



HAL
open science

Analyzing randomized search heuristics via stochastic domination

Benjamin Doerr

► **To cite this version:**

Benjamin Doerr. Analyzing randomized search heuristics via stochastic domination. Evolutionary Computation in Combinatorial Optimization - 18th European Conference, EvoCOP 2018, 2018, Parma (Italie), Italy. pp.115-137, 10.1016/j.tcs.2018.09.024 . hal-04484605

HAL Id: hal-04484605

<https://hal.science/hal-04484605>

Submitted on 4 Apr 2024

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

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

Better Runtime Guarantees Via Stochastic Domination*

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

May 2, 2019

Abstract

Apart from few exceptions, the mathematical runtime analysis of evolutionary algorithms is mostly concerned with expected runtimes. In this work, we argue that stochastic domination is a notion that should be used more frequently in this area. Stochastic domination allows to formulate much more informative performance guarantees, it allows to decouple the algorithm analysis into the true algorithmic part of detecting a domination statement and the probability-theoretical part of deriving the desired probabilistic guarantees from this statement, and it helps finding simpler and more natural proofs.

As particular results, we prove a fitness level theorem which shows that the runtime is dominated by a sum of independent geometric random variables, we prove the first tail bounds for several classic runtime problems, and we give a short and natural proof for Witt's result that the runtime of any (μ, p) mutation-based algorithm on any function with unique optimum is subdominated by the runtime of a variant of the $(1 + 1)$ EA on the ONEMAX function. As side-products, we determine the fastest unbiased $(1+1)$ algorithm for the

*Extended version of a paper that appeared at *EvoCOP 2018* [Doe18a]. This version contains as new material a section on known precise runtime distributions, a Chernoff bound for sums of independent coupon collector runtimes, several new tail bounds for classic runtime results, and a section on counter-examples.

LEADINGONES benchmark problem, both in the general case and when restricted to static mutation operators, and we prove a Chernoff-type tail bound for sums of independent coupon collector distributions.

Keywords: Evolutionary algorithms, runtime analysis.

1 Introduction

The analysis of evolutionary algorithms via mathematical means is an established part of evolutionary computation research. The subarea of *runtime analysis* aims at giving proven performance guarantees on the time an evolutionary algorithm takes to find optimal or near-optimal solutions. Traditionally, this area produces estimates for the expected runtime, which are occasionally augmented by tail bounds. A justification for the restriction to expectations was that already for very simple evolutionary algorithms and optimization problems, the stochastic processes arising from running this algorithm on this problem are so complicated that any more detailed analysis is infeasible. See the analysis how the $(1 + 1)$ evolutionary algorithm optimizes linear functions $f: \{0, 1\}^n \rightarrow \mathbb{R}; x \mapsto a_1x_1 + \dots + a_nx_n$ in [DJW02] for an example.

In this work, we shall argue that the restriction to expectations is less justified and propose stochastic domination as an alternative. It is clear that the precise distribution of the runtime of an evolutionary algorithm often is out of reach (see Section 2 for an overview on what runtime distributions are known exactly). Finding the precise distribution is maybe not even an interesting target because most likely already the result will be too complicated to be useful. What would be very useful is a possibly not absolutely tight upper bound-type statement that concerns the whole distribution of the runtime.

One way to formalize such statement is via the notion of *stochastic domination*. A real-valued random variable Y stochastically dominates another one X if and only if for each $\lambda \in \mathbb{R}$, we have

$$\Pr[X \leq \lambda] \geq \Pr[Y \leq \lambda].$$

If X and Y describe the runtimes of two algorithms A and B , then this domination statement is a very strong way of saying that algorithm A is at least as fast as B . See Section 3 for a more detailed discussion of the implication of such a domination statement.

1.1 Our Results

In this work, we shall give three main arguments for a more frequent use of domination arguments in runtime analysis and support these with several new runtime results.

Stochastic domination is often easy to show. Surprisingly, despite being a much stronger type of assertion, stochastic domination statements are often easy to obtain. The reason is that many of the classic proofs implicitly contain all necessary information, they only fail to formulate the result as a domination statement.

As an example, we prove in Section 4 a natural domination version of the classic fitness level method. In analogy to the classic result, which translates pessimistic improvement probabilities p_1, \dots, p_{m-1} into an expected runtime estimate $E[T] \leq \sum_{i=1}^{m-1} \frac{1}{p_i}$, we show that under the same assumptions the runtime is dominated by the sum of independent geometric random variables with success probabilities p_1, \dots, p_{m-1} . This statement implies the classic result, but also implies tail bound for the runtime via Chernoff bounds for geometric random variables. We note that, while our extension of the fitness level theorem is very natural, the proof is not totally obvious, which might explain why previous works were restricted to the expectation version.

Stochastic domination allows to separate the core algorithm analysis and the probability theoretic derivation of probabilistic runtime statements. The reason why stochastic domination statements are often easy to obtain is that they are close to the actions of the algorithms. When we are waiting for the algorithm to succeed in performing a particular action, then it is a geometric distribution that describes this waiting time. To make such a statement precise for a given problem, we need to understand how the algorithm achieves the particular goal. This requires a good understanding of the algorithm and the problem as well as some elementary probability theory and discrete mathematics, but usually no greater expertise in probability theory. We give in Sections 4.3 and 5 several examples, mostly using the fitness level method, but also concerning the single-source shortest paths problem, where the fitness level method is not suitable to give the best known results.

Once we have a domination statement formulated, e.g., that the runtime is dominated by a sum of independent geometric distributions, then deeper probability theoretic arguments like Chernoff-type tail bounds come into play. This part of the analysis is independent of the algorithm and only relies on the domination statement. Exemplarily, we derive tail bounds for the

runtime of the $(1 + \lambda)$ EA on ONEMAX, the $(\mu + 1)$ EA on LEADINGONES, the $(1 + 1)$ EA on jump functions, the $(1 + 1)$ EA for sorting, the multi-criteria $(1 + 1)$ EA for the single-source shortest path problem, and some more results.

That the two stages of the analysis are of a different nature is also visible in the history of runtime analysis. As discussed above, many classic proofs essentially contains all ingredients to formulate a domination statement. However, the mathematical tools to analyze sums of independent geometric random variables were developed much later (and this development is still ongoing).

This historical note also shows that from the viewpoint of research organization, it would have been profitable if previous works would have been formulated in terms of domination statements. This might have spurred a faster development of suitable Chernoff bounds and, in any case, it would have made it easier to apply the recently found Chernoff bounds (see Theorems 14 to 16) to these algorithmic problems.

Stochastic domination often leads to more natural and shorter proofs. To demonstrate this in an application substantially different from the previous ones, we regard the classic lower-bound result which, in simple words, states that solving ONEMAX via the $(1 + 1)$ EA is the fastest optimization process among many mutation-based algorithms and problems with unique optimum. This statement, of course, should be formulated via stochastic domination and this has indeed been done in previous work. However, as we shall argue in Section 6, we also have the statement that when comparing the two processes the distance of the current-best solution from the optimum for the general function dominates this distance for the ONEMAX function. This natural statement immediately implies the domination relation between the runtimes. We make this precise for the current-strongest ONEMAX-is-easiest result [Wit13]. This will shorten the previous, two-and-a-half pages long complicated proof to an intuitive proof of less than a page.

Side-results. We use our discussion of previous results on precisely known distributions of runtimes of evolutionary algorithms to (i) give a formal proof for the distribution result on the runtime of the $(1 + 1)$ EA on the LEADINGONES benchmark function stated without proof in [DJWZ13], (ii) to extend this result to a wide class of $(1+1)$ type algorithms and to time-to-target runtimes, and (iii) to determine the best possible unbiased $(1+1)$ algorithm

for this benchmark function, both in the general case and with the restriction to static mutation operators (Section 2.3).

We also use our discussion of Chernoff bounds for sums of geometric random variables to show a Chernoff bound for sums of coupon collector random variables (Section 4.2).

Related work

The use of stochastic domination is not totally new in the theory of evolutionary computation, however, the results so far appear rather sporadic than systematic.

Possibly the first to use the notion of stochastic domination was Droste, who in [Dro03, Dro04] employed it to make precise an argument often used in an informal manner, namely that some artificial random process is not faster than the process describing a run of the algorithm under investigation. Using domination as *tool in proofs* has been done subsequently in various applications, among others, binary particle swarm (PSO) algorithms [SW08], ant-colony optimization [ST12], evolutionary multi-objective optimization [BF13], optimization of monotonic functions [CDF14, Len18], non-elitist algorithms [RS14], fixed budget analyses [LS15], the simple GA [OW15], estimation-of-distribution algorithms [FKKS17, KW17], genetic programming [DKLL17], island models [DFF⁺17], and dynamic and noisy optimization [DNDD⁺18].

The first time that stochastic domination was used to *formulate results* (though without explicitly calling it stochastic domination) was in [BE08], where several result were shown of the type that when one mutation operator is better than another one (in a suitable domination sense), then the algorithm using the first one always has at least as good solutions as the second (in the sense of stochastic domination). In [Wit13], the result is shown that the runtime of an arbitrary (μ, p) mutation-based algorithm on any objective function with unique optimum dominates the runtime of the $(1 + 1)$ EA with mutation rate p and best-of- μ initialization on the ONEMAX function (see Section 6 for more details on this result).

What comes closest to this work is the paper [ZLLH12], which also tries to establish runtime analysis beyond expectations in a formalized manner. The notion proposed in [ZLLH12], called *probable computational time* $L(\delta)$, is the smallest time T such that the algorithm under investigation within the first T fitness evaluations finds an optimal solution with probability at least $1 - \delta$. In some sense, this notion can be seen as an inverse of the classic tail bound language, which makes statements of the type that the probability that the algorithm needs more than T iterations, is at most some function in T , which

is often has an exponential tail. To prove results on the probable computational time, the authors implicitly use (but do not prove) the result that the fitness level method gives a stochastic domination by a sum of independent geometric random variables (which we prove in Section 4). When using this result to obtain bounds on the probable computational time, they suffer from the fact that at that time, good tail bounds for sums of geometric random variables with different success probabilities were not yet available (these appeared only in [Wit14,DD15]). For this reason, they had to use a self-made tail bound, which unfortunately gives a non-trivial tail probability only for values higher than twice the expectation. With this limitation, tail bounds for the runtimes of the algorithms RLS, $(1 + 1)$ EA, $(\mu + 1)$ EA, MMAS* and binary PSO on the optimization problems ONEMAX and LEADINGONES are proven.

In [DL17], the complexity theoretic notion of p -Monte Carlo black-box complexity was introduced, which serves at providing more information than what expected runtimes convey for lower bounds on runtimes.

2 Exact Distributions

We said above that stochastic domination can be used to formulate runtime bounds in a distributional sense, which can give a very complete picture but without having to determine the exact runtime distribution. Before going into details, we shall summarize in this section the few situations in which it has been possible to describe the runtime distribution precisely. Not surprisingly, these all regard very simple randomized search heuristics.

This section contains a few new results, e.g., the determination of the fastest unbiased $(1+1)$ algorithm for the LEADINGONES function for static and dynamic mutation operators.

2.1 Random Search

The *random search heuristics* repeatedly samples solutions chosen uniformly at random until some termination criterion is met. Therefore, the time to sample a certain solution (or one of a given subset) follows a *geometric distribution*. We say that a random variable X is geometrically distributed with parameter $0 < p \leq 1$ and write $X \sim \text{Geom}(p)$ if

$$\Pr[X = k] = (1 - p)^{k-1}p$$

for all $k \in \mathbb{N}$. Note that this implies $\Pr[X \geq k] = (1 - p)^{k-1}$ and $E[X] = \frac{1}{p}$.

Lemma 1. Consider some search or optimization problem with finite search space \mathcal{S} . Let \mathcal{T} be a non-empty subset of \mathcal{S} . Let T be the random variable describing the first time random search samples a solution from \mathcal{T} . Then

$$T \sim \text{Geom}\left(\frac{|\mathcal{T}|}{|\mathcal{S}|}\right).$$

In particular, the time T to find a unique optimum in a problem with bit-string representation of length n is $T \sim \text{Geom}(2^{-n})$.

2.2 Randomized Local Search

The *randomized local search* (RLS) heuristic starts with a random search point and then repeats sampling a neighbor of the current solution and replacing the current solution with this offspring if it is at least as good. In the case of the bit-string representation of length n , neighbors are bit-strings which differ in exactly one bit. Hence this heuristic repeats flipping a single random bit and continues with the new solution if it is not worse than the parent. This is exactly Algorithm 1 with the 1-bit mutation operator.

Algorithm 1: The family of (1+1) search heuristics for maximizing a given objective function $f: \{0, 1\}^n \rightarrow \mathbb{R}$. When the mutation operator returns a random Hamming neighbor of x , that is, flips a random bit in x , then this algorithm is the RLS heuristic. When the mutation operator flips each bit independently with probability p , then we obtain the (1 + 1) EA with mutation rate p . The classic (1 + 1) EA uses a mutation rate of $p = \frac{1}{n}$.

```

1 Choose  $x \in \{0, 1\}^n$  uniformly at random;
2 for  $t = 1, 2, 3, \dots$  do
3    $y \leftarrow \text{mutate}(x)$ ;
4   if  $f(y) \geq f(x)$  then  $x \leftarrow y$ ;
```

When used to optimize the classic benchmark function ONEMAX: $\{0, 1\}^n \rightarrow \mathbb{R}$ defined by

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

or any other *strictly monotonic* function $f: \{0, 1\}^n \rightarrow \mathbb{R}$ (where flipping a 0-bit into a 1 always increases the fitness), the resulting optimization

process is equivalent to a coupon collector process in which each type of coupon is present initially with probability $\frac{1}{2}$, see [DD16]. The runtime of the coupon collector process with $n - k$ of the n coupons initially present is $\sum_{i=1}^k \text{Geom}(\frac{i}{n})$, where this sum is understood as sum of independent distributions. Consequently, the runtime of RLS on strictly monotonic functions can be described as follows, where we write $X \sim \text{Bin}(n, p)$ to indicate that X follows a *binomial distribution* with parameters n and $p \in [0, 1]$, that is, we have $\Pr[X = i] = \binom{n}{i} p^i (1 - p)^{n-i}$ for all $i \in [0..n]$.

Lemma 2. *Let $f: \{0, 1\}^n \rightarrow \mathbb{R}$ be any strictly monotonic function. Let T be the number of iterations¹ taken by RLS to sample the (unique) optimum of f . Then*

$$T \sim \sum_{i=0}^{n-1} 1_{X \leq i} \cdot \text{Geom}(\frac{n-i}{n}),$$

where $X \sim \text{Bin}(n, \frac{1}{2})$ and we assume that X and the arising geometric distributions are independent.

2.3 The LeadingOnes Benchmark Problem

For a surprisingly large class of algorithms the precise runtime on the LEADINGONES benchmark problem can be determined. The function LEADINGONES: $\{0, 1\}^n \rightarrow \mathbb{R}$ is defined by

$$\text{LEADINGONES}(x) = \min\{i \in [0..n] \mid \forall j \in [1..n] \cap \mathbb{R}_{\leq i} : x_j = 1\}$$

for all $x \in \{0, 1\}^n$.

In parallel independent work, the precise expected runtime of the $(1 + 1)$ EA on the LEADINGONES benchmark function was determined in [BDN10, Sud13] (note that [Sud13] is the journal version of a work that appeared at the same conference as [BDN10]). In terms of results, the two works are very similar. The work [Sud13] is also regards the $(1 + 1)$ EA with Best-of- μ initialization and shows that a lower-order runtime gain can be obtained from taking as initial solution the best of, say, n random search points. The work [BDN10] also shows that the often recommended mutation rate of $p = 1/n$ is not optimal. A runtime smaller by 16% can be obtained from taking $p = 1.59/n$ and another 12% can be gained by using a fitness-dependent mutation rate.

¹In this work, we shall call the runtime of an iterative algorithm the number of iterations taken until the optimum is generated for the first time. In comparison with the classic definition, the number of fitness evaluations until the optimum is evaluated, we thus do not count the initial search points and, in the case that per iteration $\lambda > 1$ individuals are generated, we do not count the runtime of the iteration as λ , but as one.

In terms of methods, the two works are substantially different. An advantage of the analysis in [BDN10] is that it can easily be reformulated to give the precise distribution of this runtime. This was first observed in [DJWZ13], but used several times subsequently. The arguments in [BDN10] are not specific to the $(1+1)$ EA, but are valid for any other $(1+1)$ algorithm as defined in Algorithm 1. Since already the observation in [DJWZ13] was not formally proven, we now quickly formulate the general result and prove it.

Theorem 3. *Consider the run of a $(1+1)$ algorithm optimizing the LEADINGONES function. Let T be the first time the optimum is generated. Then*

$$T \sim \sum_{i=0}^{n-1} X_i \cdot \text{Geom}(q_i),$$

where X_0, X_1, \dots, X_{n-1} are uniformly distributed binary random variables, the X_i and $\text{Geom}(q_i)$ are mutually independent, and for all $i \in [0..n-1]$ we denote by q_i the probability that the mutation operator generates from a search point of fitness exactly i a strictly better search point. Consequently, $E[T] = \frac{1}{2} \sum_{i=0}^{n-1} \frac{1}{q_i}$, where we read $\frac{1}{p_i} = \infty$ when $p_i = 0$.

Proof. For all $i \in [0..n]$, denote by T_i^0 the runtime of the algorithm when starting with a random search point of fitness exactly i . Note that if x is such a random search point, then $x_j = 1$ for $j \in [1..i]$, $x_{i+1} = 0$, and for $j \in [i+2..n]$ the x_j are independent random variables uniformly distributed in $\{0, 1\}$. Let T_i^{rand} denote the runtime of the algorithm when starting with a random search point x of fitness at least i , that is, we have $x_j = 1$ for $j \in [1..i]$ and for $j \in [i+1..n]$ the x_j are independent random variables uniformly distributed in $\{0, 1\}$.

Trivially, we have $T_n^0 = T_n^{\text{rand}} = 0$. The main technical insight is that due to the random initialization in the algorithm, we have

$$T_i^0 = \text{Geom}(q_i) + T_{i+1}^{\text{rand}} \quad (1)$$

for $i < n$. Here $\text{Geom}(q_i)$ describes the waiting time for an improvement when having a search point of fitness exactly i . At the moment when this improvement happens, the bits with index $i+2$ or higher are still independent and uniformly distributed binary random variables, since their value cannot have had any influence on the run of the algorithm. Formally speaking, we have $(x_{i+2}, \dots, x_n) = (x_{i+2}^0, \dots, x_n^0) \oplus (y_{i+2}, \dots, y_n)$, where x^0 denotes the random initial search point and $y = (y_{i+2}, \dots, y_n)$ describes which bits have changed compared to the initial bit-string. Note that y is independent of $(x_{i+2}^0, \dots, x_n^0)$. Hence since $(x_{i+2}^0, \dots, x_n^0)$ is a random binary string, also $(x_{i+2}, \dots, x_n) = (x_{i+2}^0, \dots, x_n^0) \oplus (y_{i+2}, \dots, y_n)$ is a random binary string.

With (1) we easily compute

$$\begin{aligned} T_i^{\text{rand}} &= X_i T_i^0 + (1 - X_i) T_{i+1}^{\text{rand}} \\ &= X_i (\text{Geom}(q_i) + T_{i+1}^{\text{rand}}) + (1 - X_i) T_{i+1}^{\text{rand}} \\ &= X_i \text{Geom}(q_i) + T_{i+1}^{\text{rand}}, \end{aligned}$$

where X_i is a uniformly distributed binary random variable independent of any randomness used in the other distributions. From the latter equation, a simple induction proves the claim. Note that the runtime of the algorithm is T_0^{rand} . \square

As observed in [DJWZ13] and exploited for a fixed-budget analysis, an analogous result is valid for the time $T^{\rightarrow a}$ to first reach a search point of fitness at least $a \in [0..n]$, namely $T^{\rightarrow a} \sim \sum_{i=0}^{a-1} X_i \cdot \text{Geom}(q_i)$; note that in Lemma 5 of [DJWZ13], the range of the sum starts at 1, but this clearly is a typo. Since this extension to time-to-target runtimes can be important (see [CD17]), we quickly make this point precise even if in the remainder of this work we restrict ourselves to the traditional runtime analysis target of determining the time to find the optimum.

Corollary 4. *In the notation of Theorem 3, the time $T^{\rightarrow a}$ the algorithm takes to find a solution of fitness at least a is*

$$T^{\rightarrow a} \sim \sum_{i=0}^{a-1} X_i \cdot \text{Geom}(q_i)$$

for all $a \in [0..n]$.

Proof. Let $a \in [0..n]$. Denote by T_i^0, T_i^{rand} the analogues of the times defined in the previous proof, but for the target of reaching a fitness of at least a . Then $T_a^0 = T_a^{\text{rand}} = 0$ and for $i < a$ the same relations hold as above. This proves the claim. \square

Theorem 3 allows to determine precisely the expected runtimes of many (1+1) algorithms. More importantly, it also allows to determine optimal parameters for the mutation operator.

- For the classic RLS heuristic always flipping a single random bit, we have $q_i = \frac{1}{n}$ for all $i \in [0..n-1]$. Consequently, the expected runtime is $E[T] = 0.5n^2$. As we shall show below, this is also the best static unbiased mutation operator for the LEADINGONES problem.

- Motivated by an analysis how evolutionary algorithms can solve problems with unknown solution length (that is, how many bits are actually relevant for the problem), in [DDK17, Lemma 3.4] an extension of the RLS heuristic was analyzed which, for a given sequence $p_1, p_2, \dots, p_n \in [0, 1]$ with $\sum_{i=1}^n p_i \leq 1$, as 1-bit mutation flips the i -th bit with probability exactly p_i . For this mutation operator, an expected runtime of $\frac{1}{2} \sum_{i=1}^n \frac{1}{p_i}$ was shown for instances of fixed length n . Hence for known solution length, there is no gain from flipping bits with position-dependent probabilities. This result is now an immediate consequence of Theorem 3.
- For the variant of RLS which randomly mixes the 1-bit and 2-bit flip operator, that is, with probability P one random bit is flipped and with probability $1 - P$ two (different) random bits are flipped, we have $q_i = \frac{P}{n} + 2(1 - P) \frac{n-i-1}{n(n-1)}$. From this, a slightly tedious calculation gives an expected runtime of $E[T] = \frac{1}{4(1-p)} \ln\left(\frac{2-p}{p}\right) n^2 + o(n^2)$, see [LOW17] and note that the fact that the authors consider 2-bit flips with repetition (that is, with probability $1/n$ the same bit is flipped twice resulting in a copy of the parent) has no influence on the result apart from lower-order terms.
- When flipping k random bits in a search point with fitness i , the probability $q(n, k, i)$ of obtaining a strictly better solution is $q(n, k, i) = \frac{k(n-i-1) \dots (n-i-k+1)}{n(n-1) \dots (n-k+1)}$. Consequently, $q(n, k, i) \leq q(n, k+1, i)$ if and only if $i \leq \frac{n-k}{k+1}$. This gives that the best (that is, largest) value for q_i is obtained from flipping $k(n, i) := \lfloor \frac{n}{i+1} \rfloor$ bits when the current fitness is i .² This fitness-dependent choice of the mutation operator gives a runtime of approximately $0.39n^2$. Since any unary unbiased mutation operator (see [LW12]) is the convex combination of k -bit flip operators (see [DDY16b] or [DKLW13]), this algorithm also is the fastest unbiased (1+1) algorithm for LEADINGONES.
- For the (1 + 1) EA, that is, when the mutation operator is standard-bit mutation with mutation rate p , we have $q_i = q_i(p) = (1 - p)^i p$ for all $i \in [0..n - 1]$. This yields an expected runtime of $E[T] = \frac{1}{2p^2} ((1 - p)^{1-n} - (1 - p))$. For the standard choice $p = \frac{1}{n}$, this is $E[T] =$

²In an earlier version of this work, we said that $\lfloor \frac{n+1}{i+1} \rfloor$ would be the optimal number of bits to flip in a search point with fitness i . Apart from the trivial case $i = 0$, where the previous formulation gives $n + 1$ instead of the desired value of n , both versions are correct. For all $i > 0$ such that two definitions differ, both numbers of bits to be flipped give the same probability for a fitness improvement. Thanks to Carola Doerr and Markus Wagner for pointing me to the new formula from their forthcoming work [DW18].

$\frac{1}{2}(e-1-o(1))n^2 \approx 0.86n^2$. A better (and the asymptotically best among all static rates) expected runtime is obtained from using $p \approx 1.59/n$, giving $E[T] \approx 0.77n^2$. The best (fitness-dependent) mutation rate is using $p_i = \frac{1}{i+1}$ when the current fitness is i , giving a runtime of $E[T] = \frac{e}{4}n^2 \pm O(n) \approx 0.77n^2$. All these results are from [BDN10]. Apparently unaware of this work, the expected runtime in the case $p = \frac{1}{n}$ was again determined in [DAS⁺17].

- Since the heavy-tailed mutation operator proposed in [DLMN17] applies standard-bit mutation with a randomly chosen mutation rate, the q_i are convex combinations of the $q_i(p)$ computed above. This again determines the expected runtime of the $(1+1)$ EA with this mutation operator, however, a simple closed formula as for standard-bit mutation with fixed rate is not known.
- When allowing position-dependent mutation rates for the $(1+1)$ EA, that is, the i -th bit is flipped with probability p_i independently for all $i \in [1..n]$, then an expected runtime of $E[T] = \frac{1}{2} \sum_{i=1}^n (p_i \prod_{j=1}^{i-1} (1 - p_j))^{-1}$ was shown in [DDK17]. This is minimized by taking $p_i = \frac{1}{i}$ for all i , which yields $E[T] = 0.5n^2$.

We now prove the results promised above that the 1-bit mutation operator as used by RLS is the (unique) best unbiased static mutation operator for the $(1+1)$ algorithm scheme. We say that a $(1+1)$ algorithm uses a static mutation operator if for all $x \in \{0,1\}^n$ the distribution of the offspring generated from x is the same throughout the run of the algorithm.

Theorem 5. *For any $(1+1)$ algorithm using static unbiased mutation, the expected runtime on the LEADINGONES function is at least $0.5n^2$. Moreover, RLS is the only such algorithm satisfying this bound.*

Proof. Consider a static unbiased mutation operator. As above, there are $r_0, r_1, \dots, r_n \in [0,1]$ such that $\sum_{k=0}^n r_k = 1$ and such that the mutation operator (unchanged throughout the run of the algorithm) can be written as “sample $k \in [0..n]$ with probability r_k and then flip exactly k random bits”.

Let again $q(n, k, i) = \frac{k(n-i-1)\dots(n-i-k+1)}{n(n-1)\dots(n-k+1)}$ be the probability of obtaining a strictly better solution from a search point of fitness i by flipping k bits. We apply Theorem 3 with $q_i = \sum_{k=0}^{n-1} r_k q(n, k, i)$ and obtain an expected optimization time of $E[T] = \frac{1}{2} \sum_{i=0}^{n-1} \frac{1}{q_i}$. We show that this expression, viewed as function of the r_k , has a unique minimum for $r_1 = 1$ (and hence $r_k = 0$ for all $k \neq 1$).

It is clear that flipping zero bits can never lead to an improvement. Hence moving any positive mass on r_0 to r_1 would strictly improve the algorithm. Without loss of generality, we thus assume in the following that $r_0 = 0$.

We first show that, regardless of the remaining r_k , we have $\sum_{i=0}^{n-1} q_i = 1$. This is, naturally, equivalent to saying that

$$\sum_{i=0}^{n-1} q(n, k, i) = 1 \quad (2)$$

for all $k \in [1..n]$. Since we have a precise expression for the $q(n, k, i)$, equation (2) in principle could be shown via Faulhaber's formula [Fau31]. Fortunately, a much simpler probabilistic argument can be applied. Let A_i be the event that in an application of the k -bit mutation operator the i -th bit is flipped and that no bit j with $j < i$ is flipped. Then $\Pr[A_i] = q(n, k, i - 1)$. Trivially, the A_i are disjoint events that cover the whole probability space. Hence

$$1 = \Pr[A_1 \cup \dots \cup A_n] = \sum_{i=1}^n \Pr[A_i] = \sum_{i=0}^{n-1} q(n, k, i)$$

for all $k \in [0..n - 1]$ and thus $\sum_{i=0}^{n-1} q_i = 1$.

Next, we recall that among all q_0, \dots, q_{n-1} with $\sum_{i=0}^{n-1} q_i = 1$, the sum $\sum_{i=0}^{n-1} \frac{1}{q_i}$ of the reciprocals is minimal if and only if the q_i are all equal. Since $i \mapsto q(n, k, i)$ is non-increasing, so is $i \mapsto q_i$, and the only way to have the q_i all equal is that for all $k \in [1..n]$ the function $i \mapsto r_k q(n, k, i)$ is constant. Since only for $k = 1$ the function $i \mapsto q(n, k, i)$ is constant, we have $r_k = 0$ for all $k \neq 1$, and thus $r_1 = 1$ as claimed. \square

2.4 Conclusion on Exact Runtime Distributions

The results on LEADINGONES, in particular, how the knowledge of the precise distribution easily yields optimal parameter values, show that understanding the precise runtime distribution would be highly desirable.

However, it has to be noticed that the results presented above are limited to very restricted settings and it is highly unlikely that they can be extended significantly. For example, that the exact runtime analysis of randomized local search on ONEMAX can be extended to other algorithms appears hard to believe given that for the runtime of the $(1 + 1)$ EA on ONEMAX, despite significant efforts [Rud97, DFW10, DFW11, Sud13, HPR⁺18], not even an exact value for the expectation is known. Similarly, the results for LEADINGONES depend crucially on the facts that only $(1+1)$ algorithms are regarded and that the initial individual is random. For a different initialization, none of the

results would hold and also analogous results determining the exact runtime distribution are not in sight.

For this reason, the approach discussed in the following is promising: To not try to find the exact distribution, but to ask for distributions which are upper bounds in a very strong sense, namely in the sense of stochastic domination. With this notion we will then, e.g., be able to say that the runtime distribution of the $(1+1)$ EA on ONEMAX is dominated by $\sum_{i=1}^n \text{Geom}(\frac{i}{en})$ or that the runtime distribution for an arbitrary $(1+1)$ algorithm on LEADINGONES is dominated by $\sum_{i=0}^{n-1} \text{Geom}(q_i)$ regardless of how the initial solution is chosen.

We note without further detail that another way of weakening the aim of an exact runtime distribution is to give a limiting distribution only. This has been done for the runtime of RLS and $(1+1)$ EA on ONEMAX and the NEEDLE function [GKS99] and on the LEADINGONES function [Lad05]. For example, it is shown in [Lad05] that the runtime T_n of the $(1+1)$ EA with mutation rate c/n on the n -dimensional LEADINGONES function has the property that

$$\frac{T_n - \frac{e^c - 1}{2c^2}}{n^{3/2}}$$

converges in distribution to a centered Gaussian random variable with variance

$$\sigma^2 = \frac{3(e^{2c} - 1)}{8c^3}.$$

The results on the runtime of the $(1+1)$ EA on ONEMAX and LEADINGONES have been reproven with different methods and made significantly more precise in [HPR⁺18].

3 Stochastic Domination

In this section, we recall the definition of *stochastic domination* and collect a few known properties of this notion. For an extensive treatment of various forms of stochastic orders, we refer to [MS02].

3.1 Definition of Stochastic Domination

Stochastic domination is usually defined as follows.

Definition 6 (Stochastic domination). *Let X and Y be random variables not necessarily defined on the same probability space. We say that Y stochastically dominates X , written as $X \preceq Y$, if for all $\lambda \in \mathbb{R}$ we have $\Pr[X \leq \lambda] \geq \Pr[Y \leq \lambda]$.*

If Y dominates X , then the cumulative distribution function of Y is pointwise not larger than the one of X . The definition of domination is equivalent to “ $\forall \lambda \in \mathbb{R} : \Pr[X \geq \lambda] \leq \Pr[Y \geq \lambda]$ ”, which shows more clearly why we feel that Y is at least as large as X .

Concerning nomenclature, we remark that some research communities in addition require that the inequality is strict for at least one value of λ . Hence, intuitively speaking, Y is strictly larger than X . From the mathematical perspective, this appears not very practical, and from the computer science perspective it is not clear what can be gained from this alternative definition. Consequently, our definition above is more common in computer science (though, e.g., [ZLLH12] also use the alternative definition).

One advantage of comparing two distributions via the notion of domination is that this makes a statement over the whole domain of the distributions, including “rare events” on the tails. If the runtime T_A of some algorithm A is dominated by the runtime T_B of algorithm B , then not only A is better than B in average, but also exceptionally large runtimes occur less frequent when using A .

A second advantage is that domination is invariant under monotonic re-scaling. Imagine that running an algorithm for time t incurs some cost $c(t)$. We may clearly assume that c is monotonically increasing, that is, that $t_1 \leq t_2$ implies $c(t_1) \leq c(t_2)$. Then $T_A \preceq T_B$ implies $c(T_A) \preceq c(T_B)$. Hence changing the cost measure does not change our feeling that algorithm A is better than B . Note that this is different for expectations. We may well have $E[T_A] < E[T_B]$, but $E[c(T_A)] > E[c(T_B)]$.

3.2 Properties of Stochastic Domination

We collect a few useful properties of stochastic domination. We start with three elementary observations.

Lemma 7. *If $X \preceq Y$, then $E[X] \leq E[Y]$.*

Lemma 8. *The following two conditions are equivalent.*

(i) $X \preceq Y$.

(ii) *For all monotonically increasing functions $f: \mathbb{R} \rightarrow \mathbb{R}$, we have*

$$E[f(X)] \leq E[f(Y)].$$

Lemma 9. *Let X and Y be random variables.*

(i) *If X and Y are defined on the same probability space and $X \leq Y$, then $X \preceq Y$.*

(ii) If X and Y are identically distributed, then $X \preceq Y$.

The following non-trivial lemma will be needed in our proof of the extended fitness level theorem. It was first proven in a slightly weaker form in [DHK12] and [Doe11]. The current version is from [Doe18b].

Lemma 10. *Let X_1, \dots, X_n be arbitrary discrete random variables. Let X_1^*, \dots, X_n^* be discrete random variables that are mutually independent. Assume that for all $i \in [1..n]$ and all x_1, \dots, x_{i-1} with $\Pr[X_1 = x_1, \dots, X_{i-1} = x_{i-1}] > 0$, we have*

$$\Pr[X_i \geq k \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1}] \leq \Pr[X_i^* \geq k]$$

for all $k \in \mathbb{R}$, that is, X_i^* dominates $(X_i \mid X_1 = x_1, \dots, X_{i-1} = x_{i-1})$. Then

$$\sum_{i=1}^n X_i \preceq \sum_{i=1}^n X_i^*.$$

The following result is again elementary. For discrete random variables, it is a special case of Lemma 10.

Lemma 11. *Let X_1, \dots, X_n be independent random variables defined over some common probability space. Let Y_1, \dots, Y_n be independent random variables defined over a possibly different probability space. If $X_i \preceq Y_i$ for all $i \in [1..n]$, then*

$$\sum_{i=1}^n X_i \preceq \sum_{i=1}^n Y_i.$$

3.3 Coupling

Stochastic domination is tightly connected to *coupling*. Coupling is an analysis technique that consists of suitably defining two unrelated random variables over the same probability space to ease comparing them. Let X and Y be two random variables not necessarily defined over the same probability space. We say that (\tilde{X}, \tilde{Y}) is a *coupling* of (X, Y) if \tilde{X} and \tilde{Y} are defined over a common probability space and if X and \tilde{X} as well as Y and \tilde{Y} are identically distributed. This definition itself is very weak. (X, Y) have many couplings and most of them are not interesting. So the art of coupling as a proof and analysis technique is to find a coupling of (X, Y) that allows to derive some useful information. This is often problem-specific, however, also the following general result is known. It in particular allows to couple dominating random variables. We shall use it for this purpose in Section 6.

Theorem 12. *Let X and Y be random variables. Then the following two statements are equivalent.*

(i) $X \preceq Y$.

(ii) *There is a coupling (\tilde{X}, \tilde{Y}) of (X, Y) such that $\tilde{X} \leq \tilde{Y}$.*

Without going into further detail, we note that coupling is not restricted to real-valued random variables. In the analysis of population-based evolutionary algorithms, a powerful strategy to prove lower bounds is to couple the true population of the algorithm with the population of an artificial process without selection and by this overcome the difficult dependencies introduced by the variation-selection cycle of the algorithm. This was first done in [Wit06] and [Wit08] for the analysis of the $(\mu + 1)$ EA and an elitist steady-state GA. This technique then found applications for memetic algorithms [Sud09], aging-mechanisms [JZ11], non-elitist algorithms [LY12], multi-objective evolutionary algorithms [DKV13], and the $(\mu + \lambda)$ EA [ADFH18].

4 Domination-based Fitness Level Method

In this section, we prove a version of the fitness level theorem that gives domination statements and we apply it to several classic problems.

The fitness level method, invented by Wegener [Weg01], is one of the most successful early analysis methods in the theory of evolutionary computation. It builds on the idea of partitioning the search space into levels A_i , $i = 1, \dots, m$, which contain search points of strictly increasing fitness (that is, for all $i \in [1..m - 1]$, $x \in A_i$, and $y \in A_{i+1}$, the fitness of y is better than the one of x).

We then try to show a lower bound p_i for the probability that, given that the best-so-far search point is in A_i , we generate in one iteration a search point in a higher level (due to this reference to the best-so-far search point, the fitness level theorem is mostly used for elitist algorithms). From this data, the fitness level theorem gives the estimate

$$E[T] \leq \sum_{i=1}^{m-1} \frac{1}{p_i}$$

for the time T (that is, the number of iterations taken) to find a search point in the highest level. Traditionally it is assumed that this highest level contains only optimal solutions, but this restriction can be omitted and the

theorem then gives bounds for the time needed to reach a solution having at least a certain fitness.

We also note that, in principle, there is no need to restrict the fitness level method to fitness levels. Any partition of the set of possible states of the algorithm into a sequence of subsets such that the algorithm cannot leave the current state to a state in a lower-order set would be sufficient. We analyze an application of this type in Section 4.3.3, but stick to the usual fitness level language in the remainder.

4.1 Domination-Version of the Fitness Level Theorem

We shall now show that under the same assumptions, a much stronger statement is valid, namely that the runtime T is dominated by $\sum_{i=1}^{m-1} \text{Geom}(p_i)$, that is, a sum independent random variables following geometric distributions with success probabilities p_i (see Section 2.1 for a definition of the geometric distribution). This result appears to be very natural and was used without proof in [ZLLH12], yet its proof requires the non-trivial Lemma 10.

Theorem 13 (Domination version of the fitness level method). *Consider an elitist evolutionary algorithm \mathcal{A} maximizing a function $f: \Omega \rightarrow \mathbb{R}$. Let A_1, \dots, A_m be a partition of Ω such that for all $i, j \in [1..m]$ with $i < j$ and all $x \in A_i, y \in A_j$, we have $f(x) < f(y)$. Set $A_{\geq i} := A_i \cup \dots \cup A_m$. Let p_1, \dots, p_{m-1} be such that for all $i \in [1..m-1]$ we have that if the best search point in the current parent population is contained in A_i , then independently of the past with probability at least p_i the next parent population contains a search point in $A_{\geq i+1}$.*

Denote by T the (random) number of iterations \mathcal{A} takes to generate a search point in A_m . Then

$$T \preceq \sum_{i=1}^{m-1} \text{Geom}(p_i),$$

where this sum is to be understood as a sum of independent geometric distributions.

Proof. Consider a run of the algorithm \mathcal{A} . For all $i \in [1..m]$, let T_i be the first time (iteration) when \mathcal{A} has generated a search point in $A_{\geq i}$. Then $T = T_m = \sum_{i=1}^{m-1} (T_{i+1} - T_i)$. By assumption, $T_{i+1} - T_i$ is dominated by a geometric random variable with parameter p_i regardless what happened before time T_i . Consequently, Lemma 10 gives the claim. \square

4.2 Chernoff Bounds for Sums of Independent Geometric Random Variables

By Lemma 7, the expected runtime in Theorem 13 satisfies $E[T] \leq \sum_{i=1}^{m-1} \frac{1}{p_i}$, which is the common version of the fitness level theorem [Weg01]. However, by using tail bounds for sums of independent geometric random variables, we also obtain runtime bounds that hold with high probability. This was first proposed in [ZLLH12], but did not give very convincing results due to the lack of good tail bounds at that time. We briefly present the tail bounds known by now and then give a few examples how to use them together with the new fitness level theorem.

Theorem 14. *Let X_1, \dots, X_n be independent geometric random variables with success probabilities p_1, \dots, p_n . Let $p_{\min} := \min\{p_i \mid i \in [1..n]\}$. Let $X := \sum_{i=1}^n X_i$ and $\mu = E[X] = \sum_{i=1}^n \frac{1}{p_i}$.*

(i) For all $\delta \geq 0$,

$$\Pr[X \geq (1 + \delta)\mu] \leq \frac{1}{1 + \delta} (1 - p_{\min})^{\mu(\delta - \ln(1 + \delta))} \quad (3)$$

$$\leq \exp(-p_{\min}\mu(\delta - \ln(1 + \delta))) \quad (4)$$

$$\leq \left(1 + \frac{\delta\mu p_{\min}}{n}\right)^n \exp(-\delta\mu p_{\min}) \quad (5)$$

$$\leq \exp\left(-\frac{(\delta\mu p_{\min})^2}{2n(1 + \frac{\delta\mu p_{\min}}{n})}\right). \quad (6)$$

(ii) For all $0 \leq \delta \leq 1$,

$$\Pr[X \leq (1 - \delta)\mu] \leq (1 - \delta)^{p_{\min}\mu} \exp(-\delta p_{\min}\mu) \quad (7)$$

$$\leq \exp\left(-\frac{\delta^2\mu p_{\min}}{2 - \frac{4}{3}\delta}\right) \quad (8)$$

$$\leq \exp(-\frac{1}{2}\delta^2\mu p_{\min}). \quad (9)$$

Estimates (3) and (4) are from [Jan17], bound (5) is from [Sch00], and (6) follows from the previous by standard estimates. The lower tail bound (7) is from [Jan17]. It implies (8) via standard estimates. Inequality (9) appeared in [Sch00].

It is surprising that none of these useful bounds appeared in a reviewed journal. For the case that all geometric random variables have the same success probability p , the bound

$$\Pr[X \geq (1 + \delta)\mu] \leq \exp\left(-\frac{\delta^2 n - 1}{2(1 + \delta)}\right) \quad (10)$$

appeared in [DHK11].

The bounds of Theorem 14 allow the geometric random variables to have different success probabilities, however, the tail probability depends only on the smallest of them. This is partially justified by the fact that the corresponding geometric random variable has the largest variance, and thus might be most detrimental to the desired strong concentration. If the success probabilities vary significantly, however, then this approach gives overly pessimistic tail bounds. Witt [Wit14] proves the following result, which can lead to stronger estimates.

Theorem 15. *Let X_1, \dots, X_n be independent geometric random variables with success probabilities p_1, \dots, p_n . Let $X = \sum_{i=1}^n X_i$, $s = \sum_{i=1}^n (\frac{1}{p_i})^2$, and $p_{\min} := \min\{p_i \mid i \in [1..n]\}$. Then for all $\lambda \geq 0$,*

$$\Pr[X \geq E[X] + \lambda] \leq \exp\left(-\frac{1}{4} \min\left\{\frac{\lambda^2}{s}, \lambda p_{\min}\right\}\right), \quad (11)$$

$$\Pr[X \leq E[X] - \lambda] \leq \exp\left(-\frac{\lambda^2}{2s}\right). \quad (12)$$

As we shall see, we often encounter sums of independent geometrically distributed random variables X_1, \dots, X_n with success probabilities p_i proportional to i . For this case, the following result from [DD15] gives stronger tail bounds than the previous result. Recall that the harmonic number H_n is defined by $H_n = \sum_{i=1}^n \frac{1}{i}$.

Theorem 16. *Let X_1, \dots, X_n be independent geometric random variables with success probabilities p_1, \dots, p_n . Assume that there is a number $C \leq 1$ such that $p_i \geq C \frac{i}{n}$ for all $i \in [1..n]$. Let $X = \sum_{i=1}^n X_i$. Then*

$$E[X] \leq \frac{1}{C} n H_n \leq \frac{1}{C} n (1 + \ln n), \quad (13)$$

$$\Pr[X \geq (1 + \delta) \frac{1}{C} n \ln n] \leq n^{-\delta} \text{ for all } \delta \geq 0. \quad (14)$$

A brief look at the last lines of the proof of this result in [DD15] shows that also the following, minimally stronger tail bound holds for all $\lambda \in \mathbb{N}$.

$$\Pr[X \geq \frac{1}{C} n \ln n + \lambda] \leq n \left(1 - \frac{C}{n}\right)^{\frac{1}{C} n \ln n + \lambda - 1} < \left(1 - \frac{C}{n}\right)^{\lambda - 1}$$

This immediately gives the domination result

$$X \preceq \frac{1}{C} n \ln n + \text{Geom}\left(\frac{C}{n}\right). \quad (15)$$

This domination result allows to show tail bounds for sums of independent random variables having a distribution as in Theorem 16. While we shall not need this in this work, the regular occurrence of such distributions in runtime analysis justifies showing the following result.

Theorem 17. *Let Y_1, \dots, Y_m be independent, not necessarily identically distributed, and let each one satisfy the assumptions made on X in Theorem 16. Let $Y = \sum_{i=1}^m Y_i$. Then $E[Y] \leq \frac{1}{C}nmH_n \leq \frac{1}{C}(\ln(n) + 1)nm$ and*

$$\Pr[Y \geq \frac{1}{C}(\ln(n) + 1)nm + \lambda] \leq \exp\left(-\frac{\lambda^2 C^2}{2n^2 m(1 + \frac{\lambda C}{nm})}\right).$$

Proof. Let $p = \frac{C}{n}$ and let Z be the sum of m independent $\text{Geom}(p)$ random variables. Since $Y_i \preceq \frac{1}{C}n \ln n + \text{Geom}(p)$ by (15), Lemma 11 shows $Y \preceq \frac{1}{C}mn \ln n + Z$. Hence for all $\lambda \geq 0$ we compute

$$\begin{aligned} \Pr[Y \geq \frac{1}{C}(\ln(n) + 1)nm + \lambda] &\leq \Pr[Z \geq E[Z] + \lambda] \\ &\leq \exp\left(-\frac{\lambda^2 p^2}{2m(1 + \frac{\lambda p}{m})}\right), \end{aligned}$$

where the last estimate follows from (6). □

We formulate explicitly the following corollary for sums of independent identical coupon collector distributions. In the basic coupon collector process, there are n types of coupons. In each round, we obtain a coupon of a random type. The question is how many rounds are needed to have at least one coupon of each type. As a minimal extension of the very basic version, let us assume that we start this process having already coupons of $n - k$ different types, that is, we are missing only k types. Let then D_n^k be the distribution describing the number of rounds taken until we have a coupon of each type. Clearly, $D_n^k = \sum_{i=1}^k \text{Geom}(\frac{i}{n})$. By the above theorem (taking k as n and C as $\frac{k}{n}$), we obtain the following result.

Corollary 18. *Let $m, n \in \mathbb{N}$, $k \in [1..n]$, and Y be the sum of m independent D_n^k coupon collector distributions. Then*

$$\begin{aligned} E[Y] &= mnH_k \leq mn(\ln(k) + 1), \\ \Pr[Y \geq mn(\ln(k) + 1) + \delta n] &\leq \exp\left(-\frac{\delta^2}{2m(1 + \frac{\delta}{m})}\right). \end{aligned}$$

4.3 Applications of the Fitness Level Theorem

We now present several examples where our fitness level theorem gives an upper bound in the domination sense and where this, often together with the just presented tail bounds, leads to new tail bounds for the runtime of different evolutionary algorithms.

4.3.1 General Upper Bound for Mutation-based Algorithms

Consider the optimization of an arbitrary $f: \{0, 1\}^n \rightarrow \mathbb{R}$ via any evolutionary algorithm \mathcal{A} which creates its offspring as random search points and via standard-bit mutation with mutation rate $p \leq \frac{1}{2}$ from previous search points. Then for any such offspring the probability that it is an optimal solution is at least p^n , regardless of the history of the optimization process. Consequently, the runtime T of \mathcal{A} on f satisfies

$$T \preceq \text{Geom}(p^n),$$

which implies

$$\begin{aligned} E[T] &\leq p^{-n}, \\ \Pr[T \geq \gamma p^{-n} + 1] &\leq (1 - p^n)^{\gamma p^{-n}} \leq e^{-\gamma} \text{ for all } \gamma \geq 0. \end{aligned}$$

Note that here the tail bound follows immediately from the definition of the geometric distribution. The result on the expectation was shown for the $(1 + 1)$ EA with mutation rate $p = \frac{1}{n}$ already in the seminal paper [DJW02] with essentially the same argument.

4.3.2 Performance of the $(1 + 1)$ EA and $(1 + \lambda)$ EA on OneMax

How the $(1 + 1)$ EA optimizes the ONEMAX test function (see Section 2.2 for its definition) is one of the first results in runtime analysis. Using the fitness level method with the levels $A_i := \{x \in \{0, 1\}^n \mid \text{ONEMAX}(x) = i\}$, $i = 0, 1, \dots, n$, one easily obtains that the expectation of the runtime T is at most

$$E[T] \leq enH_n \leq en(\ln(n) + 1). \quad (16)$$

For this, it suffices to compute that the probability to leave the i -th fitness level satisfies $p_i \geq \frac{n-i}{n}(1 - \frac{1}{n})^{n-1} \geq \frac{n-i}{en}$, see [DJW02] for the details. By Theorems 13 and 16, we also obtain

$$\begin{aligned} T &\preceq \sum_{i=1}^n \text{Geom}\left(\frac{i}{en}\right), \\ \Pr[T \geq (1 + \delta)en \ln n] &\leq n^{-\delta} \text{ for all } \delta \geq 0. \end{aligned}$$

The domination result, while not very deep, appears to be new, whereas the tail bound was proven before via multiplicative drift analysis [DG13].

Things become more interesting (and truly new) when regarding the $(1 + \lambda)$ EA instead of the $(1 + 1)$ EA. We denote by $d(x) = n - \text{ONEMAX}(x)$

the distance of x to the optimum. Let $t = \lfloor \frac{\ln(\lambda)-1}{2\ln\ln\lambda} \rfloor$. We partition the lowest $L = \lfloor n - \frac{n}{\ln\lambda} \rfloor$ fitness levels (that is, the search points with $\text{ONEMAX}(x) < L$) into $\lceil \frac{L}{t} \rceil$ sets each spanning at most t consecutive fitness levels. In [DK15], it was shown that the probability to leave such a set in one iteration is at least $p_0 \geq (1 - \frac{1}{e})$. The remaining search points are partitioned into sets of points having equal fitness, that is, $A_i = \{x \mid \text{ONEMAX}(x) = i\}$, $i = L, \dots, n$. For these levels, the probability to leave a set in one iteration is at least

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

For $i \leq n - \frac{en}{\lambda}$, this is at least $p_i \geq 1 - \frac{1}{e}$ by the well-known estimate $1 + r \leq e^r$ valid for all $r \in \mathbb{R}$. For $i > n - \frac{en}{\lambda}$, we use a new version of the Weierstrass product inequality (Lemma 4.8 in [Doe18b], stating that $\prod_{i=1}^n (1 - x_i) \leq 1 - S + \frac{1}{2}S^2$ for all $x_1, \dots, x_n \in [0, 1]$ and $S = \sum_{i=1}^n x_i$) to estimate $p_i \geq 1 - (1 - \frac{\lambda}{en}(n-i) + \frac{1}{2}(\frac{\lambda}{en}(n-i))^2) \geq \frac{1}{2} \frac{\lambda}{en}(n-i)$. Using the abbreviation $T_0 := \lceil \frac{L}{t} \rceil + \lceil \frac{n}{\ln\lambda} \rceil - \lceil \frac{en}{\lambda} \rceil + 1$, our fitness level theorem gives the following domination bound for the number T of *iterations* until the optimum is found.

$$T \preceq \sum_{i=1}^{T_0} \text{Geom}(1 - \frac{1}{e}) + \sum_{i=1}^{\lceil \frac{en}{\lambda} - 1 \rceil} \text{Geom}(\frac{1}{2} \frac{\lambda i}{en}). \quad (17)$$

To avoid uninteresting case distinctions, let us assume that $\lambda = \omega(1)$ and $\lambda = n^{O(1)}$. In this case, the above domination statement immediately gives

$$E[T] \leq \frac{e}{e-1} T_0 + 2e \frac{n \ln(\lceil \frac{en}{\lambda} \rceil)}{\lambda} = (1 + o(1)) \left(\frac{2e}{e-1} \frac{n \ln \ln \lambda}{\ln \lambda} + 2e \frac{n \ln(n)}{\lambda} \right),$$

which is the result of [DK15] with explicit constants. With Theorems 14 and 16, equation (17) also implies tail bounds, however, for general λ these become quite technical. For this reason, we omit any further details.

4.3.3 Performance of the $(\mu + 1)$ EA on LeadingOnes

The runtime of the $(\mu + 1)$ EA on various problems was analyzed in [Wit06]. Among them, we exemplarily regard the LEADINGONES problem (see Section 2.3 for its definition). We shall not need the assumption of a random initialization in this section.

Following the general proof idea of [Wit06], we let

$$M := \min\{\lceil n / \ln(en) \rceil, \mu\}.$$

For all $a \in [0..n]$ and $b \in [1..M]$, we say that the $(\mu + 1)$ EA optimizing LEADINGONES is in state (a, b) if the maximal fitness of an individual in the population is a and if there are exactly b individuals having this maximal fitness. We say that the algorithm is in state (a, M) also when the maximal fitness in the population is a and there are more than M individuals with this fitness. We sort these states in lexicographic order, that is, we say that (a', b') is better than (a, b) if $a' > a$ or if $a' = a$ and $b' \geq b$. From the definition of the $(\mu + 1)$ EA, it is clear that the algorithm cannot go from a better state to a worse one. Consequently, by viewing the states as fitness levels, we can apply the fitness level theorem (Theorem 13). We observe that for all $a < n$ and $b < M$, the probability of leaving the state (a, b) to a better state is at least

$$p_{ab} := \frac{b}{\mu} \left(1 - \frac{1}{n}\right)^a \geq \frac{b}{e\mu},$$

since to leave the state it suffices to select one of the b best individuals and to not flip any of the a bits contributing to the fitness (this creates a new best individual which replaces one of the inferior individuals which are still in the population). The probability to leave the state (a, M) with $a < n$ is at least

$$p_{aM} := \frac{M}{\mu} \frac{1}{n} \left(1 - \frac{1}{n}\right)^a \geq \frac{M}{en\mu}.$$

Consequently, the runtime T of the $(\mu + 1)$ EA on the LEADINGONES function satisfies

$$\begin{aligned} T &\preceq \sum_{a=0}^{n-1} \sum_{b=1}^M \text{Geom}(p_{ab}), \\ E[T] &\leq \sum_{a=0}^{n-1} \sum_{b=1}^M \frac{1}{p_{ab}} = en\mu(\ln(M-1) + 1 + \frac{n}{M}) \\ &\leq en\mu(\ln(en) + \max\{\frac{n}{\mu}, \ln(en)\}) \leq en(2\mu \ln(en) + n), \\ \Pr[T \geq E[T] + \lambda] &\leq \exp\left(-\frac{1}{4} \min\left\{\frac{\lambda^2}{e^2 n^3 \mu^2 / M^2 + \frac{\pi^2}{6} e^2 n \mu^2}, \frac{\lambda M}{en\mu}\right\}\right) \\ &= \exp\left(-\frac{1}{4} \min\left\{\frac{(1 - o(1))\lambda^2 M^2}{e^2 n^3 \mu^2}, \frac{\lambda M}{en\mu}\right\}\right). \end{aligned}$$

The tail bound follows from using Theorem 15 with

$$s = n \left(\left(\frac{en\mu}{M}\right)^2 + \sum_{b=1}^{M-1} \left(\frac{e\mu}{b}\right)^2 \right) \leq \frac{e^2 n^3 \mu^2}{M^2} + \frac{\pi^2}{6} e^2 n \mu^2 = (1 + o(1)) \frac{e^2 n^3 \mu^2}{M^2}$$

and $p_{\min} = \min\{\frac{M}{en\mu}, \frac{1}{e\mu}\} = \frac{M}{en\mu}$, noting that $M = o(n)$ and $M \leq n$.

Apart from a small improvement in the constants, stemming from slightly more careful estimates, the result on the expected runtime is the same as the one in [Wit06], which is $E[T] \leq 3en \max\{\mu \ln(en), n\}$, see Theorem 1 of [Wit06] and recall that we count the number of iterations, that is, we ignore the μ fitness evaluations of the initial individuals. The tail bound, as discussed earlier, is stronger than the one in [ZLLH12] due to the stronger Chernoff bounds for sums of geometric random variables which are available now.

4.3.4 Performance of the (1 + 1) EA on Jump Functions

For $k \in [1..n]$, the n -dimensional jump function $\text{JUMP}_k: \{0, 1\}^n \rightarrow \mathbb{R}$ was defined in [DJW02] by

$$\text{JUMP}_k(x) = \begin{cases} \text{ONEMAX}(x) + k & \text{if } \text{ONEMAX}(x) \in [0..n - k] \cup \{n\}, \\ n - \text{ONEMAX}(x) & \text{otherwise.} \end{cases}$$

Exemplarily for the runtime T_k of the (1 + 1) EA on JUMP_k , the fitness level theorem with $P_k = n^{-k}(1 - \frac{1}{n})^{n-k}$ immediately yields

$$T \preceq \sum_{i=1}^{k-1} \text{Geom}(\frac{n-i}{en}) + \sum_{i=0}^{n-k-1} \text{Geom}(\frac{n-i}{en}) + \text{Geom}(P_k),$$

where for $k \geq 2$ the last term is by far the most important one (so we do not care about the minor improvements of this bound, which are easy to obtain). Interestingly, we have a nearly matching lower bound *in the domination sense*. To avoid unnecessary technicalities, we restrict ourselves to the case $k \in [2.. \frac{n}{4}]$, but similar results could be shown for larger values of k as well.

Lemma 19. *Let $n \in \mathbb{N}$ and $k \in [2.. \frac{n}{4}]$. Let T_k be the runtime of the (1 + 1) EA on the function JUMP_k . Then with $P_k = n^{-k}(1 - \frac{1}{n})^{n-k}$ we have*

$$X \cdot \text{Geom}(P_k) \preceq T_k,$$

where X is a binary random variable with $\Pr[X = 1] = 1 - \exp(-\frac{n}{8})$.

Proof. A simple application of the additive Chernoff bound (e.g., Theorem 10.7 in [Doe18b]) shows that with probability at least $1 - \exp(-\frac{n}{8})$, the random initial search point has a ONEMAX value of at most $\frac{3}{4}n \leq n - k$. In this situation, the current search point of the (1 + 1) EA will have a ONEMAX value of at most $n - k$ as long as the optimum is not found. For

any search point with ONEMAX value $v \leq n - k$, the probability to reach the optimum in one step is $n^{-(n-v)}(1 - \frac{1}{n})^v \leq P_k$. Consequently, conditional on the initial search point having a ONEMAX value of at most $n - k$, the runtime satisfies $\text{Geom}(P_k) \preceq T$, giving the claim. \square

We note without proof that the results of this subsection can easily be extended to other elitist mutation-based algorithms, since the main argument that with high probability once the optimum has to be generated from a search point in Hamming distance at least k remains valid.

4.3.5 Performance of the $(1 + 1)$ EA on the Sorting Problem

One of the first combinatorial problems regarded in the theory of evolutionary computation is how a combinatorial $(1 + 1)$ EA sorts an array of length n . One of several setups regarded in [STW04] is modelling the sorting problem as the minimization of the number of inversions in the array. We assume that the $(1 + 1)$ EA mutates an array by first determining a number k according to a Poisson distribution with parameter $\lambda = 1$ and then performing $k + 1$ random exchanges (to ease the presentation, we do not use the jump operations also employed in [STW04], but it is easy to see that this does not significantly change things). It is easy to see that exchanging two elements that are in inverse order reduces the number of inversions by at least one. Hence if there are i inversions, then with probability $\frac{1}{e}i\binom{n}{2}^{-1}$, the $(1 + 1)$ EA inverts exactly one of the inversions and thus improves the fitness (note that $\Pr[k = 0] = \frac{1}{e}$). By our fitness level theorem, the runtime T is dominated by the independent sum $\sum_{i=1}^{\binom{n}{2}} \text{Geom}(\frac{i}{e\binom{n}{2}})$. Hence

$$E[T] \leq e \binom{n}{2} H_{\binom{n}{2}} \leq \frac{e}{2} n^2 (1 + 2 \ln n),$$

$$\Pr[T \geq (1 + \delta) e n^2 \ln n] \leq \binom{n}{2}^{-\delta},$$

where again the statement on the expectation has appeared before in [STW04] (with slightly different constants stemming from the use of a slightly different algorithm) and the tail bound (stemming from Theorem 14) is new.

We note without proof that a similar analysis would also prove a tail bound for the slightly faster $O(n^2)$ time evolutionary sorting algorithm of [DH08].

4.3.6 Further Results

In this section, we state a few more results to demonstrate the range of applicability of domination arguments. Since for each of them a detailed discussion with precise proofs would require a longer explanation of the problem and the precise evolutionary algorithm used, we just state the result and point the reader interested in the details to the literature. In all cases, the domination result follows directly from the proof of the original result on the expected runtime.

Eulerian cycles. The Eulerian cycle problem asks for a cycle in an undirected graph that traverses each edge exactly once. Ending a series [Neu08, DHN07, DKS07] of works discussing suitable representations for the Eulerian cycle problem, matchings in the adjacency lists were found to give very good results [DJ07]. Among several algorithms proposed in [DJ07], we regard the $(1 + 1)$ EA that uses perfect matchings in the adjacency lists as genotype and an edge-based mutation operator. Using a fitness level argument that regards the number of disjoint cycles in the individual, one observes that the runtime T of this algorithm is dominated by the independent sum

$$T \preceq \sum_{i=1}^{m/3} \text{Geom}\left(\frac{i}{2em}\right).$$

By Theorem 16, this implies

$$\begin{aligned} E[T] &\leq 2emH_{m/3}, \\ \Pr[T \geq 2(1 + \delta)em \ln \frac{m}{3}] &\leq \left(\frac{m}{3}\right)^{-\delta} \text{ for all } \delta \geq 0. \end{aligned}$$

Vertex covers. In an undirected graph, a vertex cover is a set of vertices such that each edge contains at least one of them. In [OHY09], it was shown that the classic $(1 + 1)$ EA using a binary representation can efficiently find minimal vertex covers on paths. Since this was shown via a fitness level argument, we easily obtain that the runtime T on a path of length n is dominated by

$$T \preceq \sum_{i=1}^{n-1} \text{Geom}\left(\Omega\left(\left(\frac{i}{n}\right)^4\right)\right) =: \bar{T},$$

which gives an expectation of

$$E[T] \leq E[\bar{T}] = O(n^4)$$

and, via Theorem 15, a tail estimate of

$$\Pr[T \geq E[\bar{T}] + \lambda] \leq \exp\left(-\Omega\left(\frac{\lambda}{n^4}\right)\right).$$

The $(1 + (\lambda, \lambda))$ genetic algorithm. The $(1 + (\lambda, \lambda))$ GA was proposed in [DDE15] as an algorithm that, despite elitism, also profits from generating search points which are inferior to the current-best solution. All analyses of this algorithm suggest to use it with a mutation rate of $p = \frac{\lambda}{n}$ and a crossover bias $c = \frac{1}{\lambda}$, where λ denotes the size of the two offspring populations used by the algorithm. We shall assume these parameters as well.

Assume that $\lambda = \omega(1)$ as this is the situation in which the $(1 + (\lambda, \lambda))$ GA can outperform the $(1 + 1)$ EA. Let $L = \lceil n \ln \ln(\lambda) / \ln(\lambda) \rceil$. Let C_1, C_2 be sufficiently large constants. Imitating the proof of the asymptotically tight runtime analysis in [DD18], we see that the runtime of the $(1 + (\lambda, \lambda))$ GA on the ONEMAX function, measured via the number of iterations, is dominated by $\sum_{i=1}^{C_1 L} \text{Geom}(p_i)$ with $p_i = C_2(1 - (1 - \frac{i}{n})^{\lambda^2/2})$ for $i \in [1..L]$ and $p_i = 1 - 1/e - o(1)$ for $i \in [L + 1..C_2 L]$.

5 Beyond the Fitness Level Theorem

Above we showed that in all situations where the classic fitness level method can be applied we immediately obtain that the runtime is dominated by a suitable sum of independent geometric distributions. We now show that the domination-by-distribution argument is not restricted to such situations. As an example, we use another combinatorial problem from [STW04], the single-source shortest path problem, and obtain a simplified and more natural proof of the currently strongest runtime bound for this problem from [DHK11]. We note without proof that similar arguments could be used in the analysis of other problems where the evolutionary algorithm builds up the optimal solution incrementally from structurally smaller solutions, such as other path problems [The09, DJ10, DHK12] or dynamic programming [DEN⁺11].

In [STW04], the single-source shortest path problem in a connected undirected graph $G = (V, E)$ with edge weights $w: E \rightarrow \mathbb{N}$ and source vertex $s \in V$ was solved via a $(1 + 1)$ EA as follows. Individuals are arrays of pointers such that each vertex different from the source has a pointer to another vertex. If, for a vertex v , following the pointers gives a path from v to s , then the length (=sum of weights of its edges) of this path is the fitness of this vertex; otherwise the fitness of this vertex is infinite. The fitness of an individual is the vector of the fitnesses of all vertices different from the source. In the selection step, an offspring is accepted if and only if all vertices have an at least as good fitness as in the parent. This is called a multi-objective formulation of the single-source shortest path problem in [STW04]. Mutating an individual means choosing a number k from a Poisson distribution with parameter $\lambda = 1$ and then changing $k + 1$ pointers to random new targets.

The main analysis argument in [STW04] is that, due to the use of the multi-objective fitness, a vertex that is connected to the source via a shortest path (let us call such a vertex *optimized* in the following) remains optimized for the remaining run of the algorithm. Hence we can perform a structural fitness level argument over the number of optimized vertices. The probability to increase this number is at least $p := \frac{1}{e(n-1)(n-2)}$ because there is at least one non-optimized vertex v for which the next vertex u on a shortest path from v to s is already optimized. Hence with probability $\frac{1}{e}$ we have $k = 0$, with probability $\frac{1}{n-1}$ we choose v , and with probability $\frac{1}{n-2}$ we rewire its pointer to u . This gives an expected runtime of at most

$$E[T] \leq (n-1)/p = e(n-1)^2(n-2).$$

To obtain better bounds for certain graph classes, the number n_i of vertices for which a shortest path with fewest edges consists of i edges is defined in [STW04]. With a fitness level argument similar to the one above, it takes an expected time of at most $en^2H_{n_1}$ to have all n_1 -type vertices optimally connected to the source. After this, an expected number of at most $en^2H_{n_2}$ iteration suffices to connect all n_2 -vertices to the source. Iterating this argument, a runtime estimate of

$$E[T] \leq en^2 \sum_{i=1}^{n-1} H_{n_i} \leq en^2 \sum_{i=1}^{n-1} (\ln(n_i) + 1)$$

is obtained in [STW04]. This expression remains of order $\Theta(n^3)$ in the worst case, but becomes, e.g., $O(n^2\ell \log(\frac{2n}{\ell}))$ when each vertex can be connected to the source via a shortest path having at most ℓ edges.

By arguments different from fitness levels and with some technical effort, this result was improved to $O(n^2 \max\{\log n, \ell\})$ in [DHK11]. We now show that this improvement could have been obtained via domination arguments in a very natural way.

Consider some vertex v different from the source. Fix some shortest path P from v to s having at most ℓ edges. Let V_v be the set of (at most ℓ) vertices on P different from s . As before, in each iteration we have a probability of at least $p = \frac{1}{e(n-1)(n-2)}$ that a non-optimized vertex of V_v becomes optimized. Consequently, the time T_v to connect v to s via a shortest path is dominated by a sum of ℓ independent $\text{Geom}(p)$ random variables. We conclude that there are random variables X_{ij} , $i \in [1..\ell]$, $j \in [1..n-1]$, such that

- (i) $X_{ij} \sim \text{Geom}(p)$ for all $i \in [1..\ell]$ and $j \in [1..n-1]$,
- (ii) for all $j \in [1..n-1]$ the variables $X_{1j}, \dots, X_{\ell j}$ are independent, and

(iii) the runtime T is dominated by $\max\{Y_j \mid j \in [1..n-1]\}$, where $Y_j := \sum_{i=1}^{\ell} X_{ij}$ for all $j \in [1..n-1]$.

It remains to deduce from this domination statement a runtime bound. Let $\delta = \max\{\frac{4\ln(n-1)}{\ell-1}, \sqrt{\frac{4\ln(n-1)}{\ell-1}}\}$. For all $j \in [1..n-1]$, by (10), we estimate

$$\begin{aligned} \Pr[Y_j \geq (1+\delta)E[Y_j]] &\leq \exp\left(-\frac{1}{2} \frac{\delta^2}{1+\delta}(\ell-1)\right) \\ &\leq \exp\left(-\frac{1}{4} \min\{\delta^2, \delta\}(\ell-1)\right) \\ &\leq \exp\left(-\frac{1}{4} \frac{4\ln(n-1)}{\ell-1}(\ell-1)\right) \\ &= \frac{1}{n-1}. \end{aligned}$$

For all $\varepsilon > 0$, again by (10), we compute

$$\begin{aligned} \Pr[Y_j \geq (1+\varepsilon)(1+\delta)E[Y_j]] &\leq \exp\left(-\frac{1}{2} \frac{(\delta+\varepsilon+\delta\varepsilon)^2}{(1+\delta)(1+\varepsilon)}(\ell-1)\right) \\ &\leq \exp\left(-\frac{1}{2} \frac{\delta^2(1+\varepsilon)^2}{(1+\delta)(1+\varepsilon)}(\ell-1)\right) \\ &\leq \exp\left(-\frac{1}{2} \frac{\delta^2}{1+\delta}(\ell-1)\right)^{1+\varepsilon} \\ &\leq (n-1)^{-(1+\varepsilon)}. \end{aligned}$$

Recall that the runtime T is dominated by $Y = \max\{Y_j \mid j \in [1..n-1]\}$, where we did not make any assumption on the correlation of the Y_j . In particular, they do not need to be independent. Let

$$T_0 = (1+\delta)E[Y_1] = (1+\delta)\frac{\ell}{p}.$$

Then

$$\Pr[Y \geq (1+\varepsilon)T_0] \leq \sum_{j=1}^{n-1} \Pr[Y_j \geq (1+\varepsilon)T_0] \leq (n-1)^{-\varepsilon}$$

by the union bound. Transforming this tail bound into an expectation via standard arguments, e.g., Corollary 6.2 (c) in [Doe18b], we obtain $E[Y] \leq (1 + \frac{1}{\ln(n-1)})T_0$. We thus have shown the following result.

Theorem 20. *Let G be an undirected graph on n vertices together with integral edge weights. Let s be a vertex. Assume that all vertices are connected*

to s via a shortest path consisting of at most ℓ edges. Let T be the runtime of the EA proposed in [STW04] to solve the single-source shortest path problem in G . Then, with $\delta = \max\{\frac{4\ln(n-1)}{\ell-1}, \sqrt{\frac{4\ln(n-1)}{\ell-1}}\}$, $p = \frac{1}{\varepsilon(n-1)(n-2)}$, and $T_0 := (1 + \delta)\frac{\ell}{p}$, we have

$$E[T] \leq \left(1 + \frac{1}{\ln(n-1)}\right) T_0,$$

$$\Pr[T \geq (1 + \varepsilon)T_0] \leq (n-1)^{-\varepsilon}$$

for all $\varepsilon \geq 0$.

This result is of same asymptotic order as the bound in [DHK11]. We see the main advantage of the proof above in the fact that it is less technical and more natural than the previous one. However, our analysis also gives a better leading constant. For $\ell \gg \log(n)$, for example, our bound on the expected runtime is $(1 + o(1))eln^2$, whereas it is $8(1 + o(1))eln^2$ in [DHK11].

6 Structural Domination

So far we have used stochastic domination to compare runtime distributions. We now show that stochastic domination can be a very useful tool also to express structural properties of the optimization process. As an example, we give a short and elegant proof for the result of Witt [Wit13] that compares the runtimes of mutation-based algorithms. The main reason why our proof is significantly shorter than the one of Witt is that we use the notion of stochastic domination also for the distance from the optimum.

To state this result, we need the notion of a (μ, p) *mutation-based algorithm* introduced in [Sud13]. This class of algorithms is called only *mutation-based* in [Sud13], but since (i) it does not include all adaptive algorithms using mutation only, e.g., those regarded in [JW06, OLN09, BDN10, BLS14, DL16, DGWY17, DWY18], (ii) it does not include all algorithms using a different mutation operator than standard-bit mutation, e.g., those in [DDY16a, DDY16b, LOW17, DLMN17], and (iii) this notion collides with the notion of unary unbiased black-box complexity algorithms (see [LW12]), which with some justification could also be called the class of mutation-based algorithms, we feel that a notion making these restrictions precise is more appropriate.

The class of (μ, p) mutation-based algorithms comprises all algorithms which first generate a set of μ search points uniformly and independently at random from $\{0, 1\}^n$ and then repeat generating new search points from any

of the previous ones via standard-bit mutation with probability p (that is, by flipping bits independently with probability p). This class includes all $(\mu + \lambda)$ and (μ, λ) EAs which only use standard-bit mutation with static mutation rate p .

Denote by $(1 + 1) \text{EA}_{\mu,p}$ the following algorithm in this class. It first generates μ random search points. From these, it selects uniformly at random one with highest fitness and then continues from this search point as a $(1 + 1)$ EA, that is, repeatedly generates a new search point from the current one via standard-bit mutation with rate p and replaces the previous one by the new one if the new one is not worse (in terms of the fitness). This algorithm was called $(1 + 1)$ EA with BestOf(μ) initialization in [dPdLDD15].

For any algorithm \mathcal{A} from the class of (μ, p) mutation-based algorithms and any fitness function $f: \{0, 1\}^n \rightarrow \mathbb{R}$, let us denote by $T(\mathcal{A}, f)$ the runtime of the algorithm \mathcal{A} on the fitness function f , that is, the number of the first individual that is an optimal solution. Usually, this will be μ plus the number of the iteration in which the optimum was generated. To cover also the case that one of the random initial individuals is optimal, let us assume that these initial individuals are generated sequentially.

In this language, Witt [Wit13] shows the following remarkable result.

Theorem 21. *For any (μ, p) mutation-based algorithm \mathcal{A} and any $f: \{0, 1\}^n \rightarrow \mathbb{R}$ with unique global optimum,*

$$T((1 + 1) \text{EA}_{\mu,p}, \text{ONEMAX}) \preceq T(\mathcal{A}, f).$$

This result significantly extends results of a similar flavor in [BE08, DJW12, Sud13]. The importance of such types of results is that they allow to prove lower bounds for the performance of many algorithm on essentially arbitrary fitness functions by just regarding the performance of the $(1 + 1) \text{EA}_{\mu,p}$ on ONEMAX.

Let us denote by $|x|_1$ the number of ones in the bit string $x \in \{0, 1\}^n$. Using similar arguments as in [DJW00, Section 5] and [DJW12, Lemma 13], Witt [Wit13] shows the following natural domination relation between offspring generated via standard-bit mutation.

Lemma 22. *Let $x, y \in \{0, 1\}^n$. Let $p \in [0, \frac{1}{2}]$. Let x', y' be obtained from x, y via standard-bit mutation with rate p . If $|x|_1 \leq |y|_1$, then $|x'|_1 \preceq |y'|_1$.*

We are now ready to give our alternate proof for Theorem 21. While it is clearly shorter than the original one in [Wit13], we also feel that it is more natural. In very simple words, it shows that $T(\mathcal{A}, f)$ dominates $T((1 + 1) \text{EA}_{\mu,p}, \text{ONEMAX})$ because the search points generated in the run

of the $(1 + 1)$ EA $_{\mu,p}$ on ONEMAX always are at least as close to the optimum (in the domination sense) as in the run of \mathcal{A} on f , and this follows from the previous lemma, a suitable coupling, and induction.

Proof of Theorem 21. As a first small technicality, let us assume that the $(1 + 1)$ EA $_{\mu,p}$ in iteration $\mu + 1$ does not choose a random optimal search point, but the last optimal search point. Since all the first μ individuals are generated independently, this modification does not change anything.

Since \mathcal{A} treats bit-positions and bit-values in a symmetric fashion, we may without loss of generality assume that the unique optimum of f is $(1, \dots, 1)$.

Let $x^{(1)}, x^{(2)}, \dots$ be the sequence of search points generated in a run of \mathcal{A} on the fitness function f . Hence $x^{(1)}, \dots, x^{(\mu)}$ are independently and uniformly distributed in $\{0, 1\}^n$ and all subsequent search points are generated from suitably chosen previous ones via standard-bit mutation with rate p . Let $y^{(1)}, y^{(2)}, \dots$ be the sequence of search points generated in a run of the $(1 + 1)$ EA $_{\mu,p}$ on the fitness function ONEMAX.

We show how to couple these random sequences of search points in a way that $|\tilde{x}^{(t)}|_1 \leq |\tilde{y}^{(t)}|_1$ for all $t \in \mathbb{N}$. We take as common probability space Ω simply the space that $(x^{(t)})_{t \in \mathbb{N}}$ is defined on and let $\tilde{x}^{(t)} = x^{(t)}$ for all $t \in \mathbb{N}$.

We define the $\tilde{y}^{(t)}$ inductively as follows. For $t \in [1.. \mu]$, let $\tilde{y}^{(t)} = x^{(t)}$. Note that this trivially implies $|\tilde{x}^{(t)}|_1 \leq |\tilde{y}^{(t)}|_1$ for these search points. Let $t > \mu$ and assume that $|\tilde{x}^{(t')}|_1 \leq |\tilde{y}^{(t')}|_1$ for all $t' < t$. Let $s \in [1..t - 1]$ be maximal such that $|\tilde{y}^{(s)}|_1$ is maximal among $|\tilde{y}^{(1)}|_1, \dots, |\tilde{y}^{(t-1)}|_1$. Let $r \in [1..t - 1]$ be such that $x^{(t)}$ was generated from $x^{(r)}$ in the run of \mathcal{A} on f . By induction, we have $|x^{(r)}|_1 \leq |\tilde{y}^{(r)}|_1$. By choice of s we have $|\tilde{y}^{(r)}|_1 \leq |\tilde{y}^{(s)}|_1$. Consequently, we have $|x^{(r)}|_1 \leq |\tilde{y}^{(s)}|_1$. By Lemma 22 and Theorem 12, there is a random $\tilde{y}^{(t)}$ (defined on Ω) such that $\tilde{y}^{(t)}$ has the distribution of being obtained from $\tilde{y}^{(s)}$ via standard-bit mutation with rate p and such that $|x^{(t)}|_1 \leq |\tilde{y}^{(t)}|_1$.

With this construction, the sequence $(\tilde{y}^{(t)})_{t \in \mathbb{N}}$ has the same distribution as $(y^{(t)})_{t \in \mathbb{N}}$. This is because the first μ elements are random and then each subsequent one is generated via standard-bit mutation from the current-best one, which is just the way the $(1 + 1)$ EA $_{\mu,p}$ is defined. At the same time, we have $|\tilde{x}^{(t)}|_1 \leq |\tilde{y}^{(t)}|_1$ for all $t \in \mathbb{N}$. Consequently, we have $\min\{t \in \mathbb{N} \mid |\tilde{y}^{(t)}|_1 = n\} \leq \min\{t \in \mathbb{N} \mid |x^{(t)}|_1 = n\}$. Since $T((1 + 1) \text{ EA}_{\mu,p}, \text{ ONEMAX})$ and $\min\{t \in \mathbb{N} \mid |\tilde{y}^{(t)}|_1 = n\}$ are identically distributed and also $T(\mathcal{A}, f)$ and $\min\{t \in \mathbb{N} \mid |x^{(t)}|_1 = n\}$ are identically distributed, we have $T((1 + 1) \text{ EA}_{\mu,p}, \text{ ONEMAX}) \preceq T(\mathcal{A}, f)$. \square

7 Counter-Examples

Stochastic domination is a strong property. Therefore, it is easy to encounter situations in which the intuitive feeling that one process is faster than a second one does not imply stochastic domination. Here are a few examples.

Random search vs. $(1 + 1)$ EA on OneMax. Let us compare the performance of the random search heuristic (cf. Section 2.1) and the $(1 + 1)$ EA on the ONEMAX benchmark function. Both intuitive considerations and the known results on the expected runtimes (2^n for random search, see Lemma 1, and at most $en(\ln(n) + 1)$ for the $(1 + 1)$ EA, see (16)) suggest that the $(1 + 1)$ EA is much more efficient than random search. Nevertheless, the runtime T_{RS} of random search does not dominate the runtime T_{EA} of the $(1 + 1)$ EA, as we show now.

Note that to have a fair comparison between these two algorithms, we here regard as runtime truly the number of individuals evaluated until the optimum is first evaluated, that is, we do count the evaluation of the initial individual for the $(1 + 1)$ EA. For either of the two algorithms, let X_1 and X_2 be the first two individuals generated (where we assume that the algorithm does not stop when X_1 is already the optimum). Denoting the optimal solution by x^* , the runtime T of this algorithm satisfies

$$\Pr[T \leq 2] = \Pr[X_1 = x^*] + \Pr[X_2 = x^*] - \Pr[X_1 = X_2 = x^*].$$

Note that X_2 also for the $(1 + 1)$ EA is uniformly distributed in $\{0, 1\}^n$. Hence

$$\begin{aligned} \Pr[T_{RS} \leq 2] &= 2^{-n+1} - 2^{-2n} = (1 + o(1))2^{-n+1}, \\ \Pr[T_{EA} \leq 2] &= 2^{-n+1} - (1 - \frac{1}{n})^n 2^{-n} = (1 + o(1))(1 - \frac{1}{2e})2^{-n+1}, \end{aligned}$$

contradicting the assertion that $T_{EA} \preceq T_{RS}$.

Fitness-proportionate selection. When optimizing objective functions with a strong fitness-distance correlation via a reasonable evolutionary algorithm, then it seems plausible that replacing individuals by better ones can only improve the performance. This is not true, for example, when fitness-proportionate selection is used, as the following example shows.

Consider the optimization of the ONEMAX function, which has the perfect fitness-distance correlation. Assume we use an algorithm which at some time selects an individual x from a population $P = \{x_1, \dots, x_\mu\}$, $\mu \geq 2$, via fitness-proportionate selection, that is, we have

$$\Pr[x = x_i] = \frac{f(x_i)}{\sum_{j=1}^{\mu} f(x_j)},$$

where $f = \text{ONEMAX}$. Now let x'_1, \dots, x'_μ be search points with $f(x'_i) \geq f(x_i)$ for all $i \in [1..\mu]$. Let $P' = \{x'_1, \dots, x'_\mu\}$ and assume that we select x' from P' via fitness-proportionate selection. Then we do not necessarily have $f(x) \preceq f(x')$.

For an extreme counter-example, let n be a multiple of 10 and let the search points x_i be such that $f(x_1) = 0.8n$ and $f(x_i) = 0$ for $i \in [2..\mu]$. Then

$$\Pr[f(x) \geq 0.8n] = \Pr[x = x_1] = 1.$$

Assume that $f(x'_i) = f(x_i) + 0.1n$ for all $i \in [1..\mu]$. Then

$$\Pr[f(x') \geq 0.8n] = \Pr[x' = x'_1] = \frac{0.8n + 0.1n}{0.8n + \mu \cdot 0.1n} = \frac{9}{\mu + 8} < 1.$$

Dependencies. So far, we mostly experienced that stochastic domination is very well-behaved. If $X \preceq Y$, then all reasonable upper bound statements on Y hold as well for X . Also, stochastic domination allowed us to compare two random variables irrespective of possible dependencies (in fact, we do not even need that they are defined over a common probability space). We now want to outline a point where some caution is required, and this is that dependencies between other random variables cannot be ignored. For example, $X_1 \preceq Y_1$ and $X_2 \preceq Y_2$ do not necessarily imply that

$$\max\{X_1, X_2\} \preceq \max\{Y_1, Y_2\}.$$

Whether such a statement is true depends on the correlation between the X_i and the correlation between the Y_i . For this reason, we needed to construct the coupling (which determines the dependencies) of the two processes in the proof of Theorem 21.

As a simple counter-example, we can again regard the processes of running random search and the $(1+1)$ EA on the ONEMAX function. Let X_0, X_1, \dots denote the search points generated by random search and let Y_0, Y_1, \dots denote the search points generated by the $(1+1)$ EA. Clearly, each X_i is uniformly distributed in $\{0, 1\}$, hence $f(X_i) \sim \text{Bin}(n, \frac{1}{2})$. For the ONEMAX process, we have $f(Y_0) \sim \text{Bin}(n, \frac{1}{2})$. For all $i \geq 0$, we also have $f(Y_0) \preceq f(Y_i)$, and hence $f(X_i) \preceq f(Y_i)$. Consequently, we have $f(X_i) \preceq f(Y_i)$ for all $i \geq 0$. However, as shown above, we do not have $\max\{f(X_0), f(X_1)\} \preceq \max\{f(Y_0), f(Y_1)\}$.

8 Conclusion

In this work, we argued that stochastic domination can be very useful in runtime analysis, both to formulate more informative results and to obtain simpler and more natural proofs. We also showed that in many situations, in particular, whenever the fitness level method is applicable, it is easily possible to describe the runtime via a domination statement.

We note however that not all classic proofs easily reveal details on the distribution. For results obtained via random walk arguments, e.g., the optimization of the short path function SPC_n [JW01], monotone polynomials [WW05], or vertex covers on paths-like graphs [OHY09], as well as for results proven via additive drift [HY01], the proofs often give little information about the runtime distribution. Note however that in [GKS99] results on the convergence of the (suitably scaled) runtime distribution could be obtained.

For results obtained via the average weight decrease method [NW07] or multiplicative drift analysis [DG13], the proofs also do not give information on the runtime distribution. However, the probabilistic runtime bound of type $\Pr[T \geq T_0 + \lambda] \leq (1 - \delta)^\lambda$ obtained from these methods implies that the runtime is dominated by $T \preceq T_0 - 1 + \text{Geom}(1 - \delta)$.

Overall, both from regarding these results and the history of the field, we suggest to more frequently formulate results via domination statements. Even in those cases where the probabilistic tools at the moment are not ready to exploit such a statement, there is a good chance future developments overcome this shortage and then it pays off if the result is readily available in a distribution form and not just as an expectation.

References

- [ADFH18] Denis Antipov, Benjamin Doerr, Jiefeng Fang, and Tangi Hetet. Runtime analysis for the $(\mu + \lambda)$ EA optimizing OneMax. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1459–1466. ACM, 2018.
- [BDN10] Süntje Böttcher, Benjamin Doerr, and Frank Neumann. Optimal fixed and adaptive mutation rates for the LeadingOnes problem. In *Parallel Problem Solving from Nature, PPSN 2010*, pages 1–10. Springer, 2010.

- [BE08] Pavel A. Borisovsky and Anton V. Eremeev. Comparing evolutionary algorithms to the (1+1)-EA. *Theoretical Computer Science*, 403:33–41, 2008.
- [BF13] Karl Bringmann and Tobias Friedrich. Parameterized average-case complexity of the hypervolume indicator. In *Genetic and Evolutionary Computation Conference, GECCO 2013*, pages 575–582. ACM, 2013.
- [BLS14] Golnaz Badkobeh, Per Kristian Lehre, and Dirk Sudholt. Unbiased black-box complexity of parallel search. In *Parallel Problem Solving from Nature, PPSN 2014*, pages 892–901. Springer, 2014.
- [CD17] Eduardo Carvalho Pinto and Carola Doerr. Discussion of a more practice-aware runtime analysis for evolutionary algorithms. In *Artificial Evolution, EA 2017*, pages 298–305, 2017.
- [CDF14] Sylvain Colin, Benjamin Doerr, and Gaspard Férey. Monotonic functions in EC: anything but monotone! In *Genetic and Evolutionary Computation Conference, GECCO 2014*, pages 753–760. ACM, 2014.
- [DAS⁺17] Yifei Du, Kenji Aoki, Makoto Sakamoto, Hiroshi Furutani, and Kunihito Yamamori. Markov chain analysis of Leading Ones problem. *Artificial Life and Robotics*, 22:443–448, 2017.
- [DD15] Benjamin Doerr and Carola Doerr. A tight runtime analysis of the $(1+(\lambda, \lambda))$ genetic algorithm on OneMax. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2015*, pages 1423–1430. ACM, 2015.
- [DD16] Benjamin Doerr and Carola Doerr. The impact of random initialization on the runtime of randomized search heuristics. *Algorithmica*, 75:529–553, 2016.
- [DD18] Benjamin Doerr and Carola Doerr. Optimal static and self-adjusting parameter choices for the $(1 + (\lambda, \lambda))$ genetic algorithm. *Algorithmica*, 80:1658–1709, 2018.
- [DDE15] Benjamin Doerr, Carola Doerr, and Franziska Ebel. From black-box complexity to designing new genetic algorithms. *Theoretical Computer Science*, 567:87–104, 2015.

- [DDK17] Benjamin Doerr, Carola Doerr, and Timo Kötzing. Unknown solution length problems with no asymptotically optimal run time. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 921–928. ACM, 2017.
- [DDY16a] Benjamin Doerr, Carola Doerr, and Jing Yang. k -bit mutation with self-adjusting k outperforms standard bit mutation. In *Parallel Problem Solving from Nature, PPSN 2016*, pages 824–834. Springer, 2016.
- [DDY16b] Benjamin Doerr, Carola Doerr, and Jing Yang. Optimal parameter choices via precise black-box analysis. In *Genetic and Evolutionary Computation Conference, GECCO 2016*, pages 1123–1130. ACM, 2016.
- [DEN⁺11] Benjamin Doerr, Anton V. Eremeev, Frank Neumann, Madeleine Theile, and Christian Thyssen. Evolutionary algorithms and dynamic programming. *Theoretical Computer Science*, 412:6020–6035, 2011.
- [DFE⁺17] Benjamin Doerr, Philipp Fischbeck, Clemens Frahnw, Tobias Friedrich, Timo Kötzing, and Martin Schirneck. Island models meet rumor spreading. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 1359–1366. ACM, 2017.
- [DFW10] Benjamin Doerr, Mahmoud Fouz, and Carsten Witt. Quasirandom evolutionary algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 1457–1464. ACM, 2010.
- [DFW11] Benjamin Doerr, Mahmoud Fouz, and Carsten Witt. Sharp bounds by probability-generating functions and variable drift. In *Genetic and Evolutionary Computation Conference, GECCO 2011*, pages 2083–2090. ACM, 2011.
- [DG13] Benjamin Doerr and Leslie Ann Goldberg. Adaptive drift analysis. *Algorithmica*, 65:224–250, 2013.
- [DGWY17] Benjamin Doerr, Christian Gießen, Carsten Witt, and Jing Yang. The $(1+\lambda)$ evolutionary algorithm with self-adjusting mutation rate. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 1351–1358. ACM, 2017. Full version available at <http://arxiv.org/abs/1704.02191>.

- [DH08] Benjamin Doerr and Edda Happ. Directed trees: A powerful representation for sorting and ordering problems. In *Congress on Evolutionary Computation, CEC 2008*, pages 3606–3613. IEEE, 2008.
- [DHK11] Benjamin Doerr, Edda Happ, and Christian Klein. Tight analysis of the (1+1)-EA for the single source shortest path problem. *Evolutionary Computation*, 19:673–691, 2011.
- [DHK12] Benjamin Doerr, Edda Happ, and Christian Klein. Crossover can provably be useful in evolutionary computation. *Theoretical Computer Science*, 425:17–33, 2012.
- [DHN07] Benjamin Doerr, Nils Hebbinghaus, and Frank Neumann. Speeding up evolutionary algorithms through asymmetric mutation operators. *Evolutionary Computation*, 15:401–410, 2007.
- [DJ07] Benjamin Doerr and Daniel Johannsen. Adjacency list matchings: an ideal genotype for cycle covers. In *Genetic and Evolutionary Computation Conference, GECCO 2007*, pages 1203–1210. ACM, 2007.
- [DJ10] Benjamin Doerr and Daniel Johannsen. Edge-based representation beats vertex-based representation in shortest path problems. In *Genetic and Evolutionary Computation Conference, GECCO 2010*, pages 759–766. ACM, 2010.
- [DJW00] Stefan Droste, Thomas Jansen, and Ingo Wegener. A natural and simple function which is hard for all evolutionary algorithms. In *IEEE International Conference on Industrial Electronics, Control, and Instrumentation, IECON 2000*, pages 2704–2709. IEEE, 2000.
- [DJW02] Stefan Droste, Thomas Jansen, and Ingo Wegener. On the analysis of the (1+1) evolutionary algorithm. *Theoretical Computer Science*, 276:51–81, 2002.
- [DJW12] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64:673–697, 2012.
- [DJWZ13] Benjamin Doerr, Thomas Jansen, Carsten Witt, and Christine Zarges. A method to derive fixed budget results from expected optimisation times. In *Genetic and Evolutionary Computation Conference, GECCO 2013*, pages 1581–1588. ACM, 2013.

- [DK15] Benjamin Doerr and Marvin Künnemann. Optimizing linear functions with the $(1+\lambda)$ evolutionary algorithm—different asymptotic runtimes for different instances. *Theoretical Computer Science*, 561:3–23, 2015.
- [DKLL17] Benjamin Doerr, Timo Kötzing, J. A. Gregor Lagodzinski, and Johannes Lengler. Bounding bloat in genetic programming. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 921–928. ACM, 2017.
- [DKLW13] Benjamin Doerr, Timo Kötzing, Johannes Lengler, and Carola Winzen. Black-box complexities of combinatorial problems. *Theoretical Computer Science*, 471:84–106, 2013.
- [DKS07] Benjamin Doerr, Christian Klein, and Tobias Storch. Faster evolutionary algorithms by superior graph representation. In *Foundations of Computational Intelligence, FOCI 2007*, pages 245–250. IEEE, 2007.
- [DKV13] Benjamin Doerr, Bojana Kodric, and Marco Voigt. Lower bounds for the runtime of a global multi-objective evolutionary algorithm. In *Congress on Evolutionary Computation, CEC 2013*, pages 432–439. IEEE, 2013.
- [DL16] Duc-Cuong Dang and Per Kristian Lehre. Self-adaptation of mutation rates in non-elitist populations. In *Parallel Problem Solving from Nature, PPSN 2016*, pages 803–813. Springer, 2016.
- [DL17] Carola Doerr and Johannes Lengler. Introducing elitist black-box models: When does elitist behavior weaken the performance of evolutionary algorithms? *Evolutionary Computation*, 25:587–606, 2017.
- [DLMN17] Benjamin Doerr, Huu Phuoc Le, Régis Makhmara, and Ta Duy Nguyen. Fast genetic algorithms. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 777–784. ACM, 2017. Full version available at <http://arxiv.org/abs/1703.03334>.
- [DNDD⁺18] Raphaël Dang-Nhu, Thibault Dardinier, Benjamin Doerr, Gautier Izacard, and Dorian Nogneng. A new analysis method for

- evolutionary optimization of dynamic and noisy objective functions. In *Genetic and Evolutionary Computation Conference, GECCO 2018*. ACM, 2018.
- [Doe11] Benjamin Doerr. Analyzing randomized search heuristics: Tools from probability theory. In Anne Auger and Benjamin Doerr, editors, *Theory of Randomized Search Heuristics*, pages 1–20. World Scientific, 2011.
- [Doe18a] Benjamin Doerr. Better runtime guarantees via stochastic domination. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2018*, pages 1–17. Springer, 2018.
- [Doe18b] Benjamin Doerr. Probabilistic tools for the analysis of randomized optimization heuristics. *CoRR*, abs/1801.06733v3, 2018.
- [dPdLDD15] Axel de Perthuis de Laillevault, Benjamin Doerr, and Carola Doerr. Money for nothing: Speeding up evolutionary algorithms through better initialization. In *Genetic and Evolutionary Computation Conference, GECCO 2015*, pages 815–822. ACM, 2015.
- [Dro03] Stefan Droste. Analysis of the (1+1) EA for a dynamically bitwise changing OneMax. In *Genetic and Evolutionary Computation Conference, GECCO 2003*, pages 909–921. Springer, 2003.
- [Dro04] Stefan Droste. Analysis of the (1+1) EA for a noisy OneMax. In *Genetic and Evolutionary Computation Conference, GECCO 2004*, pages 1088–1099. Springer, 2004.
- [DW18] Carola Doerr and Markus Wagner. Simple on-the-fly parameter selection mechanisms for two classical discrete black-box optimization benchmark problems. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 943–950. ACM, 2018.
- [DWY18] Benjamin Doerr, Carsten Witt, and Jing Yang. Runtime analysis for self-adaptive mutation rates. In *Genetic and Evolutionary Computation Conference, GECCO 2018*, pages 1475–1482. ACM, 2018.

- [Fau31] Johannes Faulhaber. *Academia Algebrae, darinnen die miraculose Inventiones, zu den höchsten Cossen weiters continuirt und profitiert werden*. 1631.
- [FKKS17] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, and Andrew M. Sutton. The compact genetic algorithm is efficient under extreme gaussian noise. *IEEE Trans. Evolutionary Computation*, 21:477–490, 2017.
- [GKS99] Josselin Garnier, Leila Kallel, and Marc Schoenauer. Rigorous hitting times for binary mutations. *Evolutionary Computation*, 7:173–203, 1999.
- [HPR⁺18] Hsien-Kuei Hwang, Alois Panholzer, Nicolas Rolin, Tsung-Hsi Tsai, and Wei-Mei Chen. Probabilistic analysis of the (1+1)-evolutionary algorithm. *Evolutionary Computation*, 26:299–345, 2018.
- [HY01] Jun He and Xin Yao. Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127:51–81, 2001.
- [Jan17] Swante Janson. Tail bounds for sums of geometric and exponential variables. *ArXiv e-prints*, arXiv:1709.08157, 2017.
- [JW01] Thomas Jansen and Ingo Wegener. Evolutionary algorithms - how to cope with plateaus of constant fitness and when to reject strings of the same fitness. *IEEE Transactions on Evolutionary Computation*, 5:589–599, 2001.
- [JW06] Thomas Jansen and Ingo Wegener. On the analysis of a dynamic evolutionary algorithm. *Journal of Discrete Algorithms*, 4:181–199, 2006.
- [JZ11] Thomas Jansen and Christine Zarges. On benefits and drawbacks of aging strategies for randomized search heuristics. *Theoretical Computer Science*, 412:543–559, 2011.
- [KW17] Martin S. Krejca and Carsten Witt. Lower bounds on the run time of the univariate marginal distribution algorithm on OneMax. In *Foundations of Genetic Algorithms, FOGA 2017*, pages 65–79. ACM, 2017.

- [Lad05] Véronique Ladret. Asymptotic hitting time for a simple evolutionary model of protein folding. *Journal of Applied Probability*, 42:39–51, 2005.
- [Len18] Johannes Lengler. A general dichotomy of evolutionary algorithms on monotone functions. *CoRR*, abs/1803.09227, 2018.
- [LOW17] Andrei Lissovoi, Pietro Simone Oliveto, and John Alasdair Warwicker. On the runtime analysis of generalised selection hyper-heuristics for pseudo-boolean optimisation. In *Genetic and Evolutionary Computation Conference, GECCO 2017*, pages 849–856. ACM, 2017.
- [LS15] Johannes Lengler and Nicholas Spooner. Fixed budget performance of the (1+1) EA on linear functions. In *Foundations of Genetic Algorithms, FOGA 2015*, pages 52–61. ACM, 2015.
- [LW12] Per Kristian Lehre and Carsten Witt. Black-box search by unbiased variation. *Algorithmica*, 64:623–642, 2012.
- [LY12] Per Kristian Lehre and Xin Yao. On the impact of mutation-selection balance on the runtime of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 16:225–241, 2012.
- [MS02] Alfred Müller and Dietrich Stoyan. *Comparison Methods for Stochastic Models and Risks*. Wiley, 2002.
- [Neu08] Frank Neumann. Expected runtimes of evolutionary algorithms for the eulerian cycle problem. *Computers & OR*, 35:2750–2759, 2008.
- [NW07] Frank Neumann and Ingo Wegener. Randomized local search, evolutionary algorithms, and the minimum spanning tree problem. *Theoretical Computer Science*, 378:32–40, 2007.
- [OHY09] Pietro Simone Oliveto, Jun He, and Xin Yao. Analysis of the (1+1)-EA for finding approximate solutions to vertex cover problems. *IEEE Transactions on Evolutionary Computation*, 13:1006–1029, 2009.
- [OLN09] Pietro Simone Oliveto, Per Kristian Lehre, and Frank Neumann. Theoretical analysis of rank-based mutation - combining exploration and exploitation. In *Congress on Evolutionary Computation, CEC 2009*, pages 1455–1462. IEEE, 2009.

- [OW15] Pietro Simone Oliveto and Carsten Witt. Improved time complexity analysis of the simple genetic algorithm. *Theoretical Computer Science*, 605:21–41, 2015.
- [RS14] Jonathan E. Rowe and Dirk Sudholt. The choice of the offspring population size in the $(1, \lambda)$ evolutionary algorithm. *Theoretical Computer Science*, 545:20–38, 2014.
- [Rud97] Günter Rudolph. *Convergence Properties of Evolutionary Algorithms*. Verlag Dr. Kováč, 1997.
- [Sch00] Christian Scheideler. *Probabilistic Methods for Coordination Problems*. University of Paderborn, 2000. Habilitation thesis. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.70.1319>.
- [ST12] Dirk Sudholt and Christian Thyssen. A simple ant colony optimizer for stochastic shortest path problems. *Algorithmica*, 64:643–672, 2012.
- [STW04] Jens Scharnow, Karsten Tinnefeld, and Ingo Wegener. The analysis of evolutionary algorithms on sorting and shortest paths problems. *Journal of Mathematical Modelling and Algorithms*, 3:349–366, 2004.
- [Sud09] Dirk Sudholt. The impact of parametrization in memetic evolutionary algorithms. *Theoretical Computer Science*, 410:2511–2528, 2009.
- [Sud13] Dirk Sudholt. A new method for lower bounds on the running time of evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 17:418–435, 2013.
- [SW08] Dirk Sudholt and Carsten Witt. Runtime analysis of binary PSO. In *Genetic and Evolutionary Computation Conference, GECCO 2008*, pages 135–142. ACM, 2008.
- [The09] Madeleine Theile. Exact solutions to the traveling salesperson problem by a population-based evolutionary algorithm. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP 2009*, pages 145–155. Springer, 2009.
- [Weg01] Ingo Wegener. Theoretical aspects of evolutionary algorithms. In *International Colloquium on Automata, Languages and Programming, ICALP 2001*, pages 64–78. Springer, 2001.

- [Wit06] Carsten Witt. Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-Boolean functions. *Evolutionary Computation*, 14:65–86, 2006.
- [Wit08] Carsten Witt. Population size versus runtime of a simple evolutionary algorithm. *Theoretical Computer Science*, 403:104–120, 2008.
- [Wit13] Carsten Witt. Tight bounds on the optimization time of a randomized search heuristic on linear functions. *Combinatorics, Probability & Computing*, 22:294–318, 2013.
- [Wit14] Carsten Witt. Fitness levels with tail bounds for the analysis of randomized search heuristics. *Information Processing Letters*, 114:38–41, 2014.
- [WW05] Ingo Wegener and Carsten Witt. On the optimization of monotone polynomials by simple randomized search heuristics. *Combinatorics, Probability & Computing*, 14:225–247, 2005.
- [ZLLH12] Dong Zhou, Dan Luo, Ruqian Lu, and Zhangang Han. The use of tail inequalities on the probable computational time of randomized search heuristics. *Theoretical Computer Science*, 436:106–117, 2012.