Analysis of a Gray-Box Operator for Vertex Cover

Samuel Baguley, Tobias Friedrich, Timo Kötzing, Xiaoyue Li, Marcus Pappik, and Ziena Zeif

{firstname.lastname}@hpi.de

Hasso Plattner Institute
University of Potsdam
Potsdam, Germany

Combinatorial optimization problems are a prominent application area of evolutionary algorithms, where the (1+1) EA is one of the most investigated. We extend this algorithm by introducing some problem knowledge with a specialized mutation operator which works under the assumption that the number of 1s of a solution is critical, as frequently happens in combinatorial optimization. This slight modification increases the chance to correct wrongly placed bits while preserving the simplicity and problem independence of the (1+1) EA.

As an application of our algorithm we examine the vertex cover problem on certain instances, where we show that it leads to asymptotically better runtimes and even finds with higher probability optimal solutions in comparison with the usual (1+1) EA. Precisely, we compare the performance of both algorithms on paths and on complete bipartite graphs of size n. Regarding the path we prove that, for a particular initial configuration, the (1+1) EA takes in expectation $\Theta(n^4)$ iterations while the modification reduces this to $\Theta(n^3)$, and present experimental evidence that such a configuration is reached. Concerning the complete bipartite graph our modification finds the optimum in polynomial time with probability $1-1/2^{\Omega(n^\xi)}$ for every positive constant $\xi < 1$, which improves the known probability of $1 - 1/\text{poly}(n)$ for the (1+1) EA.

1 Introduction

Evolutionary algorithms (EAs) are a diverse and highly versatile class of randomized search heuristics, and their effectiveness at solving combinatorial optimization problems has been recognized and widely studied since the nineties [19, Chap. 14]. Typically, EAs are implemented with naïve mutation behavior, and adapt to a particular problem only via an associated fitness function. This simplicity has lead to success in many applications, but leaves room for improvement on specific problems.

The minimum vertex cover problem is an ideal target for evolutionary algorithms because it is NP-complete [10], and appears in a broad range of practical applications. Already in 1994 Bäck and Khuri [1] defined the canonical fitness function for EAs on vertex cover and provided experimental evidence that on some graph instances EAs can produce better approximate solutions than bespoke solvers. In the years that followed runtime analysis of EAs tended to be heuristic and experimental, but recently rigorous theoretical analyses have gained prominence. The various approaches to this include proving runtime bounds for approximate solutions [9, 15], analyzing performance on best- or worst-case graph instances [13, 14], and implementing genetic mechanisms like crossover [6]. It is also common for new research to define variations of well-studied EAs, like Olivieto et al.'s (1+1) EA with an additional 'diversity maintenance mechanism' [14], or a fitness function with multiple objectives [5, 11].

In the application of evolutionary algorithms to real world problems, the full black-box view of the fitness function (ignoring all additional information about its inner workings) is at most the initial prototype; significant improvements can be made by applying problem knowledge for designing tailored variation operators [17]. The resulting view is called gray-box optimization, and research tries to understand what general structural properties can be usefully employed in evolutionary computation [20]. We are interested in combinatorial problems where a certain subset has to be selected; typically, the number of elements in the chosen set is rather important. For example, a minimum spanning tree contains exactly $n - 1$ edges, the vertex cover problem tries to minimize the number of vertices picked, and the independent set problem maximizes the selected vertices. While [12] argued that variation operators should be unbiased with respect to whether a 0 or a 1 encodes a specific property of a solution, the gray-box view of combinatorial problems wants the 1 to mean 'in' and 0 to mean 'out', and then design operators that can make use of this meaning. From the very general work of Rowe and Vose [18] we can see that only for problems which have symmetrical meaning in 1 and 0 can the operators be assumed unbiased.

In this paper we augment the (1+1) EA by allowing a second kind of mutation called a 'balanced flip', in which a bit is chosen uniformly at random, and then swapped with the bit of a neighboring vertex of opposite value, if one exists. At every iteration the augmented algorithm chooses between a balanced flip and a regular step of the (1+1) EA with equal probability. The 'balanced flip' behavior is already exhibited by the (1+1) EA on nearly-optimal bit-strings, the difference being that it takes an order of magnitude longer to perform each flip. Thus the evolutionary algorithm presented here refines the (1+1) EA by incorporating its emergent behavior on inputs close to the optimum directly into the mutation operator; this is somewhat in the spirit of Giel and Wegener's 'local (1+1) EA' [7]. The immediate benefit of our augmentation is a tighter upper runtime bound for the Balanced (1+1) EA on paths.

In Section 3 we compare the runtime of the (1+1) EA solving vertex cover on paths to that of the Balanced (1+1) EA. Our analysis of the (1+1) EA uses techniques similar to those in [7] to obtain an expected runtime of $O(n^4)$, and we then apply those same techniques to the Balanced (1+1) EA to prove an expected runtime of $O(n^3)$. We also give a partial proof of $\Omega(n^4)$ expected runtime for the (1+1) EA, assuming that the algorithm at some point produces a bit-string with a long consecutive sequence of incorrectly assigned bits. We provide experimental evidence to motivate this condition, and delay a full proof to future work.

In Section 4 we consider the complete bipartite graph $K_{L,R}$, which is a common instance class of interest when it comes to evolutionary algorithms. Under the assumption that

the ratio R/L is at least 2, the (1+1) EA is known to have exponential runtime with at least polynomial probability [5, Theorem 5]. In contrast, the Balanced (1+1) EA has far better performance on this graph instance, with exponentially small probability of not reaching the optimum in polynomial time.

2 Preliminaries

In this article we demonstrate that incorporating even small amounts of problem knowledge can significantly improve reduce expected runtime of the (1+1) EA solving vertex cover. We provide a new evolutionary algorithm that we call *Balanced (1+1) EA* (Algorithm 2) that we compare with the classical *(1+1) EA* (Algorithm 1).

We introduce briefly the graph terminology used in this paper. A graph will be denoted $G = (V, E)$, and for any subgraph H of G, we write $V(H)$ and $E(H)$ for vertices and edges of H respectively. We denote an edge $e = \{u, v\} \in E$ by uv and the neighborhood of a vertex $v \in V$ in G by $N(v) = \{u \in V \mid uv \in E\}$. Lastly, we define the length of a path $P = (V, E)$ to be number of vertices $|V|$.

A *vertex cover* for a given graph $G = (V, E)$ is a subset of vertices $V' \subseteq V$ which satisfies $V' \cap e \neq \varnothing$ for all $e \in E$. A solution to the *vertex cover problem* is a vertex cover V' which has minimal cardinality $|V'|$, in the sense that no cover of smaller cardinality exists.

Depending on the context, we shall use X to denote either the stochastic process induced by one of the algorithms or an arbitrary state of that process $X \in \{0, 1\}^n$. We make sure that the meaning of X is always clarified. For $t \in \mathbb{N}$, X_t denotes the configuration after t iterations, and X_0 corresponds to the initial configuration of the process X. Both algorithms studied in this work have search space $\{0, 1\}^n$ of bit strings of length n, where $n = |V|$. If $V = (v_1, \ldots, v_n)$ is some enumeration of the vertices of G, then a bit string $X = (x_1, \ldots, x_n) \in \{0, 1\}^n$ corresponds to a subset of $V_X \subseteq V$ via the mapping $v_i \in V_x \Leftrightarrow x_i = 1$. We call a bit string X *feasible* if V_X is a vertex cover for G, and *infeasible* otherwise. If V_X is a minimal vertex cover then X is called *optimal*. Because the relationship between X and V_X is one-to-one, we will abuse notation and use X to represent the bit string or subset of vertices interchangeably. Letting $u(X) := |\{e \in E \mid V_X \cap e = \emptyset\}|$ be the number of uncovered edges in X, we have that X is feasible if and only if $u(X) = 0$. Further, we define $|X|_1 := |V_X|$ to be the number of ones in X, $|X|_0 := n - |X|_1$ to be the number of zeroes, and $\mathrm{OPT} = |X_{\mathrm{opt}}|_1$ to be the number of ones in the minimal vertex cover.

The *fitness* of X is given by the integer $f(X)$, where

$$f : \{0, 1\}^n \to \mathbb{N} : X \mapsto |X|_1 + (n + 1)u(X)$$

defines the fitness function. The multiplicative factor $n + 1$ ensures that the (1+1) EA and Balanced (1+1) EA prioritise covering any uncovered edge over reducing the size of X, and that once covered an edge will never become uncovered. This fitness function (and slight variations of it) has been standard for evolutionary algorithms solving the vertex cover problem since [1]. With this in hand, the (1+1) EA can be defined.

Because the (1+1) EA has positive probability of sampling any point of $\{0, 1\}^n$ in step 3, regardless of the current state X, it almost surely produces a solution to the vertex cover problem in finite time. A fundamental property of interest, and the main focus of this article, is the *expected* time to find this solution. If the optimum cover is unique then a naïve estimate based on this behavior yields an extremely poor expected runtime of n^n

---
Algorithm 1: (1+1) EA
---
1 Choose $X \in \{0,1\}^n$ (if unspecified, uniformly at random)
2 **while** *stopping criterion not met* **do**
3 $Y \leftarrow$ flip each bit of X independently with probability $1/n$
4   **if** $f(Y) \leq f(X)$ **then**
5 $X \leftarrow Y$
6   **end**
7 **end**
---

for the (1+1) EA. Worst-case analysis, for example in [5, Theorem 5], suggests that no better can be expected for general G.

We now introduce a version of the (1+1) EA which chooses with equal probability between two sampling behaviors, one of which is the same as in the (1+1) EA, while the other mimics behavior of the (1+1) EA on feasible bit strings. We call this second kind of behavior a 'balanced flip', and the augmented algorithm is therefore called the Balanced (1+1) EA.

Algorithm 2: Balanced (1+1) EA
---
1 Choose $X \in \{0,1\}^n$ (if unspecified, uniformly at random)
2 **while** *stopping criterion not met* **do**
3 $p \leftarrow$ pick a number u.a.r. in the interval $(0,1)$
4   **if** $p \leq 1/2$ **then**
5 $Y \leftarrow$ flip each bit of X independently with probability $1/n$
6   **end**
7   **else**
8 $v \leftarrow$ pick a vertex u.a.r from $V(G)$
9 $N_v \leftarrow \{v' \in N(v) \mid x_v \neq x_u\}$
10   **if** $N_v \neq \varnothing$ **then**
11 $u \leftarrow$ pick a vertex u.a.r from N_v
12 $Y \leftarrow$ flip x_u and x_v
13   **end**
14   **else**
15   **go to** line 2
16   **end**
17   **end**
18   **if** $f(Y) \leq f(X)$ **then**
19 $X \leftarrow Y$
20   **end**
21 **end**

2.1 Feasible bit strings

In order to analyze the expected runtimes of the (1+1) EA and Balanced (1+1) EA, we use the following multiplicative drift theorem to bound the time it takes for X to become feasible.

Theorem 1 (Doerr et. al [4])**.** *Let* $(X_t, t \in \mathbb{N})$ *be a stochastic process on* $S \cup \{0\}$, *where* $S \subsetneq \mathbb{R}^+$ *is a set positive numbers with minimum* $s_{\min}$, *and let* $T = \inf\{t \in \mathbb{N} \mid X_t = 0\}$. *Suppose there exists a real number* $\delta > 0$ *such that for all* $s \in S$ *and* $t < T$, $\mathbb{E}[X_t - X_{t+1} \mid X_t = s] \geq \delta s$. *Then, for all* $s_0 \in S$,

$$\mathbb{E}[T \mid X_0 = s_0] \leq \frac{\ln(s_0) - \ln(s_{\min}) + 1}{\delta}.$$

Lemma 2. *Let* $G = (V, E)$ *be a graph with* $n = |V|$ *vertices, and let* T *be the first time that the (1+1) EA samples a feasible string. Then* $\mathbb{E}[T] \leq en(\ln(n) + 1/2)$ *and for* $k > 0$, $\mathbb{P}[T > 2ekn(\ln(n) + 1/2)] \leq 2^{-k}$.

Further, if S *is the first time that the Balanced (1+1) EA samples a feasible string, then* $\mathbb{E}[S] \leq 2en(\ln(n) + 1/2)$ *and for* $k > 0$, $\mathbb{P}[S > 4ekn(\ln(n) + 1/2)] \leq 2^{-k}$.

Proof. We start by analysing the (1+1) EA. Let $X_t \in \{0, 1\}^n$ denote the state of the (1+1) EA after t iterations, where X_0 is the initial random configuration. We shall apply Theorem 1 to the process $u(X_t)$, $t \geq 0$. For all $t \in \mathbb{N}$, let $F_t \subseteq E$ be the set of edges that are not covered by X_t and let $U_t := \bigcup_{e \in F_t} e$. Note that the algorithm attempts to flip exactly one bit with probability $(1 - 1/n)^{n-1} \geq 1/e$. Whenever a single bit is flipped, this bit is chosen uniformly at random. Consider the (not necessarily induced) subgraph (U_t, F_t) of G. If the flipped bit corresponds to a vertex $v \in U_t$, then $u(X_t) - u(X_{t+1})$ is equal to the degree of v in (U_t, F_t). Thus, given a uniformly random bit from U_t is flipped, the expected difference $u(X_t) - u(X_{t+1})$ is the average degree of (U_t, F_t), which is $2|F_t| / |U_t| = 2u(X_t)/|U_t|$ by the hand shake lemma. Let us denote the event that exactly one bit in U_t is flipped, and no other, by C. Since $u(X_t)$ is non-increasing, $u(X_t) - u(X_{t+1})$ is non-negative. Therefore, we have

$$\mathbb{E}[u(X_t) - u(X_{t+1}) \mid u(X_t)] \geq \mathbb{E}[(u(X_t) - u(X_{t+1}))\mathbb{1}_C \mid u(X_t)]$$
$$\geq \frac{2u(X_t)}{|U_t|}\mathbb{P}[C \mid u(X_t)]$$
$$\geq \frac{2u(X_t)}{|U_t|}\frac{|U_t|}{en} \geq \frac{2u(X_t)}{en}.$$

So the drift condition of Theorem 1 is satisfied by $u(X)$. Since $u(X) \in \{0, 1, \ldots, n^2\}$, we conclude that $\mathbb{E}[T] \leq en(\ln(n) + 1/2)$.

To bound the expected time for the Balanced (1+1) EA to find a feasible solution, note that with probability $1/2$ it performs the same actions as the (1+1) EA. Moreover, the balanced flip operations cannot increase $u(X_t)$ due to the fitness function f. Thus we can essentially use the same proof as for the (1+1) EA but with drift $\mathbb{E}[u(X_t) - u(X_{t+1}) \mid u(X_t)] \geq u(X_t)/en$, which yields an expected hitting time of at most $2en(\ln(n) + 1/2)$.

In order to prove the second part of the statement, we start once again with the (1+1) EA. We prove our claim via induction over $k \in \mathbb{N}$. The base case $k = 0$ is trivial. Now, consider some $k > 0$ and assume the claim holds for $k - 1$. Let A be the event that no feasible solution was found within the first $2e(k-1)n(\ln(n) + 1/2)$ iterations and let B be the event that no feasible solution is found within the last $2en(\ln(n) + 1/2)$ iterations. Then the desired probability can be written as $\mathbb{P}[A \text{ and } B] = \mathbb{P}[B \mid A]\mathbb{P}[A]$. By our induction hypothesis, we have $\mathbb{P}[A] \leq 2^{-(k-1)}$. Further, note that the bound on the expected number of iterations for finding a feasible solution, proven above, holds for arbitrary starting states. Thus, by Markov's inequality, we have $\mathbb{P}[B \mid A] \leq \frac{1}{2}$, which

proves the upper bound. The proof for the Balanced (1+1) EA is obtained analogously, using the respective expected hitting time. $\square$

We frequently make use of Lemma 2 throughout different parts of this paper. Note that a desirable property of the update criteria of both Algorithms 1 and 2 and the fitness f is that, once a feasible solution is found, the algorithms never leave the space of feasible solutions again.

2.2 Lower Bound Probability for Symmetric Random Walks

The technical core of our proof for a lower bound on the running time for solving vertex cover on paths is the following general lemma. It essentially lower bounds the hitting time of a process that is dominated by a symmetric random walk. There are two main difference between Lemma 3 and most bounds that can be found in the literature. Firstly, the presented bound does not only hold in expectation but at least with constant probability and, secondly, the description of the process allows for arbitrarily large jumps, as long as their probability is bounded.

Lemma 3. *Let $d \in \mathbb{N}$ with $d \geq 4\ln(10)$ and let $Z_t \in \mathbb{N}$ be a stochastic process such that, for some $p \in (0,1]$ and $q \in [0,1)$ with $q + 2p \leq 1$, it holds that*

(i) $\forall s \in \mathbb{N} \cap [0, d-1] : \mathbb{P}\left[Z_{t+1} = s+1 \mid Z_t = s\right] \leq p$,

(ii) $\forall s \in \mathbb{N} \cap [1, d-1] : \mathbb{P}\left[Z_{t+1} = s+1 \mid Z_t = s\right] \leq \mathbb{P}\left[Z_{t+1} = s-1 \mid Z_t = s\right]$ and

(iii) $\forall s \in \mathbb{N} \cap [0, d] : \mathbb{P}\left[|Z_{t+1} - Z_t| > 1 \mid Z_t = s\right] \leq q$.

Then, for $Z_0 = 0$ and $T = \inf\{t \in \mathbb{N} \mid Z_t \geq d\}$, it holds that

$$\mathbb{P}\left[T \geq \min\left\{\frac{1}{5q}, \frac{d^2(1-q)}{4\ln(10)p}\right\}\right] \geq \frac{16}{25}.$$

To prove Lemma 3, we need two additional ingredients. The first one is an elementary bound on the lower tail of a geometric random variable, which will allow us to lower bound the time until a certain event appears.

Lemma 4. *Let X be a geometric random variable with success probability $p \in (0,1)$. For all $c \in \mathbb{R}_{\geq 0}$ it holds that $\mathbb{P}\left[X \geq \frac{c}{p}\right] \geq 1 - c$.*

Proof. If $\frac{c}{p} \leq 1$ then $\mathbb{P}\left[X \geq \frac{c}{p}\right] = 1 \geq 1 - c$ holds trivially. Assume $\frac{c}{p} > 1$ and let $k = \left\lceil \frac{c}{p} \right\rceil \geq 2$. By the definition of a geometric random variable and Bernoulli's inequality, we have

$$\mathbb{P}\left[X \geq \frac{c}{p}\right] = \mathbb{P}\left[X \geq k\right] = (1-p)^{k-1} \geq 1 - p \cdot (k-1).$$

Observing that $\left\lceil \frac{c}{p} \right\rceil - 1 \leq \frac{c}{p}$ concludes the proof. $\square$

The second ingredient is a concentration result on the sum of independent geometric random variables, which is similar to the Chernoff bound. More specifically, we are interested in obtaining a lower bound on such a sum.

Theorem 5. *[Doerr [3], Theorem* $1.10.32(b)$*] Let* $X_1, \ldots X_n$ *be independent geometric random variables with common success probability* $p > 0$*. Then, for* $X = \sum_{i=1}^n X_i$*,* $\mu = \frac{n}{p}$ *and all* $\delta \in (0, 1)$ *it holds that*

$$\mathbb{P}\left[X \leq (1 - \delta)\mu\right] \leq e^{-\frac{\delta^2}{2 - \frac{4}{3}\delta}n}.$$

Given these two result, we are able to prove Lemma 3.

Proof of Lemma 3. We start by constructing a modified stochastic process $Z'_t \in D$ for $D = \mathbb{Z} \cap [-d, d+1]$ with the following transition probabilities

- $\forall s \in D \setminus \{-d, d\} : \mathbb{P}\left[Z'_{t+1} = s - 1 \mid Z'_t = s\right] = \mathbb{P}\left[Z'_{t+1} = s + 1 \mid Z'_t = s\right] = p,$

- $\forall s \in D \setminus \{-d, d\} : \mathbb{P}\left[Z'_{t+1} = d + 1 \mid Z'_t = s\right] = q$ and

- $\forall s \in \{-d, d, d+1\} : \mathbb{P}\left[Z'_{t+1} = s \mid Z'_t = s\right] = 1$

Set $Z'_0 = 0$ and $T' = \inf\{t \in \mathbb{N} \mid |Z'_t| \geq d\}$. Note that due to Lemma 3(ii), $|Z'_t|$ and Z_t can be coupled in such a way that $|Z'_t| \geq Z_t$ for all $t < T$ and $|Z'_t| \geq d$ for all $t \geq T$. In particular, $T \geq T'$ almost surely. Consequently we have for all $t \in \mathbb{N}$ that $\mathbb{P}\left[T \geq t\right] \geq \mathbb{P}\left[T' \geq t\right]$.

We proceed by analyzing T'. To this end, we start by treating cases in which Z'_t directly jumps from $s \in D$ with $|s| \leq d - 1$ to $d + 1$ separately. We define a new stopping time $T_j = \inf\{t \in \mathbb{N}_{\geq 1} \mid |Z'_{t-1}| < d$ and $Z'_t = d + 1\}$. Note that, for all $m \in \mathbb{N}$,

$$\mathbb{P}\left[T' \geq m\right] \geq \mathbb{P}\left[T' \geq m \text{ and } T_j \geq m\right]$$
$$= \mathbb{P}\left[T' \geq m \mid T_j \geq m\right] \cdot \mathbb{P}\left[T_j \geq m\right].$$

We set $m = \min\left\{\frac{1}{5q}, \frac{d^2(1-q)}{4\ln(10)p}\right\}$ and lower bound $\mathbb{P}\left[T_j \geq m\right]$ and $\mathbb{P}\left[T' \geq m \mid T_j \geq m\right]$ separately.

To obtain a lower bound on $\mathbb{P}\left[T_j \geq m\right]$, observe that, at any point in time $t \in \mathbb{N}$, Z'_t has a probability of at most q to do the desired jump. Thus, T_j dominates a geometrically distributed random variable with success probability q. As $m \leq \frac{1}{5q}$, Lemma 4 yields $\mathbb{P}\left[T_j \geq m\right] \geq \frac{4}{5}$.

For lower bounding $\mathbb{P}\left[T' \geq m \mid T_j \geq m\right]$, let $Y_t \in D$ be a stochastic process with

- $\forall s \in D \setminus \{-d, d\} :$
  $\mathbb{P}\left[Y_{t+1} = s - 1 \mid Y_t = s\right] = \mathbb{P}\left[Y_{t+1} = s + 1 \mid Y_t = s\right] = \frac{p}{1-q},$

- $\forall s \in \{-d, d, d+1\} : \mathbb{P}\left[Y_{t+1} = s \mid Y_t = s\right] = 1.$

Set $Y_0 = 0$ and $T_Y = \inf\{t \in \mathbb{N} \mid |Y_t| \geq d\}$. By defining Y in this way, and given that $T_j \geq m$, we can couple Y_t and Z'_t in such a way that $Y_t = Z'_t$ for all $t < m$. Therefore, we have $\mathbb{P}\left[T' \geq m \mid T_j \geq m\right] = \mathbb{P}\left[T_Y \geq m\right]$. Further, note that Y_t changes at most by 1 in each step and has a symmetric probability to increase or decrease. We proceed to lower bound $\mathbb{P}\left[T_Y \geq m\right]$ in two steps. Firstly, we show that, with positive probability, Y_t changes at least $\Omega\left(d^2\right)$ times to hit $|Y_t| \geq d$. Secondly, we argue that this takes at least m time steps.

Define the random set $S = \{t \in \mathbb{N}_{\geq 1} \mid Y_t \neq Y_{t-1}\}$ and let $N = |S|$. Note that N is almost surely finite, as at any time t the probability that the process increases $2d - Y_t$ times in a row is bounded below by $p^d/(1-q)^d$, after which $Y_t = d$ and the process remains constant. Let $\tau_1 < \tau_2 < \cdots < \tau_N$ denote the ordered elements of S and set

$\tau_0 = 0$. Note that, as long as $|Y_t| < d$, there is a non-zero probability for Y_t to change in each step. Thus, we have $Y_{\tau_N} = d$ and $T_Y = \tau_N$ with probability 1. Now, we define the process $Y_i' = Y_{\tau_i \wedge N}$ where $\tau_i \wedge N$ is shorthand for the minimum of τ_i and N. Y_i' is a martingale with respect to the natural filtration over Y_t, as it increases or decreases with probability $\frac{1}{2}$ at each step if $\tau_i < N$ and remains constant otherwise. Using Azuma's inequality and the fact that the step size of Y_i' is upper bounded by 1, we obtain for $i \leq \frac{d^2}{2\ln(10)}$ that $\mathbb{P}\left[Y_i' \geq d\right] \leq \mathrm{e}^{-\frac{d^2}{2i}} \leq 1/10$. As almost surely $|Y_N'| = |Y_{\tau_N}| = d$, this yields $\mathbb{P}\left[N \leq \frac{d^2}{2\ln(10)}\right] \leq 1/10$, meaning that Y_t requires to change at least $\frac{d^2}{2\ln(10)}$ times with probability at least $\frac{9}{10}$. It remains to bound the time that it takes Y_t to change this often. To this end, note that $\tau_N = \sum_{i=1}^{N}(\tau_i - \tau_{i-1})$. Recall that $T_Y = \tau_N$ almost surely. Setting $b = \frac{d^2}{2\ln(10)}$ we obtain

$$
\begin{aligned}
\mathbb{P}\left[T_Y \geq m\right] = \mathbb{P}&\left[\sum_{i=1}^{N} \tau_i - \tau_{i-1} \geq m\right] \\
\geq \mathbb{P}&\left[\sum_{i=1}^{N} \tau_i - \tau_{i-1} \geq m \text{ and } N \geq b\right] \\
= \mathbb{P}&\left[\sum_{i=1}^{N} \tau_i - \tau_{i-1} \geq m \;\middle|\; N \geq b\right] \cdot \mathbb{P}\left[N \geq b\right] \\
\geq \mathbb{P}&\left[\sum_{i=1}^{b} \tau_i - \tau_{i-1} \geq m\right] \cdot \frac{9}{10}.
\end{aligned}
$$

Next, observe that each of the differences $\tau_i - \tau_{i-1}$ independently follows a geometric distribution with success probability $\frac{p}{(1-q)}$. Using Theorem 5 with $n = b = \frac{d^2}{2\ln(10)}$, $\mu = \frac{d^2(1-q)}{2\ln(10)p}$ and $\delta = \frac{1}{2}$ we have

$$
\mathbb{P}\left[\sum_{i=1}^{b} \tau_i - \tau_{i-1} < m\right] \leq \mathbb{P}\left[\sum_{i=1}^{b} \tau_i - \tau_{i-1} < \frac{d^2(q-1)}{4ln(10)p}\right] \leq \frac{1}{10}
$$

for $d \geq 4\ln(10)$.

Therefore, we get $\mathbb{P}\left[T' \geq m \mid T_j \geq m\right] \geq \mathbb{P}\left[T_Y \geq m\right] \geq \frac{9}{10} \cdot \frac{9}{10} > \frac{4}{5}$ and $\mathbb{P}\left[T' \geq m\right] \geq \frac{4}{5} \cdot \frac{4}{5} = \frac{16}{25}$, which concludes the proof. □

3 Vertex Cover on Paths

Although the worst-case expected runtime of the (1+1) EA is known to be exponential in n, [5, Theorem 5], far better runtime bounds can be obtained on simple graph instances. In this section we take G to be a path of length $|V| = n$, with the i-th bit of X corresponding to the i-th vertex of G, numbering sequentially along the path. By considering path instances, we can precisely analyse the runtime of the (1+1) EA, and then transfer that analysis to the Balanced (1+1) EA, showing that in expectation it performs better by a linear factor. It is reasonable to expect that the Balanced (1+1) EA also improves on the runtime of the (1+1) EA on other graph instances, but rigorous analysis of those will require more general tools than those we develop below.

To achieve the lower bound for either algorithm, we require two additional assumptions, the first of which is defined in more detail in Section 3.2 and supported by experimental evidence in Fig. 1.

(A) With probability $\geq 1/2$, there exists a $t \geq 0$ such that X_t is feasible and contains a connected subpath of vertices which are not in their optimal state with size $\Theta(n)$.

(B) $|V| = n$ is odd.

Equivalently to (A), we could assume that X has an initial distribution which is supported by the subset of $\{0,1\}^n$ containing only feasible bit strings. Assumption (A) simply ensures that X hits this subset with probability $1/2$, and since X is a strong Markov process, restarting at this hitting time preserves the lower bound. Assumption (B) is needed for particular arguments in Section 3.2, but seems unlikely to be necessary.

The two main theorems of this section are as follows.

Theorem 6. *Let $P = (V, E)$ be a path of length $n = |V|$. In expectation, the (1+1) EA samples an optimal solution in $\mathcal{O}(n^4)$ iterations, and in $\Omega(n^4)$ iterations if assumptions (A) and (B) hold true.*

Theorem 7. *Let $P = (V, E)$ be a path of length $n = |V|$. In expectation, the Balanced (1+1) EA samples an optimal solution in $\mathcal{O}(n^3)$ iterations, and in $\Omega(n^3)$ iterations if assumptions (A) and (B) hold true.*

3.1 Upper Bound Running Time

By Lemma 2, we know that the (1+1) EA finds a feasible solution in expectation in $\mathcal{O}(n \log(n))$ iterations, given any starting distribution, including in particular the uniform distribution. Once a feasible solution is found, all states in following iterations are feasible as well. Hence our analysis can focus on the expected runtime given a feasible initial state. If n is odd, there exists a unique optimal vertex cover of size $\mathrm{OPT} = (n-1)/2$, and if n is even, there are multiple optimal covers of size $\mathrm{OPT} = n/2$. We define the *level* of X at time t to be $|X_t|_1 - \mathrm{OPT}$, to parametrize the deviation of X from the optimum. Our strategy is to compute the expected number of iterations that X spends at any level ℓ, and then sum over all levels to obtain an upper bound for total runtime. For this purpose, we define Y_ℓ to be the total number of iterations spent at level $\ell \geq 0$, i.e. $Y_\ell = \sum_{t=0}^{Y} \mathbf{1}\{|X_t|_1 = \ell + \mathrm{OPT}\}$, and let $Y = \sum_{i=1}^{\lceil n/2 \rceil} Y_i$ be the first hitting time of level $\ell = 0$.

The main work of this section is in proving the following bounds on the expected time spent at level ℓ - which may be zero - by the (1+1) EA and Balanced (1+1) EA respectively.

Lemma 8. *The expected number of iterations that the (1+1) EA spends at level ℓ is of order $O(n^4/\ell^2)$.*

Lemma 9. *The expected number of iterations that the Balanced (1+1) EA spends at level ℓ is of order $O(n^3/\ell^2)$.*

With these results in hand, we prove the following upper bound on the number of expected iterations before the (1+1) EA samples an optimal solution for the first time.

Theorem 10. *Let $P = (V, E)$ be a path of length $n = |V|$. Then the expected number of iterations that the (1+1) EA needs to sample an optimal solution is of order $\mathcal{O}(n^4)$, and the expected number of iterations that the Balanced (1+1) EA needs to sample an optimal solution is of order $\mathcal{O}(n^3)$.*

Proof. In either case, by Lemma 2 X becomes feasible for the vertex cover after an expected $\mathcal{O}(n \log(n))$ iterations.

At this point the algorithm has an expected remaining runtime of

$$\mathbb{E}\left[Y\right] = \mathbb{E}\left[\sum_{\ell=1}^{\lceil \frac{n}{2} \rceil} Y_\ell\right] = \sum_{\ell=1}^{\lceil \frac{n}{2} \rceil} \mathbb{E}\left[Y_\ell\right]$$

which according to Lemma 8 is of order $\mathcal{O}(n^4)$ for the (1+1) EA, and by Lemma 9 is $\mathcal{O}(n^3)$ for the Balanced (1+1) EA, since in either case $\sum_{\ell=1}^{\lceil \frac{n}{2} \rceil} 1/\ell^2$ is bounded above by a constant independent of n. $\qquad\square$

A Random Walk Coupling: The proofs of Lemmas 8 and 9 make use of a coupling argument, given below, between X and an independent random walk on a path of length $\Theta(n)$. To take advantage of that coupling we prove the following Lemma; an alternative proof of this result using drift theory can be found in Göbel et al. [8, Theorem 13].

Lemma 11. *Let $q \in (0, 1/2]$, and let Z be the symmetric Markov chain on $\{0, 1, \ldots, d\}$ with transitions $P(r, s) = \mathbb{P}\left[Z_{t+1} = r | Z_t = s\right]$ defined by*

$$P(s, s-1) = \begin{cases} q & \text{if } 1 \leq s \leq d-1, \\ 2q & \text{if } s = d, \\ 0 & \text{if } s = 0, \end{cases} \qquad P(s, s+1) = \begin{cases} q & \text{if } 1 \leq s \leq d-1, \\ 0 & \text{if } s \in \{0, d\}, \end{cases}$$

and $P(s, s) = 1 - P(s, s-1) - P(s, s+1)$. This is the symmetric random walk with reflecting barrier at d and absorbing barrier at 0, with an additional probability to stay put at any vertex. Suppose $Z_0 = d$ and $T = \min\{t \geq 0 : Z_t = 0\}$ be the first time that Z hits 0. Then $\mathbb{E}\left[T\right] \leq d^2/2q$.

Proof. Let Y be the symmetric random walk on $\{0, 1, \ldots, 2d\}$ started at $Y_0 = d$, with absorbing barriers 0 and $2d$. It is well known - see for example Cox and Miller [2, Exercise 2.12] - that the expected hitting time of $\{0, 2d\}$ by Y is $\mathbb{E}\left[\min\{t \geq 0 : Y_t \in \{0, 2d\}\right] = d^2$.

If instead Y moves left or right with probability q and stays put with probability $1 - 2q$ then Y sits for an independent geometrically distributed time S at each vertex before moving, so Wald's equation gives $\mathbb{E}\left[\min\{t \geq 0 : Y_t \in \{0, 2d\}\right] = d^2\mathbb{E}\left[S\right] = d^2/2q$. Since Z has the same law as $d - |Y|$, it has expected runtime $\mathbb{E}\left[T\right] = d^2/2q$. $\qquad\square$
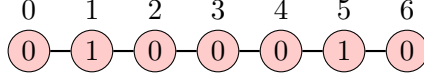
The unique optimal solution on the odd-length path - which will be the terminal value of X - has alternate vertices selected, starting with the second and ending with the $(n-1)$-th. Similarly, both optima of the even-length path are similarly alternating, apart from a possible pair of adjacent ones. Let us suppose that X is at level ℓ and is feasible. We shall begin by supposing that X is irreducible, and bound the time it takes to become reducible. Since X is feasible, no two neighboring vertices are unselected, and since X is irreducible, there is no subpath of consecutive ones of length more than 2. An example for $n = 11$ is given below.
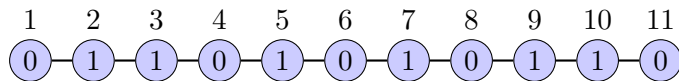
Here X is at level $\ell = 1$, although in this example most nodes (from 3 to 9 inclusive) are not in their optimal state. Now we can define a dual process $\tilde{X}$ on the path of length $2 + (n-1)/2$, which is coupled to X, in the following way: for $i \in \{1, \ldots, (n-1)/2\}$,

$$\tilde{X}(i) = X(2i) \cdot \mathbb{1}\{\text{at least one of } X(2i-1) \text{ and } X(2i+1) \text{ is a } 1\}$$
$$= X(2i)(X(2i-1) + X(2i+1) - X(2i-1)X(2i+1)),$$

and $\tilde{X}(0) = X(1)$, $\tilde{X}(n+1/2) = X(n)$. Then the dual state to the one above is

$$\begin{array}{ccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 \\ \boxed{0} - \boxed{1} - \boxed{0} - \boxed{0} - \boxed{0} - \boxed{1} - \boxed{0} \end{array}$$

If X is the (1+1) EA, then conditional on X flipping at most two bits at a time,[1] the 'particles' at positions 1 and 5 move independently like symmetric random walks (with reflective boundaries) until the time that they are adjacent, at which point X is reducible. If X is the Balanced (1+1) EA, then the same random walk behavior is observed, but with differing transition probabilities. In this example there are only two such particles, since X is at level $\ell = 1$, but in general there will be 2ℓ. As long as X is not optimal, we define the minimum distance (measured by number of zero nodes) between any two particles in $\tilde{X}$ to be the longest subpath of consecutive zeroes (corresponding to subpaths of alternating bits in X). This is itself a stochastic process on $\{0, \ldots, d\}$, where $d \leq 2\ell$. For ease of notation we denote this process by H.

The (1+1) EA - Proof of Lemma 8: Take X to be the (1+1) EA, and for the moment, let us condition on the event that the algorithm only flips 2 bits at every iteration. Under this conditioning, the process H can only make steps of size one in either direction. It need not be symmetric, but decreases (that is, $H_{t+1} = H_{t-1}$) with probability $\geq 2p^2(1-p)^{n-2}$, and increases with probability $\leq 2p^2(1-p)^{n-2}$. Thus the expected number of iterations taken for H to hit zero is less than the expected hitting time T of zero by the independent process Z on $\{0, 1, \ldots, d\}$ satisfying the conditions of Lemma 11 with $q = 2p^2(1-p)^{n-2}$. Recalling that $p = 1/n$ and $d \leq n/2\ell$, the following corollary is immediate.

Corollary 12. $\mathbb{E}[T] \leq d^2/4p^2(1-p)^{n-2} \leq n^4/\ell^2$.

Thus we have an upper bound on the expected time for X to become reducible, conditional on the (1+1) EA flipping no more than 2 bits at every step. Now we want to remove that conditioning. To do so, we modify the transitions of Z so that it returns to position d with probability $4p^4$ at every step; this is an upper bound for the probability that $(H_{t+1} - H_t) \geq 2$, that is, that H increases by 2 or more in a single step. Let us denote this modified version of Z by $\tilde{Z}$. The transitions of $\tilde{Z}$ are

$$\tilde{Z}_{t+1} = \begin{cases} \tilde{Z}_t + 1 & \text{with probability } 2p^2(1-p)^{n-2}, \\ \tilde{Z}_t - 1 & \text{with probability } 2p^2(1-p)^{n-2}, \\ d & \text{with probability } 4p^4, \\ \tilde{Z}_t & \text{otherwise}, \end{cases}$$

when $\tilde{Z}_t \in \{1, \ldots, d-1\}$, and similar in the obvious way for $\tilde{Z}_t = d$. In essence, $\tilde{Z}$ is 'restarted' every time H increases by 2 or more, and so the expectation of the hitting

---

[1] Since X is feasible and irreducible, it cannot flip only one bit in a step.

time $\tilde{T} = \min\{t \geq 0 : \tilde{Z}_t = 0\}$ is greater than the expected time it takes for H to hit zero - that is, for X to become reducible.

The question to answer now is how many times $\tilde{Z}$ will have to restart; we shall show that it is constant in expectation. Each 'run' of $\tilde{Z}$, starting at position d and ending when it hits zero or is restarted, is independent from the others. They are also all indentically distributed, with the same law as T from Corollary 12. The number of restarts is therefore a geometric random variable, say R, and to calculate its mean we need only calculate the probability of success, that is, of $\tilde{Z}$ hitting zero before it is sent back to d.

The 'restart time' of $\tilde{Z}$ is itself a geometric random variable S, independent of the position of $\tilde{Z}$, with success probability $4p^4$. Therefore, for any $r \in \mathbb{N}^+$,

$$\mathbb{P}\left[T < S\right] \geq \mathbb{P}\left[T \leq r < S\right] = \mathbb{P}\left[T \leq r\right](1 - 4p^4)^r.$$

We saw in Corollary 12 that $\mathbb{E}\left[T\right] \leq n^4/\ell^2$, and it follows from Markov's inequality that $\mathbb{P}\left[T \leq 2\mathbb{E}\left[T\right]\right] \geq 1/2$. Therefore, taking $r = 2n^4/\ell^2$, we have that

$$\mathbb{P}\left[T < S\right] \geq \frac{(1 - 4n^{-4})^{2n^4/\ell^2}}{2} \to c \in (0,1) \quad \text{as } n \to \infty.$$

Then Wald's equation yields that

$$\mathbb{E}\left[\tilde{T}\right] \leq \mathbb{E}\left[T\right]\mathbb{E}\left[R\right] = \frac{n^4}{\ell^2\mathbb{P}\left[T \geq S\right]} = \frac{n^4}{(1-c)\ell^2}.$$

This gives an upper bound for the expected number of iterations it takes for H to hit zero - that is, X to become reducible - given that X started in level ℓ.

It remains to note that when X is in a reducible state, it has a positive probability of taking a step to an irreducible state but remaining at level ℓ. However, the probability of this occurring before X is reduced is bounded above by a constant,[2] and so an argument similar to above using Wald's equation and restarting gives that $\mathbb{E}\left[Y_\ell\right] \leq kn^4/\ell^2$ for some constant $k > 1$; this proves Lemma 8.

The Balanced (1+1) EA - Proof of Lemma 9: The same arguments as for the (1+1) EA hold for the Balanced (1+1) EA, with the difference that H decreases with larger probability. More precisely, with probability $1/2$ the algorithm chooses a bit uniformly at random, and attempts to swap it with a neighbor of a different value. There are 4 balanced flips out of at most $2n$ which cause H to decrease; conditioning the algorithm on flipping no more than 2 bits in a single step, this leads to a decrease probability of

$$q \geq \frac{1}{2}\left(\frac{4}{2n} + 4p^2(1-p)^{(n-2)}\right),$$

which again via Lemma 11 leads to a runtime bound of $d^2/2q \leq n^3/\ell^2$. The same restarting arguments as above then yield the result of Lemma 9.

3.2 Lower Bound Running Time

In this section, we prove the lower bound parts of Theorems 6 and 7. That is, given a suitably chosen starting configuration, the (1+1) EA and Balanced (1+1) EA require at least $\Omega\left(n^4\right)$ iterations to find the minimum vertex cover on a path of odd length n.

---

[2]This requires two simultaneous bit flips, and roughly has probability ℓ/n^2, compared to $1/n$ for simply reducing. Markov's inequality gives the constant.

Again we note that such instances G only have one minimal vertex cover, namely the set of vertices at even positions, which simplifies our analysis.

We start by introducing some additional notation and terminology that will come in handy for stating and proving the main result of the section. Let P be a path of odd length n with vertices $v_1, \ldots, v_n$ and let $X \in \{0,1\}^n$ be a bit string $x_1, \ldots, x_n$, representing a solution candidate for the vertex cover problem.

Similarly to Section 3.1, a major role is played in these proofs by connected subpaths of X in which all vertices are not in their optimal state, but they appear in a different form, and so we introduce some new notation. Let $B(X) \subseteq [n]^2$ denote the set of all tuples (i,j) with $i \leq j$ that corresponds to endpoints of subpath $\{v_i, \ldots, v_j\} \subseteq V$ such that all vertices v_k for $i \leq k \leq j$ are not in their optimal state. That is, every x_k for $i \leq k \leq j$ has value 0 if and only if k is even. We shall call such subpaths 'bad paths'. Then X represents a minimum vertex cover if and only if X is feasible and $B(X) = \emptyset$. Assumption (A) is equivalent to assuming that for initial configuration X_0, $B(X_0)$ contains exactly on tuple (i,j) with $j - i \in \Theta(n)$.

The main results of this section are the following two statements.

Theorem 13. *Let $P = (V, E)$ be a path of odd length $n = |V|$. Assume a feasible initial configuration X_0 with $B(X_0) = \{(i,j)\}$ such that $j - i \in \Theta(n)$. With constant positive probability the (1+1) EA requires at least $\Omega\left(n^4\right)$ iterations to find the optimal solution. Consequently, the expected number of iterations is in $\Omega(n^4)$.*

Theorem 14. *Let $P = (V, E)$ be a path of odd length $n = |V|$. Assume a feasible initial configuration X_0 with $B(X_0) = \{(i,j)\}$ such that $j - i \in \Theta(n)$. With constant positive probability the Balanced (1+1) EA requires at least $\Omega\left(n^3\right)$ iterations to find the optimal solution. Consequently, the expected number of iterations is in $\Omega(n^3)$.*

The (1+1) EA - Proof of Theorem 13: As long as $B(X_t) \neq \emptyset$, let $l(X_t)$ and $r(X_t)$ denote the left and right endpoint of the bad path in P with respect to X_t (i.e., $B(X_t) = \{(l(X_t), r(X_t))\}$). Let $m = l(X_0) + (r(X_0) - l(X_0))/2$ and set $l(X_t) = r(X_t) = m$ for all $t \in \mathbb{N}$ with $B(X_t) = \emptyset$. Define $b(X_t) = r(X_t) - l(X_t)$ and observe that $T = \inf\{t \in \mathbb{N} \mid b(X_t) \leq 0\}$ is a lower bound on the required number of iterations to reach the optimum.

We proceed by showing that $T \in \Omega(n^4)$ with positive probability. To this end, define two new stopping times $T_l = \inf\{t \in \mathbb{N} \mid l(X_t) \geq m - 4\}$ and $T_r = \inf\{t \in \mathbb{N} \mid r(X_t) \leq m + 4\}$. Observe that $T \geq \min\{T_l, T_r\}$. Thus, lower bounding T_l and T_r with positive probability suffices to obtain the desired result for T.

The following lemma characterizes the transition probabilities of $l(X_t)$ and $r(X_t)$. Applying it in combination with Lemma 3, we are able to prove Theorem 13.

Lemma 15. *For all $t < \min\{T_l, T_r\}$ it holds that*

- $\mathbb{P}\left[|l(X_{t+1}) - l(X_t)| > 2 \mid X_t\right] \leq 2/n^4$,

- $\mathbb{P}\left[l(X_{t+1}) = l(X_t) + 2 \mid X_t\right] \leq 1/n^2$ *and*

- $\mathbb{P}\left[l(X_{t+1}) = l(X_t) + 2 \mid X_t\right] = \mathbb{P}\left[l(X_{t+1}) = l(X_t) - 2 \mid X_t\right]$ *whenever $l(X_t) \geq 3$*

and analogously

- $\mathbb{P}\left[|r(X_{t+1}) - r(X_t)| > 2 \mid X_t\right] \leq 2/n^4$,

- $\mathbb{P}\left[r(X_{t+1}) = r(X_t) - 2 \mid X_t\right] \leq 1/n^2$ *and*

- $\mathbb{P}\left[r(X_{t+1}) = r(X_t) + 2 \mid X_t\right] = \mathbb{P}\left[r(X_{t+1}) = r(X_t) - 2 \mid X_t\right]$ *whenever* $r(X_t) \leq n - 2$.

Proof. First note that we assume X_0 to be feasible. Thus, X_t is also feasible for all $t \in \mathbb{N}$. Moreover, in the initial configuration X_0 we have exactly one bad path from position $l(X_0)$ to $r(X_0)$. This means, for every $k \in [n] \cap [l(X_0), r(X_0)]$, the bit at position k is 1 in X_0 if and only if k is odd. On the other hand, for all $k \in [n] \setminus [l(X_0), r(X_0)]$, the bit at position k is 1 in X_0 if and only if k is even. Thus, the only bit flips that the (1+1) EA can do is

- flip an even length bit sequence either starting at $l(X_0)$ or ending at $l(X_0) - 1$,

- flip an even length bit sequence either ending at $r(X_0)$ or starting at $r(X_0) + 1$

- or do both of the above at the same time.

As long as this leads to a state X_1 with $B(X_1) \neq \emptyset$, these properties remain in place. Thus, inductively, the same holds for all $t < \min\{T_l, T_r\}$.

Given this, we prove the claimed transition probabilities for $l(X_t)$. The transitions of $r(X_t)$ are analyzed analogously. Note that for $t < \min\{T_l, T_r\}$ we have $r(X_t) - l(X_t) \geq 8$. Thus, changing $l(X_t)$ by more than 2 requires at least flipping either the bits $l(X_t) - 4$ to $l(X_t) - 1$ or $l(X_t)$ to $l(X_t) + 3$, each of which happens with probability at most $\frac{1}{n^4}$. By union bound we get $P[|l(X_{t+1}) - l(X_t)| > 2|X_t] \leq 2/n^4$. Similarly, in order to increase $l(X_t)$ by exactly 2, the bits $l(X_t)$ and $l(X_t) + 1$ need to be flipped at once. Thus, it holds that $P[l(X_{t+1}) = l(X_t) + 2|X_t] \leq 1/n^2$. Finally, consider X_t such that $l(X_t) \geq 3$ and assume a set of bits $S \subseteq [n]$ is flipped, such that $l(X_t)$ changes by exactly 2. If this succeeds, either $\{l(X_t) - 2, l(X_t) - 1\} \subseteq S$ or $\{l(X_t), l(X_t) + 1\} \subseteq S$ but not both. Assume $\{l(X_t) - 2, l(X_t) - 1\} \subseteq S$, meaning that $l(X_t)$ decreases by 2, and construct the set $S' = (S \setminus \{l(X_t) - 2, l(X_t) - 1\}) \cup \{l(X_t), l(X_t) + 1\}$. Note that, if flipping the bit set S succeeds, flipping S' would also succeed as both result in the same fitness value. Moreover, we have $|S| = |S'|$, which means that both operations have the same probability. Finally, observe that the construction of S' from S is reversible. Thus, the overall probability of decreasing $l(X_t)$ by 2 is the same as increasing it by 2. $\qquad\square$

Using the characterization of the transition probabilities given in Lemma 15 we use Lemma 3 to prove Theorem 13.

Proof of Theorem 13. Let $l(X_t), r(X_t), b(X_t), m, T, T_l$ and T_r be defined as above. As discussed earlier, T is a lower bound on the number of iterations that the (1+1) EA requires to get to the optimum. Furthermore, it holds that $T \geq \min\{T_l, T_r\}$. Thus, for every $t \in \mathbb{N}$, it holds that

$$\mathbb{P}\left[T \geq t\right] \geq 1 - \mathbb{P}\left[T_l < t \text{ or } T_r < t\right] \geq 1 - \mathbb{P}\left[T_l < t\right] - \mathbb{P}\left[T_r < t\right].$$

We proceed by using Lemma 3 to prove that, if $r(X_0) - l(X_0) \in \Theta(n)$, then there is some $\tau \in \Theta(n^4)$ such that $\mathbb{P}\left[T_l \geq \tau\right] \geq \frac{16}{25}$ and $\mathbb{P}\left[T_r \geq \tau\right] \geq \frac{16}{25}$.

To this end, we consider the process $Z_t = \frac{l(X_t) - l(X_0)}{2} \cdot \mathbb{1}_{\{l(X_t) \geq l(X_0)\}}$ and set $d = \left\lfloor \frac{r(X_0) - l(X_0)}{4} \right\rfloor - 2$. Note that $d \in \Theta(n)$, which implies $d \geq 4\ln(10)$ for n sufficiently large. By Lemma 15, we know that Z_t satisfies the requirements of Lemma 3 for $p = \frac{1}{n^2}$ and $q = \frac{2}{n^4}$. Thus, for $\tau = \min\left\{\frac{2}{5}n^4, \frac{d^2 n^2}{8\ln(10)}\left(1 - \frac{2}{n^4}\right)\right\} \in \Theta(n^4)$ and $T_d = \inf\{t \in \mathbb{N} \mid Z_t \geq d\}$,

Lemma 3 yields $\mathbb{P}[T_d \geq \tau] \geq \frac{16}{25}$. Now, note that for our choice of Z_t and d, $l(X_t) \geq m - 4$ implies $Z_t \geq d$. Therefore, we have $\mathbb{P}[T_l \geq \tau] \geq \frac{16}{25}$.

For $r(X_t)$, we can argue analogously that $\mathbb{P}[T_r \geq \tau] \geq \frac{16}{25}$ for the same $d \in \Theta(n)$ and $\tau \in \Theta(n^4)$ as above and using the process $Z_t = \frac{r(X_0) - r(X_t)}{2} \cdot \mathbb{1}_{\{r(X_t) \leq r(X_0)\}}$. Thus, we obtain $\mathbb{P}[T \geq \tau] \geq \frac{7}{25}$, which proves the first part of the statement. For the second part, note that by $T \geq 0$ and the law of total expectation $\mathbb{E}[T] \geq \tau \cdot \frac{7}{25} \in \Theta(n^4)$. By monotonicity of the expectation, this carries over to the expected number of iterations. $\quad\square$

The Balanced (1+1) EA - Proof of Theorem 14: The proof of Theorem 14 works analogously to that of Theorem 13. The only substantial difference is in the transition probabilities in Lemma 15. In particular, the Balanced (1+1) EA satisfies

$$P[l(X_{t+1}) = l(X_t) + 2 | X_t] = 1/2(1/n^2 + 1/n), \text{ and}$$
$$P[r(X_{t+1}) = r(X_t) - 2 | X_t] = 1/2(1/n^2 + 1/n).$$

Clearly $1/2(1/n^2 + 1/n) \leq 1/n$, which gives a lower bound of $\Omega(n^3)$ for the (1+1) EA.

Experimental Evidence for Long Bad Paths: So far we have argued that given a feasible initial configuration with exactly one bad path of length in $\Theta(n)$, the (1+1) EA requires with constant probability $\Omega(n^4)$ iterations to reach the optimum. Analogously, we have shown that the Balanced (1+1) EA requires $\Omega(n^3)$ steps.

We are now going to present experimental evidence that, starting from a uniformly chosen configuration, the (1+1) EA and the Balanced (1+1) EA both reach a state $X \in \{0,1\}^n$ with $f(X) = \mathrm{OPT} + 1$ with high probability. Once such a state is reached, our experiments suggest that the bad path of X has linear length with at least constant probability.

We did experiments on paths of odd length $n \in \{51 + 10 \cdot k \mid k \in \mathbb{N} \cap [0, 15]\}$ with 100 runs for each n and each algorithm. In every iteration, the algorithm starts with a uniformly random initial configuration. Once a state X with $f(X) = \mathrm{OPT} + 1$ is reached, we record the length of the bad path in X, divided by n. Since X has only one vertex more than the optimal cover, there will be only one bad path, and all other vertices will have their optimal value.

In the case that the algorithm never reaches such a state (i.e. it jumps directly from fitness level $\ell > 1$ to the optimum), the length of the bad path is set to zero, which is a worst-case. Interestingly, this behaviour wasn't observed in even a single simulation, which suggests that such a jump is difficult for either algorithm to perform.

The results are given in Fig. 1. Each box indicates span from the first to the third quartile (i.e. the inner 50% of the data) and the whiskers indicate minimum and maximum values that were observed. The result for both algorithms look very alike. Apart from minor deviations, most of the iterations resulted in similar relative lengths of the bad paths, independent of n. The medians seem to fluctuate around a value of $1/3$ and never drop below $1/5$, which behavior is consistent throughout all tested values of n.

Our experiments suggest that, with constant probability, both algorithms enter the last fitness level with a bad path of linear length. Therefore, when starting from a uniformly random initial configuration, with constant probability a state is reached to which Theorems 13 and 14 apply. This suggests that, even when starting from a uniform random initialization, the expected number of iterations to find the optimum are in $\Omega(n^4)$ for the (1+1) EA and in $\Omega(n^3)$ for the Balanced (1+1) EA, due to the appearance of such linear bad paths.
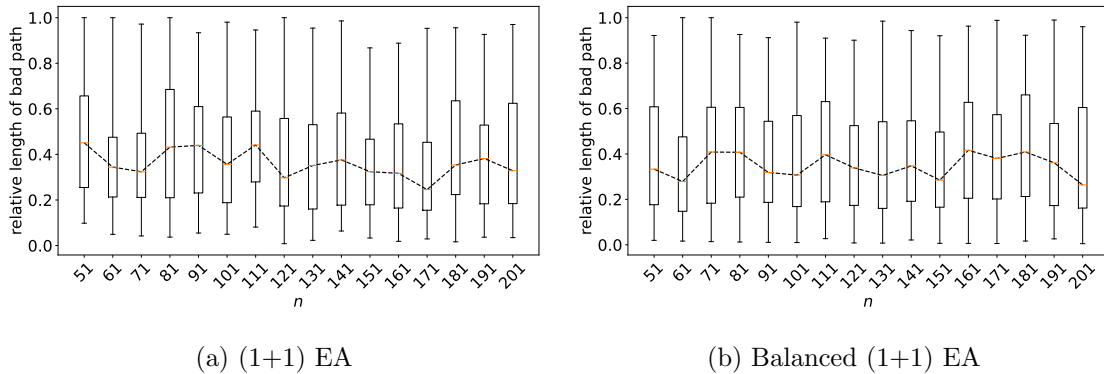
(a) (1+1) EA   (b) Balanced (1+1) EA

Figure 1: Relative length of the bad path at last fitness level for (1+1) EA and Balanced (1+1) EA

4 Vertex Cover on Complete Bipartite Graphs

A common instance class of interest when it comes to evolutionary algorithms for the vertex cover problem are complete bipartite graphs. In this section, we investigate basic properties of the Balanced (1+1) EA in this setting. Specifically, for $L, R \in \mathbb{N} \setminus \{0\}$, let $K_{L,R} = (V, E)$ be a graph, such that V can be partitioned into sets V_L and V_R with $|V_L| = L$ and $|V_R| = R$, and $E = \{\{v, u\} \mid v \in V_L \text{ and } u \in V_R\}$. We call V_L the left partition and V_R the right partition. A bit string $X \in \{0, 1\}^{L+R}$ represents a feasible solution for the vertex cover problem on $K_{L,R}$ if and only if $V_L \subseteq X$ or $V_R \subseteq X$.

Now assume $R/L = c$ for some $c > 1$, and note that in this case $X = V_L$ is the unique optimum of the vertex cover problem on $K_{L,R}$. On the other hand, the state $X = V_R$ represents a local optimum with respect to the fitness function f. Once this local optimum is reached, X must flip at least $2L$ bits in one step to escape. This causes the (1+1) EA to need an exponential number of iterations in expectation, and local search heuristics like Random Local Search (RLS) to fail completely. The probability of reaching the optimal solution in polynomial time is usually closely related to the probability of never getting too close to that local optimum, which makes complete bipartite graphs to an interesting instance class to study.

In this section, we specifically consider the following setting. We assume the ratio $c := R/L > 0$ to be fixed and investigate the probability for finding the optimum in polynomial time asymptotically in L. Our main result is that the probability of not finding the optimum in polynomial time decays exponentially in L as long as $c > 2$. As Table 1 shows, this is quite different from the RLS, where the probability is constant for any fixed c, and the normal (1+1) EA, where it only decays polynomialy in L. Note all asymptotic behavior here is stated in terms of L. This can be easily translated to asymptotics in terms of the absolute number of vertices $n = L + R = (c + 1)L$. In particular, as long as the ratio c is considered a fixed constant, this only changes linear factors, leading to the results that are presented in the abstract.

The main technical result that we use to prove this bound on the number of iterations can informally be stated as follows. As long as V_R is more than twice as big as V_L, the probability of selecting at least $R - L$ vertices in V_R before selecting all vertices in V_L decays exponentially in the size of L.

---

[3]In fact it is part of the proof of Theorem 5 in [5], but not part of the main statement.

15

RLS	(1+1) EA	Balanced (1+1) EA
$\leq \frac{c}{c+1}$	$\leq 1 - \frac{1}{\text{poly}(L)}$	$\geq 1 - 2^{-\Omega(L^{\xi})}$ for all $0 \leq \xi < 1$ if $c > 2$
[5, Theorem 4]	[5, Theorem 5] [3]	see Theorem 17

Table 1: Probability of finding the minimum vertex cover on $\boldsymbol{K_{L,R}}$ with $\boldsymbol{c := R/L > 1}$ in time poly $(\boldsymbol{L})$.

Lemma 16. *Let $K_{L,R}$ be a complete bipartite graph with vertex partitions V_L of size L and V_R of size R, and set $c := R/L$. Let $X_t \in \{0,1\}^{L+R}$ be the sequence of states of the Balanced (1+1) EA with fitness function f on $K_{L,R}$ and let X_0 be chosen uniformly at random. Define $T_L = \inf\{t \in \mathbb{N} \mid |V_L \cap X_t| = L\}$ and let $T'_R = \inf\{t \in \mathbb{N} \mid |V_R \cap X_t| \geq R - L\}$. If $c > 2$, then, for all positive constants $\xi < 1$ and L sufficiently large, it holds that $\mathbb{P}[T_L \geq T'_R] \leq 2^{-\Omega(L^{\xi})}$.*

The following probabilistic bound on the number of iterations for finding the minimum vertex cover on $K_{L,R}$ can be derived from Lemma 16.

Theorem 17. *Let $K_{L,R}$ be a complete bipartite graph with vertex partitions V_L of size L and V_R of size R, and set $c := R/L$. If $c > 2$ then, for all positive constants $\xi < 1$, it holds that the Balanced (1+1) EA with fitness function f, starting from a uniformly random initial configuration, requires at most $\mathcal{O}\left((c+1)L^2 \log((c+1)L)\right)$ iterations to find the minimum vertex cover of $K_{L,R}$ with probability at least $1 - 2^{-\Omega(L^{\xi})}$.*

Proof. We prove this statement in two steps. At first, we argue that, with sufficiently high probability, we reach a state where all vertices in V_L are selected and more than L vertices from V_R are not selected in at most $\mathcal{O}\left((c+1)L^2 \log((c+1)L)\right)$ iterations. In the second step, we argue that, once such a state is reached, we go to the optimum in at most $\mathcal{O}\left((c+1)L^2 \log((c+1)L)\right)$ iterations with high probability.

For the first part, let X_t denote the states of the Balanced (1+1) EA and let T_L and T'_R be defined as in Lemma 16. Formally, we want to lower bound $\mathbb{P}[T_L \leq \tau$ and $T_L < T'_R]$ for some suitably chosen $\tau \in \mathcal{O}\left((c+1)L^2 \log((c+1)L)\right)$. Let $T = \inf\{t \in \mathbb{N} \mid X_t \text{ is feasible}\}$ and observe that, if $T_L < T'_R$, then $T = T_L$. Thus, we have

$$\mathbb{P}\left[T_L \leq \tau \text{ and } T_L < T'_R\right] = \mathbb{P}\left[T \leq \tau \text{ and } T_L < T'_R\right].$$

Using De Morgan's laws and union bound, we get

$$\mathbb{P}\left[T \leq \tau \text{ and } T_L < T'_R\right] = 1 - \mathbb{P}\left[T > \tau \text{ or } T_L \geq T'_R\right]$$
$$\geq 1 - \mathbb{P}\left[T > \tau\right] - \mathbb{P}\left[T_L \geq T'_R\right].$$

By Lemma 16, we have $\mathbb{P}[T_L \geq T'_R] \leq 2^{-\Omega(L^{\xi})}$. Setting $\tau = 4\mathrm{e}L(L+R)(\ln(L+R)+1/2)$, Lemma 2 yields $\mathbb{P}[T \geq \tau] \leq 2^{-L}$. Thus, we get that the Balanced (1+1) EA reaches a state with all vertices in V_L selected and more than L vertices in V_R are not selected in at most $\tau \in \mathcal{O}\left((c+1)L^2 \log((c+1)L)\right)$ iterations.

Now, assume we start in a state X_0 with $V_L \subseteq X_0$ and $|X_0 \cap V_R| < R - L$. Note that adding vertices from V_R requires removing at least as many vertices from V_L. However, to remain feasible, removing any vertex from V_L, requires adding all vertices from V_R that are not selected yet. As there are at lest $L + 1$ such vertices and at most L vertices in V_L that could be removed, the Balanced (1+1) EA can never add a vertex from V_R or

16

remove a vertex from V_L in such a state. Thus, all we need to bound is the required time to remove all vertices in $X_0 \cap V_R$ to also bound the time for reaching the optimum. First, observe that for each fixed vertex in $X_0 \cap V_R$, the probability of not being removed after $\tau' \in \mathbb{N}$ steps is at most $\left(1 - \frac{1}{L+R}\right)^{\tau'} \leq \mathrm{e}^{-\frac{\tau'}{L+R}} = \mathrm{e}^{-\frac{\tau'}{(c+1)L}}$. Thus, using union bound, the probability for not reaching the optimum after τ' steps is at most

$$|X_0 \cap V_R| \mathrm{e}^{-\frac{\tau'}{(c+1)L}} < (R - L)\mathrm{e}^{-\frac{\tau'}{(c+1)L}} = (c - 2)L\mathrm{e}^{-\frac{\tau'}{(c+1)L}}.$$

Choosing $\tau' = (c+1)L\left(\ln((c-2)L) + L\right) \in o\left((c+1)L^2\ln((c+1)L)\right)$, this probability is upper bounded by e^{-L}.

Observe that, if the first phase needs τ iterations and the second phase needs τ' iterations, then the total number of iterations for finding the optimum is $\tau + \tau' \in \mathcal{O}\left((c+1)L^2\ln((c+1)L)\right)$. The probability that at least one of both phases fails is at most $2^{-\Omega(L^\xi)} + \mathrm{e}^{-L} \leq 2^{-\Omega(L^\xi)}$, which concludes the proof. $\qquad\square$

We proceed by proving Lemma 16. To simplify notation, we define the processes $L_t = V_L \cap X_t$ and $R_t = V_R \cap X_t$. A central step in our proof of Lemma 16 is to argue that T_R' is super polynomial with high probability. Once this is done, it remains to show that, with sufficiently high probability, all vertices in V_L are selected beforehand.

To prove the first part, we apply the following negative drift theorem.

Theorem 18 (Oliveto et. al [16])**.** *Let* $(Y_t)_{t\in\mathbb{N}}$ *be random variables over an interval I adapted to a filtration $(\mathcal{F}_t)_{t\in\mathbb{N}}$, and let $b \in I$ such that $Y_0 \geq b$. Suppose that there exists an interval $[a, b] \subseteq I$, two constants $\delta, \varepsilon > 0$, and, possibly depending on $d := b - a$, a function $r(d)$ satisfying $1 \leq r(d) \in o(d/\log(d))$ such that, for all $t \in N$, it holds that*

(a) $\mathbb{E}\left[(Y_{t+1} - Y_t) \cdot \mathbb{1}_{\{a < Y_t < b\}} \mid \mathcal{F}_t\right] \geq \varepsilon \cdot \mathbb{1}_{\{a < Y_t < b\}}$, *and*

(b) for all $j \in \mathbb{N}$, it holds that

$$\mathbb{P}\left[|Y_{t+1} - Y_t| \cdot \mathbb{1}_{\{Y_t > a\}} \geq j \cdot \mathbb{1}_{\{Y_t > a\}} \mid \mathcal{F}_t\right] \leq \frac{r(d)}{(1 + \delta)^j}.$$

Then there is a constant $z > 0$ and a function $m(d) \in \Omega(d/r(d))$ such that, for the hitting time $T := \inf\{t \in \mathbb{N} \mid Y_t \leq a\}$, it holds that

$$\mathbb{P}\left[T \leq 2^{\frac{zd}{r(d)}}\right] = 2^{-m(d)}.$$

An obvious choice would be apply the above theorem directly to $|R_t|$. However, this would lead to rather complicated transition probabilities. Instead, we define the following slightly modified process $S_t \in \{0, 1\}^R$ based on X_t:

- Initially, $S_0 = R_0$.

- Whenever X_t attempts to flip bits for some set of vertices $A \subseteq V$, S_t flips all bits in $(V_R \cap A) \setminus S_t$ (in set notation, this is $S_{t+1} = S_t \cup (V_R \cap A)$).

- Whenever X_t attempts a balanced flip with initial vertex $v \in S_t$, the corresponding bit in S_t is set to 0 (i.e., $S_{t+1} = S_t \setminus \{v\}$).

The following relationship between R_t and S_t will come in handy throughout our analysis.

Lemma 19. *Let T_L, T'_R and S_t be defined as above. Then $R_t \subseteq S_t$ for all $t \leq \min\{T_L, T'_R\}$.*

Proof. We prove the statement via induction over t. For the base case $t = 0$, note that $S_0 = R_0$ by definition. Consider some $0 < t \leq \min\{T_L, T'_R\}$ and assume $R_{t-1} \subseteq S_{t-1}$. If the Balanced (1+1) EA attempts to flip the bits that correspond to a the set of vertices $A \subseteq V$ in iteration t, then $R_t \subseteq R_{t-1} \cup (A \cap V_R) \subseteq S_{t-1} \cup (A \cap V_R) = S_t$.

Now, assume the Balanced (1+1) EA attempts to do a balanced flip in iteration t and selects some $v \in V$ as starting vertex. Note that $t \leq T'_R$ implies that $t - 1 < T'_R$ and therefore $|R_{t-1}| < R - L$, which means that a balanced flip can never add a vertex in the right partition. If $v \in V_L$, then either a vertex in the right partition is removed or $R_t = R_{t-1}$. In any case, $R_t \subseteq R_{t-1} \subseteq S_{t-1} = S_t$. For $v \in V_R$, we distinguish between three different cases. If $v \in V_R \setminus S_{t-1}$, then also $v \in V_R \setminus R_{t-1}$ and nothing happens. On the other hand, if $v \in R_{t-1}$, then $t \leq T_L$ implies $t - 1 < T_L$ and $L \setminus L_t \neq \emptyset$. Thus, the bit corresponding to v gets flipped to 0 and a vertex from the left partition is added to the solution. It holds that $R_t = R_{t-1} \setminus \{v\} \subseteq S_{t-1} \setminus \{v\} = S_t$. Finally, assume that $v \in S_{t-1} \setminus R_{t-1}$. Note that the existence of such a vertex v implies $R_{t-1} \subseteq S_{t-1} \setminus \{v\}$. In this case, the Balanced (1+1) EA does nothing and v is removed from S_{t-1}. Thus, we have $R_t = R_{t-1} \subseteq S_{t-1} \setminus \{v\} = S_t$. In all cases, $R_t \subseteq S_t$, which concludes the proof. $\square$

We are aiming to use the stopping time $T_S = \inf\{t \in \mathbb{N} \mid |S_t| \geq R - L\}$ as a lower bound for T'_R. The following lemma justifies this.

Lemma 20. *For T'_R and T_S as above it holds that $T'_R \geq T_S$.*

Proof. We distinguish between two cases. If $T_L \geq T'_R$, then Lemma 19 implies that $R_t \subseteq S_t$ for all $t \leq T'_R$. Thus, $|R_t| \leq |S_t|$ for all such t and especially $R - L \leq \left|R_{T'_R}\right| \leq \left|S_{T'_R}\right|$, which implies $T'_R \geq T_S$. Now, assume $T_L < T'_R$. Note that this implies that X_{T_L} is feasible and has a fitness value of $f(X_{T_L}) < R$. Observe that increasing the number of selected vertices in the right partition would require removing vertices from the left partition. However, to remain feasible, all vertices from the right partition must be added, resulting in a fitness value of at least R. Thus, we have for all $t > T_L$ that $|R_t| \leq |R_{T_L}| < R - L$. Therefore, $T'_R = \infty$ and $T_S \leq T'_R$ holds trivially. $\square$

Based on Lemma 20, the following statement is essentially derived by applying a negative drift argument to $|S_t|$.

Lemma 21. *Let T'_R be defined as above. If $c > 2$, then, for every positive constant $\xi < 1$ and all $L \geq (\max\{2, 4/(c-2)\})^{\frac{1}{1-\xi}}$, there is a constant $\alpha > 0$ and a function $g(L) \in \Omega(L^\xi)$ such that*

$$\mathbb{P}\left[T'_R \leq 2^{\alpha L^\xi}\right] \leq 2^{-g(L)},$$

given the initial configuration X_0 is chosen uniformly at random.

Proof. We start by assuming that the initial configuration is chosen such that $|R_0| \leq R - L - L^\xi$. By Lemma 20 it is sufficient to prove the statement for T_S instead. To simplify notation, set $s_t = |S_t|$. Observe that s_t has the following transition probabilities:

- $\mathbb{P}\left[s_{t+1} = s_t - 1 \mid s_t\right] = \frac{1}{2} \cdot \frac{s_t}{L+R}$

- for all $j \in \mathbb{N} \cap [1, R - r_t]$:
  $\mathbb{P}\left[s_{t+1} = s_t + j \mid s_t\right] = \frac{1}{2} \cdot \binom{R-r_t}{j} \cdot \left(\frac{1}{L+R}\right)^j \cdot \left(1 - \frac{1}{L+R}\right)^{R-r_t-j}$

- with the remaining probability mass, $s_{t+1} = s_t$

Note that the second case (i.e., s_t increases) can be described more intuitively as follows. With a probability of $\frac{1}{2}$, we draw a binomial random variable Q with $R - r_t$ trials and success probability $\frac{1}{L+R}$, and $s_{t+1} = s_t + Q$.

To bound T_S, we apply Theorem 18 to the transformed stochastic process $Y_t = R - L - s_t$ with its natural filtration $\mathcal{F}_t$ and stopping time $T_Y = \inf\{t \in \mathbb{N} \mid Y_t \leq 0\} = T_S$. Note that, for $s_t \leq R - L - L^\xi$, we have $Y_t \geq L^\xi$. To this end, we argue that, for any fixed constant $c > 2$ and all $L \geq (\frac{4}{c-2})^{\frac{1}{1-\xi}}$, the requirements (a) and (b) of Theorem 18 are satisfied for $a = 0$, $b = L^\xi$, $r(d) = e^{1/3} \geq 1$, $\delta = e^{1/3} - 1 > 0$ and $\epsilon = \frac{c-2}{4(c+1)} > 0$.

Theorem 18 condition (a): Note that this condition is trivially satisfied for $Y_t \leq a = 0$ and $Y_t \geq b = L^\xi$ due to the indicator functions. To simplify notation, we assume from now on that $Y_t \in (0, L^\xi)$ and omit those indicator functions. By the definition of Y_t, it holds that

$$
\begin{aligned}
\mathbb{E}\left[Y_{t+1} - Y_t \mid \mathcal{F}_t\right] &= -\mathbb{E}\left[s_{t+1} - s_t \mid \mathcal{F}_t\right] \\
&= -\frac{1}{2} \cdot \left(\frac{R - s_t}{R + L} - \frac{s_t}{R + L}\right) \\
&= \frac{2s_t - R}{2(R + L)}.
\end{aligned}
$$

Substituting $s_t = R - L - Y_t$ and $R = cL$, and upper bounding Y_t by L^ξ yields

$$
\mathbb{E}\left[Y_{t+1} - Y_t \mid \mathcal{F}_t\right] = \frac{R - 2L - 2Y_t}{2(R + L)} \geq \frac{(c-2)}{2(c+1)} - \frac{L^{\xi-1}}{c+1}.
$$

Thus, for $L \geq \left(\frac{4}{c-2}\right)^{\frac{1}{1-\xi}}$ we have $\mathbb{E}\left[Y_{t+1} - Y_t \mid \mathcal{F}_t\right] \geq \frac{c-2}{2(c+1)} = \epsilon$.

Theorem 18 condition (b): Again, we omit the indicator functions and assume $Y_t > a = 0$. We do a case distinction based on $j \in \mathbb{N}$. First, observe that for $j \in \{0, 1\}$ it trivially holds that

$$
\mathbb{P}\left[|Y_{t+1} - Y_t| = j \leq 1\right] \leq \frac{e^{1/3}}{e^{j/3}} = \frac{r(d)}{(1 + \delta)^j}
$$

Thus, the condition is satisfied. For $j \geq 2$, note that such jumps can only appear when Y_t decreases. Therefore, we focus on the distribution $Y_t - Y_{t+1} = s_{t+1} - s_t$. As discussed earlier, given s_t, $s_{t+1} - s_t$ is dominated by a binomial random variable Q with $R - s_t$ trials and success probability $\frac{1}{L+R}$. Q in turn is dominated by a binomial random variable Q' with the same success probability but $L + R$ trials. By Chernoff's bound (see Theorem 1.10.1 in [3]) and the fact that $\mathbb{E}\left[Q'\right] = 1$ we obtain

$$
\mathbb{P}\left[|Y_{t+1} - Y_t| = j \mid Y_t\right] \leq \mathbb{P}\left[Q' \geq j\right] \leq e^{-\frac{(j-1)}{3}} = \frac{r(d)}{(1 + \delta)^j}.
$$

To save space, we write λ for $R - L - L^\xi$. Applying Theorem 18, we obtain that there is a constant $z > 0$ and a function $m(d) \in \Omega(d)$ such that $\mathbb{P}\left[T_Y \leq 2^{\frac{zd}{e^{1/3}}} \mid R_0 \leq \lambda\right] \leq 2^{-m(d)}$. Noting that $d = L^\xi$ and setting $\alpha = \frac{z}{e^{1/3}}$ proves that, for $g'(L) = m(L^\xi)$ we obtain

$$
\mathbb{P}\left[T'_R \leq 2^{\alpha L^\xi} \mid |R_0| \leq \lambda\right] \leq \mathbb{P}\left[T_S \leq 2^{\alpha L^\xi} \mid |R_0| \leq \lambda\right] \leq 2^{-g'(L)}.
$$

19

Next, assume we start with an initial configuration X_0 that is chosen uniform at random. Observe that $|R_0|$ follows a binomial distribution with success probability $\frac{1}{2}$ and R trials. Next, observe that for $L \geq 2^{\frac{1}{1-\xi}}$ and $c > 2$ it holds that

$$\left(1 + \frac{1}{2}\right) \geq \left(2 - \frac{1}{c}\right) \geq \left(2 - \frac{2 - 2L^{\xi-1}}{c}\right) = \frac{cL - L - L^{\xi}}{\frac{c}{2}L}.$$

As $\mathbb{E}[|R_0|] = \frac{1}{2}R = \frac{c}{2}L$, we have $\left(1 + \frac{1}{2}\right)\mathbb{E}[|R_0|] \geq cL - L - L^{\xi}$ and by Chernoff's bound (see Theorem 1.10.1 in [3])

$$\mathbb{P}[|R_0| > \lambda] \leq \mathbb{P}\left[|R_0| > \left(1 + \frac{1}{2}\right)\mathbb{E}[|R_0|]\right] \leq e^{-\frac{c}{24}L}.$$

By using the law of total probability we obtain

$$\begin{aligned}
&\mathbb{P}\left[T_R' \leq 2^{\alpha L^{\xi}}\right] \\
&= \mathbb{P}\left[T_R' \leq 2^{\alpha L^{\xi}} \text{ and } |R_0| \leq \lambda\right] + \mathbb{P}\left[T_R' \leq 2^{\alpha L^{\xi}} \text{ and } |R_0| > \lambda\right] \\
&\leq \mathbb{P}\left[T_R' \leq 2^{\alpha L^{\xi}} \;\middle|\; |R_0| \leq \lambda\right] + \mathbb{P}[|R_0| > \lambda] \\
&\leq 2^{-g'(L)} + e^{-\frac{c}{24}L} \\
&\leq 2^{-g(L)}
\end{aligned}$$

for some function $g(L) \in \Omega\left(L^{\xi}\right)$, which proves the claim. $\qquad\square$

Having Lemma 21 at hand, we are ready to prove Lemma 16 and hence our main theorem of this section (Theorem 17).

Proof of Lemma 16. By Lemma 21 we know that for L sufficiently large and $c > 2$ there is a constant $\alpha > 0$ and a function $g(L) \in \Omega(L^{\xi})$ such that $\mathbb{P}\left[T_R' \leq 2^{\alpha L^{\xi}}\right] \leq 2^{-g(L)}$.

We know that $\mathbb{P}[T_L < T_R'] \geq \mathbb{P}[T_L \leq \tau \text{ and } T_R' > \tau]$ holds for every $\tau \in \mathbb{N}$. Let $T = \inf\{t \in \mathbb{N} \mid X_t \text{ is feasible }\}$ and observe that, if $T_R' > \tau$, then $T_L \leq \tau$ if and only if $T \leq \tau$. Consequently, we have $\mathbb{P}[T_L \leq \tau \text{ and } T_R' > \tau] = \mathbb{P}[T \leq \tau \text{ and } T_R' > \tau]$. Moreover, using union bound, we get

$$\begin{aligned}
\mathbb{P}\left[T \leq \tau \text{ and } T_R' > \tau\right] &= 1 - \mathbb{P}\left[T > \tau \text{ or } T_R' \leq \tau\right] \\
&\geq 1 - \mathbb{P}[T > \tau] - \mathbb{P}\left[T_R' \leq \tau\right].
\end{aligned}$$

By choosing $\tau = 4e(c + 1)L^2\left(\ln((c+1)L) + 1/2\right)$, Lemma 2 yields $\mathbb{P}[T > \tau] \leq 2^{-L}$. Moreover, for L sufficiently large, we have $\tau \leq 2^{\alpha L^{\xi}}$. Thus, $\mathbb{P}[T_R' \leq \tau] \leq 2^{-g(L)}$ for a function $g(L) \in \Omega(L^{\xi})$. Consequently, $\mathbb{P}[T_L < T_R'] \geq 1 - 2^{-\Omega(L^{\xi})}$ and $\mathbb{P}[T_L \geq T_R'] \leq 2^{-\Omega(L^{\xi})}$, which concludes the proof. $\qquad\square$

Besides Theorem 17, a variety of other properties that might be of independent interest can be derived from Lemma 16. For example, the following corollary shows that the probability of X filling the larger side before all vertices in the smaller partition are selected, and the probability to ever select the entire larger side at all, both decay exponentially in L as well.

Corollary 22. *Consider the setting of Lemma 16 and let $T_R = \inf\{t \in \mathbb{N} \mid V_R \subseteq X_t\}$. If $c > 2$, then, for all positive constants $\xi < 1$ and L sufficiently large, it holds that*

(1) $\mathbb{P}\left[T_L \geq T_R\right] \leq 2^{-\Omega\left(L^{\xi}\right)}$

(2) $\mathbb{P}\left[T_R < \infty\right] \leq 2^{-\Omega\left(L^{\xi}\right)}$.

Proof. Part (1) follows trivially from $T_R \geq T_R'$. For part (2), note that, as soon as all vertices in V_L are selected, adding vertices to the right partition requires removing at least as many vertices on the left. To remain feasible, however, all free vertices in the right partition must be added at once. If there are more than L vertices in the right partition not selected, this is not possible as there are at most L vertices in the left partition in total that could be removed. Thus, from this point on, no more vertex from the right partition can be added and $T_L < T_R'$ implies $T_R = \infty$. $\square$

References

[1] Thomas Bäck and Sami Khuri. An evolutionary heuristic for the maximum independent set problem. *Proceedings of the First IEEE Conference on Evolutionary Computation. IEEE World Congress on Computational Intelligence*, pages 531–535 vol.2, 1994.

[2] D.R. Cox and H.D. Miller. *The Theory of Stochastic Processes*. Routledge, 1965.

[3] Benjamin Doerr. Probabilistic tools for the analysis of randomized optimization heuristics. *Theory of Evolutionary Computation*, page 1–87, Nov 2019.

[4] Benjamin Doerr, Daniel Johannsen, and Carola Winzen. Multiplicative drift analysis. *Algorithmica*, 64(4):673–697, 2012.

[5] Tobias Friedrich, Jun He, Nils Hebbinghaus, Frank Neumann, and Carsten Witt. Approximating covering problems by randomized search heuristics using multi-objective models. *Evolutionary Computation*, 18(4):617–633, 2010.

[6] Tobias Friedrich, Timo Kötzing, Martin S. Krejca, Samadhi Nallaperuma, Frank Neumann, and Martin Schirneck. Fast building block assembly by majority vote crossover. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*, GECCO '16, page 661–668. Association for Computing Machinery, 2016.

[7] Oliver Giel and Ingo Wegener. Evolutionary algorithms and the maximum matching problem. In *STACS 2003*, pages 415–426. Springer Berlin Heidelberg, 2003.

[8] Andreas Göbel, Timo Kötzing, and Martin S Krejca. Intuitive analyses via drift theory. *arXiv preprint arXiv:1806.01919*, 2018.

[9] Thomas Jansen, Pietro S Oliveto, and Christine Zarges. Approximating vertex cover using edge-based representations. In *Proceedings of the twelfth workshop on Foundations of genetic algorithms XII*, pages 87–96, 2013.

[10] Richard M. Karp. Reducibility among combinatorial problems. *Complexity of Computer Computations*, pages 85–103, 1972.

[11] Stefan Kratsch and Frank Neumann. Fixed-parameter evolutionary algorithms and the vertex cover problem. *Algorithmica*, 65(4):754–771, 2013.

[12] Per Kristian Lehre and Carsten Witt. Black-box search by unbiased variation. *Algorithmica*, 64(4):623–642, 2012.

[13] Pietro S Oliveto, Jun He, and Xin Yao. Evolutionary algorithms and the vertex cover problem. In *2007 IEEE Congress on Evolutionary Computation*, pages 1870–1877. IEEE, 2007.

[14] Pietro S Oliveto, Jun He, and Xin Yao. Analysis of population-based evolutionary algorithms for the vertex cover problem. In *2008 IEEE Congress on Evolutionary Computation (IEEE World Congress on Computational Intelligence)*, pages 1563–1570. IEEE, 2008.

[15] Pietro S 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(5):1006–1029, 2009.

[16] Pietro S Oliveto and Carsten Witt. Simplified drift analysis for proving lower bounds in evolutionary computation. *Algorithmica*, 59(3):369–386, 2011.

[17] Jannik Peters, Daniel Stephan, Isabel Amon, Hans Gawendowicz, Julius Lischeid, Lennart Salabarria, Jonas Umland, Felix Werner, Martin S. Krejca, Ralf Rothenberger, Timo Kötzing, and Tobias Friedrich. Mixed integer programming versus evolutionary computation for optimizing a hard real-world staff assignment problem. In J. Benton, Nir Lipovetzky, Eva Onaindia, David E. Smith, and Siddharth Srivastava, editors, *Proceedings of the Twenty-Ninth International Conference on Automated Planning and Scheduling, ICAPS 2018, Berkeley, CA, USA, July 11-15, 2019*, pages 541–554. AAAI Press, 2019.

[18] Jonathan E. Rowe and Michael D. Vose. Unbiased black box search algorithms. In Natalio Krasnogor and Pier Luca Lanzi, editors, *13th Annual Genetic and Evolutionary Computation Conference, GECCO 2011, Proceedings, Dublin, Ireland, July 12-16, 2011*, pages 2035–2042. ACM, 2011.

[19] Ruhul Sarker, Masoud Mohammadian, and Xin Yao. *Evolutionary Optimization*. Kluwer Academic Publishers, 2002.

[20] L. Darrell Whitley, Francisco Chicano, and Brian W. Goldman. Gray box optimization for mk landscapes (NK landscapes and max-ksat). *Evol. Comput.*, 24(3):491–519, 2016.