all $\kappa \geq \kappa_0$ the following holds. If the initial state (x_0, λ_0) is κ-safe with $|x_0|_1 < 3n/4$ then with probability at least $1 - \gamma^{-\Omega(\kappa)}$ all states during the next $\min\{\gamma^\kappa, T_{3n/4}\}$ generations are weakly κ-safe.

Proof. Let (x_0, λ_0) denote the initial state of the self-adjusting $(1, \lambda)$ EA. If $|x_0|_1 \geq 3n/4$ the statement is trivial, hence we assume $|x_0|_1 < 3n/4$. As in the proof of Lemma 3.7 in [17], we are setting up to apply the negative drift theorem [27, 28].

The value of γ will be determined later on, ensuring $1 < \gamma \leq 2$. Choosing $\kappa_0 \geq 324F^{1/s}$ and recalling that the initial state is κ-safe, Lemma 4.14 states that, with probability at least $1 - 2^{-2\kappa}$, $\lambda_1 \leq 2^{-\kappa/2} \cdot (\kappa/2)!$ as long as the number of ones is smaller than $3n/4$. By induction and a union bound, this holds for the first $\min\{\gamma^\kappa, T_{3n/4}\}$

generations with probability at least $1 - 2^{-2\kappa} \cdot \min\{\gamma^\kappa, T_{3n/4}\} \geq 1 - 2^{-2\kappa} \cdot \gamma^\kappa \geq 1 - \gamma^{-k}$ as $\gamma \leq 2$, unless the algorithm jumps back to the first slope. We assume in the following that the bound on λ_t always applies, while the algorithm remains on the second slope (and has not found a search point with at least $3n/4$ ones yet).

Let $\alpha := \frac{se}{e-1} \log_F \left(8e + \log_{\frac{e}{e-1}} (2/\varepsilon) F^{1/s}\right) = O(1)$ abbreviate the maximum difference between the potential h and the number of ones, then we start with a potential of at least $2n/3+\kappa-\alpha$. We apply the negative drift theorem [27, 28] to an interval $[a, b] := [2n/3 + \kappa/2, 2n/3+\kappa-\alpha]$ with respect to the current potential. By choosing $\kappa_0 \geq 6\alpha$, we can ensure that $b-a = \kappa-\alpha-\kappa/2 = \kappa/3+\kappa/6-\alpha \geq \kappa/3$.

We pessimistically assume that the number-of-ones component of h can only increase by at most 1. Lemma 4.13 has already shown that, even under this assumption, the drift is at least a positive constant. This implies the first condition of Theorem 2 in [28]. For the second condition, we need to bound transition probabilities for the potential. Owing to our pessimistic assumption, the number of ones can only increase by at most 1.

For jumps decreasing the number of ones, we need to argue more carefully. Let $i = |x_t|_1 \geq a$ be the current number of ones and let $p_{i,j}$ be the probability that $|x_{t+1}|_1 = i - j$. Note that a jump back to the first slope it is sufficient that one offspring has at most $2n/3$ ones. A necessary requirement is that j bits flip, which has probability at most $1/(j!)$. By a union bound over λ offspring, $p_{i,j} \leq \lambda/(j!) \leq 2^{-\kappa/2} \cdot (\kappa/2)!/(j!)$ using our bound on λ. For $j \geq \kappa/2$ (as $\kappa \geq 2$), we have $j! \geq (\kappa/2)! \cdot 2^{j-\kappa/2}$ and $p_{i,j} \leq 2^{-j}$. This implies for all $i \geq a$ and all j with $i - j \leq 2n/3$:

$$\Pr\left(|x_t|_1 - |x_{t+1}|_1 \geq j\right) \leq \sum_{j' \geq j} 2^{-j'} \leq 2 \cdot 2^{-j}.$$

In particular, $p^{\uparrow}_{i,\lambda} \leq \sum_{j=\kappa/2}^{n-i} p_{i,j} \leq \sum_{j=\kappa/2}^{\infty} 2 \cdot 2^{-j} = 4 \cdot 2^{-\kappa/2}$.

For $i - j > 2n/3$ the number of ones only decreases by j if *all* offspring decrease their number of ones by at least j, or if there is one offspring on the first slope. The probability of all offspring decreasing their number of ones by j is bounded by the probability that the first offspring decreases its number of ones by j. This is bounded by the probability of j bits flipping, which is at most $1/(j!) \leq 2/2^j$. Hence,

$$\forall i - j > 2n/3\colon \Pr\left(|x_t|_1 - |x_{t+1}|_1 \geq j\right) \leq 2 \cdot 2^{-j} + p^{\uparrow}_{i,\lambda} \leq 6 \cdot 2^{-j}.$$

The possible penalty in the definition of h changes by at most $\max\left(\frac{se}{e-1}, \frac{se}{e-1} \cdot \frac{1}{s}\right) = \frac{e}{e-1} < 1$. Hence, for all t,

$$\Pr\left(|h(X_{t-1}) - h(X_t)| \geq j + 1 \mid h(X_t) > a\right) \leq \frac{12}{2^{j+1}},$$

which meets the second condition of Theorem 2 in [28] when choosing $\delta := 1$ and $r(\ell) := 12$.

The negative drift theorem [27, 28] now implies that there exists a constant c^* such that the probability of the number of ones dropping below a in $2^{c^*(b-a)/12} \geq 2^{c^*\kappa/36}$ generations (or reaching a search point with at least $3n/4$ ones) is $2^{-\Omega((b-a)/12)} = 2^{-\Omega(\kappa)}$. Choosing $\gamma := \min\{2^{c^*/36}, 2\}$, this is at least γ^κ generations and a probability of $\gamma^{-\Omega(\kappa)}$ as claimed. Taking a union bound over this failure probability and that from Lemma 4.14 proves the claim. □

Now we show that with probability $\Omega(1)$ a search point with at least $3n/4$ ones is reached, without returning to the first slope and without resetting λ. Thus, with the claimed probability the algorithm behaves as the self-adjusting $(1, \lambda)$ EA from [17] on OneMax throughout this part of the run.

Lemma 4.16. *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. Assume the conditions from Lemma 4.15 hold for constants γ and $\kappa := \frac{\log \log n}{\log \log \log n}$. Then with probability $\Omega(1)$ a search point with at least $3n/4$ ones is reached within $O(n)$ generations.*

Moreover, with the claimed probability the algorithm does not go back to the first slope and does not reset λ before a search point with at least $3n/4$ ones is reached.

Proof. The statement of Lemma 4.15 satisfies the preconditions of Lemma 4.13 for $d = \kappa$. Then Lemma 4.13 implies a positive drift

$$\mathrm{E}\left(h(X_{t+1}) - h(X_t) \mid X_t\right) \geq \frac{1}{2e} - \frac{s}{2(e-1)} =: \delta$$

for the next γ^κ generations, unless a search point with at least $3n/4$ ones has been reached. In the latter case we are done, hence we assume that the drift is bounded from below as stated through the next $t_\kappa := \min\{\gamma^\kappa, n/\delta\}$ generations. By the additive drift theorem [15], the expected time to increase the potential by $\delta/12 \cdot \min\{\gamma^\kappa, n/\delta\}$, while the drift bound holds, is at most $\frac{\delta/12 \cdot \min\{\gamma^\kappa, n/\delta\}}{\delta} = \min\{\gamma^\kappa/12, n/(12\delta)\}$. By Markov's inequality, the probability that after t_κ steps the potential has not increased by $\delta/12 \cdot \min\{\gamma^\kappa, n/\delta\} = \min\{\delta/12 \cdot \gamma^\kappa, n/12\}$ is at most $1/12$.

Assuming that the potential has increased by $\min\{\delta/12 \cdot \gamma^\kappa, n/12\}$, by Definition 4.13 the number of ones has increased by $\min\{\delta/12 \cdot \gamma^\kappa, n/12\} - O(1)$. Since we start with at least $2n/3 + \kappa = 2n/3 + \omega(1)$ ones, there is a generation amongst the next γ^κ generations in which the number of ones is at least $\min\{2n/3 + \delta/12 \cdot \gamma^\kappa, 3n/4\}$.

Let $\kappa_0 := \kappa = \frac{\log \log n}{\log \log \log n}$ and define $\kappa_i := \delta/12 \cdot \gamma^{\kappa_{i-1}}$ for $i > 0$. Note that we have just showed that we have found a search point with at least $\min\{\kappa_1, 3n/4\}$ ones. If the number of ones is less than $3n/4$, the current state (x_t, λ_t) is κ_1-safe as it is weakly κ_0-safe and so $\lambda_t \leq 2^{-\kappa_0/2} \cdot (\kappa_0/2)! \leq 2^{-\kappa_1/2} \cdot (\kappa_1/2)!$.

Iterating the above argument with κ_1 instead of κ_0, we find a search point with at least $\min\{2n/3 + \kappa_2, 3n/4\}$ ones within the next $\min\{\gamma^{\kappa_1}, n/\delta\}$ generations, with probability at least $1 - 1/12 - \gamma^{-\Omega(\kappa_1)}$. We again iterate the argument with κ_2 and once again with κ_3. We claim that $t_{\kappa_2} := \min\{\gamma^{\kappa_2}, n/\delta\} = n/\delta$ and show this by bounding κ_1, κ_2 and κ_3 from below.

$$\kappa_1 = \frac{\delta}{12} \gamma^{\frac{\log \log n}{\log \log \log n}} = 2^{\log(\gamma) \cdot \frac{\log \log n}{\log \log \log n} + \log(\delta/12)}$$
$$\geq 2^{\log\left(\frac{2}{\log(\gamma)} \log \log n\right)} = \frac{2}{\log(\gamma)} \log \log n.$$

Now, κ_2 is at least

$$\kappa_2 = \frac{\delta}{12} \gamma^{\kappa_1} \geq \frac{\delta}{12} \gamma^{\frac{2}{\log(\gamma)} \log \log n} = \frac{\delta}{12} \log^2 n \geq \frac{2}{\log(\gamma)} \log n$$

for n large enough. Likewise,

$$\kappa_3 = \frac{\delta}{12} \gamma^{\kappa_2} \geq \frac{\delta}{12} \cdot \gamma^{\frac{2}{\log(\gamma)} \log n} = \frac{\delta}{12} \cdot n^2.$$

Together, we have that within $\min\{\gamma^{\kappa_0}, n/\delta\} + \min\{\gamma^{\kappa_1}, n/\delta\} + \min\{\gamma^{\kappa_2}, n/\delta\} = O(n)$ generations, with probability at least $1 - \frac{3}{12} - \gamma^{-\Omega(\kappa_0)} - \gamma^{-\Omega(\kappa_1)} - \gamma^{-\Omega(\kappa_2)} \geq \frac{3}{4} - 3\gamma^{-\Omega(\kappa_0)} = \Omega(1)$ we have reached a search point with at least $3n/4$ ones, without going back to the first slope. The probability of a reset during $O(n)$ expected generations is exponentially small by Lemma 4.3 and (9), hence this failure probability can be absorbed in the $\Omega(1)$ probability bound. This completes the proof. □

4.4 Finding the global optimum

Once the self-adjusting $(1, \lambda)$ EA moves far away from the cliff, the probability of jumping back up the cliff is reduced, and the next part of the optimisation resembles OneMax. The algorithm still can reset λ to 1. Such a steep decrease of λ would typically make the algorithm slip down the second slope until λ recovers to large enough values that support hill climbing. Hence, resets would break the runtime analysis made in [17]. We show in Lemma 4.17 that, with probability $\Omega(1)$, the algorithm neither jumps back up the cliff nor resets λ during the last part of the optimisation. This allows us to apply the previous analysis from [17] on OneMax.

Lemma 4.17. *Consider the self-adjusting $(1, \lambda)$ EA as in Theorem 4.1. For any initial $\lambda_0 \leq \lambda_{\max}$ with $\lambda_0 = O(n \log n)$ and any initial search point x_0 with $|x_0|_1 \geq 3n/4$ the probability that the self-adjusting $(1, \lambda)$ EA creates the optimum without jumping back up the cliff or resetting λ to 1 is at least $1 - \frac{1}{e-1} - O\left(\frac{\log^3(n)}{n}\right)$.*

Proof. As long as the self-adjusting $(1, \lambda)$ EA does not jump back up the cliff, the self-adjusting $(1, \lambda)$ EA behaves as the self-adjusting $(1, \lambda)$ EA on OneMax. Additionally, if it does not have an unsuccessful generation with $\lambda = \lambda_{\max}$ it will never reset to 1, behaving as the self-adjusting $(1, \lambda)$ EA studied in [17].

From [17, Theorem 3.1 and Theorem 3.5] we know that the self-adjusting $(1, \lambda)$ EA solves OneMax in expected $O(n)$ generations and $O(n \log n)$ evaluations. Therefore, within these expected times our algorithm either finds the global optimum, jumps back up the cliff or resets 1. We show that the with probability $\Omega(1)$, a global optimum is reached.

In order for λ to reset at the same time as there is a jump back up the cliff, at least one offspring must flip $n/3$ one-bits and all other offspring must not increase their fitness. The probability of flipping $n/3$ bits is $n^{-\Omega(n)}$, hence the probability of both events happening at the same time is at most $n^{-\Omega(n)}$.

By Lemma 3.7 in [17] if the initial search point x_0 has $|x_0|_1 \geq 3n/4$, with probability $1 - O(1/n)$ the number of one-bits will never drop below $3n/4 - O(\log n)$ before finding the optimum. This means that, for the algorithm to jump back up the cliff at least one offspring must flip a linear amount of bits. The probability that one offspring flips a linear amount of bits is $n^{-\Omega(n)}$. By (9), the probability that this happens during $O(n \log n)$ expected evaluations is still $n^{-\Omega(n)}$. In the following we assume that we never return to the first slope.

To show that there is never an unsuccessful generation with $\lambda = \lambda_{\max}$ (i.e. λ never resets to 1) we divide the optimisation in two phases. The first phase ends the first time a state (x_t, λ_t) is found with $\lambda_t \geq 4 \log n$ and $|x_t|_1 \geq n - 3 \ln n$ or the optimum is found, and the second phase ends when the optimum is found.

During the first phase, since $\lambda_{\max} > 4 \log n$, we can only reach $\lambda = \lambda_{\max}$ if $|x|_1 < n - 3 \ln n$ otherwise we would start phase two. In order to reach $\lambda = \lambda_{\max}$ at least one generation with $\lambda \geq en$ must be unsuccessful. The probability of an unsuccessful generation with $\lambda \geq en$ is at most

$$1 - p_{i,\lambda}^{\rightarrow} + p_{i,\lambda}^{\uparrow} \leq \left(1 - \frac{n-i}{en}\right)^{en} \leq \left(1 - \frac{3 \ln n}{en}\right)^{en} \leq e^{-3 \ln n} = n^{-3}.$$

Given that the optimum is found after $O(n)$ expected generations, by (9) the probability of reaching $\lambda = \lambda_{\max}$ during the first phase is $O(1/n^2)$.

For the second phase we first argue that the current fitness does not decrease, with high probability. The second phase starts with $\lambda \geq 4 \log n$ and by Lemma 3.7 in [17] while $\lambda_t \geq 4 \log n$ the fitness is not reduced before reaching the optimum with probability $1 - O(1/n)$. We now show that $\lambda \geq 4 \log n$ throughout the remainder of the run with high probability.

By Lemma 3.6 in [17] from $|x|_1 \geq n - 3 \ln n$ in expectation the optimum will be reached in $O(\log n)$ generations. To reduce λ to a value smaller than $4 \log n$, a generation with $\lambda < 4F \log n$ must be successful. This event has a probability of at most

$$1 - \left(\frac{i}{n}\right)^{\lambda} \leq 1 - \left(1 - \frac{3 \ln n}{n}\right)^{4F \log n} \leq \frac{12F \log^2 n}{n \log e} = O\left(\frac{\log^2(n)}{n}\right).$$

By (9) the probability that λ is reduced to a value less than $4 \log n$ during the next $O(\log n)$ generations is $O\left(\frac{\log^3(n)}{n}\right)$. Accounting for both failures with probability $1 - O\left(\frac{\log^3(n)}{n}\right)$ each fitness value is left at most once.

Now we can calculate the probability of resetting λ by considering at most one generation with $\lambda = \lambda_{\max}$ per fitness value. We only have a reset of λ if one such generation is unsuccessful. Thus, the probability of resetting λ during the second phase is at most

$$\begin{aligned}
\sum_{i=n-3\ln n}^{n-1} \left(1 - p_{i,\lambda_{\max}}^{+}\right) &\leq \sum_{i=n-3\ln n}^{n-1} \left(1 - \frac{n-i}{en}\right)^{enF^{1/s}} \\
&\leq \sum_{i=n-3\ln n}^{n-1} e^{-F^{1/s}(n-i)} = \sum_{j=1}^{3\ln n} e^{-F^{1/s} j} \leq \sum_{j=1}^{\infty} e^{-F^{1/s} j} \\
&= \frac{1}{1 - e^{-F^{1/s}}} - 1 = \frac{1}{\exp(F^{1/s}) - 1} \leq \frac{1}{e-1}.
\end{aligned}$$

Adding up all failure probabilities completes the proof. □

4.5 Putting Things Together

Now we are able to prove the claimed bounds of $O(n)$ expected generations and $O(n \log n)$ expected evaluations from Theorem 4.1.

Proof of Theorem 4.1. From any initial state, by Lemma 4.2 we reach a solution x_t with $|x_t|_1 \geq 2n/3$ in expected $O(n)$ generations and $O(n + \lambda_{\max})$ evaluations.

Then, by Lemma 4.9, the algorithm reaches a κ-safe state (for $\kappa := \frac{\log \log n}{\log \log \log n}$) or a search point with at least $3n/4$ ones in $O(\log(\lambda_{\max}) \log n)$ expected generations and $O(\lambda_{\max} \log n)$ expected evaluations.

Together, along with $\lambda_{\max} = \Omega(n)$ and $\lambda_{\max} = \text{poly}(n)$, the total time to reach a κ-safe state or a search point with at least $3n/4$

ones from an arbitrary initial state is $O(\log(\lambda_{\max})\log n + n) = O(n)$ generations and $O(\lambda_{\max}\log n)$ evaluations.

By Lemma 4.16 with probability $\Omega(1)$ we reach a search point with at least $3n/4$ ones within $O(n)$ generations, without going back to the first slope or resetting λ. During this time, the algorithm behaves like the self-adjusting $(1,\lambda)$ EA without resetting on OneMax and we obtain an upper bound of $O(n\log n)$ evaluations from [17]. Hence, the expected number of evaluations until we return to the first slope, reset λ or reach a search point with $3n/4$ ones is $O(n\log n)$. In expectation, a constant number of trials suffices to find a search point with at least $3n/4$ ones. Hence, from any initial state, in expectation in $O(n)$ generations and $O(\lambda_{\max}\log n)$ evaluations we reach a search point with at least $3n/4$ ones.

Likewise, from a search point with at least $3n/4$ ones, by Lemma 4.17 with probability $\Omega(1)$ we find the optimum without resetting λ or returning to the first slope, and hence the analysis from [17] still applies. Thus, in expected $O(n)$ generations and $O(n\log n)$ evaluations we either reach the global optimum, return to the first slope or reset λ, and the probability of reaching the optimum is $\Omega(1)$. Iterating this argument an expected constant number of times proves the claimed bound. □

5 CONCLUSIONS

The usefulness of parameter control has so far mainly been demonstrated for elitist EAs on relatively easy problems. For the more difficult multimodal problem Cliff we showed that the self-adjusting $(1,\lambda)$ EA using success-based rules and a reset mechanism can find the global optimum in $O(n)$ expected generations and $O(n\log n)$ expected evaluations. This is a speedup of order $\Omega(n^{2.9767}/\log n)$ over the expected optimisation time with the best fixed value of λ.

The latter conclusion was obtained by refining the previous bounds on the expected optimisation time of the $(1,\lambda)$ EA on Cliff from [18], $O(e^{5\lambda}) = O(148.413^{\lambda})$ and $\min\{n^{n/4}, e^{\lambda/4}\}/3 = \min\{n^{n/4}, 1.284^{\lambda}\}/3$, towards bounds of $\Omega(\xi^{\lambda})$ and $O(\xi^{\lambda}\log n)$, for $\xi \approx 6.196878$, revealing the degree of the polynomial in the expected runtime of the $(1,\lambda)$ EA with the best fixed λ as $\eta \approx 3.976770136$.

Our results demonstrate the power of parameter control for the multimodal Cliff problem and that drastic performance improvement can be obtained. Several open questions remain, for instance, in how far our results generalise to Cliff functions where the position of the cliff is chosen differently from $2n/3$ ones and how the considered algorithm would perform on other problems.

REFERENCES

[1] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. 2014. Unbiased Black-Box Complexity of Parallel Search. In *Proc. of PPSN XIII*. Springer, 892–901.

[2] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. 2015. Black-box Complexity of Parallel Search with Distributed Populations. In *Proc. of FOGA*. ACM, 3–15.

[3] Süntje Böttcher, Benjamin Doerr, and Frank Neumann. 2010. Optimal Fixed and Adaptive Mutation Rates for the LeadingOnes Problem. In *Proc. of PPSN XI*, Vol. 6238. Springer, 1–10.

[4] Brendan Case and Per Kristian Lehre. 2020. Self-Adaptation in Nonelitist Evolutionary Algorithms on Discrete Problems with Unknown Structure. *IEEE Trans. Evol. Comput.* 24, 4 (2020), 650–663.

[5] Dogan Corus, Pietro S. Oliveto, and Donya Yazdani. 2020. When Hypermutations and Ageing Enable Artificial Immune Systems to Outperform Evolutionary Algorithms. *Theor. Comput. Sci.* 832 (2020), 166–185.

[6] Duc-Cuong Dang and Per Kristian Lehre. 2016. Self-adaptation of Mutation Rates in Non-elitist Populations. In *Proc. of PPSN XIV*. Springer, Cham, 803–813.

[7] Benjamin Doerr. 2020. *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization.* Springer, Chapter Probabilistic Tools for the Analysis of Randomized Optimization Heuristics, 1–87.

[8] Benjamin Doerr and Carola Doerr. 2018. Optimal Static and Self-Adjusting Parameter Choices for the (1+(λ,λ)) Genetic Algorithm. *Algorithmica* 80, 5 (2018), 1658–1709.

[9] Benjamin Doerr and Carola Doerr. 2020. *Theory of Evolutionary Computation: Recent Developments in Discrete Optimization.* Springer, Chapter Theory of Parameter Control for Discrete Black-box Optimization: Provable Performance Gains Through Dynamic Parameter Choices, 271–321.

[10] Benjamin Doerr, Carola Doerr, and Franziska Ebel. 2015. From Black-Box Complexity to Designing New Genetic Algorithms. In *Theor. Comput. Sci.*, Vol. 567. 87–104.

[11] Benjamin Doerr, Carola Doerr, and Johannes Lengler. 2019. Self-Adjusting Mutation Rates with Provably Optimal Success Rules. In *Proc. of GECCO*. ACM.

[12] Benjamin Doerr, Christian Gießen, Carsten Witt, and Jing Yang. 2019. The (1+ λ) Evolutionary Algorithm with Self-Adjusting Mutation Rate. *Algorithmica* 81, 2 (2019), 593–631.

[13] Benjamin Doerr, Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2018. On the Runtime Analysis of Selection Hyper-Heuristics with Adaptive Learning Periods. In *Proc. of GECCO*. ACM, 1015–1022.

[14] Benjamin Doerr, Carsten Witt, and Jing Yang. 2021. Runtime Analysis for Self-adaptive Mutation Rates. *Algorithmica* 83, 4 (2021), 1012–1053.

[15] Jun He and Xin Yao. 2004. A Study of Drift Analysis for Estimating Computation Time of Evolutionary Algorithms. *Nat. Comput.* 3, 1 (2004), 21–35.

[16] Mario Alejandro Hevia Fajardo and Dirk Sudholt. 2020. On the Choice of the Parameter Control Mechanism in the $(1 + (\lambda, \lambda))$ Genetic Algorithm. In *Proc. of GECCO*. ACM, 832–840.

[17] Mario Alejandro Hevia Fajardo and Dirk Sudholt. 2021. Self-Adjusting Population Sizes for Non-Elitist Evolutionary Algorithms: Why Success Rates Matter. In *Proc. of GECCO*. ACM, 1151–1159.

[18] Jens Jägersküpper and Tobias Storch. 2007. When the Plus Strategy Outperforms the Comma Strategy and When Not. In *Proc. of IEEE FOCI*. IEEE, 25–32.

[19] Thomas Jansen and Dirk Sudholt. 2010. Analysis of an Asymmetric Mutation Operator. *Evol. Comput.* 18, 1 (2010), 1–26.

[20] Jörg Lässig and Dirk Sudholt. 2011. Adaptive Population Models for Offspring Populations and Parallel Evolutionary Algorithms. In *Proc. of FOGA*. ACM, 181–192.

[21] Per Kristian Lehre and Carsten Witt. 2012. Black-Box Search by Unbiased Variation. *Algorithmica* 64, 4 (2012), 623–642.

[22] Johannes Lengler. 2020. A General Dichotomy of Evolutionary Algorithms on Monotone Functions. *IEEE Trans. Evol. Comput.* 24, 6 (2020), 995–1009.

[23] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2019. On the Time Complexity of Algorithm Selection Hyper-Heuristics for Multimodal Optimisation. In *Proc. of AAAI*, Vol. 33. 2322–2329.

[24] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2020. How the Duration of the Learning Period Affects the Performance of Random Gradient Selection Hyper-Heuristics. In *Proc. of AAAI*, Vol. 34. 2376–2383.

[25] Andrei Lissovoi, Pietro S. Oliveto, and John Alasdair Warwicker. 2020. Simple Hyper-heuristics Control the Neighbourhood Size of Randomised Local Search Optimally for LeadingOnes. *Evol. Comput.* 28, 3 (2020), 437–461.

[26] Andrea Mambrini and Dirk Sudholt. 2015. Design and Analysis of Schemes for Adapting Migration Intervals in Parallel Evolutionary Algorithms. *Evol. Comput.* 23, 4 (2015), 559–582.

[27] Pietro S. Oliveto and Carsten Witt. 2011. Simplified Drift Analysis for Proving Lower Bounds in Evolutionary Computation. *Algorithmica* 59, 3 (2011), 369–386.

[28] Pietro S. Oliveto and Carsten Witt. 2012. Erratum: Simplified Drift Analysis for Proving Lower Bounds in Evolutionary Computation. *ArXiv e-prints* (2012). arXiv:1211.7184

[29] Tiago Paixão, Jorge Pérez Heredia, Dirk Sudholt, and Barbora Trubenová. 2017. Towards a Runtime Comparison of Natural and Artificial Evolution. *Algorithmica* 78, 2 (2017), 681–713.

[30] Amirhossein Rajabi and Carsten Witt. 2020. Evolutionary Algorithms with Self-adjusting Asymmetric Mutation. In *Proc. of PPSN XVI*. Springer, 664–677.

[31] Amirhossein Rajabi and Carsten Witt. 2020. Self-Adjusting Evolutionary Algorithms for Multimodal Optimization. In *Proc. of GECCO*. ACM, 1314–1322.

[32] Amirhossein Rajabi and Carsten Witt. 2021. Stagnation Detection with Randomized Local Search. In *EvoCOP*. Springer, 152–168. Full version available at http://arxiv.org/abs/2101.12054.

[33] Jonathan E. Rowe and Dirk Sudholt. 2014. The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theor. Comput. Sci.* 545 (2014), 20–38.

[34] Abraham Wald. 1944. On Cumulative Sums of Random Variables. *Ann. Math. Stat.* 15, 3 (1944), 283 – 296.

[35] Carsten Witt. 2013. Tight Bounds on the Optimization Time of a Randomized Search Heuristic on Linear Functions. *Comb. Probab. Comput.* 22, 2 (2013), 294–318.