

Upper Bounds on the Running Time of the Univariate Marginal Distribution Algorithm on OneMax

Carsten Witt¹ 

Received: 13 June 2017 / Accepted: 2 June 2018 / Published online: 11 June 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract The Univariate Marginal Distribution Algorithm (UMDA) is a randomized search heuristic that builds a stochastic model of the underlying optimization problem by repeatedly sampling λ solutions and adjusting the model according to the best μ samples. We present a running time analysis of the UMDA on the classical OneMax benchmark function for wide ranges of the parameters μ and λ . If $\mu \geq c \log n$ for some constant $c > 0$ and $\lambda = (1 + \Theta(1))\mu$, we obtain a general bound $O(\mu n)$ on the expected running time. This bound crucially assumes that all marginal probabilities of the algorithm are confined to the interval $[1/n, 1 - 1/n]$. If $\mu \geq c' \sqrt{n} \log n$ for a constant $c' > 0$ and $\lambda = (1 + \Theta(1))\mu$, the behavior of the algorithm changes and the bound on the expected running time becomes $O(\mu \sqrt{n})$, which typically holds even if the borders on the marginal probabilities are omitted. The results supplement the recently derived lower bound $\Omega(\mu \sqrt{n} + n \log n)$ by Krejca and Witt (Proceedings of FOGA 2017, ACM Press, New York, pp 65–79, 2017) and turn out to be tight for the two very different choices $\mu = c \log n$ and $\mu = c' \sqrt{n} \log n$. They also improve the previously best known upper bound $O(n \log n \log \log n)$ by Dang and Lehre (Proceedings of GECCO '15, ACM Press, New York, pp 513–518, 2015) that was established for $\mu = c \log n$ and $\lambda = (1 + \Theta(1))\mu$.

Keywords Randomized search heuristics · Estimation-of-distribution algorithms · UMDA · Running time analysis

An extended abstract of this article appeared in the proceedings of the 2017 Genetic and Evolutionary Computation Conference (GECCO 2017) [28].

✉ Carsten Witt
cawi@dtu.dk

¹ DTU Compute, Technical University of Denmark, Kongens Lyngby, Denmark

1 Introduction

Estimation-of-distribution algorithms (EDAs, [15]) are randomized search heuristics that have emerged as a popular alternative to classical evolutionary algorithms like Genetic Algorithms. In contrast to the classical approaches, EDAs do not store explicit populations of search points but develop a probabilistic model of the fitness function to be optimized. Roughly, this model is built by sampling a number of search points from the current model and updating it based on the structure of the best samples.

Although many different variants of EDAs (cf. [11]) and many different domains are possible, theoretical analysis of EDAs in discrete search spaces often considers running time analysis over $\{0, 1\}^n$. The simplest of these EDAs have no mechanism to learn correlations between bits. Instead, they store a Poisson binomial distribution, i. e., a probability vector \mathbf{p} of n independent probabilities, each component p_i denoting the probability that a sampled bit string will have a 1 at position i .

The first theoretical analysis in this setting was conducted by Droste [7], who analyzed the *compact Genetic Algorithm* (cGA), an EDA that only samples two solutions in each iteration, on linear functions. Papers considering other EDAs [2–5] followed. Also iteration-best *Ant Colony Optimization* (ACO), historically classified as a different type of search heuristic, can be considered as an EDA and analyzed in the same framework [21].

Recently, the interest in the theoretical running time analysis of EDAs has increased [6, 8, 9, 14, 26]. Most of these works derive bounds for a specific EDA on the popular ONEMAX function, which counts the number of 1s in a bit string and is considered to be one of the easiest functions with a unique optimum [25, 27]. In this paper, we follow up on recent work on the *Univariate Marginal Distribution Algorithm* (UMDA [20]) on ONEMAX.

The UMDA is an EDA that samples λ solutions in each iteration, selects $\mu < \lambda$ best solutions, and then sets the probability p_i (hereinafter called *frequency*) to the relative occurrence of 1s among these μ individuals. The algorithm has already been analyzed some years ago for several artificially designed example functions [2–5]. However, none these papers considered the most fundamental benchmark function in theory, the ONEMAX function. In fact, the running time analysis of the UMDA on the simple ONEMAX function has turned out to be rather challenging; the first such result, showing the upper bound $O(n \log n \log \log n)$ on its expected running time for certain settings of μ and λ , was not published until 2015 [6].

In a recent related study, Wu et al. [29] present the first running time analysis of the cross-entropy method (CE), which is a generalization of the UMDA and analyze it on ONEMAX and another benchmark function. Using $\mu = n^{1+\epsilon} \log n$ for some constant $\epsilon > 0$ and $\lambda = \omega(\mu)$, they obtain that the running time of CE on ONEMAX is $O(\lambda n^{1/2+\epsilon/3} / \rho)$ with overwhelming probability, where ρ is a parameter of CE. Hence, if $\rho = \Omega(1)$, including the special case $\rho = 1$ where CE collapses to UMDA, a running time bound of $O(n^{3/2+(4/3)\epsilon} \log n)$ holds, i. e., slightly above $n^{3/2}$. Technically, Wu et al. use concentration bounds such as Chernoff bounds to bound the effect of so-called genetic drift, which is also considered in the present paper, as well as anti-concentration results, in particular for the Poisson binomial distribution, to obtain their statements.

All bounds can hold with high probability only since CE is formulated without so-called borders on the frequencies.

Very recently, these upper bounds were supplemented by a general lower bound of the kind $\Omega(\mu\sqrt{n} + n \log n)$ [14], proving that the UMDA cannot be more efficient than simple evolutionary algorithms on this function, at least if $\lambda = (1 + \Theta(1))\mu$. As the upper bounds due to [6] and the recent lower bounds were apart by a factor of $\Theta(\log \log n)$, it was an open problem to determine the asymptotically best possible running time of the UMDA on ONEMAX.

In this paper, we close this gap and show that the UMDA can optimize ONEMAX in expected time $O(n \log n)$ for two very different, carefully chosen values of μ , always assuming that $\lambda = (1 + \Theta(1))\mu$. In fact, we obtain two general upper bounds depending on μ . If $\mu \geq c\sqrt{n} \log n$, where c is a sufficiently large constant, the first upper bound is $O(\mu\sqrt{n})$. This bound exploits that all p_i move more or less steadily to the largest possible value and that with high probability there are no frequencies that ever drop below $1/4$. Around $\mu = \Theta(\sqrt{n} \log n)$, there is a phase transition in the behavior of the algorithm. With smaller μ , the stochastic movement of the frequencies is more chaotic and many frequencies will hit the lowest possible value during the optimization. Still, the expected optimization time is $O(\mu n)$ for $\mu \geq c' \log n$ and a sufficiently large constant $c' > 0$ if all frequencies are confined to the interval $[1/n, 1 - 1/n]$, as typically done in EDAs. If frequencies are allowed to drop to 0, the algorithm will typically have infinite optimization time below the phase transition point $\mu \sim \sqrt{n} \log n$, whereas it typically will be efficient above.

Interestingly, Dang and Lehre [6] used $\mu = \Theta(\ln n)$, i. e., a value below the phase transition to obtain their $O(n \log n \log \log n)$ bound. This region turns out to be harder to analyze than the region above the phase transition, at least with our techniques. However, our proof also follows an approach being widely different from [6]. There the so-called level based theorem, a very general upper bound technique, is applied to track the stochastic behavior of the best-so-far ONEMAX-value. While this gives a rather short and elegant proof of the upper bound $O(n \log n \log \log n)$, the generality of the technique does not give much insight into how the probabilities p_i of the individuals bits develop over time. We think that it is crucial to understand the working principles of the algorithm thoroughly and present a detailed analysis of the stochastic process at bit level, as also done in many other running time analyses of EDAs [8, 9, 14, 26].

This paper is structured as follows: in Sect. 2, we introduce the setting we are going to analyze and summarize some tools from probability theory that are used throughout the paper. In particular, a new negative drift theorem is presented. It generalizes previous formulations by making milder assumptions on steps in the direction of the drift than on steps against the drift. In this section, we also give a detailed analysis of the update rule of the UMDA, which results in a bias of the frequencies p_i towards higher values. These techniques are presented for the ONEMAX-case, but contain some general insights that may be useful in analyses of different fitness functions. In Sect. 3, we prove the upper bound for the case of μ above the phase transition point $\Theta(\sqrt{n} \log n)$. The case of μ below this point is dealt with in Sect. 4. We finish with some conclusions. The “Appendix” gives a self-contained proof of the new drift theorem.

Independent, related work Very recently, Lehre and Nguyen [16] independently obtained the upper bound $O(\lambda n)$ for $c \log n \leq \mu = O(\sqrt{n})$ and $\lambda = \Omega(\mu)$ using a refined application of the so-called level-based method. Our approach also covers larger μ (but requires $\lambda = \mu(1 + \Theta(1))$) and is technically different.

2 Preliminaries

We consider the so-called *Univariate Marginal Distribution Algorithm* (UMDA [20]) in Algorithm 1 that maximizes the pseudo-Boolean function f . Throughout this paper, we have $f := \text{ONEMAX}$, where, for all $x = (x_1, \dots, x_n) \in \{0, 1\}^n$,

$$\text{ONEMAX}(x) = \sum_{i=1}^n x_i.$$

Note that the unique maximum is the all-ones bit string. However, a more general version can be defined by choosing an arbitrary optimum $a \in \{0, 1\}^n$ and defining, for all $x \in \{0, 1\}^n$, $\text{ONEMAX}_a(x) = n - d_H(x, a)$, where $d_H(x, a)$ denotes the Hamming distance of the bit strings x and a . Note that ONEMAX_{1^n} is equivalent to the original definition of ONEMAX . Our analyses hold true for any function ONEMAX_a , with $a \in \{0, 1\}^n$, due to symmetry of the UMDA’s update rule.

Algorithm 1: Univariate Marginal Distribution Algorithm (UMDA); algorithm UMDA* is obtained if the line indexed [R] is omitted.

```

t ← 0, pt,1 ← pt,2 ← ⋯ ← pt,n ← 1/2;
while termination criterion not met do
    Pt ← ∅;
    for j ∈ {1, …, λ} do
        for i ∈ {1, …, n} do
            xt,i(j) ← 1 with prob. pt,i and xt,i(j) ← 0 with prob. 1 - pt,i;
        Pt ← Pt ∪ {xt(j)};
    Sort individuals in P descending by fitness (such that f(xt(1)) ≥ ⋯ ≥ f(xt(μ))), breaking ties
    uniformly at random;
    for i ∈ {1, …, n} do
        pt+1,i ← (∑j=1μ xt,i(j)) / μ;
        [R] Restrict pt+1,i to be within [1/n, 1 - 1/n];
    t ← t + 1;

```

We call bit strings *individuals* and their respective ONEMAX -values *fitness*.

The UMDA does not store an explicit population but does so implicitly, as usual in EDAs. For each of the n different bit positions, it stores a rational number p_i , which we call *frequency*, determining how likely it is that a hypothetical individual would have a 1 at this position. In other words, the UMDA stores a probability distribution over $\{0, 1\}^n$. The starting distribution samples according to the uniform distribution, i. e., $p_i = 1/2$ for $i \in \{1, \dots, n\}$.

In each so-called generation t , the UMDA samples λ individuals such that each individual has a 1 at position i , where $i \in \{1, \dots, n\}$ with probability $p_{t,i}$, independent of all the other frequencies. Thus, the number of 1s is sampled according to a Poisson binomial distribution with probability vector $\mathbf{p}_t = (p_{t,i})_{i \in \{1, \dots, n\}}$.

After sampling λ individuals, μ of them with highest fitness are chosen, breaking ties uniformly at random (so-called *selection*). Then, for each position, the respective frequency is set to the relative occurrence of 1s in this position. That is, if the chosen μ best individuals have x 1s at position i among them, the frequency p_i will be updated to x/μ for the next iteration. Note that such an update allows large jumps like, e.g., from $(\mu - 1)/\mu$ to $1/\mu$.

If a frequency is either 0 or 1, it cannot change anymore since then all values at this position will be either 0 or 1. To prevent the UMDA from getting stuck in this way, we narrow the interval of possible frequencies down to $[1/n, 1 - 1/n]$ and call $1/n$ and $1 - 1/n$ the *borders* for the frequencies. Hence, there is always a chance of sampling 0s and 1s for each position. This is a common approach used by other EDAs as well, such as the cGA or ACO algorithms (cf. the references given in the introduction). We also consider a variant of the UMDA called UMDA* where the borders are not used. That algorithm will typically not have finite expected running time; however, it might still be efficient with high probability if it is sufficiently unlikely that frequencies get stuck at bad values.

Overall, we are interested in upper bounds on the UMDA's expected number of *function evaluations* on ONEMAX until the optimum is sampled; this number is typically called *running time* or *optimization time*. Note that this equals λ times the expected number of generations until the optimum is sampled.

In all of our analyses, we assume that $\lambda = (1 + \beta)\mu$ for some arbitrary constant $\beta > 0$ and use μ and λ interchangeably in asymptotic notation. Of course, we could also choose $\lambda = \omega(\mu)$ but then each generation would be more expensive. Choosing $\lambda = \Theta(\mu)$ lets us basically focus on the minimal number of function evaluations per generation, as μ of them are at least needed to make an update.

2.1 Useful Tools from Probability Theory

We will see that the number of 1s sampled by the UMDA at a certain position is binomially distributed with the frequency as success probability. In our analyses, we will therefore often have to bound the tail of binomial and related distributions. To this end, many classical techniques such as Chernoff-Hoeffding bounds exist. The following version, which includes the knowledge of the variance, is particularly handy to use.

Lemma 1 [18] *If X_1, \dots, X_n are independent, and $X_i - E(X_i) \leq b$ for $i \in \{1, \dots, n\}$, then for $X := X_1 + \dots + X_n$ and any $d \geq 0$ it holds that*

$$\Pr(X - E(X) \geq d) \leq e^{-\frac{d^2}{2\sigma^2(1+\delta/3)}},$$

where $\sigma^2 := \text{Var}(X)$ and $\delta := bd/\sigma^2$.

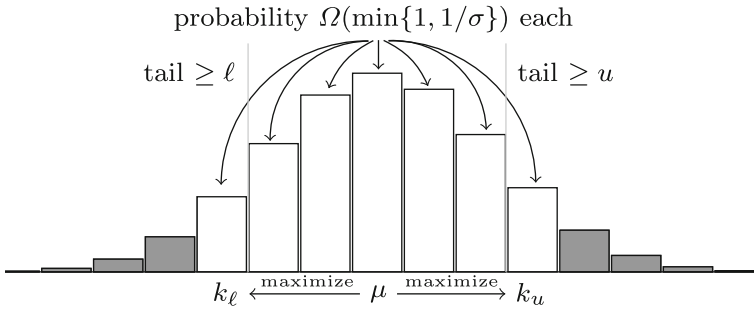


Fig. 1 Illustration of Lemma 2

The following lemma describes a result regarding the Poisson binomial distribution which we find very intuitive. However, as we did not find a sufficiently related result in the literature, we give a self-contained proof here. Roughly, the lemma considers a chunk of the distribution around the expected value whose joint probability is a constant less than 1 and then argues that every point in the chunk has a probability that is at least inversely proportional to the variance. See Fig. 1 for an illustration.

Lemma 2 *Let X_1, \dots, X_n be independent Poisson trials. Denote $p_i = \Pr(X_i = 1)$ for $i \in \{1, \dots, n\}$, $X := \sum_{i=1}^n X_i$, $\mu := E(X) = \sum_{i=1}^n p_i$ and $\sigma^2 := \text{Var}(X) = \sum_{i=1}^n p_i(1 - p_i)$. Given two constants $\ell, u \in (0, 1)$ such that $\ell + u < 1$, let $k_\ell := \min\{i \mid \Pr(X \leq i) \geq \ell\}$ and $k_u := \max\{i \mid \Pr(X \geq i) \geq u\}$. Then it holds that $\Pr(X = k) = \Omega(\min\{1, 1/\sigma\})$ for all $k \in \{k_\ell, \dots, k_u\}$, where the Ω -notation is with respect to n .*

Proof To begin with, we note that $k_\ell \leq k_u$. This holds since by assumption $\Pr(X < k_\ell) < \ell$ and $\Pr(X > k_u) < u$, hence $\Pr(k_\ell \leq X \leq k_u) \geq 1 - \ell - u > 0$, using $\ell + u < 1$. If $k_\ell > k_u$ happened, we would obtain a contradiction.

We first handle the case $\sigma = o(1)$ separately. This implies $\Pr(X - E(X) \geq 1/2) = o(1)$ and analogously $\Pr(X - E(X) \leq -1/2) = o(1)$. Namely, if we had $\Pr(X - E(X) \geq 1/2) = \Omega(1)$, then $E((X - E(X))^2) = \Omega(1)$, contradicting the assumption $\sigma = o(1)$; analogously for the other inequality.

Let $[E(X)]$ be the integer closest to $E(X)$, which is unique since, as argued in the previous paragraph, $E(X) - [E(X)] = 1/2$ would contradict $\sigma = o(1)$. We assume $[E(X)] = \lceil E(X) \rceil$ and note that the case $[E(X)] = \lfloor E(X) \rfloor$ is analogous. From the previous paragraph, we now obtain that $\Pr(X \leq [E(X)]) = o(1)$ and therefore $\Pr(X \geq [E(X)]) = 1 - o(1)$. Moreover, again using $\sigma = o(1)$, we also obtain $\Pr(X \geq [E(X)] + 1) = o(1)$. Since ℓ and u are positive constants less than 1, it immediately follows that $k_\ell = k_u = [E(X)]$ and $\Pr(X = k_\ell) = 1 - o(1) = \Omega(1)$.

In the following, we assume $\sigma = \Omega(1)$. We use that $k_\ell \leq [E(X)]$ or $k_u \geq [E(X)]$ (or both) must hold; otherwise, since $[E(X)] \leq [E(X)] + 1$, we would contradict the fact $k_\ell \leq k_u$. Hereinafter we consider the case $k_\ell \leq [E(X)]$ and note that other case is symmetrical. We start by proving $\Pr(X = k_\ell) = \Omega(1)$. To this end, we recall the unimodality of the Poisson binomial distribution function, more precisely $\Pr(X = i) \leq \Pr(X = i + 1)$ for $i \leq [E(X)] - 1$ and $\Pr(X = i) \geq \Pr(X = i + 1)$ for

$i \geq \lceil E(X) \rceil$ [24]. Hence, denoting $\alpha := \Pr(X = k_\ell)$, we have $\Pr(X = i) \leq \alpha$ for all $i \leq k_\ell$. It follows that $\Pr(X \leq k_\ell - \ell/(2\alpha)) \geq \ell/2 = \Omega(1)$ since $\Pr(X \leq k_\ell) \geq \ell$ by definition. We remark (but do not use) that this also implies a lower bound on k_ℓ .

If $\alpha = o(1/\sigma)$ held, the fact that $\Pr(X \leq k_\ell - \ell/(2\alpha)) = \Omega(1)$ would imply $\sqrt{\text{Var}(X)} = \Omega(1/\alpha) = \omega(\sigma)$, contradicting our assumption $\sqrt{\text{Var}(X)} = \sigma$. Hence, $\Pr(X = k_\ell) = \Omega(1/\sigma)$. Again using the monotonicity of the Poisson binomial distribution, we have $\Pr(X = i) = \Omega(1/\sigma)$ for all $i \in \{k_\ell, \dots, \lfloor E(X) \rfloor\}$. If $k_u \leq \lfloor E(X) \rfloor$, this already proves $\Pr(X = i) = \Omega(1/\sigma)$ for all $i \in \{k_\ell, \dots, k_u\}$ and nothing is left to show. Otherwise the bound follows for the remaining $i \in \{\lceil E(X) \rceil, \dots, k_u\}$ by a symmetrical argument, more precisely by first showing that $\Pr(X = k_u) = \Omega(1/\sigma)$ and then using that $\Pr(X = i) \geq \Pr(X = i + 1)$ for $i \geq \lceil E(X) \rceil$. \square

As mentioned, we will study how the frequencies associated with single bits evolve over time. To analyze the underlying stochastic process, the following theorem will be used. It generalizes the so-called *simplified drift theorem with scaling* from [22]. The crucial relaxation is that the original version demanded an exponential decay w.r.t. jumps in both directions, more precisely the second condition below was on $\Pr(|X_{t+1} - X_t| \geq jr)$. We now only have sharp demands on jumps in the undesired direction while there is a milder assumption (included in the first item) on jumps in the desired direction. Roughly speaking, if constants in the statements do not matter, the previous version of the drift theorem is implied by the current one as long as $r = O(1)$.

The theorem uses the notation $E(X \mid \mathcal{F}_t; A)$ for filtrations \mathcal{F}_t and events A to denote the expected value $E(X \mid \mathcal{F}_t)$ in the conditional probability space on event A . If A is not a null event, then $E(X \mid \mathcal{F}_t; A) \geq \epsilon$ is equivalent to $E(X - \epsilon; A \mid \mathcal{F}_t) \geq 0$, where the notation “ $;$ A ” just denotes the multiplication with $\mathbb{1}\{A\}$; in fact the notation $E(X; A \mid \mathcal{F}_t)$ is often used in the literature, e.g., by Hajek [10]. Additionally $X \preceq Y$ denotes that X is stochastically at most as large as Y . The proof of the theorem is given in the “Appendix”.

Theorem 3 (Generalizing [22]) *Let $X_t, t \geq 0$, be real-valued random variables describing a stochastic process over some state space, adapted to a filtration \mathcal{F}_t . Suppose there exist an interval $[a, b] \subseteq \mathbb{R}$ and, possibly depending on $\ell := b - a$, a drift bound $\epsilon := \epsilon(\ell) > 0$, a typical forward jump factor $\kappa := \kappa(\ell) > 0$, a scaling factor $r := r(\ell) > 0$ as well as a sequence of functions $\Delta_t := \Delta_t(X_{t+1} - X_t)$ satisfying $\Delta_t \leq X_{t+1} - X_t$ such that for all $t \geq 0$ the following three conditions hold:*

1. $E(\Delta_t \cdot \mathbb{1}\{\Delta_t \leq \kappa\epsilon\} \mid \mathcal{F}_t; a < X_t < b) \geq \epsilon$,
2. $\Pr(\Delta_t \leq -jr \mid \mathcal{F}_t; a < X_t) \leq e^{-j}$ for all $j \in \mathbb{N}$,
3. $\lambda\ell \geq 2 \ln(4/(\lambda\epsilon))$, where $\lambda := \min\{1/(2r), \epsilon/(17r^2), 1/(\kappa\epsilon)\}$.

Then for $T^ := \min\{t \geq 0 \mid X_t \leq a\}$ it holds that $\Pr(T^* \leq e^{\lambda\ell/4} \mid \mathcal{F}_0; X_0 \geq b) = O(e^{-\lambda\ell/4})$.*

To derive upper bounds on hitting times for an optimal state, drift analysis is used, in particular in scenarios where the drift towards the optimum is not state-homogeneous. Such a drift is called *variable* in the literature. A clean form of a variable drift theorem, generalizing previous formulations from [12] and [19], was presented in [23]. The following formulation has been proposed in [17].

Theorem 4 (Variable drift, upper bound) *Let $(X_t)_{t \in \mathbb{N}_0}$, be a stochastic process, adapted to a filtration \mathcal{F}_t , over some state space $S \subseteq \{0\} \cup [x_{\min}, x_{\max}]$, where $x_{\min} > 0$. Let $h(x) : [x_{\min}, x_{\max}] \rightarrow \mathbb{R}^+$ be a monotone increasing function such that $1/h(x)$ is integrable on $[x_{\min}, x_{\max}]$ and $E(X_t - X_{t+1} \mid \mathcal{F}_t) \geq h(X_t)$ if $X_t \geq x_{\min}$. Then it holds for the first hitting time $T := \min\{t \mid X_t = 0\}$ that*

$$E(T \mid \mathcal{F}_0) \leq \frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^{X_0} \frac{1}{h(x)} dx.$$

Finally, we need the following lemma in our analysis of the impact of the so-called 2nd-class individuals in Sect. 2.2. Its statement is very specific and tailored to our applications. Roughly, the intuition is to show that $E(\min\{C, X\})$ is not much less than $\min\{C, E(X)\}$ for $X \sim \text{Bin}(D, p)$ and $D \geq C$. Here and in the following, we write $\text{Bin}(a, b)$ to denote the binomial distribution with parameters a and b .

Lemma 5 *Let $X \sim \text{Bin}(D, p)$. Let $C \in \{1, \dots, D\}$. Then*

$$E(\min\{C, X\}) \geq Cp + \frac{1}{4}p(1 - p) \min\{C, D - C\}.$$

Proof We start by deriving a general lower bound on the expected value of $\min\{C, X\}$. The idea is to decompose the random variable X , which is a sum of D independent trials, into the the first C and the remaining $D - C$ trials. Let $Y \sim \text{Bin}(C, p)$ and $Z \sim \text{Bin}(D - C, p)$. Hence, $X = Y + Z$ and, since Z is independent of Y , we have $\min\{C, X\} = \min\{C, Y\} + \min\{C - Y, Z\} = Y + \min\{C - Y, Z\}$. We also note that $E(Y) = Cp$ and $E(Z) = (D - C)p$.

Assume that for some $k < C$ and some $p^* > 0$, we know that $\Pr(Y \leq k) \geq p^*$. Then by the law of total probability

$$\begin{aligned} E(\min\{C, X\}) &\geq (E(Y \mid Y \leq k) + E(\min\{C - k, Z\})) \Pr(Y \leq k) \\ &\quad + E(Y \mid Y > k) \Pr(Y > k) \\ &= E(Y) + E(\min\{C - k, Z\}) \Pr(Y \leq k) \\ &\geq E(Y) + p^* \cdot E(\min\{C - k, Z\}). \end{aligned} \tag{1}$$

In the following, we will distinguish between two cases with respect to p , in which appropriately chosen pairs (k, p^*) imply the lemma.

Case 1 $p \leq 1 - 2/C$. Hence $C(1 - p) \geq 2$, which implies $\lfloor C(1 - p)/2 \rfloor \geq 1$. Therefore, $\lceil E(Y) \rceil \leq E(Y) + 1 \leq Cp + \lfloor C(1 - p)/2 \rfloor$. We apply the bound $\Pr(Y \leq \lceil E(Y) \rceil) \geq 1/2$, which is equivalent to the well-known bound $\Pr(A \geq \lfloor E(A) \rfloor) \geq 1/2$ that holds for all binomially distributed random variables A [13]. Hence,

$$\Pr(Y \leq Cp + \lfloor C(1 - p)/2 \rfloor) \geq \frac{1}{2}.$$

Using (1) with $p^* := 1/2$ and $k := Cp + \lfloor C(1 - p)/2 \rfloor$ we conclude

$$\begin{aligned} E(\min\{C, X\}) &\geq Cp + \frac{1}{2} \cdot E(\min\{C(1 - p) - \lfloor C(1 - p)/2 \rfloor, Z\}) \\ &= Cp + \frac{1}{2} \cdot E(\min\{\lceil C(1 - p)/2 \rceil, \text{Bin}(D - C, p)\}) \\ &\geq Cp + \frac{1}{2} \cdot E(\text{Bin}(\min\{\lceil C(1 - p)/2 \rceil, D - C\}, p)) \\ &= Cp + \frac{1}{2}p \min\{\lceil C(1 - p)/2 \rceil, D - C\} \\ &\geq Cp + \frac{1}{2}p \min\{C(1 - p)/2, (D - C)(1 - p)/2\} \\ &= Cp + \frac{1}{4}p(1 - p) \min\{C, D - C\}, \end{aligned}$$

where the second inequality exploits that $\text{Bin}(A, p)$ is stochastically larger than $\text{Bin}(B, p)$ for all $B \leq A$, and clearly $\text{Bin}(B, p) \leq B$. The third inequality uses that $(1 - p)/2 \leq 1$ and the final equality exploits that $1 - p$ is non-negative. Hence, the lemma holds in this case.

Case 2 $p > 1 - 2/C$. The aim is to show that $\Pr(Y \leq C - 1) \geq C(1 - p)/3$. In the subcase that $C \leq 3$, we clearly have $\Pr(Y \leq C - 1) \geq 1 - p \geq C(1 - p)/3$. If $C \geq 4$, we work with $q := 1 - p \leq 2/C \leq 2$ and note that $\Pr(Y \leq C - 1) = 1 - p^C = 1 - (1 - q)^C$. Now,

$$1 - (1 - q)^C \geq 1 - e^{-qC} \geq 1 - \left(1 - \frac{qC}{3}\right) = \frac{qC}{3} = \frac{C(1 - p)}{3},$$

where the first inequality uses $e^x \geq 1 + x$ for $x \in \mathbb{R}$ and the second $e^{-x} \leq 1 - \frac{x}{3}$ for $x \leq 2$. Hence, $\Pr(Y \leq C - 1) \geq C(1 - p)/3$ for all $C \in \{1, \dots, D\}$. Using (1) with $k = C - 1$ and $p^* := \frac{C(1-p)}{3}$ and proceeding similarly to Case 1, we obtain

$$\begin{aligned} E(\min\{C, X\}) &\geq Cp + \frac{C(1 - p)}{3} \cdot E(\min\{1, Z\}) \\ &\geq Cp + \frac{C(1 - p)}{3} E(\text{Bin}(\min\{1, D - C\}, p)) \\ &= Cp + \frac{C(1 - p)}{3} p \min\{1, D - C\} \\ &= Cp + \frac{p(1 - p)}{3} \min\{C, C(D - C)\} \\ &\geq Cp + \frac{p(1 - p)}{3} \min\{C, D - C\}, \end{aligned}$$

where the last inequality used that $C \geq 1$ and the previous equality used that C is non-negative. This concludes Case 2 and proves the lemma. □

2.2 On the Stochastic Behavior of Frequencies

To bound the expected running time of UMDA and UMDA*, it is crucial to understand how the n frequencies associated with the bits evolve over time. The symmetry of the fitness function ONEMAX implies that each frequency evolves in the same way, but not necessarily independently of the others. Intuitively, many frequencies should be close to their upper border for making it sufficiently likely to sample the optimum, i. e., the all-1s string.

To understand the stochastic process on a frequency, it is useful to consider the UMDA without selection for a moment. More precisely, assume that each of the λ offspring has the same probability of being selected as one of the μ individuals determining the frequency update. Then the frequency describes a random walk that is a martingale, i. e., in expectation it does not change. With ONEMAX, individuals with higher values are more likely to be among the μ updating individuals. However, since only the accumulated number of 1-bits per individual matters for selection, a single frequency may still decrease even if the step leads to an increase of the best-so-far seen ONEMAX value. We will spell out the bias due to selection in the remainder of this section.

In the following, we consider an arbitrary but fixed bit position j and denote by $p_t := p_{t,j}$ its frequency at time t . Moreover, let X_t , where $0 \leq X_t \leq \mu$, be the number of ones at position j among the μ offspring selected to compute p_t . Then $p_t = \text{cap}_{1/n}^{1-1/n}(X_t/\mu)$, where $\text{cap}_\ell^h(a) := \max\{\min\{a, h\}, \ell\}$ caps frequencies at their borders.

2.2.1 Ranking, 1st-Class Individuals, 2nd-Class Individuals and Candidates

Consider the fitness of all individuals sampled during one generation of the UMDA w. r. t. $n - 1$ bits, i. e., all bits but bit j . Assume that the individuals are sorted in levels decreasingly by their fitness; each individual having a unique rank, where ties are broken arbitrarily. Level $n - 1$ is called the topmost, and level 0 the lowermost. Let w^+ be the level of the individual with rank μ , and let w^- be the level of the individual with rank $\mu + 1$. Since bit j has not been considered so far, its ONEMAX-value can potentially increase each individual's level by 1. Now assume that $w^+ = w^- + 1$. Then, individuals from level w^- can end up with the same fitness as individuals from level w^+ , once bit j has been sampled. Thus, individuals from level w^+ were still prone to selection.

Among the μ individuals chosen during selection, we distinguish between two different types: 1st-class and 2nd-class individuals. 1st-class individuals are those which have so many 1s at the $n - 1$ other bits such that they had to be chosen during selection no matter which value bit j has. The remaining of the μ selected individuals are the 2nd-class individuals; they had to compete with other individuals for selection. Therefore, their bit value at position j is biased towards 1 compared to 1st-class individuals. Note that 2nd-class individuals can only exist if $w^+ \leq w^- + 1$, since in this case, individuals from level w^- can still be as good as individuals from level w^+ after sampling bit j .

Given X_t , let C_{t+1}^* denote the number of 2nd-class individuals in generation $t + 1$. Note that the total number of 1s at position j in the 1st-class individuals during generation $t + 1$ follows a binomial distribution with success probability $p_t = X_t/\mu$, assuming X_t/μ is within the interval $[1/n, 1 - 1/n]$. Since we have $\mu - C_{t+1}^*$ 1st-class individuals, the distribution of the number of 1s in these follows $\text{Bin}(\mu - C_{t+1}^*, X_t/\mu)$.

We proceed by analyzing the number of 2nd-class individuals and how they bias the number of 1s, leading to the Lemmas 6–9 below. The underlying idea is that both the number of 2nd-class individuals is sufficiently large and that at the same time, these 2nd-class individuals were selected from an even larger set to allow many 1s to be gained at the considered position j . This requires a careful analysis of the level where the rank- μ individual ends up.

For $i \in \{0, \dots, n - 1\}$, let C_i denote the cardinality of level i , i. e., the number of individuals in level i during an arbitrary generation of the UMDA, and let $C_{\geq i} = \sum_{a=i}^{n-1} C_a$. Let M denote the index of the first level from the top such that the number of sampled individuals is greater than μ when including the following level, i. e.,

$$M := \max\{i \mid C_{\geq i-1} > \mu\}.$$

Note that M can never be 0, and only if $M = n - 1$, C_M can be greater than μ . Due to the definition of M , if $M \neq n - 1$, level $M - 1$ contains the individual of rank $\mu + 1$, so level $M - 1$ contains the cut where the best μ out of λ offspring are selected. Individuals in levels at least $M + 1$ are definitely 1st-class individuals since they still will have rank at least μ even if the bit j sampled last turns out to be 0. 2nd-class individuals, if any, have to come from levels M , $M - 1$ and $M - 2$ (still in terms of the ranking before sampling bit j). Individuals from level M may still be selected (but may also not) for the μ updating individuals even if bit j turns out as 0. Individuals from level $M - 2$ have to sample a 1 at bit j to be able to compete with the individuals from levels M and $M - 1$; still it is not sure that they will end up in the μ updating individuals.

To obtain a pessimistic bound on the bias introduced by 2nd-class individuals, we concentrate on level $M - 1$. Note that all individuals from level $M - 1$ sampling bit j as 1 will certainly be selected unless the $\mu - C_{\geq M}$ remaining slots for the μ best are filled up. We call the individuals from level $M - 1$ 2nd-class candidates and denote their number by $D_{t+1}^* := C_{M-1}$. By definition, $D_{t+1}^* = C_{\geq M-1} - C_{\geq M}$. We also introduce the notation $C_{t+1}^{**} := \mu - C_{\geq M}$ and note that $C_{t+1}^* \geq C_{t+1}^{**}$ since 2nd-class individuals also may come from levels M and $M - 2$, in addition to level $M - 1$. Our definition of D_{t+1}^* only covers the candidates for 2nd-class individuals that come from level $M - 1$; so the candidates from levels $M - 2$ and M are not part of our notation.

In the following, we often drop the index $t + 1$ from C_{t+1}^* , C_{t+1}^{**} , and D_{t+1}^* if there is no risk of confusion.

2.2.2 Illustration

Figure 2 illustrates the concepts we have introduced so far. On the left-hand side, $\lambda = 14$ individuals are ranked with respect to their ONEMAX-value, ignoring bit j . Level M is the last level from the top such that at most μ individuals are sampled in

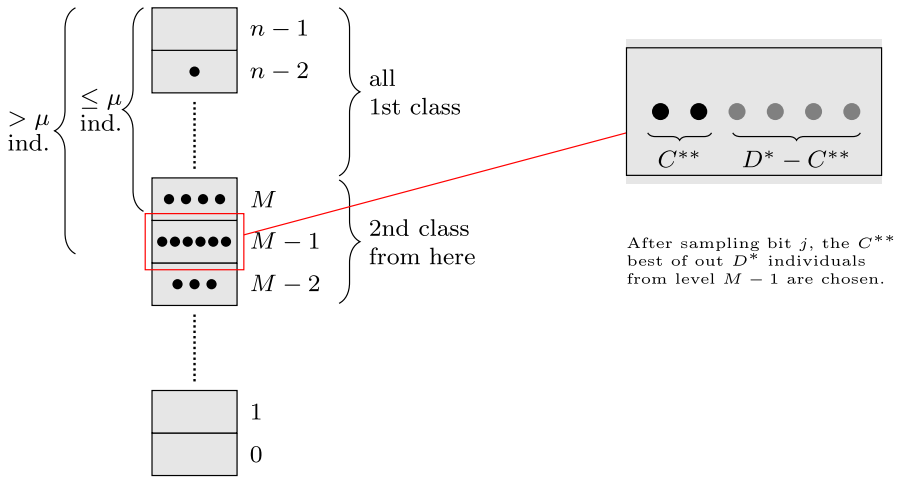


Fig. 2 Illustration of the ranking of the individuals after sampling $n - 1$ bits for $\lambda = 14$ and $\mu = 7$. Finally, $C^{**} = 2$ individuals out of $D^* = 6$ from level $M - 1$ will definitely be selected. Some individuals from levels M and $M - 2$ may also be 2nd-class, in which case $C^* > C^{**}$ holds

level M or above. The individuals from level $M + 1$ and above will be selected for sure even if bit j turns out as 0 and are therefore 1st-class individuals. Level $M - 1$ contains the individual of rank μ and possibly further individuals. In general, if there are D^* individuals in level $M - 1$ and $C_{\geq M}$ in higher levels, then these D^* individuals are the 2nd-class candidates. After finally bit j has been sampled, selection will take the best $C^{**} = \mu - C_{\geq M}$ out of these D^* candidates. They become 2nd-class individuals. Recall that $C^* \geq C^{**}$ as the latter lacks possible 2nd-class individuals from levels $M - 2$ and M .

In the following, the crucial idea is to show that D^* is expected to be larger than C^{**} . That is, we expect to have more 2nd-class candidates (in level $M - 1$) than can actually be selected as 2nd-class individuals. This is dealt with in Lemma 6 below. Roughly speaking, it shows that the number of 2nd-class individuals is stochastically as least as large as if it was sampled from a binomial distribution with parameters $\Theta(\mu)$ and $\Theta(1/\sigma_t)$, where $\sigma_t^2 := \sum_{i=1}^n p_{t,i}(1 - p_{t,i})$ is the sampling variance of the UMDA. This result can be interpreted as follows. It is well known that the Poisson binomial distribution with vector \mathbf{p}_t has a mode of $O(1/\sigma_t)$ [1]. Hence, if we just look at the number of individuals that has a certain number k of 1s at some position, then this is determined by a binomial distribution with parameters λ and $O(1/\sigma_t)$. Lemmas 6 and Lemma 7 together show that essentially the same holds even if we consider the individuals from C^* , i. e., specific individuals from level $M - 1$, each of which is drawn from the complicated conditional distribution induced by the definition of level M . Also, it establishes a similar probabilistic bound for $D^* - C^{**}$, the difference between the number of 2nd-class individuals from level $M - 1$ and the number of 2nd-class candidates, since this difference is responsible for the selection bias. By definition, it always holds that $D^* - C^{**} \geq 1$.

For our analysis, knowledge of the sheer number of 2nd-class individuals and candidates is not yet sufficient. Therefore, afterwards Lemma 8 deals with the number of 1s sampled in the 2nd-class individuals and candidates. This result is then finally used to obtain Lemma 9, which quantifies the bias due to selecting 2nd-class individuals in a drift statement. More precisely, the expected value of X_{t+1} , the number of 1s at position j in the μ selected individuals at time $t + 1$, is bounded from below depending on X_t . This statement is also formulated with respect to the expected success probabilities $E(p_{t+1} | p_t)$ in the lemma.

We are now ready to state the first of the above-mentioned four lemmas. It shows that the number of 2nd-class individuals from level $M - 1$ follows a binomial distribution $\text{Bin}(\mu, q)$, the second parameter of which will be analyzed in the subsequent Lemma 7. In addition, the lemma establishes a similar result for $D_{t+1}^* - C_{t+1}^{**}$, the overhang in 2nd-class candidates w. r. t. the number of 2nd-class individuals that can come from level $M - 1$. Again Lemma 7 will analyze the second parameter of the respective distribution.

Lemma 6 For all $t \geq 0$,

1. $C_{t+1}^* \geq C_{t+1}^{**}$ and $C_{t+1}^{**} \sim \text{Bin}(\mu, q)$ for some random $q \leq 1$.
2. $D_{t+1}^* - C_{t+1}^{**} \sim 1 + \text{Bin}(\lambda - \mu - 1, q')$ for some random $q' \leq 1$.

Proof We first prove the first statement in a detailed manner and then show that the second one can be proven similarly. Hence, we now concentrate on the distribution of $C^{**} = C_{t+1}^{**} = \mu - C_{\geq M}$, which, as outlined above, is a lower bound on C_{t+1}^* . To this end, we carefully investigate and then reformulate the stochastic process generating the λ individuals (before selection), restricted to $n - 1$ bits. Each individual is sampled by a Poisson binomial distribution for a vector of probabilities $\mathbf{p}'_t = (p'_{t,1}, \dots, p'_{t,n-1})$ obtained from the frequency vector of the UMDA by leaving the entry belonging to bit j out (i. e., $\mathbf{p}'_t = (p_{t,1}, \dots, p_{t,j-1}, p_{t,j+1}, \dots, p_{t,n})$). Counting its number of 1s, each of the λ individuals then falls into some level i , where $0 \leq i \leq n - 1$, with some probability q_i depending on the vector \mathbf{p}'_t . Since the individuals are created independently, the number of individuals in level i is binomially distributed with parameters λ and q_i .

Next, we take an alternative view on the process of putting individuals into levels, using the principle of deferred decisions. We imagine that the process first samples all individuals in level 0 (through λ trials, all of which either hit the level or not), then (using the trials which did not hit level 0) all individuals in level 1, ..., up to level $n - 1$.

The number of individuals C_0 in level 0 is still binomially distributed with parameters λ and q_0 . However, after all individuals in level 0 have been sampled, the distribution changes. We have $\lambda - C_0$ trials left, each of which can hit one of the levels 1 to $n - 1$. In particular, such a trial will hit level 1 with probability $q_1/(1 - q_0)$, by the definition of conditional probability since level 0 is excluded. This holds independently for all of the $\lambda - C_0$ trials so that C_1 follows a binomial distribution with parameters $\lambda - C_0$ and $q_1/(1 - q_0)$. Inductively, also all C_i for $i > 1$ are binomially distributed; e. g., C_{n-1} is distributed with parameters $\lambda - C_{n-2} - \dots - C_0$ and 1. Note that this model of the sampling process can also be applied for any other permutation of the levels; we will make use of this fact.

Recall that our aim is to analyze C^{**} . Formally, by applying the law of total probability, its distribution looks as follows for $k \in \{0, \dots, \lambda\}$:

$$\Pr(C^{**} \geq k) = \sum_{i=1}^{n-1} \Pr(M = i) \cdot \Pr(\mu - C_{\geq i} \geq k \mid M = i). \tag{2}$$

To prove the first item of the lemma, it is now sufficient to bound $\Pr(\mu - C_{\geq i} \geq k \mid M = i)$ by the distribution function belonging to a binomial distribution for all $i \in \{1, \dots, n - 1\}$ (recalling that $M = 0$ is impossible).

We reformulate the underlying event appropriately. Here we note that

$$\{\mu - C_{\geq i} \geq k\} \cap \{M = i\}$$

is equivalent to

$$\{C_{\leq i-1} \geq \lambda - \mu + k\} \cap \{M = i\},$$

where $C_{\leq i} = \sum_{j=0}^i C_j$, and, using the definition of M , this is also equivalent to

$$\{C_{\leq i-1} \geq \lambda - \mu + k\} \cap \{C_{\leq i-2} < \lambda - \mu\}.$$

We now use the above-mentioned view on the stochastic process and assume that levels 0 to $i - 2$ have been sampled and a number of experiments in a binomial distribution is carried out to determine the individuals from level $i - 1$. Hence, considering the outcome of $C_{\leq i-2}$ and using that $C_{i-1} = C_{\leq i-1} - C_{\leq i-2}$, our event is equivalent to that the event

$$E^* := \{C_{i-1} \geq (\lambda - \mu - a) + k\} \cap \{C_{\leq i-2} = a\}$$

happens for some $a < \lambda - \mu$. Recall from our model that C_{i-1} follows a binomial distribution with $\lambda - a$ trials and with a certain success probability s . The number of trials left after having sampled levels $0, \dots, i - 2$ is at least μ since $a < \lambda - \mu$. The probability of E^* is determined by conditioning on that $C_{\leq i-2} = a$, i.e., a samples have fallen into levels $0, \dots, i - 2$ and that afterwards $i - 1$ has already been hit by $\lambda - \mu - a$ samples. Then μ trials are left that still may sample within C_{i-1} . Altogether, we proven that C_{i-1} , conditioning on $M = i$, follows a binomial distribution with parameters μ and q , where the value of q depends on the random M . This proves the first item of the lemma.

We now use a dual line of argumentation to prove the second item of the lemma. While the item is concerned with $D_{i+1}^* - C_{i+1}^{**}$, it is more convenient to analyze $C_{\geq M-1}$ and then exploit that

$$D_{i+1}^* - C_{i+1}^{**} = (C_{\geq M-1} - C_{\geq M}) - (\mu - C_{\geq M}) = C_{\geq M-1} - \mu, \tag{3}$$

which follows directly from the definition of C_{i+1}^{**} and D_{i+1}^* .

We claim that

$$C_{\geq M-1} \sim \mu + 1 + \text{Bin}(\lambda - \mu - 1, q'), \tag{4}$$

for some probability q' depending on the outcome of M . To show this, we take the same view on the stochastic process as above but imagine now that the levels are sampled in the order from $n - 1$ down to 0. Conditioning on that levels $n - 1, \dots, M$ have been sampled, there are at least $\lambda - \mu$ trials are left to populate level $M - 1$ since by definition less than μ samples fall into levels $n - 1, \dots, M$. However, by definition of M , at least $\mu + 1 - C_{\geq M}$ of these trials must fall into level $M - 1$. Afterwards, there are $\lambda - \mu - 1$ trials left, each of which may hit level $M - 1$ or not. This proves (4). \square

As announced, the purpose of the following lemma is to analyze the second parameters of the binomial distributions that appear in Lemma 6. Roughly speaking, up to an exponentially small failure probability, we obtain $\Omega(1/\sigma_t)$ as success probability of the binomial distribution.

Lemma 7 *Let $\sigma_t^2 := \sum_{i=1}^n p_{t,i}(1 - p_{t,i})$ be the sampling variance of the UMDA. Consider $C_{t+1}^{**} \sim \text{Bin}(\mu, q)$ and $D_{t+1}^* - C_{t+1}^{**} \sim 1 + \text{Bin}(\lambda - \mu - 1, q')$ as defined in Lemma 6. There is an event E^* with $\Pr(E^*) = 1 - e^{-\Omega(\mu)}$ such that for all $t \geq 0$ the following holds:*

1. *Conditioned on E^* , it holds that $q = \Omega(1/\sigma_t)$. Hence $E(C_{t+1}^{**} \mid \sigma_t) = \Omega(\mu/\sigma_t)$.*
2. *Conditioned on E^* , it holds that $q' = \Omega(1/\sigma_t)$. Hence $E(D_{t+1}^* - C_{t+1}^{**} \mid \sigma_t) = 1 + \Omega((\lambda - \mu - 1)/\sigma_t)$.*

Proof We recall the first statement from Lemma 6 and the stochastic model used in its proof. Let X be the number of 1s at the considered position j in a single individual sampled in the process of creating the λ offspring (without conditioning on certain levels being hit). The aim is to derive bounds on q using X . By our stochastic model, q denotes the probability to sample an individual with $M - 1$ 1s, given that it cannot have less than $M - 1$ 1s. By omitting this condition, we clearly do not increase the probability. Hence, we pessimistically assume that $q = \Pr(X = i - 1)$, given $M = i$. The latter probability heavily depends on M . We will now concentrate on the values of i where $\Pr(M = i)$ is not too small.

The random variable X follows a Poisson binomial distribution with vector \mathbf{p}'_j as defined in the Proof of Lemma 6. Clearly, the variance of this distribution, call it $\tilde{\sigma}_i^2$, is smaller than σ_i^2 since bit j is left out. Still, since $\sigma_i^2 \geq n \cdot (1/n)(1 - 1/n) = 1 - 1/n$ due to the borders on the frequencies, we obtain $\tilde{\sigma}_i^2 \geq 1 - 1/n - 1/4$ and $\tilde{\sigma}_i^2 = \Theta(\sigma_i^2)$.

We define

$$L := \min \left\{ i \mid \Pr(X \leq i) \geq \frac{1}{2 + 2\beta} \right\}$$

and

$$U := \max \left\{ i \mid \Pr(X \geq i) \geq \frac{\beta}{2 + 2\beta} \right\},$$

where β is still the constant from our assumption $\lambda = (1 + \beta)\mu$. By Chernoff bounds, both the number of individuals sampled above U is less than $\frac{1}{1+\beta}\lambda = \mu$ and the number of individuals sampled below L is less than $\frac{\beta}{1+\beta}\lambda = \lambda - \mu$ with probability $1 - e^{-\Omega(\lambda)}$. Then the μ th ranked individual will be within $Z := [L, U]$ with probability at least $1 - e^{-\Omega(\mu)}$, which means that

$$\Pr(M \in Z) = 1 - e^{-\Omega(\mu)}. \tag{5}$$

Note that $M \in Z$ is the event E^* mentioned in the statement of the lemma.

We now assume $M \in Z$ and apply Lemma 2, using $\ell := 1/(2 + 2\beta)$ and $u := \beta/(2 + 2\beta)$, in accordance with the above definition of L and U . Hence, every level in Z , in particular level $M - 1$, is hit with probability $\Omega(\min\{1, 1/\tilde{\sigma}_t\}) = \Omega(1/\sigma_t)$.

Hence with probability $1 - e^{-\Omega(\mu)}$ we have that $q = \Omega(1/\sigma_t)$ and therefore $C_{t+1}^{**} \sim \text{Bin}(\mu, \Omega(1/\sigma_t))$. Using the properties of the binomial distribution and the law of total probability, we obtain $E(C^{**}) = \Omega(\mu/\sigma_t)$, which proves the first item of the lemma.

The second item is proven similarly. We recall from Lemma 6 that $D_{t+1}^* - C_{t+1}^{**} = 1 + \text{Bin}(\lambda - \mu - 1, q')$ where q' is the probability of hitting level $M - 1$, assuming levels $n - 1, \dots, M$ have been sampled. Hence with probability $1 - e^{-\Omega(\mu)}$ according to (5), we have $q' = \Omega(1/\sigma_t)$. Altogether, we obtain $E(D_{t+1}^* - C_{t+1}^{**}) = 1 + \Omega((\lambda - \mu - 1)/\sigma_t)$, which concludes the proof of the second item of the lemma. \square

As mentioned above, we now know much about the distribution of the number of 2nd-class individuals and candidates. The next step is to bound the number of 1s sampled at position j in these individuals.

Lemma 8 *For all $t \geq 0$*

$$X_{t+1} \geq \text{Bin}(\mu - C_{t+1}^{**}, X_t/\mu) + \min\{C_{t+1}^{**}, \text{Bin}(D_{t+1}^*, X_t/\mu)\}.$$

Proof We essentially show that the expected overhang in 2nd-class candidates from level $M - 1$ compared to 2nd-class individuals from this level allows a bias of the frequency towards higher values, as detailed in the following. We recall that we ignore 2nd-class individuals stemming from levels $M - 2$ and M . These might introduce a bias that is only larger, which is why the statement of the lemma only establishes a lower bound on X_{t+1} .

In each of the D_{t+1}^* 2nd-class candidates from level $M - 1$, bit j is sampled as 1 with probability X_t/μ . Only a subset of the candidates, namely the C_{t+1}^{**} 2nd-class individuals from this level, is selected for the best μ offspring determining the next frequency. As observed above in Sect. 2.2, the number of 1s at position j in the 1st-class individuals is binomially distributed with parameters $\mu - C_{t+1}^*$ and X_t/μ . We have $C_{t+1}^* \geq C_{t+1}^{**}$ and recall that the distribution of X_{t+1} becomes stochastically smallest when equality holds. Hence, we obtain for the number of 1s (at position j) in the μ selected offspring that

$$X_{t+1} \geq \text{Bin}(\mu - C_{t+1}^{**}, X_t/\mu) + \min\{C_{t+1}^{**}, \text{Bin}(D_{t+1}^*, X_t/\mu)\},$$

which is what we wanted to show. □

Finally, based on the preceding two lemmas, we can quantify the bias of the frequencies due to selection in a simple drift statement. The following lemma is crucially used in the drift analyses that prove Theorems 10 and 12.

Lemma 9 *Let $\mu = \omega(1)$. Then for all $t \geq 0$,*

$$E(X_{t+1} \mid X_t, \sigma_t) = X_t + \Omega((\mu/\sigma_t)(X_t/\mu)(1 - X_t/\mu)).$$

If $p_t \leq 1 - c/n$, where $c > 0$ is a sufficiently large constant, then

$$E(p_{t+1} \mid p_t, \sigma_t) = p_t + \Omega(p_t(1 - p_t)/\sigma_t).$$

Proof We start with the bound

$$X_{t+1} \geq \text{Bin}(\mu - C, X_t/\mu) + \min\{C, \text{Bin}(D, X_t/\mu)\} \tag{6}$$

from Lemma 8, where we write $C := C_{t+1}^{**}$ and $D := D_{t+1}^*$ for notational convenience. We will estimate the expected value of X_{t+1} based on this stochastic lower bound. By Lemma 5, the expected value of the minimum is at least as large as the minimum of

$$C \frac{X_t}{\mu} + \frac{1}{4}(D - C) \frac{X_t}{\mu} \left(1 - \frac{X_t}{\mu}\right)$$

and

$$C \frac{X_t}{\mu} + \frac{1}{4}C \frac{X_t}{\mu} \left(1 - \frac{X_t}{\mu}\right),$$

where C and D are still random.

Taking the expected value in (6), we obtain

$$\begin{aligned} E(X_{t+1} \mid X_t, \sigma_t, C, D) &\geq E(\text{Bin}(\mu - C, X_t/\mu) \mid X_t, \sigma_t, C, D) + C \frac{X_t}{\mu} \\ &\quad + \frac{1}{4} \frac{X_t}{\mu} \left(1 - \frac{X_t}{\mu}\right) \min\{C, D - C\} \\ &\geq X_t + \frac{1}{4} \frac{X_t}{\mu} \left(1 - \frac{X_t}{\mu}\right) \min\{C, D - C\}, \end{aligned} \tag{7}$$

where the last inequality computed the expected value of the binomial distribution. We also note that C and D are independent of X_t as X_t counts the number of ones in bit j , which is not used to determine C and D .

Recall that the overall aim is to bound $E(X_{t+1} \mid X_t, \sigma_t) = E(E(X_{t+1} \mid X_t, \sigma_t, C, D) \mid X_t, \sigma_t)$ from below. Hence, inspecting the last bound from (7), we are left with the task to prove a lower bound on

$$E(\min\{C, D - C\} \mid X_t, \sigma_t) = E(\min\{C, D - C\} \mid \sigma_t),$$

using that C and $D - C$ are independent of X_t . We recall from Lemmas 6 and 7 that $C \geq \text{Bin}(\mu, q)$ and $D - C \geq 1 + \text{Bin}(\lambda - \mu - 1, q')$, where $q = \Omega(1/\sigma_t)$ and $q' = \Omega(1/\sigma_t)$ with probability $1 - e^{-\Omega(\mu)}$. Also, Lemma 7 yields $E(C \mid \sigma_t) = \Omega(\mu/\sigma_t)$ and $E(D - C \mid \sigma_t) \geq 1 + \Omega((\lambda - \mu - 1)/\sigma_t) = 1 + \Omega(\mu/\sigma_t)$, where the last equality used our assumption that $\lambda = (1 + \Theta(1))\mu$.

We distinguish between two cases. If $\mu/\sigma_t \leq \kappa$ for an arbitrary constant $\kappa > 0$ (chosen later) then $1/\sigma_t = O(1/\mu) = o(1)$ by our assumption on μ . Working with the lower bound 1 on $D - C$, we get

$$\begin{aligned} E(\min\{C, D - C\} \mid \sigma_t) &\geq E(C \cdot \mathbb{1}\{C \leq 1\} \mid \sigma_t) \\ &\geq \Pr(C = 1) = \binom{\mu}{1} \Omega((1/\sigma_t)(1 - o(1))) = \Omega(\mu/\sigma_t). \end{aligned}$$

If $\mu/\sigma_t \geq \kappa$ and κ is chosen sufficiently large but constant then Chernoff bounds yield that each of the events

$$C \geq \frac{E(C)}{2}$$

and

$$D \geq \frac{E(D - C)}{2}$$

fail to happen with probability at most $1/4$, so by a union bound both events happen simultaneously with probability at least $1/2$. This proves

$$E(\min\{C, D - C\} \mid \sigma_t) = \Omega(\mu/\sigma_t)$$

also in this case. Plugging this back into (7), we obtain

$$E(X_{t+1} \mid X_t) = X_t + \Omega((\mu/\sigma_t)(X_t/\mu)(1 - X_t/\mu)),$$

which proves the statement on $E(X_{t+1} \mid X_t, \sigma_t)$ from this lemma.

To conclude on the expected value of p_{t+1} , we recall from Algorithm 1 that $p_{t+1} := \text{cap}_{1/n}^{1-1/n}(X_{t+1}/\mu)$. Using our assumption $p_t \leq 1 - c/n$ we get $1 - X_t/\mu \geq c/n$. Hence, as hitting the upper border changes the frequency by only at most $1/n$ and the lower border can be ignored here, we also obtain

$$E(p_{t+1} \mid p_t) \geq p_t + \Omega((1/\sigma_t)p_t(1 - p_t))$$

if c is large enough to balance the implicit constant in the Ω . □

We note that parts of the analyses behind Lemmas 6–9 are inspired by [14]; in particular, the modeling of the stochastic process and the definition of M follow that paper closely. However, as [14] is concerned with lower bounds on the running time, it bounds the number of 2nd-class individuals from above and needs a very different argumentation in the core of its proofs.

3 Above the Phase Transition

We now prove our main result for the case of large λ . It implies an $O(n \log n)$ running time behavior if $\mu = c\sqrt{n} \log n$.

Theorem 10 *Let $\lambda = (1 + \beta)\mu$ for an arbitrary constant $\beta > 0$, let $\mu \geq c\sqrt{n} \log n$ for some sufficiently large constant $c > 0$ as well as $\mu = n^{O(1)}$. Then with probability $\Omega(1)$, the optimization time of both UMDA and UMDA* on ONEMAX is bounded from above by $O(\lambda\sqrt{n})$. For UMDA, also the expected optimization time is bounded in this way.*

The Proof of Theorem 10 follows a well-known approach that is similar to techniques partially independently proposed in several previous analyses of EDAs and of ant colony optimizers [8, 21, 26]. Here we show that the approach also works for the UMDA. Roughly, a drift analysis is performed with respect to the sum of frequencies. In Lemma 9, we have already established a drift of frequencies towards higher values. Still, there are random fluctuations (referred to as *genetic drift* in [26]) of frequencies that may lead to undesired decreases towards 0. The Proof of Theorem 10 uses that under the condition on μ , typically all frequencies stay sufficiently far away from the lower border; more precisely, no frequency drops below $1/4$. Then the drift is especially beneficial.

The following lemma formally shows that, if μ is not too small, the positive drift along with the fine-grained scale implies that the frequencies will generally move to higher values and are unlikely to decrease by a large distance. Using the lemma, we will obtain a failure probability of $O(n^{-cc'})$ within $n^{cc'}$ generations, which can subsume any polynomial number of steps by choosing c large enough.

Lemma 11 *Consider an arbitrary bit and let p_t be its frequency at time t . Suppose that $\mu \geq c\sqrt{n} \log n$ for a sufficiently large constant $c > 0$. For $T := \min\{t \geq 0 \mid p_t \leq 1/4\}$ it then holds that $\Pr(T \leq e^{cc' \log n}) = O(e^{-cc' \log n})$, where c' is another positive constant.*

Proof The aim is to apply Theorem 3. We consider the frequency $p_t := p_{t,i}$ associated with the considered bit i and its distance $X_t := \mu p_t$ from the lower border. By initialization of the UMDA, $X_0 = \mu/2$. Note that X_t for $t > 0$ is a process on $\{\mu/n, 1, 2, \dots, \mu - 1, \mu(1 - 1/n)\}$.

In the notation of the drift theorem, we set $[a, b] := [\mu/4, \mu/2]$, hence $\ell = \mu/4$. Next we establish the three conditions. First, we observe that $E(X_{t+1} - X_t \mid X_t) = \Omega(X_t(1 - X_t/\mu)/\sqrt{n})$ (Lemma 9 along with the trivial bound $\sigma_t = O(\sqrt{n})$) for $X_t \in \{1, \dots, \mu\}$. The bound is $\Omega(\mu/\sqrt{n})$ for $X_t \in [a, b]$. Since $\mu \geq c\sqrt{n} \log n$ by assumption, we will set $\epsilon' := cc_1 \log n$ for some constant $c_1 > 0$. Hereinafter, we will omit the conditions $X_t; a < X_t < b$ from the expected values. To establish the first condition of the drift theorem, we need to show that the expected value is “typical”; formally, we will find a not too large κ such that

$$E((X_{t+1} - X_t) \cdot \mathbb{1}\{X_{t+1} - X_t \leq \kappa\epsilon'\}) \geq \frac{\epsilon'}{2}. \quad (8)$$

Then the first condition is established with $\epsilon := \epsilon'/2 = cc_1(\log n)/2$.

In the following, we show the claim that

$$E((X_{t+1} - X_t) \cdot \mathbb{1}\{X_{t+1} - X_t > \kappa'\epsilon\}) \leq \frac{\epsilon'}{2}$$

if κ is chosen sufficiently large. Obviously, this implies (8). We recall from Lemma 8 that X_{t+1} is stochastically at least as large as the sum of two random variables $Z_1 \sim \text{Bin}(\mu - C^{**}, X_t/\mu)$ and $Z_2 \sim \min\{C^{**}, \text{Bin}(D^*, X_t/\mu)\}$ for some random variables C^{**} and D^* , which are related to binomial distributions according to Lemma 6. Theorem 3 allows us to deal with a stochastic lower bound $\Delta_t(X_{t+1} - X_t)$ on the drift. We make the drift stochastically smaller in several ways. First, we assume that $\mu = c\sqrt{n} \log n$ and that $\sigma_t = \Omega(\sqrt{n})$, i.e., the pessimistic bounds used to estimate ϵ' above hold actually with equality. This is possible since the lower bound on the drift derived in the analysis from Lemma 9, using the insights from Lemmas 6–8, becomes stochastically smaller by decreasing the two parameters.

Second, we assume that $X_{t+1} = Z_1 + Z_2$ (essentially ignoring 2nd-class individuals coming from levels M and $M - 2$). But now the X_{t+1} obtained in this way is stochastically also bounded from above by $Z_1 + \text{Bin}(D^*, X_t/\mu)$, by omitting the minimum in Z_2 . We conclude that

$$\begin{aligned} X_{t+1} &\leq \text{Bin}(\mu - C^{**}, X_t/\mu) + \text{Bin}(D^*, X_t/\mu) \\ &= \text{Bin}(\mu, X_t/\mu) + \text{Bin}(D^* - C^{**}, X_t/\mu). \end{aligned}$$

By Lemma 6, $D^* - C^{**} - 1$ is binomially distributed with known number of successes $\lambda - \mu - 1$, which is $\Theta(\sqrt{n} \log n)$ by our assumptions, yet random success probability. We note that the lower bounds on $E(X_{t+1} \mid X_t, \sigma_t)$ from Lemma 9 stem from the case that the success probability is $\Omega(1/\sigma_t)$, which actually happens with probability $1 - e^{-\Omega(\mu)}$. So, without decreasing the drift, we can assume that the success probability is actually fixed at $\Theta(1/\sigma_t) = \Theta(1/\sqrt{n})$. These estimations lead to the bound

$$X_{t+1} \leq \text{Bin}(\mu, X_t/\mu) + \text{Bin}(\Theta(\sqrt{n} \log n), \Theta(1/\sqrt{n})).$$

To ease notation, we write $Z'_1 \sim \text{Bin}(\mu, X_t/\mu)$ and $Z'_2 \sim \text{Bin}(c_2\sqrt{n} \log n, c_3/\sqrt{n})$ for some unknown constants $c_2, c_3 > 0$ and obtain $X_{t+1} = Z_1 + Z_2 \leq Z'_1 + Z'_2$. We also recall that $X_t/\mu \in [1/2, 3/4]$. By Chernoff bounds, $\Pr(Z'_1 \geq E(Z'_1) + c_4(\log n)\sqrt{\mu}) \leq e^{-\Omega(c_4 \log^2 n)} \leq 1/\mu^2$ if c_4 is chosen large enough but constant. Similarly, $\Pr(Z'_2 \geq E(Z'_2) + c_5 \log n) \leq 1/\mu^2$ if c_5 is a sufficiently large constant. We now look into the event $Z'_1 + Z'_2 \geq E(Z'_1) + c_4(\log n)\sqrt{\mu} + E(Z'_2) + c_5 \log n$. A necessary condition for this to happen is that at least one of the two events $Z'_1 \geq E(Z'_1) + c_4(\log n)\sqrt{\mu}$ and $Z'_2 \geq E(Z'_2) + c_5 \log n$ happens. A union bound yields that

$$\Pr(Z'_1 + Z'_2 \geq E(Z'_1 + Z'_2) + c_4(\log n)\sqrt{\mu} + c_5 \log n) \leq \frac{2}{\mu^2}$$

and therefore clearly

$$\Pr(Z'_1 + Z'_2 \geq \mathbb{E}(Z'_1 + Z'_2) + c_6(\log n)\sqrt{\mu}) \leq \frac{2}{\mu^2}$$

for some constant $c_6 > 0$. Since $X_{t+1} \leq Z'_1 + Z'_2$, we conclude that

$$\Pr(X_{t+1} \geq \mathbb{E}(Z'_1 + Z'_2) + c_6(\log n)\sqrt{\mu}) \leq \frac{2}{\mu^2}.$$

Since $X_{t+1} \leq \mu$ and $\mathbb{E}(Z'_1 + Z'_2) \leq X_t + c_7 \log n$ for $c_7 = c_2 \cdot c_3$, we obtain

$$\mathbb{E}(X_{t+1} \cdot \mathbb{1}\{X_{t+1} \geq X_t + c_7 \log n + c_6(\log n)\sqrt{\mu}\}) \leq \frac{2}{\mu}$$

and, since $X_t \geq 0$, clearly also

$$\mathbb{E}((X_{t+1} - X_t) \cdot \mathbb{1}\{X_{t+1} - X_t \geq (c_6 + c_7)(\log n)\sqrt{\mu}\}) \leq \frac{2}{\mu} \leq \frac{\epsilon'}{2}.$$

Note that $(c_6 + c_7)(\log n)\sqrt{\mu} = \Theta(\sqrt{\mu}\epsilon')$. This proves the claim for $\kappa = \Theta(\sqrt{\mu})$ and establishes the first condition of the drift theorem.

To show the second condition, recall from Sect. 2.2 that X_{t+1} stochastically dominates $\text{Bin}(\mu, X_t/\mu)$. Hence, to analyze steps where $X_{t+1} < X_t$, we may pessimistically assume the martingale case, where X_{t+1} follows this binomial distribution, and obtain

$$\zeta^2 := \text{Var}(X_t - X_{t+1} \mid X_t) = \mu \frac{X_t}{\mu} \left(1 - \frac{X_t}{\mu}\right) \leq \frac{\mu}{4},$$

so $\zeta = O(\sqrt{\mu})$. Using Lemma 1 with $d = j\zeta$ and $b = 1$, we get $\Pr(X_{t+1} - X_t \leq -j\zeta) \leq e^{-\Omega(\min\{j^2, j\})}$. Hence, we can work with some $r = c_7\sqrt{\mu}$ for some sufficiently large constant $c_7 > 0$ and satisfy the second condition on jumps that decrease the state.

The third condition is also easily verified. We recall that $\ell = \Theta(\mu)$, $\epsilon = \Theta(\log n)$, $r = \Theta(\sqrt{\mu})$ and $\kappa = \Theta(\sqrt{\mu})$. We note that $\epsilon/r^2 = \Theta((\log n)/\mu)$, $1/r = \Theta(1/\sqrt{\mu})$ and $1/\epsilon = \Theta(1/\log n)$. Hence, the λ from the drift theorem (not to be confused with the λ of the UMDA) equals $\epsilon/(17r^2)$ and $\lambda\ell = \Omega(c \log n)$ by our assumption $\mu \geq c\sqrt{n} \log n$ from the lemma. Recalling that $\epsilon \geq cc_1(\log n)/2$ and $\ell = \mu/4$, we obtain by choosing c large enough that

$$\frac{\epsilon\ell}{17r^2} \geq 2 \ln \left(\frac{4r^2}{17\epsilon^2} \right)$$

since the right-hand side is at most $2 \ln(\Theta(\mu/\log^2 n)) = \Theta(\ln n)$ due to $\mu = n^{O(1)}$. Thereby, we satisfy the third condition. Hence, the drift theorem implies that the

first hitting time of states less than a , starting from above b is at least $e^{cc' \log n}$ with probability at least $1 - e^{-cc' \log n}$, if n is large enough and c' is chosen as a sufficiently small positive constant independent of c . \square

We now ready to prove the main theorem from this section.

Proof (Proof of Theorem 10) We use a similar approach and partially also similar presentation of the ideas as in [26]. Following [21, Theorem 3] we show that, starting with a setting where all frequencies are at least $1/2$ simultaneously, with probability $\Omega(1)$ after $O(\sqrt{n})$ generations either the global optimum has been found or at least one frequency has dropped below $1/4$. In the first case we speak of a success and in the latter of a failure. The expected number of generations until either a success or a failure happens is $O(\sqrt{n})$.

With respect to UMDA, we can use the success probability $\Omega(1)$ to bound the expected optimization time. We choose a constant $\gamma > 3$. According to Lemma 11, the probability of a failure in altogether n^γ generations is at most $n^{-\gamma}$, provided the constant c in the condition $\mu \geq c\sqrt{n} \log n$ is large enough. In case of a failure we wait until all frequencies simultaneously reach values at least $1/2$ again and then repeat the arguments from the preceding paragraph. It is easy to show via additive drift analysis for the UMDA (not the UMDA*) that the expected time for one frequency to reach the upper border is always bounded by $O(n^{3/2})$, regardless of the initial probabilities. This holds since by Lemma 9 there is always an additive drift of $\Omega(p_{t,i}(1 - p_{t,i})/\sigma_t) = \Omega(1/(n\sigma_t)) = \Omega(1/n^{3/2})$. By standard arguments on independent phases, the expected time until *all* frequencies have reached their upper border at least once is $O(n^{3/2} \log n)$. Once a frequency reaches the upper border, we apply a straightforward modification of Lemma 11 to show that the probability of a frequency decreasing below $1/2$ in time n^γ is at most $n^{-\gamma}$ (for large enough c). The probability that there is a frequency for which this happens is at most $n^{-\gamma+1}$ by the union bound. If this does not happen, all frequencies attain value at least $1/2$ simultaneously, and we apply our above arguments again. As the probability of a failure is at most $n^{-\gamma+1}$, the expected number of restarts is $O(n^{-\gamma+1})$ and the expected time until all bits recover to values at least $1/2$ only leads to an additional term of $n^{-\gamma+1} \cdot O(n^{3/2} \log n) \leq o(1)$ (as $n^{-\gamma} \leq n^{-3}$) in the expectation. We now only need to show that after $O(\sqrt{n})$ generations without failure the probability of having found the all-ones string is $\Omega(1)$.

In the rest of this proof, we consider the potential function $\phi_t := n - 1 - \sum_{i=1}^n p_{t,i}$, which denotes the total distance of the frequencies from the upper border $1 - 1/n$. For simplicity, for the moment we assume that no frequency is greater than $1 - c/n$, where c is the constant from Lemma 9. Using Lemma 9 and the linearity of expectation, we obtain for some constant $\gamma > 0$ the drift

$$\begin{aligned} E(\phi_t - \phi_{t+1} \mid \phi_t) &= \sum_{i=1}^n (p_{t+1,i} - p_{t,i}) \\ &= \sum_{i=1}^n (p_{t,i} + \gamma p_{t,i}(1 - p_{t,i})/\sigma_t - p_{t,i}) = \gamma \sigma_t, \end{aligned}$$

since $\sum_{i=1}^n p_{t,i}(1 - p_{t,i}) = \sigma_t^2$. Using our assumption $p_{t,i} \geq 1/4$, we obtain the lower bound

$$E(\phi_t - \phi_{t+1} \mid \phi_t) \geq \gamma \sqrt{\sum_{i=1}^n p_{t,i}(1 - p_{t,i})} \geq \gamma \sqrt{\sum_{i=1}^n (1 - p_{t,i})/4} = \gamma \frac{\sqrt{\phi_t}}{2}. \tag{9}$$

The preceding analysis ignored frequencies above $1 - c/n$. To take these into account, we now consider an arbitrary but fixed frequency being greater than $1 - c/n$. We claim that in each of the μ selected offspring the underlying bit (say, bit j) is set to 0 with probability at most c/n . Clearly, each of the λ offspring (before selection) has bit j set to 0 with probability at most c/n . Let us again model the sampling of offspring as a process where bit j is sampled last after the outcomes of the other $n - 1$ bits have been determined. Since setting bit j to 0 leads to a lower ONEMAX-value than setting it to 1, the probability that bit j is 0 in a selected offspring cannot be larger than c/n .

By linearity of expectation, an expected number of at most $\mu c/n$ out of μ selected offspring set bit j to 0. Hence, the expected next value of frequency of bit j must satisfy

$$E(p_{t+1,j} \mid p_{t,j}; p_{t,j} > 1 - c/n) \geq \frac{(\mu - \mu c/n) \cdot 1 + (\mu c/n) \cdot 0}{\mu} = 1 - \frac{c}{n},$$

so $E(p_{t,j} - p_{t+1,j} \mid p_{t,j}; p_{t,j} > 1 - c/n) \geq c/n$. Again by linearity of expectation, the frequencies greater than $1 - c/n$ contribute to the expected change $E(\phi_t - \phi_{t+1} \mid \phi_t)$ an amount of no less than $-c$.

Combining this with (9), we bound the drift altogether by

$$E(\phi_t - \phi_{t+1} \mid \phi_t) \geq \frac{\gamma \sqrt{\phi_t}}{2} - c.$$

If ϕ_t is above a sufficiently large constant, more precisely, if $\sqrt{\phi_t} \geq 4c/\gamma$ (equivalent to $c \leq \gamma \sqrt{\phi_t}/4$), the bound is still positive and only by a constant factor smaller than the drift bound stemming from (9). We obtain

$$E(\phi_t - \phi_{t+1} \mid \phi_t; \sqrt{\phi_t} \geq 4c/\gamma) \geq \frac{\gamma \sqrt{\phi_t}}{2} - c \geq \frac{\gamma \sqrt{\phi_t}}{2} - \frac{\gamma \sqrt{\phi_t}}{4} = \frac{\gamma \sqrt{\phi_t}}{4}.$$

We now set $h(\phi_t) := \frac{\gamma \sqrt{\phi_t}}{4}$ and apply the variable drift theorem (Theorem 4) with drift function $h(\phi_t)$, maximum n and minimum $x_{\min} = 16c^2/\gamma^2$ (since $\sqrt{\phi_t} \geq 4c/\gamma$ is required). Hence, the expected number of generations until the ϕ -value is at most $16c^2/\gamma^2$ is at most

$$\frac{x_{\min}}{h(x_{\min})} + \int_{x_{\min}}^n \frac{dx}{h(x)} \leq \frac{16c^2/\gamma^2}{\gamma \sqrt{16c^2/\gamma^2}/4} + \int_{16c^2/\gamma^2}^n \frac{dx}{\gamma \sqrt{x}/4} = O(\sqrt{n})$$

since both c and γ are constant. Hence, by Markov’s inequality, $O(\sqrt{n})$ generations, amounting to $O(\lambda\sqrt{n})$ function evaluations, suffice with probability $\Omega(1)$ to reach $\phi_t \leq 16c^2/\gamma^2 = O(1)$. It is easy to see that $\phi_t = O(1)$ implies an at least constant probability of sampling the all-ones string (assuming that all $p_{t,i}$ are at least $1/4$). Hence, the optimum is sampled in $O(\sqrt{n})$ generations with probability $\Omega(1)$, which, as outlined above, proves the first statement of the lemma and also the statement on UMDA’s expected running time. \square

The $\Omega(1)$ bound in Theorem 10 on the probability of UMDA* finding the optimum (without stagnating at wrong borders) results from two factors: the application of Markov’s inequality and the probability $\Omega(1)$ of sampling the optimum after $\sqrt{\phi_t} \leq 4c/\gamma$ has been reached from the first time. It is not straightforward to improve this bound to higher values (e. g., probability $1 - o(1)$) since one would have to analyze the subsequent development of potential in the case that the algorithm fails to sample the optimum at $\sqrt{\phi_t} \leq 4c/\gamma$.

4 Below the Phase Transition

Theorem 10 crucially assumes that $\mu \geq c\sqrt{n} \log n$ for a large constant $c > 0$. As described above, the UMDA shows a phase transition between unstable and stable behavior at the threshold $\Theta(\sqrt{n} \log n)$. Above the threshold, the frequencies typically stay well focused on their drift towards the upper border and do not drop much below $1/2$. The opposite is the case if $\mu < c'\sqrt{n} \log n$ for a sufficiently small constant $c' > 0$. Krejca and Witt [14] have shown for this regime that with high probability $n^{\Omega(1)}$ frequencies will walk to the lower border before the optimum is found, resulting in a coupon collector effect and therefore the lower bound $\Omega(n \log n)$ on the running time. It also follows directly from their results (although this was not made explicit) that UMDA* will in this regime with high probability have infinite optimization time since $n^{\Omega(1)}$ frequencies will get stuck at 0. Hence, in the regime $\mu = \Theta(\sqrt{n} \log n)$, the UMDA* turns from efficient with at least constant probability to inefficient with overwhelming probability.

Interestingly, the value $\Theta(\sqrt{n} \log n)$ has also been derived in [26] as an important parameter setting w. r. t. the update strengths called K and $1/\rho$ in the simple EDAs cGA and 2-MMAS_{ib}, respectively. Below the threshold value, lower bounds are obtained through a coupon collector argument, whereas above the threshold, the running time is $O(K\sqrt{n})$ (and $O((1/\rho)\sqrt{n})$, respectively) since frequencies evolve smoothly towards the upper border. The UMDA and UMDA* demonstrate the same threshold behavior, even at the same threshold points.

The EDAs considered in [26] use borders $1/n$ and $1 - 1/n$ for the frequencies in the same way as the UMDA. The only upper bounds on the running time are obtained for update strengths greater than $c\sqrt{n} \log n$. Below the threshold, no conjectures on upper bounds on the running time are stated; however, it seems that the authors do not see any benefit in smaller settings of the parameter since they recommend always to choose values above the threshold. Surprisingly, this does not seem to be necessary if the borders $[1/n, 1 - 1/n]$ are used. With respect to the UMDA, we will show that even for logarithmic μ it has polynomial expected running time, thanks to the borders,

while we already know that UMDA* will fail. We also think that a similar effect can be shown for the EDAs in [26].

We now give our theorem for the UMDA with small μ . If $\mu = \Omega(\sqrt{n} \log n)$, it is weaker than Theorem 10, again underlining the phase transition. The proof is more involved since it has to carefully bound the number of times frequencies leave a border state.

Theorem 12 *Let $\lambda = (1 + \beta)\mu$ for an arbitrary constant $\beta > 0$ and $\mu \geq c \log n$ for a sufficiently large constant $c > 0$ as well as $\mu = O(n^{1-\epsilon})$ for some constant $\epsilon > 0$. Then the expected optimization time of UMDA on ONEMAX is $O(\lambda n)$. For UMDA*, it is infinite with high probability if $\mu < c' \sqrt{n} \log n$ for a sufficiently small constant $c' > 0$.*

Before we prove the theorem, we state two lemmas that work out properties of frequencies that have reached the upper border. In a nutshell, the following lemma show that with high probability such frequencies will stay in the vicinity of the border afterwards, assuming a sufficiently large drift stemming from $\sigma_t = O(1)$. Moreover, it return to the border very quickly afterwards with high probability; in fact, it will only spend a constant number of steps at a non-border value with high probability. Since we are dealing with a Markov chain, the analysis does not change if a certain amount of time has elapsed. In the following lemma, we therefore w.l.o.g. assume that the time where the border is hit equals 0.

Lemma 13 *In the setting of Theorem 12, consider the frequency $p_t, t \geq 0$, belonging to an arbitrary but fixed bit and suppose that $p_0 = 1 - 1/n$ and $\sigma_t = O(1)$ for $t \geq 1$, where $\sigma_t^2 = \sum_{i=1}^n p_{t,i}(1 - p_{t,i})$. Then for every constant $c_1 > 0$ there is a constant $c_2 > 0$ such that*

- $p_1 \geq 1 - c_2/\mu$ with probability at least $1 - n^{-c_1}$.
- There is a constant $r > 0$ with the following properties: if $p_1 \geq 1 - c_2/\mu$ then with probability $1 - e^{-\Omega(c \log n)}$, $p_t \geq 1 - c_2/\mu$ for all $t \in \{2, \dots, r - 1\}$ and finally $p_r = 1 - 1/n$.

Proof By assumption, $p_0 = 1 - 1/n$. We analyze the distribution of p_1 . Since the μ best individuals are biased towards one-entries, the number N of 0s sampled at the bit among the μ best is stochastically smaller than $\text{Bin}(\mu, 1/n)$. Since $\mu = O(n^{1-\epsilon})$ according to the assumptions from Theorem 12, we obtain for any $k > 0$

$$\Pr(N \geq k) \leq \binom{\mu}{k} \left(\frac{1}{n}\right)^k \leq \frac{1}{k!} \left(\frac{\mu}{n}\right)^k \leq n^{-c'k}$$

for some constant $c' > 0$. If $N = k > 0$, then clearly $p_1 = 1 - k/\mu$ by the update rule of the UMDA (using that $\mu = o(n)$ so that p_1 is not capped at the border). We can now choose a constant k such that $c'k = c_1$ and establish the first claim with $c_2 = k$.

To establish the second claim, we consider the distance $X_t := \mu p_t$ of the frequency from the lower border for $t \geq 1$. By assumption, $X_1 = \mu - c_2$. The aim is to show the following: a phase of $O(1)$ steps will consist of increasing steps only such that the upper border is finally reached, all with probability $1 - e^{-\Omega(c \log n)}$. For technical reasons, it is not straightforward to apply Theorem 3 to show this. Instead, we use a

more direct and somewhat simpler argumentation based on the analysis of the sampling process.

Recall from Sect. 2.2, in particular Lemma 8, that X_{t+1} is obtained by summing up the number of 1s at position j sampled in both the 1st-class individuals and the 2nd-class individuals at generation $t + 1$. Denote by D^* the number of 2nd-class candidates and by C^{**} the number of 2nd-class individuals as in Lemma 6. From Lemma 7 we obtain that $D^* - C^{**} - 1$ with probability $1 - e^{-\Omega(\mu)}$ follows a binomial distribution with parameters $\Omega(\mu)$ and $\Omega(1/\sigma_t)$; analogously for C^{**} . We assume both to happen. Note that $\Omega(1/\sigma_t) = \Omega(1)$ by assumption, which implies $\mu/\sigma_t = \Omega(\mu) = \Omega(c \log n)$, where we used the constant c from Theorem 12.

By Chernoff bounds, the probability that $D^* - C^{**} \geq c_3\mu$ for some constant $c_3 > 0$ is at least $1 - e^{-\Omega(\mu)} = 1 - e^{-\Omega(c \log n)}$; similarly, $C^{**} \geq c_3\mu$ with probability $1 - e^{-\Omega(c \log n)}$. We assume both these event also to happen. The number of 1s sampled in the D^* candidates is at least $D^*(1 - c_4)X_t/\mu$ with probability $1 - e^{-\Omega(c \log n)}$ by Chernoff bounds (using that $X_t/\mu = \Theta(1)$), where c_4 is a constant that can be chosen small enough. Also by Chernoff bounds, the number of 1s sampled in the $\mu - C^{**}$ 1st-class individuals (pessimistically assuming that $C^* = C^{**}$ and that $\mu - C^{**} = \Omega(\mu)$) is at least $(\mu - C^{**})(1 - c_4)X_t/\mu$ with probability $1 - e^{-\Omega(c \log n)}$.

Assuming all this to happen, we have $X_{t+1} \geq X_t + (D^* - C^{**})X_t/\mu - c_4(\mu + D^* - C^{**})X_t/\mu$ (or X_{t+1} takes its maximum μ anyway) with probability $1 - e^{-\Omega(c \log n)}$. If c_4 is chosen sufficiently small compared to c_3 , we obtain $X_{t+1} \geq X_t + c_4\mu/2$ with probability $1 - e^{-\Omega(c \log n)}$. Note that the constants in the exponent of $e^{-\Omega(c \log n)}$ may be different; we use the smallest ones from these estimations and union bounds to obtain the bound $1 - e^{-\Omega(c \log n)}$ on the final probability.

If $X_{t+1} \geq X_t + c_4\mu/2$ for a constant number $r \leq 2c_2/c_4$ of iterations, then the frequency p_t has raised from $1 - c_2/\mu$ to its upper border in this number of steps. By a union bound over these many iterations, the probability that any of the events necessary for this happens to fail is most $O(1) \cdot e^{-\Omega(c \log n)} = e^{-\Omega(c \log n)}$ as claimed above. □

In the following lemma, we show a statement on the expected value of the frequency. To obtain this value, it is required that the frequency does not drop below $1 - O(1/\mu)$ and quickly returns to $1 - 1/n$, which can be satisfied by means of the previous lemma.

Lemma 14 *In the setting of Theorem 12, consider the frequency $p_t, t \geq 0$, belonging to an arbitrary but fixed bit and suppose that*

- $p_0 = 1 - 1/n$,
- $p_t \geq 1 - O(1/\mu)$ for all $t \geq 0$, and
- there is a constant $r > 0$ such that the following holds: for all t' where $p_{t'} = 1 - 1/n$ there is some $s \in \{1, \dots, r\}$ such that $p_{t'+s} = 1 - 1/n$.

Then for all $t \geq 0$ it holds that $E(p_t) = 1 - O(1/n)$.

Proof By assumption, $p_0 = 1 - 1/n$. We will analyze the distribution and expected value of p_t for $t \geq 1$. We will show that for all $t \geq 0$, we have $\Pr(p_t = 1 - 1/n) \geq 1 - c_1\mu/n$ for some constant $c_1 > 0$. By assumption, in any case $p_t \geq 1 - O(1/\mu)$.

Using the law of total probability to combine the cases $p_t = 1 - 1/n$ and $p_t < 1 - 1/n$, we obtain

$$E(p_t) \geq \left(1 - \frac{c_1\mu}{n}\right) \left(1 - \frac{1}{n}\right) + \frac{c_1\mu}{n} \frac{\mu - O(1)}{\mu} = 1 - O(1/n).$$

We are left with the claim $\Pr(p_t = 1 - 1/n) = 1 - O(\mu/n)$. Recall that on $p_t < 1 - 1/n$ we have $p_1 = 1 - O(1/\mu)$ and that again $p_t = 1 - 1/n$ after $t \leq r$ steps. We take now a simpler view by means of a two-state Markov chain, where state 0 corresponds to frequency $1 - 1/n$ and state 1 to the rest of the reachable states. Time is considered in blocks of r steps, which will be justified in the final paragraph of this proof. The transition probability from 0 to 1 is at most $O(\mu/n)$ (the expected number of 0s sampled in the μ best individuals) and the transition probability from 1 to 0 is 1; the remaining probabilities are self-loops. Now, it is easy to analyze the steady-state probabilities, which are $1 - O(\mu/n)$ for state 0 and $O(\mu/n)$ for state 1. Moreover, since the chain starts in state 0, simple calculations of occupation probabilities over time yield for state 0 a probability of $1 - O(\mu/n)$ for all points of time $t \geq 0$. More precisely, at the transition from time t to time $t + 1$ the occupation probability of state 0 can only decrease by $O(\mu/n)$. When state 1 exceeds an occupation probability of $c_2\mu/n$ for a sufficiently large constant $c_2 > 0$, the process goes to state 0 with probability at least $c_2\mu/n$, which is less than the decrease of the occupation probability for state 0 for c_2 large enough. Hence, the occupation probability of state 0 cannot drop below $1 - O(\mu/n)$.

Finally, we argue why we may consider phases of length r in the Markov chain analysis. Note that only every r th step a transition from state 1 to 0 is possible (in our pessimistic model), however, in fact every step can transit from state 0 to 1. Formally, we have to work in these additional steps in our two-state model. We do so by increasing the probability of leaving state 0 by a factor of r , which vanishes in the $O(\mu/n)$ bound used above. □

Having proved these two preparatory lemmas, we can give the proof of the main theorem from this section.

Proof (Proof of Theorem 12) The second statement can be derived from [14], as discussed above. We now focus on the first claim, basically reusing the potential function $\phi_t = n - 1 - \sum_{i=1}^n p_{t,i}$ from the Proof of Theorem 10. Let k denote the number of frequencies below $1 - c/n$ for the c from Lemma 9, w.l. o. g., these are the frequencies associated with bits $1, \dots, k$. The last $n - k$ bits are actually at $1 - 1/n$ since $1/\mu = \omega(1/n)$ by assumption. They are set to 0 with probability at most $1/n$ in each of the selected offspring, amounting to a total expected loss of at most 1. Similarly as in the Proof of Theorem 10, we compute the drift

$$\begin{aligned} E(\phi_t - \phi_{t+1} \mid \phi_t) &\geq \sum_{i=1}^k (p_{t+1,i} - p_{t,i}) - (n - k) \frac{1}{n} \\ &\geq \sum_{i=1}^k (p_{t,i} + \gamma p_{t,i}(1 - p_{t,i})/\sigma_t - p_{t,i}) - 1 \end{aligned}$$

$$\begin{aligned}
 &= \frac{\gamma \sum_{i=1}^k p_{t,i}(1 - p_{t,i})}{\sqrt{((n - k)/n)(1 - 1/n) + \sum_{i=1}^k p_{t,i}(1 - p_{t,i})}} - 1 \\
 &\geq \frac{\gamma \sum_{i=1}^k p_{t,i}(1 - p_{t,i})}{\sqrt{1 + \sum_{i=1}^k p_{t,i}(1 - p_{t,i})}} - 1 \tag{10}
 \end{aligned}$$

where γ is the implicit constant in the Ω -notation from Lemma 9, and the last equality just used the definition of σ_t . We now distinguish two cases depending on $V^* := \sum_{i=1}^k p_{t,i}(1 - p_{t,i})$, the total variance w. r. t. the bits not at the upper border.

Case 1 If $V^* \geq c'$ for some sufficiently large constant $c' > 0$, we obtain

$$\frac{\gamma V^*}{\sqrt{1 + V^*}} \geq 2,$$

and therefore

$$E(\phi_t - \phi_{t+1} \mid \phi_t) \geq \frac{\gamma V^*}{\sqrt{1 + V^*}} - 1 \geq 1$$

from (10). If $V^* < c'$, we will show by advanced arguments that the bits that have reached the upper border can almost be ignored and that the drift with respect to the other bits is still in the order $\Omega(V^*/\sqrt{1 + V^*})$. Using this (to be proved) statement, we apply variable drift (Theorem 4) with $x_{\min} = 1/\mu$ (since each $p_{t,i} = i/\mu$ for some $i \in \{1, \dots, \mu - 1\}$ if it is not at a border) and

$$h(x) := \min\{1, c''x/\sqrt{1 + x}\}$$

for some constant c'' . Let x^* be the point where $1 = c''x^*/\sqrt{1 + x^*}$ and note that x^* is some constant bigger than 1 if c'' is small enough. We obtain the upper bound

$$\frac{x_{\min}}{h(x_{\min})} + \int_{1/\mu}^n \frac{1}{h(x)} dx = \frac{(1/\mu)\sqrt{1 + 1/\mu}}{c''/\mu} + \int_{1/\mu}^{x^*} \frac{\sqrt{1 + x}}{c''x} dx + \int_{x^*}^n \frac{dx}{1} \tag{11}$$

on the expected number of generations until all frequencies have hit the upper border at least once. The anti-derivative of $\sqrt{1 + x}/x$ is $2\sqrt{1 + x} + \ln(\sqrt{1 + x} - 1) - \ln(\sqrt{1 + x} + 1)$. Hence,

$$\int_{1/\mu}^{x^*} \frac{\sqrt{1 + x}}{c''x} dx \leq O(1) - \frac{\ln(\sqrt{1 + 1/\mu} - 1)}{c''} = O(\log \mu)$$

since $\ln(\sqrt{1 + 1/\mu} - 1) \geq \ln(1/(2\mu) - 1/(8\mu^2)) = -\ln(\mu)/2 + \Theta(1)$ by a Taylor expansion. Hence, the whole bound (11) can be simplified to

$$O(1) + O(\log \mu) + O(n) = O(n)$$

using $\mu = o(n)$. When the potential has reached its minimum value, the optimum is sampled with probability $\Omega(1)$. If this fails, the argumentation can be repeated. The expected number of repetitions is $O(1)$. This corresponds to an expected running time of $O(\lambda n)$.

Case 2 We still have to show that we have a drift of $\Omega(V^*/\sqrt{1+V^*})$ if $V^* \leq c'$. Actually, we will consider a phase of κn generations for some sufficiently large constant $\kappa > 0$ and show that the claim holds with high probability throughout the phase. We then show that under this assumption the optimum is still sampled with probability $\Omega(1)$ in the phase. In case of a failure, we repeat the argumentation and get an expected number of $O(1)$ repetitions, altogether an expected running time of $O(\lambda n)$.

We have seen above that frequencies at the upper border may contribute negatively to the drift of the ϕ_t -value. Hence, to show the claim that the potential also is expected to decrease when $V^* \leq c'$, we will analyze an arbitrary but fixed frequency and use the above-proven fact that it is likely to stay in the vicinity of the upper border once having been there. We work under the assumption that we have $\sigma_t = O(1)$ within a phase of κn generations. We claim that this actually happens with probability at least $1 - O(1/n)$, a proof of which will be given below. In case of a failure, we repeat the argumentation and are done within an expected number of $O(1)$ repetitions.

We invoke Lemma 13 for any frequency that has hit the upper border and note that, unless a failure event of probability $e^{-\Omega(c \log n)} + n^{-c_1}$ happens, we can apply Lemma 14. We assume that the constant c (stemming from the assumption $\mu \geq c \log n$) as well as c_1 have been chosen appropriately such that the failure probability is $O(1/n^2)$. Hereinafter, we assume that no failure occurs. Hence, we have an expected frequency of $1 - O(1/n)$. This expected value is just the probability of sampling a 1 at the underlying bit. Consequently, if there are ℓ bits that have been at the upper border at least once, the probability of sampling only 1s at all these bits is at least

$$\prod_{i=1}^{\ell} \left(1 - O\left(\frac{1}{n}\right) \right) = \Omega(1).$$

This still allows the optimum to be sampled with probability $\Omega(1)$ after the potential on the bits that never have hit the border so far has decreased below c' .

Finally, we have to justify why with high probability $\sigma_{t'} = O(1)$ for any $t' \geq t$ within κn steps after the first time t where $\phi_t = O(1)$. The frequencies that never have been at the upper border contribute at most $c' = O(1)$ to ϕ_t by assumption and, since $\sigma_t \leq \phi_t$, also to σ_t . Frequencies that are at the upper border leave this state only with probability $O(\mu/n)$ and have a value of $1 - O(1/\mu)$ afterwards according to Lemma 13 with high probability. In every step only an expected number of $n \cdot O(\mu/n) = O(\mu)$ frequencies leave the upper border. By Chernoff bounds, the number is $O(\mu)$ even with probability $1 - e^{-\Omega(\mu)}$. Finally, since every frequency that has left the upper border returns to it within r steps with high probability, there are within the phase only $O(\mu)$ such frequencies with high probability. Their contribution to σ_t is therefore $O(\mu) \cdot O(1/\mu) = O(1)$ with high probability. The failure probability is again $O(1/n^2)$.

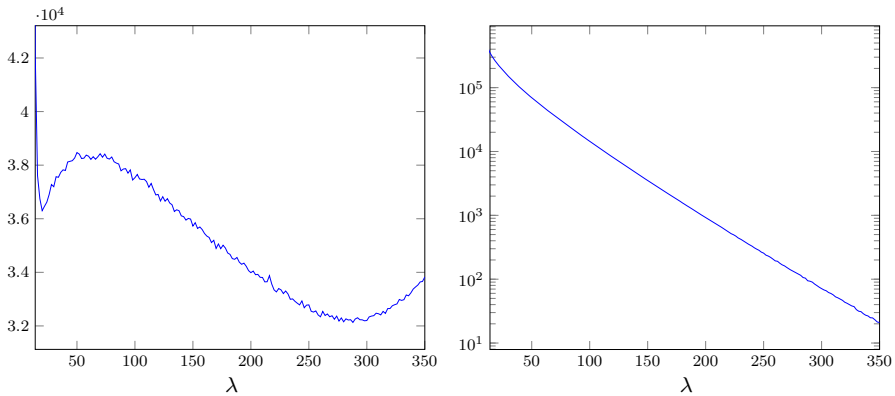


Fig. 3 Left-hand side: empirical running time of UMDA on ONEMAX, right-hand side: number of hits of lower border; for $n = 2000$, $\lambda \in \{14, 16, \dots, 350\}$, $\mu = \lambda/2$, and averaged over 3000 runs

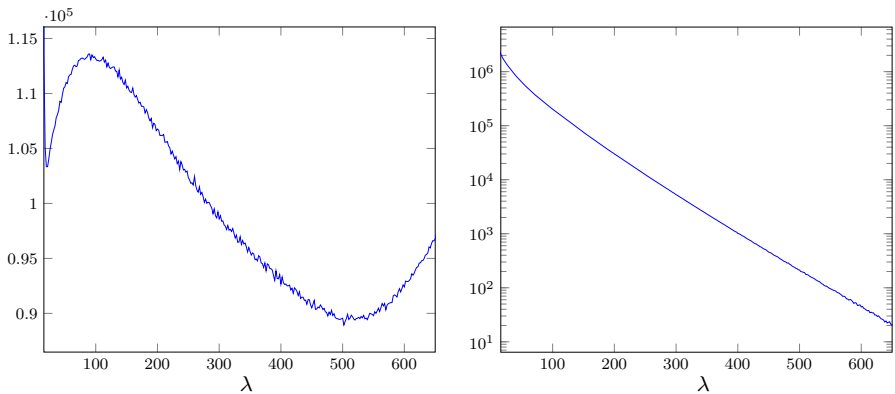


Fig. 4 Left-hand side: empirical running time of UMDA on ONEMAX, right-hand side: number of hits of lower border; for $n = 5000$, $\lambda \in \{14, 16, \dots, 650\}$, $\mu = \lambda/2$, and averaged over 3000 runs

as the constants in Lemma 13 and the c in the assumption $\mu \geq c \log n$ were chosen sufficiently large. Hence, the probability of a failure in $O(\kappa n)$ steps is $O(1/n)$.

We finally note that, although the state of the algorithm may switch between the cases $V^* \leq c'$ and $V^* > c'$ more than once, the drift argument can always be applied since we have established a drift of $\Omega(V^*/\sqrt{1 + V^*})$ for the ϕ_t -value regardless of the case. \square

We have now concluded the Proof of Theorem 12. As mentioned before, we can extract from this theorem a second value of μ that gives the $O(n \log n)$ running time bound, namely $\mu = c' \log n$. We also believe that values $\mu = o(\log n)$ will lead to a too coarse-grained frequency scale and exponential lower bounds on the running time, which can be regarded as another phase transition in the behavior. We do not give a proof here but only mention that such a phase transition from polynomial to

exponential running time is known from ACO algorithms and non-elitist $(1, \lambda)$ EAs when a parameter crosses $\log n$ [21, 23].

5 Experiments

We have carried out experiments for UMDA on ONEMAX to gain some empirical insights into the relationship between λ and the average running time. The algorithm was implemented in the C programming language using the PCG32 random number generator. The problem size was set to $n = 2000$, λ was increased from 14 to 350 in steps of size 2, μ was set to $\lambda/2$, and, due to the high variance of the runs especially for small λ , an average was taken over 3000 runs for every setting of λ . The left-hand side of Fig. 3 demonstrates that the average running time in fact shows a multimodal dependency on λ . Starting out from very high values, it takes a minimum at $\lambda \approx 20$ and then increases again up to $\lambda \approx 70$. Thereafter it falls again up to $\lambda \approx 280$ and finally increases rather steeply for the rest of the range. The right-hand side, a semi-log plot, also illustrates that the number of times the lower border $1/n$ of the frequencies is hit seems to decrease exponentially with λ . The phase transition where the behavior of frequencies turns from chaotic into stable is empirically located somewhere between 250 and 300.

Similar results are obtained for $n = 5000$, see Fig. 4. The location of the maximum does not seem to increase linearly with n .

6 Conclusions

We have analyzed the UMDA on ONEMAX and obtained the upper bounds $O(\mu\sqrt{n})$ and $O(\mu n)$ on its expected running time in different domains for μ , more precisely if $\mu \geq c\sqrt{n} \log n$ and $\mu \geq c' \log n$, respectively, where c, c' are positive constants. This implies an expected running time of $O(n \log n)$ for two asymptotic values of μ , closing the previous gap between the lower bound $\Omega(\mu\sqrt{n} + n \log n)$ and the upper bound $O(n \log n \log \log n)$. In our proofs, we provide detailed tools for the analysis of the stochastic processes at single frequencies in the UMDA. We hope that these tools will be fruitful in future analyses of EDAs.

We note that all our results assume $\lambda = O(\mu)$. However, we do not think that larger λ can be beneficial; if $\lambda = \alpha\mu$, for $\alpha = \omega(1)$, the progress due to 2nd-class individuals can be by a factor of at most α bigger; however, also the computational effort per generation would grow by this factor. A formal analysis of other ratios between μ and λ is open, as is the case of sublogarithmic μ . Moreover, we do not have lower bounds matching the upper bounds from Theorem 10 if μ is in the regime where both $\mu = \omega(\log n)$ and $\mu = o(\sqrt{n} \log n)$.

Acknowledgements Financial support by the Danish Council for Independent Research (DFF-FNU 4002–00542) is gratefully acknowledged.

Appendix

Proof of Theorem 3

We will use Hajek’s drift theorem to prove Lemma 3. As we are dealing with a stochastic process, we implicitly assume that the random variables $X_t, t \geq 0$, are adapted to some filtration \mathcal{F}_t such as the natural filtration $X_0, \dots, X_t, t \geq 0$.

We do not formulate the theorem using a potential/Lyapunov function g mapping from some state space to the reals either. Instead, we w.l.o.g. assume the random variables X_t as already obtained by the mapping.

The following theorem follows immediately from taking Conditions D1 and D2 in [10] and applying Inequality (2.8) in a union bound over $L(\ell)$ time steps.

Theorem 15 [10] *Let $X_t, t \geq 0$, be real-valued random variables describing a stochastic process over some state space, adapted to a filtration \mathcal{F}_t . Pick two real numbers $a(\ell)$ and $b(\ell)$ depending on a parameter ℓ such that $a(\ell) < b(\ell)$ holds. Let $T(\ell)$ be the random variable denoting the earliest point in time $t \geq 0$ such that $X_t \leq a(\ell)$ holds. If there are $\lambda(\ell) > 0$ and $p(\ell) > 0$ such that the condition*

$$\mathbb{E}(e^{-\lambda(\ell) \cdot (X_{t+1} - X_t)} \mid \mathcal{F}_t; a(\ell) < X_t < b(\ell)) \leq 1 - \frac{1}{p(\ell)} \tag{*}$$

holds for all $t \geq 0$ then for all time bounds $L(\ell) \geq 0$

$$\Pr(T(\ell) \leq L(\ell) \mid X_0 \geq b(\ell)) \leq e^{-\lambda(\ell) \cdot (b(\ell) - a(\ell))} \cdot L(\ell) \cdot D(\ell) \cdot p(\ell),$$

where $D(\ell) = \max\{1, \mathbb{E}(e^{-\lambda(\ell) \cdot (X_{t+1} - b(\ell))} \mid \mathcal{F}_t; X_t \geq b(\ell))\}$.

Proof (Proof of Theorem 3) We will apply Theorem 15 for suitable choices of its variables, some of which might depend on the parameter $\ell = b - a$ denoting the length of the interval $[a, b]$. The following argumentation is also inspired by Hajek’s work [10].

By assumption $\Delta_t(X_{t+1} - X_t) \leq X_{t+1} - X_t$. Clearly, for the process $X'_t = X_0 + \sum_{j=0}^{t-1} \Delta_j(X_{j+1} - X_j)$ we have $X'_t \leq X_t$. Hence, the hitting time T^* for state less than a of the original process X_t is stochastically at least as big as the corresponding hitting time of the process X'_t . In the following, we will therefore without further mention analyze X'_t instead of X_t and bound the tail of its hitting time. We work with $\Delta := \Delta_t$, which equals $X'_{t+1} - X'_t$. We still use the old notation X_t instead of X'_t .

The aim is to bound the moment-generating function (mgf.) from Condition (*). In this analysis, we for notational convenience often omit the filtration \mathcal{F}_t . First we observe that it is sufficient to bound the mgf. of $\Delta \cdot \mathbb{1}\{\Delta \leq \kappa\epsilon\}$ since

$$\begin{aligned} \mathbb{E}(e^{-\lambda\Delta}) &= \mathbb{E}(e^{-\lambda\Delta\mathbb{1}\{\Delta \leq \kappa\epsilon\} - \lambda\Delta\mathbb{1}\{\Delta > \kappa\epsilon\}}) \\ &= \mathbb{E}(e^{-\lambda\Delta\mathbb{1}\{\Delta \leq \kappa\epsilon\}} e^{-\lambda\Delta\mathbb{1}\{\Delta > \kappa\epsilon\}}) \leq \mathbb{E}(e^{-\lambda\Delta\mathbb{1}\{\Delta \leq \kappa\epsilon\}}), \end{aligned}$$

using $\Delta \mathbb{1}\{\Delta > \kappa\epsilon\} > 0$ and hence $e^{-\lambda\Delta \mathbb{1}\{\Delta > \kappa\epsilon\}} \leq 1$. In the following, we omit the factor $\mathbb{1}\{\Delta \leq \kappa\epsilon\}$ but implicitly multiply Δ with it all the time. The same goes for $\mathbb{1}\{a < X_t < b\}$.

To establish Condition (*), it is sufficient to identify values $\lambda := \lambda(\ell) > 0$ and $p(\ell) > 0$ such that

$$E(e^{-\lambda\Delta} \mathbb{1}\{a < X_t < b\}) \leq 1 - \frac{1}{p(\ell)}.$$

Using the series expansion of the exponential function, we get

$$\begin{aligned} E(e^{-\lambda\Delta} \mathbb{1}\{a < X_t < b\}) &= 1 - \lambda E(\Delta) + \sum_{k=2}^{\infty} \frac{(-\lambda)^k}{k!} E(\Delta^k) \\ &= 1 - \lambda E(\Delta) + \sum_{k=2}^{\infty} \frac{(-\lambda)^k}{k!} \left(E(\Delta^k \mathbb{1}\{\Delta \geq 0\}) + E(\Delta^k \mathbb{1}\{\Delta < 0\}) \right). \end{aligned}$$

We first concentrate on the positive steps in the direction of the expected value, more precisely, we consider for any odd $k \geq 3$

$$M_k := \frac{\lambda^k}{k!} E(\Delta^k \mathbb{1}\{\Delta \geq 0\}) - \frac{\lambda^{k+1}}{(k+1)!} E(\Delta^{k+1} \mathbb{1}\{\Delta \geq 0\}).$$

Since we implicitly multiply with $\mathbb{1}\{\Delta \leq \kappa\epsilon\}$, we have $\Delta^k \mathbb{1}\{\Delta \geq 0\} \leq (\kappa\epsilon)^k$ and hence $|E(\Delta^{k+1} \mathbb{1}\{\Delta \geq 0\})/E(\Delta^k \mathbb{1}\{\Delta \geq 0\})| \leq \kappa\epsilon$. By choosing $\lambda \leq 1/(\kappa\epsilon)$, we have

$$M_k \geq \frac{\lambda^k}{k!} E(\Delta^k \mathbb{1}\{\Delta \geq 0\}) - \frac{\lambda^k}{\kappa\epsilon(k+1)!} \kappa\epsilon E(\Delta^k \mathbb{1}\{\Delta \geq 0\}) \geq 0,$$

for $k \geq 3$ since $(1/k!)/(1/(k+1)!) = k$. Hence,

$$\begin{aligned} E(e^{-\lambda\Delta}) &\leq 1 - \lambda E(\Delta) + \frac{\lambda^2}{2} E(\Delta^2 \mathbb{1}\{\Delta \geq 0\}) \\ &\leq 1 - \lambda E(\Delta) + \frac{\lambda^2}{2} E(\Delta \cdot \kappa\epsilon \cdot \mathbb{1}\{\Delta \geq 0\}) \\ &\leq 1 - \lambda E(\Delta) + \lambda \frac{1}{2\kappa\epsilon} \cdot \kappa\epsilon \cdot E(\Delta) \leq 1 - \lambda\epsilon/2 \end{aligned}$$

where the first inequality used that $\Delta^2 \leq \Delta\kappa\epsilon$ due to our implicit multiplication with $\mathbb{1}\{\Delta \leq \kappa\epsilon\}$ everywhere and the second used again $\lambda \leq 1/(\kappa\epsilon)$. So, we have estimated the contribution of all the positive steps by $1 - \lambda E(\Delta)/2$.

We proceed with the remaining terms. We overestimate the sum by using $\Delta' := |\Delta \cdot \mathbb{1}\{\Delta < 0\}|$ and bounding $(-\lambda^k) \leq \lambda^k$ in all terms starting from $k = 2$. Incorporating

the contribution of the positive steps, we obtain for all $\gamma \geq \lambda$

$$\begin{aligned} E(e^{-\lambda\Delta}) &\leq 1 - \frac{\lambda}{2}E(\Delta) + \frac{\lambda^2}{\gamma^2} \sum_{k=2}^{\infty} \frac{\gamma^k}{k!} E(\Delta^k) \\ &\leq 1 - \frac{\lambda}{2}E(\Delta) + \frac{\lambda^2}{\gamma^2} \sum_{k=0}^{\infty} \frac{\gamma^k}{k!} E(\Delta^k) \leq 1 - \frac{\lambda}{2}\epsilon + \lambda^2 \underbrace{\frac{E(e^{\gamma\Delta'})}{\gamma^2}}_{=:C(\gamma)}, \end{aligned}$$

where the last inequality uses the first condition of the theorem, i. e., the bound on the drift.

Given any $\gamma > 0$, choosing $\lambda := \min\{1/(\kappa\epsilon), \gamma, \epsilon/(4C(\gamma))\}$ results in

$$E(e^{-\lambda\Delta} \mathbb{1}\{a < X_t < b\}) \leq 1 - \frac{\lambda}{2}\epsilon + \lambda \cdot \frac{\epsilon}{4C(\gamma)} \cdot C(\gamma) = 1 - \frac{\lambda\epsilon}{4} = 1 - \frac{1}{p(\ell)}$$

with $p(\ell) := 4/(\lambda\epsilon)$.

The aim is now to choose γ in such a way that $E(e^{\gamma\Delta'})$ is bounded from above by a constant. We get

$$E(e^{\gamma\Delta'}) \leq \sum_{j=0}^{\infty} e^{\gamma(j+1)r} \Pr(\Delta \leq -jr) \leq \sum_{j=0}^{\infty} e^{\gamma(j+1)r} e^{-j}$$

where the inequality uses the second condition of the theorem.

Choosing $\gamma := 1/(2r)$ yields

$$E(e^{\gamma\Delta'}) \leq \sum_{j=0}^{\infty} e^{(j+1)/2-j} = e^{1/2} \sum_{j=0}^{\infty} e^{-j/2} = e^{1/2} \frac{1}{1 - e^{-1/2}} \leq 4.2.$$

Hence, $C(\gamma) \leq 4.2/\gamma^2$ and therefore $\lambda \leq \epsilon/(4 \cdot 4.2r^2) < \epsilon/(17r^2)$. From the definition of λ , we altogether have $\lambda = \min\{1/(2r), \epsilon/(17r^2), 1/(\kappa\epsilon)\}$. Since $p(\ell) = 4/(\lambda\epsilon)$, we know $p(\ell) = O(r/\epsilon + r^2/\epsilon^2 + \kappa)$. Condition (*) of Theorem 15 has been established along with these bounds on $p(\ell)$ and $\lambda = \lambda(\ell)$.

To bound the probability of a success within $L(\ell)$ steps, we still need a bound on $D(\ell) = \max\{1, E(e^{-\lambda(X_{t+1}-b)} \mid X_t \geq b)\}$. If 1 does not maximize the expression then

$$\begin{aligned} D(\ell) &= E(e^{-\lambda(X_{t+1}-b)} \mid X_t \geq b) \leq E(e^{-\lambda|\Delta|} \mid X_t \geq b) \\ &\leq 1 + E(e^{\gamma\Delta'} \mid X_t \geq b), \end{aligned}$$

where the first inequality follows from $X_t \geq b$ and the second one from $\gamma \geq \lambda$ along with the bound + 1 for the positive terms as argued above. The last term can be bounded as in the above calculation leading to $E(e^{\gamma\Delta'}) = O(1)$ since that estimation

uses only the second condition, which holds conditional on $X_t > a$. Hence, in any case $D(\ell) = O(1)$. Altogether, we have

$$\begin{aligned} e^{-\lambda(\ell)\cdot\ell} \cdot D(\ell) \cdot p(\ell) &\leq e^{-\lambda\ell} \cdot \frac{4}{\lambda\epsilon} \\ &= e^{-\lambda\ell\epsilon + \ln(4/(\lambda\epsilon))} \end{aligned}$$

By the third condition, we have $\lambda\ell \geq 2 \ln(4/(\lambda\epsilon))$, which finally means that

$$e^{-\lambda(\ell)\cdot\ell} \cdot D(\ell) \cdot p(\ell) \leq O(e^{-\lambda\ell\epsilon/2})$$

Choosing $L(\ell) = e^{\lambda\ell/4}$, Theorem 15 yields

$$\Pr(T(\ell) \leq L(\ell)) \leq L(\ell) \cdot O(e^{-\lambda\ell/4}) = O(e^{-\lambda\ell/4}),$$

which proves the theorem. □

References

- Baillon, J.-B., Cominetti, R., Vaisman, J.: A sharp uniform bound for the distribution of sums of Bernoulli trials. *Comb. Probab. Comput.* **25**(3), 352–361 (2016)
- Chen, T., Tang, K., Chen, G., Yao, X.: On the analysis of average time complexity of estimation of distribution algorithms. In: *Proceedings of CEC '07*, pp. 453–460 (2007)
- Chen, T., Lehre, P.K., Tang, K., Yao, X.: When is an estimation of distribution algorithm better than an evolutionary algorithm? In: *Proceedings of CEC '09*, pp. 1470–1477 (2009)
- Chen, T., Tang, K., Chen, G., Yao, X.: Rigorous time complexity analysis of univariate marginal distribution algorithm with margins. In: *Proceedings of CEC '09*, pp. 2157–2164 (2009)
- Chen, T., Tang, K., Chen, G., Yao, X.: Analysis of computational time of simple estimation of distribution algorithms. *IEEE Trans. Evol. Comput.* **14**(1), 1–22 (2010)
- Dang, D.-C., Lehre, P.K.: Simplified runtime analysis of estimation of distribution algorithms. In: *Proceedings of GECCO '15*, pp. 513–518. ACM Press, New York (2015)
- Droste, S.: A rigorous analysis of the compact genetic algorithm for linear functions. *Nat. Comput.* **5**(3), 257–283 (2006)
- Friedrich, T., Kötzing, T., Krejca, M.S., Sutton, A.M.: The benefit of recombination in noisy evolutionary search. In: *Proceedings of ISAAC '15*, pp. 140–150. Springer, Berlin (2015)
- Friedrich, T., Kötzing, T., Krejca, M.S.: EDAs cannot be balanced and stable. In: *Proceedings of GECCO '16*, pp. 1139–1146. ACM Press, New York (2016)
- Hajek, B.: Hitting-time and occupation-time bounds implied by drift analysis with applications. *Adv. Appl. Probab.* **14**, 502–525 (1982)
- Hauschild, M., Pelikan, M.: An introduction and survey of estimation of distribution algorithms. *Swarm Evol. Comput.* **1**(3), 111–128 (2011)
- Johannsen, D.: Random combinatorial structures and randomized search heuristics. Ph.D. Thesis, Universität des Saarlandes, Germany (2010). <http://scidok.sulb.uni-saarland.de/volltexte/2011/3529>
- Kaas, R., Buhman, J.M.: Mean, median and mode in binomial distributions. *Stat. Neerl.* **34**, 13–18 (1980)
- Krejca, M.S., Witt, C.: Lower bounds on the run time of the univariate marginal distribution algorithm on OneMax. In: *Proceedings of FOGA 2017*, pp. 65–79. ACM Press, New York (2017)
- Larrañaga, P., Lozano, J.A. (eds.): *Estimation of Distribution Algorithms: A New Tool for Evolutionary Computation*, Volume 2 of *Genetic Algorithms and Evolutionary Computation*. Springer, Berlin (2002)
- Lehre, P.K., Nguyen, P.T.H.: Improved runtime bounds for the univariate marginal distribution algorithm via anti-concentration. In: *Proceedings of GECCO '17*, pp. 414–434. ACM Press, New York (2017)

17. Lehre, P.K., Witt, C.: Concentrated hitting times of randomized search heuristics with variable drift. In: Proceedings of ISAAC '14, pp. 686–697. Springer, Berlin (2014). Full technical report at [arXiv:1307.2559](https://arxiv.org/abs/1307.2559)
18. McDiarmid, C.: Concentration. In: Habib, M., McDiarmid, C., Ramirez-Alfonsin, J., Reed, B. (eds.) Probabilistic Methods for Algorithmic Discrete Mathematics, pp. 195–247. Springer, Berlin (1998)
19. Mitavskiy, B., Rowe, J.E., Cannings, C.: Theoretical analysis of local search strategies to optimize network communication subject to preserving the total number of links. *Int. J. Intell. Comput. Cybern.* **2**(2), 243–284 (2009)
20. Mühlenbein, H., Paass, G.: From recombination of genes to the estimation of distributions I. Binary parameters. In: Proceedings of PPSN IV, pp. 178–187. Springer, Berlin (1996)
21. Neumann, F., Sudholt, D., Witt, C.: A few ants are enough: ACO with iteration-best update. In: Proceedings of GECCO '10, pp. 63–70. ACM Press, New York (2010)
22. Oliveto, P.S., Witt, C.: Improved time complexity analysis of the simple genetic algorithm. *Theor. Comput. Sci.* **605**, 21–41 (2015)
23. Rowe, J.E., Sudholt, D.: The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theor. Comput. Sci.* **545**, 20–38 (2014)
24. Samuels, S.M.: On the number of successes in independent trials. *Ann. Math. Stat.* **36**(4), 1272–1278 (1965)
25. Sudholt, D.: A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Trans. Evol. Comput.* **17**(3), 418–435 (2013)
26. Sudholt, D., Witt, C.: Update strength in EDAs and ACO: how to avoid genetic drift. In: Proceedings of GECCO '16, pp. 61–68. ACM Press, New York (2016)
27. Witt, C.: Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Comb. Probab. Comput.* **22**(2), 294–318 (2013)
28. Witt, C.: Upper bounds on the runtime of the univariate marginal distribution algorithm on OneMax. In: Proceedings of GECCO '17, pp. 1415–1422 (2017)
29. Wu, Z., Michael, K., Möhring, R.H.: Stochastic runtime analysis of the cross-entropy algorithm. *IEEE Trans. Evolut. Comput.* **21**(4), 616–628 (2017)