



This is a repository copy of *On steady-state evolutionary algorithms and selective pressure: Why inverse rank-based allocation of reproductive trials is best.*

White Rose Research Online URL for this paper:  
<https://eprints.whiterose.ac.uk/167459/>

Version: Accepted Version

---

**Article:**

Corus, D., Lissovoi, A., Oliveto, P. et al. (1 more author) (2021) On steady-state evolutionary algorithms and selective pressure: Why inverse rank-based allocation of reproductive trials is best. *ACM Transactions on Evolutionary Learning and Optimization*, 1 (1). 2. ISSN 2688-299X

<https://doi.org/10.1145/3427474>

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

**Reuse**

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

**Takedown**

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing [eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk) including the URL of the record and the reason for the withdrawal request.



[eprints@whiterose.ac.uk](mailto:eprints@whiterose.ac.uk)  
<https://eprints.whiterose.ac.uk/>

# On Steady-State Evolutionary Algorithms and Selective Pressure: Why Inverse Rank-Based Allocation of Reproductive Trials is Best

DOGAN CORUS, The University of Sheffield

ANDREI LISSOVOI, The University of Sheffield

PIETRO S. OLIVETO, The University of Sheffield

CARSTEN WITT, Technical University of Denmark

We analyse the impact of the selective pressure for the global optimisation capabilities of steady-state EAs. For the standard bimodal benchmark function TwoMAX we rigorously prove that using uniform parent selection leads to exponential runtimes with high probability to locate both optima for the standard  $(\mu+1)$  EA and  $(\mu+1)$  RLS with any polynomial population sizes. On the other hand, we prove that selecting the worst individual as parent leads to efficient global optimisation with overwhelming probability for reasonable population sizes. Since always selecting the worst individual may have detrimental effects for escaping from local optima, we consider the performance of stochastic parent selection operators with low selective pressure for a function class called TRUNCATEDTwoMAX where one slope is shorter than the other. An experimental analysis shows that the EAs equipped with inverse tournament selection, where the loser is selected for reproduction and small tournament sizes, globally optimise TwoMAX efficiently and effectively escape from local optima of TRUNCATEDTwoMAX with high probability. Thus they identify both optima efficiently while uniform (or stronger) selection fails in theory and in practice. We then show the power of inverse selection on function classes from the literature where populations are essential by providing rigorous proofs or experimental evidence that it outperforms uniform selection equipped with or without a restart strategy. We conclude the paper by confirming our theoretical insights with an empirical analysis of the different selective pressures on standard benchmarks of the classical MaxSat and Multidimensional Knapsack Problems.

CCS Concepts: • **Theory of computation** → **Random search heuristics**.

Additional Key Words and Phrases: randomised search heuristics,  $(\mu+1)$  EA, parent selection, diversity, TwoMAX

## ACM Reference Format:

Dogan Corus, Andrei Lissovoi, Pietro S. Oliveto, and Carsten Witt. 2020. On Steady-State Evolutionary Algorithms and Selective Pressure: Why Inverse Rank-Based Allocation of Reproductive Trials is Best. 1, 1 (October 2020), 38 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

## 1 INTRODUCTION

In natural evolution, the better individuals adapt to the environment, the more likely it is that they are selected for reproduction – that they generate offspring [Darwin 1859]. Indeed, canonical *generational* genetic algorithms (GAs) rely on this concept to efficiently evolve populations of

---

Authors' addresses: Dogan Corus, The University of Sheffield, [d.corus@sheffield.ac.uk](mailto:d.corus@sheffield.ac.uk); Andrei Lissovoi, The University of Sheffield, [a.lissovoi@sheffield.ac.uk](mailto:a.lissovoi@sheffield.ac.uk); Pietro S. Oliveto, The University of Sheffield, [p.oliveto@sheffield.ac.uk](mailto:p.oliveto@sheffield.ac.uk); Carsten Witt, Technical University of Denmark, [cawi@imm.dtu.dk](mailto:cawi@imm.dtu.dk).

---

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

© 2020 Association for Computing Machinery.

XXXX-XXXX/2020/10-ART \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

candidate solutions [Goldberg 1989; Holland 1992]. Generational evolutionary algorithms (EAs) model natural evolution by using the current population to create a completely new population of the same size in each generation (i.e., all individuals die by the end of their generation). The more individuals of the current population are fit (i.e., have high objective function value), the more likely they are to generate offspring for the next generation. These offspring will be part of the next population independently of their fitness but the probability they will reproduce, in turn, depends on how fit they are compared to the rest of the population. All the commonly used selection operators in generational EAs and GAs (such as e.g., fitness proportional, tournament, ranking and comma selection) rely on this described property of natural evolution to increase the average fitness of the population over time. Roughly speaking, the probability that fit individuals are selected for reproduction and produce a surviving offspring in generational EAs and GAs is commonly referred to as *selective pressure*. The higher the selective pressure, the faster diversity is lost due to the best individuals taking over the population (i.e., *takeover time*). On the other hand, it is well understood that sufficiently high selective pressure is required for the population of a generational EA to evolve effectively [Corus et al. 2018a]. In particular, if the selective pressure is not high enough, then generational EAs will require exponential time to optimise any function with unique optimum [Lehre 2010].

More recently introduced evolutionary algorithms, typically referred to as *steady-state GAs* [Syswerda 1989; Whitley 1989] (or *steady-state EAs* according to whether they use crossover or not [Witt 2006]), rely on *elitist* selection (often referred to also as *cut*, *truncation* or *plus* selection) between consecutive generations. Differently to canonical ones, steady-state algorithms (EAs or GAs) evolve overlapping populations. They create an offspring population of smaller size (i.e.  $\lambda$ ) compared to the parent population (of size  $\mu > \lambda$ ) and the worst individuals from the union of both populations (i.e. of total size  $\mu + \lambda$ ) are removed to obtain a new population of the same size as that of the parent population for the next generation. This considerable change in the evolutionary process means that the best genotypes do not disappear from the population in the next generation even if they do not reproduce. Such a difference implies that, if elitism is used, it is not *crucial* that better individuals have higher probabilities of reproducing for a population to evolve effectively. In particular, from the point of view of the *survival of the fittest* [Darwin 1868; Spencer 1864], elitism makes the choice of how parents are selected for reproduction irrelevant (i.e., the probability that the best  $\mu - \lambda$  individuals will be in the next generation is 1). From this point of view, the choice of the parent selection mechanism becomes arbitrary. While when steady-state evolutionary algorithms were originally introduced in the genetic algorithm community individuals were ranked and selected with a bias towards the top [Syswerda 1989; Whitley 1989], indeed nowadays parents are commonly selected uniformly at random for reproduction in steady-state EAs and GAs<sup>1</sup> [Corus and Oliveto 2018; Dang et al. 2018; Friedrich et al. 2009; Lengler 2018; Sutton 2018; Witt 2006].

From a different perspective, the parent selection operator may considerably impact the overall performance of the algorithm (i.e., the expected time for the global optimum to be found). In particular, we argue that the selective pressure commonly used in steady-state EAs is too high because, the higher the probability of selecting better individuals, the quicker the population is taken over by the best one, hence the faster *premature convergence* [Goldberg 1989; Sudholt 2019] occurs. Such a problem was already argued shortly after the introduction of steady-state GAs to the community [Goldberg and Deb 1991]. However, the only solutions considered to avoid premature

---

<sup>1</sup>This is probably due to the fact that the German *evolution strategies* community, where continuous optimisation though was the main focus, used uniform parent selection in  $(\mu + \lambda)$  ESs [Bäck et al. 1991]. While for the mentioned reasons this clearly makes sense, very successful steady-state GAs may still apply some parent selective pressure such as the binary tournament selection used in Chu and Beasley's GA. [Chu and Beasley 1998].

convergence with overlapping populations were either to increase the population size or to use multiple populations (niching).

In this paper we analyse the impact of parent selection on the performance of steady-state EAs for global optimisation. To this end, we consider the performance of the well studied  $(\mu+1)$  EA for the standard bimodal benchmark function TwoMAX. The function has been used in the theory of evolutionary computation to evaluate the exploration capabilities of EAs in discrete search spaces because, since the optima are located very far away from each other, it is very challenging to evolve a population that contains both of them [Friedrich et al. 2009; Oliveto et al. 2018; Osuna and Sudholt 2017; Sudholt 2019]. In particular, TwoMAX has often been used to evaluate the performance of diversity enhancing mechanisms for the  $(\mu+1)$  EA because, without such mechanisms, the algorithm has been proven to be unable to identify both optima of the function in polynomial time with high probability if the population size is at most sublinear (i.e.  $\mu = o(n/\log n)$ ) [Friedrich et al. 2009].

However, no proof is available that the algorithm would not be efficient using larger population sizes. Our first result is the proof that it is not. In particular, we show that the  $(\mu+1)$  EA and the  $(\mu+1)$  RLS with standard uniform parent selection cannot identify both optima of TwoMAX in polynomial time with  $1 - o(1)$  probability for any population of polynomial size. The previous analysis from the literature assumed pessimistically that the population has nearly optimised both branches when the takeover occurs. Our analysis reveals that the takeover occurs much earlier: not long after initialisation. Hence, the negative effects of premature convergence are emphasised: while TwoMAX displays local and global optima at opposite extremities of the search space, since the takeover occurs very early it may negatively affect the algorithm on much larger classes of multimodal functions.

On the other hand, we also prove that it is not necessary to use diversity enforcing mechanisms for the two algorithms to identify both optima of TwoMAX: decreasing the selective pressure suffices! As a proof of concept we decrease it to a minimum by always selecting the worst individual of the population as parent (breaking ties uniformly at random). We prove that such a decrease leads the algorithms to the efficient global optimisation of TwoMAX for reasonable population sizes i.e., the expected runtime is  $O(\mu n \log n)$  with  $\mu = \Omega(n^{1+\epsilon})$  for the  $(\mu+1)$  RLS and  $\mu = \Omega(n^{1+2\epsilon})$  for the  $(\mu+1)$  EA.

However, while we are arguing for low selective pressures, in general we do not recommend selecting the worst individual for reproduction as it may have serious detrimental effects for the global optimisation capabilities of steady-state EAs. For example, if the worst individual of the population is stuck on a local optimum, then it will be selected for reproduction until it manages to escape. For some local optima this may be prohibitive, and even require super-polynomial time. To tackle this problem we consider the usage of stochastic parent selection operators that do not deterministically select the worst individual of the population. To this end, we consider tournament selection with (small) tournament sizes where the *loser* of the tournament is chosen for reproduction<sup>2</sup>. Thus, we increase the selective pressure while still keeping it lower than that of uniform selection. To evaluate the performance of the algorithm at escaping from local optima, we modify TwoMAX such that one slope is shorter than the other and call the resulting function class TRUNCATEDTWOMAX<sub>k</sub> where the parameter  $k$  defines the height of the local optima.

A corollary of our first result shows that the  $(\mu+1)$  EA with standard uniform selection will only identify one of the peaks with high probability and this will be the lower one with probability at least  $1/2 - o(1)$ . An experimental analysis confirms that *Inverse Elitist* selection is very effective for the global optimisation of TwoMAX and that for TRUNCATEDTWOMAX with a probability of approximately  $1/2$  it gets stuck on the lower branch unless an unlikely takeover happens before.

<sup>2</sup>We use inverse tournament for simplicity of implementation. The same results can be achieved with Inverse Rank Selection.

On the other hand, the experiments reveal that *Inverse Tournament Selection* is able to escape from the local optima, hence is efficient for both TWOMAX and TRUNCATEDTWOMAX with approximately linear population sizes.

Although the TWOMAX benchmark function is commonly used to evaluate the exploration capabilities of population based EAs, a single trajectory algorithm with a restart strategy suffices to optimise the function very efficiently. Hence, we also evaluate the performance of the *Inverse Elitist* and *Inverse Tournament Selection* on benchmark functions that have previously been used in the literature to highlight problem characteristics where populations are essential. In particular, problem classes where single trajectory EAs fail with overwhelming probability (w.o.p.), hence restart strategies are also ineffective. To this end, we pick the multimodal function SUFSAMP (which we generalise and prefer to call RIDGEWITHBRANCHES) previously used in the literature to highlight fitness landscape characteristics where the use of offspring populations is essential [Jansen et al. 2005] and a bimodal function class, that we call TWOGRADIENTS, which has been used in the literature to highlight circumstances where parent populations are essential [Witt 2006]. For both functions a single trajectory (1+1) EA gets stuck on a local optimum w.o.p. and restarts will not help to reduce the runtime to polynomial in the problem size. On the other hand, the respective population based EAs (i.e., (1+ $\lambda$ ) EA for RIDGEWITHBRANCHES and the ( $\mu$ +1) EA for TWOGRADIENTS) can efficiently identify the global optimum by avoiding the local optima. However, if the fitness values of the local and global optima are swapped appropriately, then the opposite effects can be shown: the single trajectory algorithms are effective while the population-based algorithms get stuck w.o.p. Our contributions are the following. Concerning RIDGEWITHBRANCHES, we rigorously prove that the ( $\mu$ +1) EA with *Inverse Tournament Selection* efficiently identifies all of the many optima of the function, hence also the global one independent of which optimum is set to have the maximal value. Concerning TWOGRADIENTS we rigorously prove that the ( $\mu$ +1) EA with inverse binary tournaments can find the global optimum efficiently just like the ( $\mu$ +1) EA with uniform selection. Additionally, we show experimentally that, while the (1+1) EA w.o.p. gets trapped in one optimum and the uniform selection ( $\mu$ +1) EA on the other, the *Inverse Tournament Selection* algorithms identify both optima with reasonable probability, hence can optimise the function efficiently independent of which local optimum is defined to be the optimal one.

We conclude the paper by comparing experimentally the inverse selection operators against standard uniform selection for well studied benchmark instances of two NP-Hard problems, MaxSat and the Multidimensional knapsack problem (MKP). While for the MaxSat instances, independent runs of single trajectory algorithms outperform the population based EAs, such a strategy is detrimental for MKP<sup>3</sup>. On the other hand, for both problems the inverse selection algorithms outperform uniform selection with statistical significance i.e., better solutions are found with higher probability.

The rest of the paper is structured as follows. In the next section we introduce the ( $\mu$ +1) EA and ( $\mu$ +1) RLS together with the different selection operators and the other algorithms that we will consider in the paper. We also present the TWOMAX and TRUNCATEDTWOMAX<sub>k</sub> function classes and some tools for the analyses. In Section 3 we prove that standard uniform selection or higher selective pressure does not allow the efficient global optimisation of TWOMAX independent of the population size. In Section 4 we prove that the algorithm becomes efficient if the worst individual is selected for reproduction. In Section 5 we present an experimental analysis providing evidence that inverse tournament selection with small tournament sizes is preferable to the previously considered selection operators for the optimisation of TWOMAX and TRUNCATEDTWOMAX<sub>k</sub>. In Section 6 we

<sup>3</sup>For MaxSat we implemented a ( $\mu$ +1) EA with deterministic crowding where offspring only compete with their parent. The difference in performance with  $\mu$  parallel (1+1) EAs should be minimal and not noticeable.

**Algorithm 1** ( $\mu+1$ ) EA

---

Sample  $\mu$  search points from  $\{0, 1\}^n$  independently and uniformly at random and put them into the multiset  $P$ .

**for**  $t \leftarrow 1, 2, \dots$  **do**

Pick  $x \in P$  according to the parent selection operator.

Create  $y$  by flipping each bit of  $x$  independently with probability  $1/n$ .

Remove  $\arg \min\{f(z) \mid z \in P \cup \{y\}\}$  from  $P$ , breaking ties randomly.

**end for**

---

show the superiority of the inverse selection operators over uniform selection for problems where populations are essential and where restart strategies of single trajectory algorithms fail w.o.p. In Section 7 we confirm our theoretical intuitions with an experimental analysis of the considered algorithms for the two NP-Hard problems. We finish with a discussion and some conclusions.

## 2 PRELIMINARIES

We will consider the well-studied ( $\mu+1$ ) EA described in Algorithm 1. The algorithm starts by selecting  $\mu$  individuals chosen uniformly at random. Then at each iteration an individual is selected from the population (parent selection) and an offspring is generated by mutating each bit of the selected individual with probability  $p$  (in this paper we will use the standard mutation rate  $p = 1/n$ ). Then, the worst individual is removed from the population breaking ties randomly. We will also consider its random local search variant called ( $\mu+1$ ) RLS that flips exactly one bit per mutation (instead of using the mutation rate  $p$ ).

Our aim is to analyse how the parent selection operator affects the global optimisation performance of the algorithm. Different definitions have been provided in the literature to refer to the *global optimisation* of multimodal optimisation problems. Some researchers use the term to refer to the goal of identifying one single global optimum, while others refer to the identification of multiple optimal solutions and even to a collection of global and local optima. We refer to [Preuss 2015] for an overview of the available definitions in the literature. While the insights provided in this paper apply to any of these definitions, in this paper we refer to global optimisation as the task of identifying one global optimum. In particular, all the benchmark problems considered in this paper have only one global optimum (for TwoMAX either optimum may be designated to be the global one but, in order for an algorithm to identify it with high probability, it is necessary to visit both). We will consider the following operators enumerated from highest to lowest selective pressure:

- (1)  **$K$ -Tournament selection:** Select  $K$  individuals uniformly at random from the population and choose the winner (i.e., best) of the tournament as parent;
- (2) **Uniform Selection:** choose the parent uniformly at random from the population;
- (3) **Inverse  $K$ -Tournament selection:** Select  $K$  individuals uniformly at random and choose the loser (i.e., worst) of the tournament as parent;
- (4) **Inverse Elitist Selection:** choose the parent with worst fitness from the population breaking ties uniformly at random.

Apart from comparing different selective pressures for the ( $\mu+1$ ) EA, we will also consider the use of diversity mechanisms from the literature for which good global exploration capabilities have been proven. In particular, if the use of problem specific knowledge is not allowed, then *deterministic crowding* is the best diversity mechanism that has been analysed for the TwoMAX benchmark problem we consider in the first part of this paper. The ( $\mu+1$ ) EA using deterministic

**Algorithm 2** ( $\mu+1$ ) EA with deterministic crowding

---

Sample  $\mu$  search points from  $\{0, 1\}^n$  independently and uniformly at random and put them into the multiset  $P$ .

**for**  $t \leftarrow 1, 2, \dots$  **do**

Pick  $x \in P$  uniformly at random.

Create  $y$  by flipping each bit of  $x$  independently with probability  $1/n$ .

**If**  $f(y) \geq f(x)$  **then** remove  $x$  from  $P \cup \{y\}$ , **else** keep  $P$ .

**end for**

---

crowding is described in Algorithm 2. It optimises TwoMAX in  $O(\mu n \log n)$  fitness evaluations with probability at least  $1 - 2^{-\mu}$ . The difference with the standard  $(\mu+1)$  EA is that offspring only compete with their parent rather than with the whole population. The algorithm closely resembles  $\mu$  parallel  $(1+1)$  EAs since the individuals explore the landscape independently. Indeed a  $(1+1)$  EA with a simple restart strategy optimises TwoMAX with the same probability if  $\mu$  is the number of performed restarts or parallel runs. The  $(1+1)$  EA is obtained from Algorithm 1 by setting  $\mu = 1$  and preferring the offspring to the parent when breaking ties. Other theoretically studied diversity mechanisms are either ineffective (genotype diversity), detrimental (fitness diversity) or require problem knowledge to be made efficient (fitness sharing and clearing<sup>4</sup>) for TwoMAX [Friedrich et al. 2009; Oliveto et al. 2018; Osuna and Sudholt 2017; Sudholt 2019] hence we do not consider them in this paper.

We are interested in the optimisation (here: maximisation) of the TwoMAX function, which is defined as  $\text{TwoMAX} : \{0, 1\}^n \rightarrow \mathbb{R}, x \mapsto |n/2 - |x|_1|$ . This function has two global optima, the all-ones and the all-zeros string. Leading to each optimum, there is a so-called *branch* of perfect fitness-distance correlation: if  $|x|_1 > n/2$  then the fitness is  $n/2$  minus the Hamming distance to the all-ones string and we say that  $x$  is on the right-hand branch; analogously for  $|x|_1 < n/2$ , the Hamming distance to the all-zeros string is considered and we say that  $x$  is on the so-called left-hand branch. A search point with  $n/2$  ones, having fitness 0 is said to be in the *valley*. The goal of our analyses is to evaluate whether the EAs can identify both optima of TwoMAX in expected polynomial *time* (i.e., the number of fitness function evaluations is polynomial with respect to the problem size  $n$ ). The function is depicted in Figure 1.

We also consider the following generalisation of TwoMAX:

$$\text{TRUNCATEDTwoMAX}_k(x) = \begin{cases} \text{TwoMAX}(x) & \text{if } |x|_1 \geq n/2 - k, \\ 0 & \text{otherwise.} \end{cases}$$

The parameter  $k$  defines the height of the left branch, from 0 (i.e., no left branch) to  $n/2$  (i.e., TwoMAX). The lower the parameter value, the easier is the optimisation of the function because the likelihood that individuals on the local optima take over the population decreases. The function is depicted in Figure 1.

In our analyses we will need the following useful estimates for the probability function as well as the density of the binomial distribution with parameters  $n$  and  $p$ .

<sup>4</sup>Fitness sharing has only been proven to be efficient for TwoMAX if the knowledge that it is a unitation function is exploited by the sharing function [Friedrich et al. 2009; Oliveto et al. 2018]. Without problem knowledge, hence using the Hamming distance for the sharing function, it has been experimentally shown that fitness sharing does not make the  $(\mu+1)$  EA efficient for TwoMAX [Oliveto et al. 2018]. Similar argumentations also hold for Clearing which requires large population sizes and an appropriate niching radius to optimise the function efficiently [Osuna and Sudholt 2017].

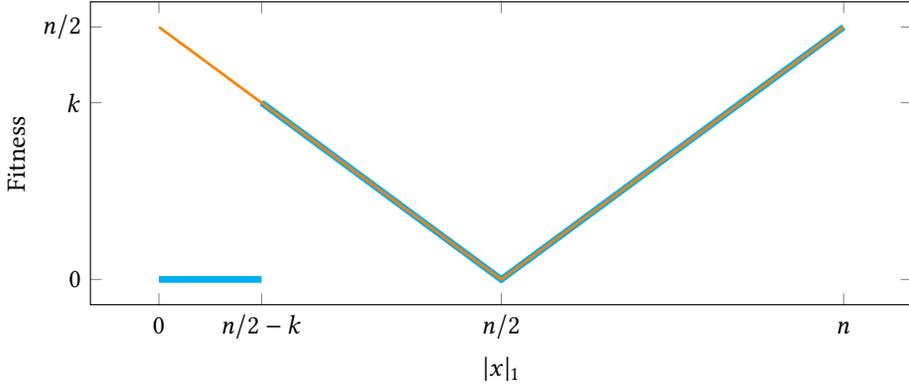


Fig. 1. TwoMAX (orange) and TRUNCATEDTWO MAX<sub>k</sub> (thick blue) functions.

LEMMA 1 (THEOREMS 1.2, 1.5 AND 1.6 IN [BOLLOBAS 1985]). Let  $X \sim \text{Bin}(n, p)$ ,  $q = 1 - p$  and suppose  $np \geq 1$ . For all  $h \geq 3/(nq)$  it holds that

$$\Pr(X = np + h) \leq \frac{1}{\sqrt{2\pi pqn}} e^{-\frac{h^2}{2pqn} + \frac{h}{qn} + \frac{h^3}{p^2n^2}}$$

Also, for all  $h \geq 0$  where  $np + h$  is an integer less than  $n$ , it holds that

$$\Pr(X = np + h) \geq \frac{1}{\sqrt{2\pi pqn}} e^{-\frac{h^2}{2pqn} - \frac{h^3}{2q^2n^2} - \frac{h^4}{3p^3n^3} - \frac{h}{2pn} - \frac{1}{12(np+h)} - \frac{1}{12(nq-h)}}$$

Finally, if  $pqn = \omega(1)$ ,  $h > 0$ ,  $h = \omega(\sqrt{pqn})$ ,  $h = o((pqn)^{2/3})$  as well as  $np + h \leq n$  then

$$\Pr(X \geq np + h) = (1 \pm o(1)) \frac{\sqrt{pqn}}{h\sqrt{2\pi}} e^{-\frac{h^2}{2pqn}}$$

As a corollary from the previous lemma, we obtain the following asymptotically tight estimate for the case of the symmetrical binomial distribution, i. e.,  $p = 1/2$ :

COROLLARY 1. Let  $X \sim \text{Bin}(n, 1/2)$ . Then for all  $h = o(n^{2/3})$  such that  $n/2 + h \in \{n/2, \dots, n\}$  it holds that

$$\Pr(X = n/2 + h) = (1 \pm o(1)) \frac{\sqrt{2}}{\sqrt{\pi n}} e^{-\frac{2h^2}{n}}$$

as well as

$$\Pr(X \geq n/2 + h) = (1 \pm o(1)) \frac{\sqrt{n}}{h\sqrt{8\pi}} e^{-\frac{2h^2}{n}}$$

In particular, the corollary will be used to analyse the maximum order statistic with respect to the number of one-bits in a uniform population of  $\mu$  individuals. Related results can be found in [de Perthuis de Laillevault et al. 2015]; however, we find the bounds based on [Bollobas 1985] presented above more convenient to use.

The following two theorems, an additive variant of the classical Chernoff-Hoeffding inequality and the artificial fitness levels theorem with tail probabilities, will be used in Section 4 for the analysis of the  $(\mu+1)$  EA.

THEOREM 1 (TH. 1.11 IN [DOERR 2011]). Let  $X_1, \dots, X_n$  be independent random variables taking values in  $[0, 1]$ . Let  $X = \sum_{i=1}^n X_i$ . Then for all  $\lambda \geq 0$ ,  $\Pr(X \geq E[X] + \lambda) \leq \exp(-\frac{2\lambda^2}{n})$  and  $\Pr(X \leq E[X] - \lambda) \leq \exp(-\frac{2\lambda^2}{n})$ .

Finally, we make use of the following artificial fitness levels theorem with tail probabilities.

**THEOREM 2 (THEOREM 2 IN [WITT 2014]).** *Consider an algorithm  $\mathcal{A}$  maximising some function  $f$  and a partition of the search space into non-empty sets  $A_1, \dots, A_m$ . Assume that the sets form an  $f$ -based partition, i. e., for  $1 \leq i < j \leq m$  and all  $x \in A_i, y \in A_j$  it holds  $f(x) < f(y)$ . We say that  $\mathcal{A}$  is in  $A_i$  or on level  $i$  if the best search point created so far is in  $A_i$ .*

*If  $p_i$  is a lower bound on the probability that a step of  $\mathcal{A}$  leads from level  $i$  to some higher level, independently of previous steps, then for any  $\delta > 0$ , the first hitting time of  $A_m$ , starting from level  $k$ , is at most*

$$\sum_{i=k}^{m-1} \frac{1}{p_i} + \delta$$

*with probability at least  $1 - e^{-\frac{\delta}{4} \cdot \min\{\frac{\delta}{s}, h\}}$ , for any finite  $s \geq \sum_{i=k}^{m-1} \frac{1}{p_i}$  and  $h = \min\{p_i \mid i = k, \dots, m-1\}$ .*

For the remainder of the paper we will use *with overwhelming probability* (w.o.p.) to mean with probability  $1 - 2^{-n^{\Omega(1)}}$ .

### 3 TOO HIGH SELECTIVE PRESSURE FAILS

It is known from previous work that the standard  $(\mu+1)$  EA (using uniform parent selection) with high probability fails to find both optima of TwOMAX if  $\mu = o(n/\log n)$  [Friedrich et al. 2009]. More precisely, with probability  $1 - o(1)$  there will be within any polynomial amount of time no step where the current population contains both optima. The ideas behind the analysis in [Friedrich et al. 2009] can briefly be summarized as follows:

- With overwhelming probability the initial population contains none of the two optima.
- As every generation creates exactly one new individual, there is a first point in time where exactly one of the optima, say the all-ones string, is created and added to the population.
- The take-over time for the all-ones string is  $O(\mu \log \mu) = o(n)$  if the all-zeros string is not found in between. With probability  $1 - o(1)$  the take-over time is still  $o(n)$ .
- Pessimistically assuming all individuals from the branch leading to the all-zeros string to have only a single-one bit, the probability of creating the all-zeros string by mutation is still at most  $1/n$ . With high probability this step does not happen during the take-over time.
- Hence, with high probability after  $o(n)$  steps the whole population consists of copies of one optimum only, and in no step before has the other optimum been in the population. The probability of creating the opposite optimum by mutation is  $1/n^n$ , which w.o.p. does not happen within any polynomial number of steps.

Our analysis will considerably improve the result from [Friedrich et al. 2009]. First of all, we show that even with arbitrarily large polynomial population sizes  $\mu$  the algorithm will fail to find both optima. Second, we show that the whole population will be on one branch of TwOMAX in very early stages of the optimisation, i. e., soon after initialisation. Hence, diversity is lost much earlier than in the final improving steps that are pessimistically considered in [Friedrich et al. 2009].

In a nutshell, the proof idea for the following theorem is to consider the uniform initialization procedure of the  $(\mu+1)$  EA and to analyse the probability of so-called outliers. More formally, an outlier is an individual  $x^*$  from the initial population that has the largest distance  $\|x^*\|_1 - n/2$  from the points in the valley  $n/2$  and also has a significantly larger distance to the valley than all other individuals. Again, an argument about the take-over time then shows that with high probability the whole population will be descendants of the outlier after  $O(\mu \log \mu n^\epsilon)$  steps. Here we use the family tree technique by Witt [Witt 2006] to show that offspring of other individuals would have

to make an exceptionally large progress within the take-over time to catch up to the level of the best. After the population has been taken over by descendants of the outlier, steps that mutate individuals to accepted individuals from the opposite branch are superpolynomially unlikely due to the large distance of the outlier's fitness to the valley.

**THEOREM 3.** *Let  $\mu \leq n^k$  for an arbitrary constant  $k \geq 0$  as well as  $\mu = \omega(1)$  and consider the  $(\mu+1)$  EA on TwoMAX. Let  $\epsilon > 0$  be an arbitrary constant. Then with probability  $1 - o(1)$ , all individuals are on one branch before the TwoMAX-value reaches  $n/2 + c\sqrt{n}(\log n + \log \mu)$ , for some sufficiently large constant  $c > 0$ , and the worst TwoMAX-value in the population is  $n/2 + \Omega(n^{1/2-\epsilon})$  at the time one branch dies out. This implies time  $2^{\Omega(n^{1/2-\epsilon})}$  with probability at least  $1 - 2^{-\Omega(n^{1/2-\epsilon})}$  to create an individual from that branch again.*

To prove Theorem 3, we will essentially look into random variables  $X_i$ , where  $1 \leq i \leq \mu$ , following the Binomial distribution with parameters  $n$  and  $1/2$ . These variables describe the number of one-bits in the  $\mu$  uniformly initialized individuals. Since we are only interested in the distance of this number to the valley  $n/2$ , we introduce  $D_i := ||x_i|_1 - n/2|$  for  $i \in \{1, \dots, \mu\}$  and relate these to  $X_i$ . The probability distribution of  $D_i$  equals the one of  $X_i - n/2$ , conditioned on  $X_i \geq n/2$ . Also,  $1/2 \leq \Pr(X_i \geq n/2) \leq 1/2 + O(1/\sqrt{n})$ , so for  $a \geq 0$

$$(2 - o(1)) \Pr(X_i \geq n/2 + a) \leq \Pr(D_i \geq a) \leq (2 + o(1)) \Pr(X_i \geq n/2 + a).$$

In the following, we will analyse the probability of an outlier, i. e., an individual that has a  $D$ -value significantly larger than the standard deviation  $\sigma = \sqrt{n/4}$  of the binomial distribution  $\text{Bin}(n, 1/2)$ . Let  $D_{(1)} \geq \dots \geq D_{(\mu)}$  be the order statistics of the  $D_i$  (for convenience, index (1) denotes the largest order statistic here, in contrast to the customary use in the literature). We shall also use the random variables  $(i)$ , where  $i \in \{1, \dots, \mu\}$ , to denote the index of the  $i$ th order statistic.

We recall the assumptions  $\mu \leq n^k$  for some constant  $k \in \mathbb{N}$  and  $\mu = \omega(1)$  from the main theorem, where the latter is mostly for convenience of proof. Our first result shows a lower bound on  $D_{(1)}$ , the largest order statistic, that holds with high probability.

**LEMMA 2.** *There is a  $\ell^* = \sqrt{n((1/2 - 1/(8k)) \ln \mu - \Theta(\log \log \mu))}$  such that  $\Pr(D_{(1)} \geq \ell^*) = 1 - o(1)$ .*

**PROOF.** Let  $\ell^*$  be the solution for  $h$  of the equality

$$\frac{\sqrt{n}}{h\sqrt{8\pi}} e^{-2h^2/n} = \frac{\mu^{-1+1/(4k)}}{2}.$$

We first prove the asymptotic expression given for  $\ell^*$  in the statement of the lemma. We write

$$\ell^* = \sqrt{n((1/2 - 1/(8k)) \ln \mu - c(n) \ln \ln \mu)}$$

for a function  $c(n)$ , i. e.,  $c(n)$  stands for the  $\Theta(\log \log \mu)$ -term in the claimed expression divided by  $\ln \ln \mu$ . Then,

$$\begin{aligned} \frac{\sqrt{n}}{\ell^* \sqrt{8\pi}} e^{-2(\ell^*)^2/n} &= \frac{\sqrt{n} \mu^{-1+1/(4k)} e^{2c(n) \ln \ln \mu}}{\ell^* \sqrt{8\pi}} = \frac{\sqrt{n} (\ln \mu)^{2c(n)}}{\sqrt{n((1/2 - 1/(8k)) \ln \mu - c(n) \ln \ln \mu)} \sqrt{8\pi}} \mu^{-1+1/(4k)} \\ &= \frac{1}{\sqrt{8\pi \left( \left( \frac{1}{2} - \frac{1}{8k} \right) (\ln \mu)^{1-4c(n)} - \frac{c(n) \ln \ln \mu}{(\ln \mu)^{4c(n)}} \right)}} \mu^{-1+1/(4k)} \end{aligned}$$

Now, the fraction equals  $1/2$  for some choice of  $c(n)$  that converges to the constant  $1/4$ , as can be seen from solving

$$\sqrt{8\pi \left( \frac{1}{2} - \frac{1}{8k} \right) (\ln \mu)^{1-4c(n)}} = 2$$

for  $c(n)$  and exploiting that  $\frac{c(n) \ln \ln \mu}{(\ln \mu)^{4c(n)}} = o(1)$  thanks to our assumption  $\mu = \omega(1)$ .

By Lemma 1, we have for any of the random variables  $X_i \sim \text{Bin}(n, 1/2)$  that  $\Pr(X_i \geq n/2 + \ell^*) \geq (1 - o(1))\mu^{-1+1/(4k)}/2$  and therefore  $\Pr(D_i \geq \ell^*) \geq (1 - o(1))\mu^{-1+1/(4k)}$ . Using the independence of the random variables, the probability at there is at least one  $i \in \{1, \dots, \mu\}$  with this property is therefore at least

$$1 - \left(1 - \mu^{-1+1/(4k)}\right)^\mu = 1 - o(1).$$

□

While the previous lemma established the existence of a high extremal value for the  $D_i$ , the next lemma shows that the maximum order statistic with high probability actually is an outlier, more precisely, that it is almost  $\sqrt{n}$  larger than the second-largest  $D$ -value.

**LEMMA 3.** *With probability  $1 - o(1)$ , there exists  $\delta^* = O(n^{1/2-\epsilon})$  for an arbitrary constant  $\epsilon > 0$  such that  $D_{(1)} - D_{(2)} \geq \delta^*$ .*

To prove Lemma 3, we need to show for all  $D_i \neq D_{(1)}$  that they are significantly smaller than  $D_{(1)}$ . However, this does not imply the lemma yet since there may be more than one variable taking the maximum value. Hence, we first prove the following auxiliary statement, showing that the largest order statistic is unique with high probability.

**LEMMA 4.** *Assume  $D_{(1)} \geq \ell^*$  and let  $i$  be an index such that  $D_i = D_{(1)}$ . Then with probability  $1 - o(1)$  it holds that  $D_j < D_i$  for all  $j \neq i$ .*

**PROOF.** The main idea of the proof is to exploit that the first bound (on the probability value) and the second bound (on the distribution) in Corollary 1 are within a ratio of  $\Theta(h/n)$ , making it unlikely for any other individual to hit exactly the value of  $D_{(1)}$ . For a moment, we consider the random  $D_i$  without the condition  $D_{(1)} \geq \ell^*$ . According to Lemma 2, we have  $\ell^* = \Theta(\sqrt{n \log \mu})$ . We already know that  $\Pr(D = \ell^*) = (2 + o(1))\Pr(X_i \geq n/2 + \ell^*)$  and obtain from the first inequality in Corollary 1 that

$$\Pr(D_i = \ell^*) = O\left(\frac{\text{polylog } \mu}{\mu^{1-1/(4k)}\sqrt{n}}\right) = O\left(\frac{\text{polylog } \mu}{\mu n^{1/4}}\right),$$

where we used  $\mu^{1/(4k)} \leq n^{1/4}$ . Also, for values  $k \geq \ell^*$  we clearly have  $\Pr(D = k) \leq \Pr(D = \ell^*)$  by the monotonicity of the binomial distribution.

We now consider the generation of the  $\mu$  random  $D_i$ -values as a stochastic process, assuming that  $D_i$  is generated at time  $i$ . Note that our condition  $D_{(1)} \geq \ell^*$  is equivalent to that there exists a (random)  $i \in \{1, \dots, \mu\}$  such that  $D_i \geq \ell^*$ . Let  $T$  denote the smallest such  $i$ . For  $j < T$  is clearly holds  $D_j < \ell^* \leq D_{(1)}$ . After time  $T$ , the condition  $D_{(1)} \geq \ell^*$  is satisfied so that each of the  $D_{T+1}, \dots, D_\mu$  is conditioned on being at most  $D_{(1)}$ . Since even  $\Pr(D_j \leq \ell^*) = 1 - o(1)$ , the conditional probability is only by a factor  $1 + o(1)$  larger than the unconditional one. Hence, it is sufficient to show that the values  $D_{T+1}, \dots, D_\mu$  unconditionally with high probability are less than  $D_{(1)}$ .

By the above calculations, for all  $i > T$ , we get

$$\Pr(D_i = D_{(1)}) = O\left(\frac{\text{polylog } \mu}{\mu n^{1/4}}\right).$$

Now a union bound over the  $\mu - T$  variables yields the desired result. □

We now prove the outlier property of the largest order statistic.

PROOF OF LEMMA 3. We recall that  $\ell^* = \Theta(\sqrt{n \log \mu})$ . By using the second inequality of Corollary 1, we note that for  $\delta^* = O(n^{1/2-\epsilon})$  and any  $i \in \{1, \dots, \mu\}$  it holds that  $C \Pr(D_i \geq \ell^*) = (1 \pm o(1)) \Pr(D_i \geq \ell^* - \delta^*)$ , where

$$C = \left(1 - \frac{\delta^*}{\ell^*}\right)^{-1} \exp\left(\frac{4\ell^*\delta}{n} - \frac{2\delta^2}{n}\right) = 1 + O(n^{-\epsilon}) .$$

Now,

$$\begin{aligned} \Pr(D_{(1)} - \delta \leq D_i < D_{(1)} \mid D_{(1)} \geq \ell^* \wedge D_i < D_{(1)}) &\leq \Pr(D_i \geq \ell^* - \delta \wedge D_i < \ell^* \mid D_i < D_{(1)}) \\ &\leq (1 - o(1))^{-1} (1 - \Pr(D_i \geq \ell^*) - (1 - \Pr(D_i \geq \ell^* - \delta))) \\ &= (1 + o(1))(C - 1) \Pr(D_i \geq \ell^*) = O(1/\mu) \mu^{1/(4k)} n^{-\epsilon} = O(\mu^{-1} n^{-\epsilon/2}) , \end{aligned}$$

where the first inequality follows from exploiting the condition that  $D_{(1)} \geq \ell^*$  and the fact that  $\Pr([x - a, x])$  is decreasing in  $x$  for any positive constant  $a$  and  $x - a$  being larger than the mode. In the second inequality we used  $\Pr(A \mid B) \leq \Pr(A)/\Pr(B)$  and

$$\Pr(D < D_{(1)}) \geq \Pr(D_{(1)} \geq \ell^*) \Pr(D < D_{(1)} \mid D_{(1)} \geq \ell^*) = 1 - o(1) ,$$

which follows from Lemma 2 and Lemma 4. The equalities then follow from bound on  $C$  derived above, the fact  $\Pr(D_i \geq \ell^*) \leq \mu^{-1+1/(4k)}$  and sufficiently large  $k$ .

Hence, we have that the unconditional probability

$$\Pr(D_{(1)} - \delta \leq D < D_{(1)}) \leq \frac{o(1/\mu)}{(1 - o(1))}$$

Now, by a union bound, the probability that there exists some  $j \neq (1)$  such that  $D_{(1)} - \delta \leq D_j < D_{(1)}$  is  $o(1)$ . In particular, the probability of  $D_{(2)}$  being in this interval is  $o(1)$ .  $\square$

Lemma 3, which establishes the outlier in the initial population is now crucial to show that the outlier will take over the whole population with high probability before the rest of the population can catch up to the outlier. The following lemma makes this take-over effect formal.

LEMMA 5. *If the best individual in a population optimising TwoMAX has a fitness value  $\Omega(n^\epsilon)$  larger than the second best for some  $\epsilon = \Omega(1)$ , then the population will be taken over by individuals from the best solution's branch whose fitness values are all at least as good as the initial best fitness in  $n^{\epsilon/4} \mu \log \mu$  generations w.o.p.*

PROOF. Using a standard analysis of take-over times (e. g., [Friedrich et al. 2009; Witt 2006]) it takes an expected number of  $\Theta(\mu \log \mu)$  generations until the population is taken over by solutions from the best individual's branch, unless an equal or better search point on the opposite branch is sampled. Using Markov's inequality and analyzing the probability of failing to take over within  $n^{\epsilon/4}$  phases of length  $\Theta(\mu \log \mu)$  each, we obtain that the take-over time is  $O(n^{\epsilon/4} \mu \log \mu)$  w.o.p.

We now use the well-known family tree technique [Witt 2006] for the analysis of the  $(\mu+1)$  EA to show that the offspring of any other individual than the initially best with high probability cannot catch up to the fitness of the initially best. The family tree of any initial solution will not reach a depth larger than  $n^{\epsilon/2}$  in  $n^{\epsilon/4} \mu \log \mu$  generations w.o.p. according to Lemma 2 in [Witt 2006]. Let us call a path with any node which has Hamming distance at least  $3n^{\epsilon/2}$  to its root a *bad path*. The probability of observing a particular bad path of length  $\ell$  in a tree of depth smaller than  $n^{\epsilon/2}$  is at most  $(1/\mu)^\ell e^{-3n^{\epsilon/2}}$ . Using a union bound over all paths of length at most  $n^{\epsilon/2}$  that can be observed

in  $n^{\epsilon/4} \mu \log \mu$  generations of the  $(\mu+1)$  EA, we upper bound the probability of sampling a solution which has  $3n^{\epsilon/2}$  more 1-bits than its root by,

$$\sum_{\ell=1}^{n^{\epsilon/2}} \binom{n^{\epsilon/4} \mu \log \mu}{\ell} (1/\mu)^\ell e^{-3n^{\epsilon/2}} \leq n^{\epsilon/2} e^{n^{\epsilon/3}} e^{-3n^{\epsilon/2}} = 2^{-\Omega(n)}.$$

The probability that a bad path will not be observed in the family tree of any solution in the initial population can be shown to be exponentially small by using a union bound once again. Hence, no offspring of the individuals different from the best one will reach the fitness of the best with high probability.  $\square$

We put everything together to prove our main result:

**PROOF OF THEOREM 3.** According to Lemma 3,  $D_{(1)} - D_{(2)} = \Omega(n^{1/2-\epsilon})$  with probability  $1 - o(1)$ . We assume this to happen. Hence, Lemma 5 can be applied, such that after  $n^{1/8} \mu \log \mu$  generations one branch will have died with probability  $1 - o(1)$ . Moreover, the lemma shows that the whole population after this time will be at least as good as the initially best one, which proves that the worst TwoMAX-value will be at least  $n/2 + D_{(1)} = n/2 + \Omega(n^{1/2-\epsilon})$  at the time one branch dies out.

Finally, using the family tree technique in the same way as for Lemma 5, we note that the best TwoMAX-value in the population will not increase by more than  $n^{1/8} \log \mu$  in  $n^{1/8} \mu \log \mu$  generations with probability  $1 - o(1)$ . Using that  $D_{(1)} = O(\sqrt{n}(\log \mu + \log n))$  with probability  $1 - o(1)$  (which follows from simple Chernoff and union bounds), we obtain that with probability  $1 - o(1)$ , the best TwoMAX-value is  $O(n/2 + c\sqrt{n}(\log n + \log \mu))$  at the time one branch dies out, as claimed in the theorem statement.  $\square$

We can straightforwardly generalize Theorem 3 to TRUNCATEDTWO MAX if the maximum fitness of the truncated branch lies well above the best fitness reached after a branch has taken over. The theorem considers a TwoMAX-value of  $n/2 + c\sqrt{n} \log(\mu n)$ , so in principle we only have to choose  $k$  slightly larger than the offset  $c\sqrt{n} \log(\mu n)$ . For convenience, we formulate our corollary with  $k = \omega(\sqrt{n} \log n)$ .

**COROLLARY 2.** *Let  $\mu \leq n^{k'}$  for an arbitrary constant  $k' \geq 0$  as well as  $\mu = \omega(1)$  and consider the  $(\mu+1)$  EA on TRUNCATEDTWO MAX with  $k = \omega(\sqrt{n} \log n)$ . Let  $\epsilon > 0$  be an arbitrary constant. Then with probability  $1/2 - o(1)$  the whole population is on the truncated branch before the global optimum is reached, and conditional on this event the optimisation time is  $2^{\Omega(n^{1/2-\epsilon})}$  w. o. p.*

### 3.1 RLS Mutation

We can obtain a result analogous to Theorem 3 also for the  $(\mu+1)$  RLS which differs from the  $(\mu+1)$  EA only in the mutation operator. Basically, the only difference is in the analysis of take-over times, which originally refer to the time until the whole population descends from a certain individual and is at least as good as it. For the  $(\mu+1)$  EA using standard-bit mutation, the  $O(\mu \log \mu)$  bound is based on the fact that an individual can produce a copy of itself with probability  $(1 - 1/n)^n \approx e^{-1}$ . This is not the case for the one-bit mutation of the  $(\mu+1)$  RLS, which either improves or degrades individuals' fitness. Nevertheless, we can use the same type of analysis if the best and second-best individual are sufficiently large apart and the offspring of the best do not become worse than the offspring of the second-best within a certain period of time.

More formally, let  $x^*$  be the initially best individual and let after  $t$  generations the set  $D_t$  denote the individuals in the population at time  $t$  that are offspring of  $x^*$ . Given  $|D_t| = k < \mu$ , the probability of choosing from  $D_t$  is  $k/\mu$  and, if this happens and the offspring is accepted, it holds  $|D_{t+1}| = k + 1$ . Hence, assuming all these offspring to be accepted, the time until the whole population descends

from  $x^*$  has expected value  $O(\mu \log \mu)$ . Note that this does not imply that all descendants are at least as good as  $x^*$  in fitness.

Let  $R^*$  be the other  $\mu - 1$  individuals from the initial population. Offspring of  $x^*$  are guaranteed to be accepted in the proposed time of expected  $O(\mu \log \mu)$  generations if

- no offspring of individuals in  $R^*$  reaches fitness at least  $(D_{(1)} + D_{(2)})/2$
- no offspring of  $x^*$  has fitness less than  $(D_{(1)} + D_{(2)})/2$

For either event to fail, the progress must be at least  $(D_{(1)} - D_{(2)})/2$  in at least one lineage of one of the  $\mu$  family trees stemming from the individuals of the initial population. Using Lemma 5 to bound the depth of each family tree by  $n^{\epsilon/2}$  and observing that the progress along any lineage due to one-bit mutation is bounded by the length of the lineage, we obtain a progress of at most  $n^{\epsilon/2}$  with probability  $1 - o(1)$  during the takeover time. This is clearly smaller than the bound  $(D_{(1)} - D_{(2)})/2 = n^{\epsilon/2}$  that results from uniform initialization according to Lemma 3 with probability  $1 - o(1)$ . Hence, we have obtained the following result:

**THEOREM 4.** *Let  $\mu \leq n^k$  for an arbitrary constant  $k \geq 0$  as well as  $\mu = \omega(1)$  and consider the  $(\mu+1)$  RLS on TwoMAX. With probability  $1 - o(1)$ , all individuals are on one branch before the TwoMAX-value reaches  $n/2 + c\sqrt{n}(\log n + \log \mu)$ , for some sufficiently large constant  $c > 0$ , and the worst TwoMAX-value in the population is  $n/2 + \Omega(n^{1/2-\epsilon})$  at the time one branch dies out. If this happens, individuals from the other branch cannot be generated any more.*

#### 4 LOW SELECTIVE PRESSURE SUCCEEDS

In this section we will show that reducing the selection pressure to its minimum by always picking the worst parent for reproduction makes the  $(\mu+1)$  EA efficient at locating both optima of TwoMAX with reasonable population sizes. In the following section we will argue why such a strategy is not sufficient for efficient global optimisation in general, and propose necessary corrections. We begin by analysing the standard bit mutation version, the  $(\mu+1)$  EA, and afterwards we consider the RLS version. At the core of both proofs is the analysis of the stochastic process  $D_t$  which keeps track of the difference between the number of individuals with more 1-bits and more 0-bits. The crucial observation is that, with inverse elitist parent selection,  $D_t$  performs a fair random walk which starts close to 0 and, for large enough  $\mu$ , cannot reach w.o.p. any of the two boundaries ( $\mu$  or  $-\mu$ ) before both optima are found.

**THEOREM 5.** *The  $(\mu+1)$  EA with  $\mu = \Omega(n^{1+2\epsilon})$  and inversely elitist parent selection finds both optima of TwoMAX in  $O(\mu n^{1+\epsilon})$  fitness function evaluations with probability  $1 - e^{-\Omega(n^\epsilon)}$  for any positive constant  $\epsilon$  such that  $1 - \epsilon = \Theta(1)$ .*

**PROOF.** We consider the population  $P_t$  at time (i. e., iteration)  $t$  and its mutually exclusive sub-populations  $P_0^t$  ( $P_1^t$ ) which consist of individuals with more 0-bits (1-bits) in their bitstrings. Let  $D_t := |P_0^t| - |P_1^t|$  denote the difference between the sizes of the sub-populations at time  $t$  and  $\mathcal{T}$  be the smallest  $t$  such that  $P_t$  consists only of  $1^n$  or  $0^n$  bit-strings. We note here that  $|D_t| = \mu$  if and only if  $P_t = P_0^t$  or  $P_t = P_1^t$ . Consequently,  $|D_{\mathcal{T}}| < \mu$  implies that both optima are found at time  $\mathcal{T}$ .

First of all, we will prove the upper bound  $\mathcal{T} < \mu n^{1+\epsilon}$  which holds with probability at least  $1 - e^{-\Omega(n^\epsilon)}$  for any constant  $\epsilon > 0$ . Then, in order to show that  $|D_t| < \mu$  for any  $t \in [\mu n^{1+\epsilon}]$ , we will first show that,  $|D_0| < \mu^{\frac{1+\epsilon}{2}}$  with probability at least  $1 - e^{-\Omega(\mu^{2\epsilon})}$ , and then finally show that for any  $t \in [\mu n^{1+\epsilon}]$  the probability that  $|D_t| - |D_0| > \mu - \mu^{\frac{1+\epsilon}{2}}$  is at most  $e^{-\Omega(n^\epsilon)}$ . Using a union bound over (1) the probability that  $\mathcal{T} > \mu n^{1+\epsilon}$ , (2) the probability that the initial difference  $|D_0| > \mu^{\frac{1+\epsilon}{2}}$  (3) all the probabilities  $|D_t| - |D_0| > \mu - \mu^{\frac{1+\epsilon}{2}}$  for  $t \in [\mu n^{1+\epsilon}]$ ; we will obtain our claim.

Given that the best individual in the population has fitness value  $i$ , the probability that a solution with fitness at least  $i + 1$  will be sampled in the next iteration is at least  $(n - i)/(\mu n)$ . Since

$\sum_{i=n/2}^{n-1} p_i^{-1} \leq \sum_{i=n/2}^{n-1} \frac{\mu \epsilon n}{n-i} \leq \mu \epsilon n \ln n$ , we can use Theorem 2 to bound the probability that  $\mathcal{T}$  is at least  $\mu \epsilon n^{1+\epsilon}$  from below.

$$\begin{aligned} \Pr(\mathcal{T} > \mu \epsilon n^{1+\epsilon}) &\leq \Pr\left(\mathcal{T} > \sum_{i=n/2}^{n-1} p_i^{-1} + \mu \epsilon n^{1+\epsilon} - \mu \epsilon n \ln n\right) \\ &\leq \exp\left(-\frac{\mu \epsilon n (n^\epsilon - \ln n)}{n}\right) = \exp(-\Omega(\mu \epsilon n^\epsilon)) \end{aligned}$$

From now on, we will assume that  $\mathcal{T} \leq \mu \epsilon n^{1+\epsilon}$  and account for the case of  $\mathcal{T} > \mu \epsilon n^{1+\epsilon}$  as part of the overall failure probability  $e^{-\Omega(n^\epsilon)}$ .

Next, we will show that, with probability  $e^{-\Omega(n^\epsilon)}$ ,  $\mu \epsilon n^{1+\epsilon}$  iterations will not be enough for one of the branches to take over the population.  $D_t$  changes only when a solution from  $P_0^t$  ( $P_1^t$ ) is selected for mutation and then an individual from the opposite sub-population  $P_1^t$  ( $P_0^t$ ) is selected for removal. Since the parent and the environmental selections are identical and independent, the probability  $p^+$  of selecting from  $P_0^t$  as parent and from  $P_1^t$  for removing is equal to the probability  $p^-$  of selecting the parent from the  $P_1^t$  and then removing from  $P_0^t$ . Moreover, all parent candidates have the same probability of being improved and the same probability of being copied, thus,

$$\Pr(D_{t+1} = D_t + 2 \mid D_{t+1} \neq D_t) = \Pr(D_{t+1} = D_t - 2 \mid D_{t+1} \neq D_t) = 1/2$$

Let a relevant step be any  $t$  such that  $D_t \neq D_{t+1}$  and  $\pi_k \in [T]$  denote the actual iteration number of the  $k$ th relevant step. Then,  $D_{\pi_{k+1}} - D_{\pi_k} \sim 4 \cdot \text{Bernoulli}(\frac{1}{2}) - 2$  and consequently  $D_{\pi_k} - D_0 \sim 4 \cdot \text{Bin}(k, \frac{1}{2}) - 2 \cdot k$ . Since the initial population consists of individuals sampled uniformly at random, the initial difference  $D_0 \sim 2 \cdot \text{Bin}(\mu, 1/2) - \mu$  and  $\Pr(D_0 > \mu^{\frac{1+\epsilon}{2}}) = \exp(-\Omega(\mu^\epsilon))$ . Symmetrically we also have  $\Pr(D_0 < -\mu^{\frac{1+\epsilon}{2}}) = \exp(-\Omega(\mu^\epsilon))$  and can use the union bound to obtain  $\Pr(|D_0| < \mu^{\frac{1+\epsilon}{2}}) = \exp(-\Omega(\mu^\epsilon))$ . Therefore, for any relevant step  $k$  and the random variable  $X \sim \text{Bin}(k, \frac{1}{2})$ ;

$$\begin{aligned} \Pr(|D_{\pi_k}| = \mu \mid |D_0| < \mu^{\frac{1+\epsilon}{2}}) &\leq \Pr(|D_{\pi_k}| - |D_0| \geq \mu - \mu^{\frac{1+\epsilon}{2}}) \leq \Pr(D_{\pi_k} - D_0 \geq \mu - \mu^{\frac{1+\epsilon}{2}}) \\ &= \Pr(4 \cdot X - 2 \cdot k \geq \mu - \mu^{\frac{1+\epsilon}{2}}) = \Pr\left(X \geq \frac{\mu - \mu^{\frac{1+\epsilon}{2}} + 2k}{4}\right) \\ &= \Pr\left(X \geq \frac{\mu - \mu^{\frac{1+\epsilon}{2}}}{4} + E[X]\right) \leq \exp\left(-\frac{2 \cdot \left(\mu - \mu^{\frac{1+\epsilon}{2}}\right)^2}{16k}\right) \end{aligned}$$

where in the last step we used Theorem 1. Taking the assumptions  $1 - \epsilon = \Theta(1)$ ,  $\mu = \Omega(n^{1+2\epsilon})$  and  $k \leq \mu \epsilon n^{1+\epsilon}$ , it holds that

$$\exp\left(-\frac{2 \cdot \left(\mu - \mu^{\frac{1+\epsilon}{2}}\right)^2}{16k}\right) = \exp\left(-\Omega\left(\frac{\mu^2}{k}\right)\right) = \exp\left(-\Omega\left(\frac{\mu}{n^{1+\epsilon}}\right)\right) = \exp(-\Omega(n^\epsilon)).$$

Note that the above probability is the probability that a single  $k \in [\mu \epsilon n^{1+\epsilon}]$  would have  $|D_{\pi_k}| \geq \mu$ . To get our final results, we will use the union bound to aggregate the probabilities of failure for all

On Steady-State Evolutionary Algorithms and Selective Pressure: Why Inverse Rank-Based Allocation of Reproductive Trials is Best

15

such  $k$  along with the probability that  $\mathcal{T} > e\mu n^{1+\epsilon}$  and that  $|D_0| > \mu^{\frac{1+\epsilon}{2}}$ .

$$\begin{aligned} \Pr(|D_{\mathcal{T}}| \neq \mu) &> 1 - e\mu n^{1+\epsilon} \exp(-\Omega(n^\epsilon)) - \exp(-\Omega(n^\epsilon)) - \exp(-\Omega(\mu n^\epsilon)) \\ &= 1 - \exp(-\Omega(n^\epsilon)). \end{aligned}$$

□

**THEOREM 6.** *For any arbitrarily small constant  $\epsilon > 0$ , the  $(\mu+1)$  RLS with inversely elitist parent selection and  $\mu = \Omega(n^{1+\epsilon})$  finds both optima of TwOMAX in  $O(\mu n \log n)$  fitness function evaluations with probability  $1 - e^{-\Omega(n^\epsilon)}$ .*

**PROOF.** We will follow the same proof idea and the definition of relevant step as in Theorem 5, however, since RLS cannot create copies, the number  $\mathcal{K}$  of relevant steps in a run is bounded from above by  $\mu \cdot n$  because a solution is improved in each relevant step and each individual can be improved at most  $n$  times. We will denote the actual iteration number of the  $k$ th relevant step as  $\pi_k \in [T]$ . At every relevant step  $k$ , the difference  $D_k := |P_0^{\pi_k}| - |P_1^{\pi_k}|$  either increases or decreases by two with equal probability  $1/2$ . Formally,  $D_{k+1} - D_k \sim 2 \cdot \text{Bernoulli}(\frac{1}{2}) - 1$  and consequently  $D_k - D_0 \sim 2 \cdot \text{Bin}(k, \frac{1}{2}) - k$ . As in the  $(\mu+1)$  EA, the initial  $D_0$  satisfies  $\Pr(|D_0| < \mu^{\frac{1+\epsilon}{2}}) = \exp(-\Omega(\mu^{2\epsilon}))$ . Therefore, for any relevant step  $k$ ,

$$\begin{aligned} \Pr(|D_k| = \mu \mid |D_0| < \mu^{\frac{1+\epsilon}{2}}) &\leq \Pr(|D_k| - |D_0| \geq \mu - \mu^{\frac{1+\epsilon}{2}}) \leq \Pr(D_k - D_0 \geq \mu - \mu^{\frac{1+\epsilon}{2}}) \\ &= \Pr(2 \cdot X - k \geq \mu - \mu^{\frac{1+\epsilon}{2}}) = \Pr\left(X \geq \frac{\mu - \mu^{\frac{1+\epsilon}{2}} + k}{2}\right) = \Pr\left(X \geq \frac{\mu - \mu^{\frac{1+\epsilon}{2}}}{2} + E[X]\right) \\ &\leq \exp\left(-\frac{2 \cdot \left(\mu - \mu^{\frac{1+\epsilon}{2}}\right)^2}{\frac{k}{2}}\right) = \exp\left(-\Omega\left(\frac{(n^{1+\epsilon})^2}{n^{1+\epsilon}}\right)\right) = \exp(-\Omega(n^\epsilon)) \end{aligned}$$

where  $X \sim \text{Bin}(k, \frac{1}{2})$  and in the last step we used the observation that  $k \leq \mu \cdot n$  and the assumption that  $\mu = \Omega(n^{1+\epsilon})$ . Using the above bound and the union bound over  $k \in [\mathcal{K}]$  we get our final result:

$$\Pr(\forall k \in [\mathcal{K}] \mid D_k| \neq \mu \mid |D_0| < \mu^{\frac{1+\epsilon}{2}}) \geq 1 - \mu \cdot n \cdot \exp(-\Omega(n^\epsilon)) = 1 - \exp(-\Omega(n^\epsilon))$$

Thus, with probability  $1 - \exp(-\Omega(n^\epsilon))$ ,  $|D_k| < \mu$  for all  $k$  and the population is not taken over by any of the branches before all the solutions reach the optimal fitness. □

## 5 NOT TOO HIGH, NOR TOO LOW PRESSURE: EXPERIMENTAL SUPPLEMENTS

It is easy to see that the Inverse Elitist selection operator considered in the previous section may struggle to identify the global optimum if the worst individual is trapped on a low quality local optimum. In such a case it will be selected until it escapes which may require prohibitive time.

In this section we present an experimental analysis of the  $(\mu+1)$  EA using the selection operators considered in the previous sections (i.e., too high and extremely low selective pressure) together with a range of intermediate selective pressures that are higher than Inverse Elitist but still considerably lower than Uniform selection. To achieve such pressures we implement *Inverse Tournament selection* as defined in Section 2 with small tournament sizes<sup>5</sup>. Our aim is to show that the latter operator is preferable to both Uniform and Inversely Elitist selection because it has low premature convergence for TwOMAX and it can effectively escape from local optima as experiments for TRUNCATEDTwOMAX<sub>k</sub> will show.

<sup>5</sup>Naturally the same selection probabilities can be achieved with rank selection.

16

Dogan Corus, Andrei Lissovoi, Pietro S. Oliveto, and Carsten Witt

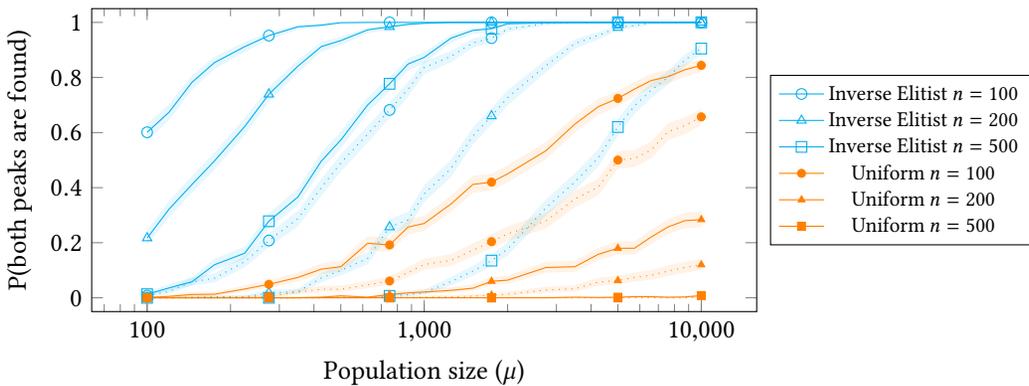


Fig. 2. Proportion of independent runs finding both peaks on TwoMax for various combinations of parent selection operators, mutation operators (solid:  $RLS_1$ , dotted: standard bit mutation), and problem sizes.

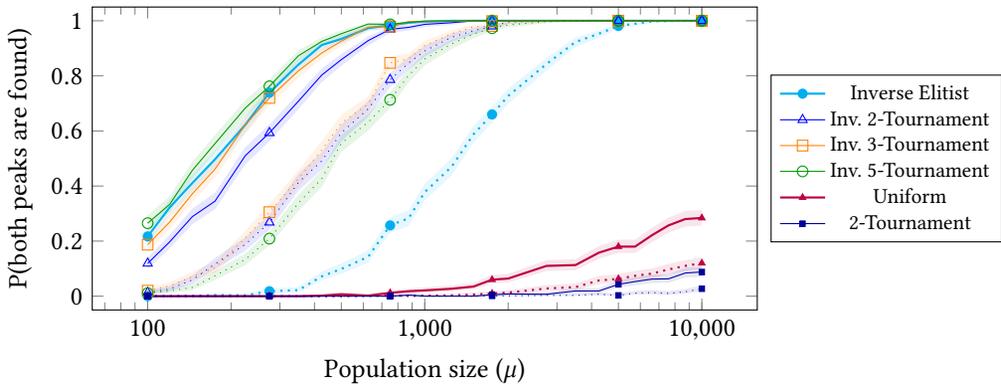


Fig. 3. Proportion of independent runs finding both peaks on TwoMax ( $n = 200$ ) for various combinations of parent selection and mutation operators (solid:  $RLS_1$ , dotted: standard bit mutation).

The results presented in this section are based on performing 1000 independent runs for each combination of problem size and algorithm parameters, and observing whether the  $(\mu+1)$  EA is able to construct both optima during the search process. The presented figures estimate the probability that the algorithms find both optima by the proportion of the independent runs which do so. Additionally, a 95% confidence interval is displayed as a shaded region around the estimated probabilities.

We start by examining the success probability of the analysed selection operators on TwoMax for small problem sizes. Figure 2 compares the performance of Inverse Elitist and Uniform parent selection operators for various population and problem sizes. As expected, Uniform parent selection requires a larger population size than Inverse Elitist parent selection in order to achieve the same probability of finding both peaks; for the considered problem sizes,  $\mu = n^{1.5}$  appears sufficient for the combination of Inverse Elitist parent selection and  $RLS_1$  mutation to achieve a high probability of discovering both peaks, while SBM mutation requires a slightly larger population size, and Uniform parent selection requires a significantly larger population size.

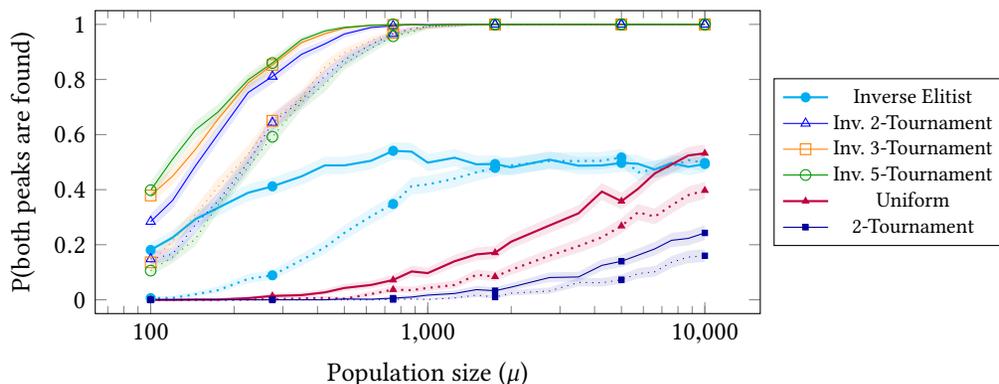


Fig. 4. Proportion of independent runs finding both peaks on  $\text{TRUNCATEDTWO MAX}_k$  ( $n = 200$ ,  $k = 80$ ) for various combinations of parent selection operators and mutation operators (solid:  $\text{RLS}_1$ , dotted: standard bit mutation).

Figure 3 compares the performance of Tournament and Inverse Tournament parent selection operators with that of Inverse Elitist and Uniform parent selection operators for a fixed problem size and varying population sizes. Generally, modest Inverse Tournament sizes (such as  $k = 3$ ) are sufficient to achieve similar-or-better probability of discovering both optima compared to Inverse Elitist parent selection when using  $\text{RLS}$  mutation, and can even outperform Inverse Elitist parent selection when using standard bit mutation. As expected, the higher selective pressure of the traditional Tournament selection results in faster takeover by a single branch, and hence a lower probability of finding both optima.

For  $\text{TRUNCATEDTWO MAX}_k$ , where the left branch terminates at an earlier local optimum, Figure 4 shows that with sufficiently high population sizes, Inverse Elitist parent selection is only able to find both peaks with probability at most  $1/2$ , while the other parent selection mechanisms are able to achieve higher probabilities of finding both peaks compared to the non-truncated setting of Figure 3. Additionally, Inverse Tournament Selection suffers a runtime penalty in this setting: when the time comes to remove from the population the local optimum on the truncated branch, larger tournament sizes result in higher waiting times to select a suitable parent to mutate. This is particularly pronounced when the population size only provides a moderate (rather than high) probability of finding both peaks. If the population is almost entirely on the truncated branch when that local optimum is the lowest-fitness individual in the population, selecting a parent from the other branch could require up to  $O(\mu^s)$  iterations, where  $s$  is the employed Inverse Tournament size, and would need to be repeated many times to restore a reasonable probability of an iteration making progress toward the global optimum. We conclude that small (i.e., constant) inverse tournament sizes should be used in practice.

For the problems considered up to now, the use of single trajectory algorithms coupled with restart strategies suffice for efficient optimisation. In the following sections we consider problems from the literature where restart strategies are ineffective.

## 6 WHERE POPULATIONS ARE ESSENTIAL BECAUSE RESTART STRATEGIES FAIL

In the previous sections we have shown that decreasing the selective pressure below uniform, increases the exploration capabilities of the  $(\mu+1)$  EA. Essentially lower selective pressure increases the takeover time, effectively allowing the population to explore the two slopes of  $\text{TWO MAX}$ .

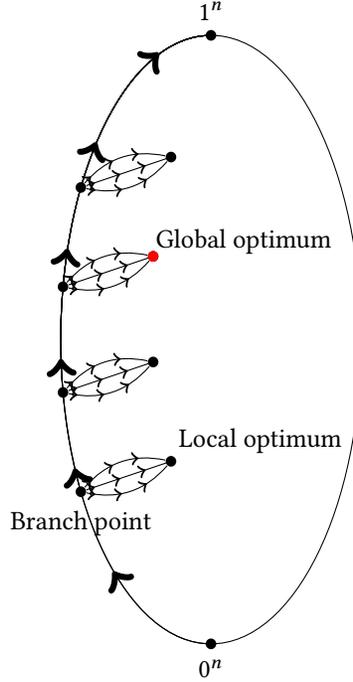
However, when multiple slopes can be identified with high probability at initialisation as for TwoMAX, then single trajectory algorithms with appropriate restart strategies can be just as, or even more, effective.

In this section we consider two benchmark function classes from the literature originally designed to highlight circumstances where populations are essential and single trajectory algorithms fail w.o.p. The two function classes have different characteristics. The first one, SUFSAMP, was introduced to show situations when offspring populations outperform single trajectory algorithms [Jansen et al. 2005]. The function consists of a main steep path leading to the global optimum with many branching points located at regular intervals. At each such point a branch leading towards a different local optimum may be followed. While single trajectory algorithms are tricked into following one of the branches, algorithms using large enough offspring populations can identify the steeper slope by sampling many solutions around each branching point. We will refer to this function using a more descriptive name: RIDGEWITHBRANCHES. The other function class, TWOGRADIENTS, is inspired by one designed to highlight when the use of parent populations is essential [Witt 2006]. It consists of two different paths: an easier one for hillclimbers leading to local optima and a slower one which leads to the global optimum. While single trajectory algorithms quickly optimise the easier path, the slower progress of parent populations on the easy path allows the  $(\mu+1)$  EA to follow the not-so-steep gradient to the global optimum. Since for both function classes, single trajectory algorithms are attracted w.o.p. to the local optima, even restarts do not allow them to identify the global optimum in polynomial time. Apart from proving that the  $(\mu+1)$  EA with the inverse selection operators optimises the problems efficiently, in this section we will also show, either through rigorous proofs or by experimental evidence, that it outperforms the same algorithm using uniform selection: the inverse selection algorithms can identify all the optima with high probability, hence optimise the functions efficiently independently of which local optimum is designated to be the global one. On the other hand, single trajectory algorithms and the uniform selection  $(\mu+1)$  EA can be trapped w.o.p. according to which optimum is the global one and even restarts cannot help them to optimise the functions efficiently.

## 6.1 A Function Class with Many Slopes

In this section we aim to highlight an important exploration characteristic of the  $(\mu+1)$  EA with inverse tournament selection that single trajectory algorithms using restarts or population-based EAs with deterministic crowding do not have. In the latter algorithms each individual explores the search space independently and no communication occurs between individuals that explore different trajectories. Differently, inverse selection does not completely isolate the individuals in the population from each other. If individuals are located on different slopes, it is more likely that slopes of lower fitness are explored first. Once the local optima of the slopes are identified, individuals on slopes of higher fitness will eventually take over and the search may proceed in other unexplored directions. This property allows an EA with inverse selection to exhaust the improvement potential of many of the gradients it encounters with good probability. In particular, it allows the algorithm to explore gradients which more greedy search strategies (e.g., using offspring populations or parent populations with low take-over times) are likely to ignore, thus identify optima that popular algorithms will struggle to find. To demonstrate this behaviour, we consider the RIDGEWITHBRANCHES function which has many slopes leading to different local optima. It was originally introduced to highlight problem characteristics where offspring populations are useful [Jansen et al. 2005]. It is defined as follows:

Fig. 5. Visual representation of RIDGEWITHBRANCHES. Arrows indicate the direction of increasing fitness while multiple edges represent different ways the branch points can be improved towards the local optima.



DEFINITION 1. For  $n \in \mathbb{N}$  we define  $k := \lfloor \sqrt{n} \rfloor$ . We use  $|x| = \text{ONEMAX}(x)$  and define the function  $g : \{0, 1\}^n \rightarrow \mathbb{N}$  for all  $x \in \{0, 1\}^n$  by

$$g(x) := \begin{cases} (i + 3)n + |x| & \text{if } (x = 0^{n-i}1^i \text{ with } 0 \leq i \leq n) \text{ or} \\ & (x = y0^{n-i-k}1^i \text{ with} \\ & i \in \{k, 2k, \dots, (k-2)k\}, \\ & y \in \{0, 1\}^k) \\ 0 & \text{otherwise,} \end{cases}$$

for all  $x = x_1x_2 \cdots x_n \in \{0, 1\}^n$  with  $x' := x_1x_2 \cdots x_{m'} \in \{0, 1\}^{m'}$  and  $x'' := x_1x_2 \cdots x_{m''} \in \{0, 1\}^{m''}$ .

We will use the function  $g$  to give RIDGEWITHBRANCHES its characteristic structure. It is essentially a Ridge function (where for  $i \in \{0, \dots, n\}$  Ridge solutions of the form  $0^{n-i}1^i$  are assigned fitness values which monotonically increase with  $i$ ) with special branch points of the form  $0^{n-j \cdot k}1^{j \cdot k}$  for any  $j \in \{1, \dots, k-1\}$  on its path at every  $k = \lfloor \sqrt{n} \rfloor$  additional trailing ones. These branch points can either be improved by increasing the ONEMAX value of its prefix (thus, obtaining an off-branch solution of the form  $y0^{n-(j+2)k}1^{(j+1) \cdot k}$  where  $y \neq 0^k$ ) or by flipping the first zero bit before its trailing ones. While improving the ONEMAX part of the function is asymptotically more likely, flipping the Ridge bit increases the function value more than it is possible by maximising the ONEMAX part in the prefix. Moreover, flipping the particular Ridge bit allows the algorithm to follow the gradient to the next branch point. In its current form  $g$  does not provide any guidance

to the algorithm for finding the first point on the RIDGE. We add this extra gradient to obtain the RIDGEWITHBRANCHES function as defined in [Jansen et al. 2005].

**DEFINITION 2.** For  $n \in \mathbb{N}$  we define  $m' := \lfloor n/2 \rfloor$  and  $m'' := \lceil n/2 \rceil$ , and  $RIDGEWITHBRANCHES : \{0, 1\}^n \rightarrow \mathbb{N}$

$$RIDGEWITHBRANCHES(x) := \begin{cases} n - OM(x'') & \text{if } x' \neq 0^{m'} \wedge x' \neq 0^{m''} \\ 2n - OM(x') & \text{if } x' \neq 0^{m'} \wedge x' = 0^{m''} \\ g(x'') & \text{if } x' = 0^{m'} \end{cases}$$

for all  $x = x_1 x_2 \cdots x_n \in \{0, 1\}^n$  with  $x' := x_1 x_2 \cdots x_{m'} \in \{0, 1\}^{m'}$  and  $x'' := x_1 x_2 \cdots x_{m''} \in \{0, 1\}^{m''}$ .

Since it is more likely to flip a ONEMAX bit rather than the Ridge bit at each branching point, single trajectory individuals, such as RLS and the (1+1) EA, get stuck on a local optimum w.o.p., hence restarts cannot help either. On the other hand, an offspring population algorithm such as the (1+ $\lambda$ ) EA with sufficiently large  $\lambda$ , samples the Ridge bit at each branching point w.o.p., hence efficiently follows the slope towards the global optimum.

We generalise the function class to make it considerably harder for evolutionary algorithms by including  $k - 1$  different functions, each one differing by the position of the global optimum. In particular, we introduce a parameter  $j \in \{1, \dots, k - 1\}$  and for  $RIDGEWITHBRANCHES_j$  the optimal fitness value is assigned to the local optimum with  $j \cdot k$  trailing ones. The function class is illustrated in Fig. 5.

**DEFINITION 3.** For  $k \in \mathbb{N}$ ,  $m' = m'' = k^2$ ,  $n = m' + m''$  and  $j \in \{1, \dots, k - 1\}$ , we define  $RIDGEWITHBRANCHES_j(x) : \{0, 1\}^n \rightarrow \mathbb{N}$

$$RIDGEWITHBRANCHES_j(x) := \begin{cases} n^3 & \text{if } x = 0^{m'} 1^k 0^{m'' - (j+1) \cdot k} 1^{j \cdot k} \\ RIDGEWITHBRANCHES(x) & \text{otherwise} \end{cases}$$

As a result we obtain a function class that none of the traditional EAs, the restart based approaches or the inverse elitist selection ( $\mu+1$ ) EA can optimise efficiently. The (1 +  $\lambda$ ) EA with an offspring population of size  $\Omega(n \log n)$  ignores the milder gradients and with probability arbitrarily close to 1 finds a local optimum (i.e. the global optimum of the original RIDGEWITHBRANCHES) [Jansen et al. 2005]. When the offspring population size is slightly increased to  $\Omega(n^{1+\epsilon})$  for any arbitrarily small constant  $\epsilon$ , then the probability that (1 +  $\lambda$ ) EA follows any of the milder gradients becomes exponentially small and the global optimum of  $RIDGEWITHBRANCHES_j$  cannot be found in polynomial time for any  $j \neq k - 1$  w.o.p. On the other hand, EAs with small populations (whether it is a single run or a restart based implementation) have a high probability of following any of the mild gradients they encounter and thus have an exponentially small probability of finding the optima that require a super-constant number of mild gradients to be ignored. If the population is large enough to ignore the mild gradients with high probability, then the ( $\mu+1$ ) EA is very likely to end up on the original global optimum at the end of the Ridge. Noting that  $RIDGEWITHBRANCHES_{k-1}$  is identical to the original function, we first extend the negative result presented in [Jansen et al. 2005] with a proof sketch that adapts it to  $RIDGEWITHBRANCHES_j$ .

**THEOREM 7.** The probability that the (1+1) EA optimises the function  $RIDGEWITHBRANCHES_j$  for any  $j > k/2$  within  $n^{O(1)}$  function evaluations is bounded above by  $2^{-\Omega(\sqrt{n})}$ .

We follow the proof of Theorem 7 in [Jansen et al. 2005] to establish that at every branch point, the (1+1) EA samples a locally optimal solution with  $k$  leading ones before it samples a new ridge solution with probability  $1 - O(1/\sqrt{n})$ . Since  $j > k/2$  implies that the best fitness value of

RIDGEWITHBRANCHES<sub>j</sub> is assigned to a local optimum with at least  $m''/2$  trailing ones, we multiply the failure probabilities of only the first  $\sqrt{n}/4$  branch points (rather than all of them as in [Jansen et al. 2005]) and obtain the same asymptotic failure probability  $2^{-\Omega(\sqrt{n} \log n)}$ . This allows us to report the same overall failure probability as in [Jansen et al. 2005] which stems from the probability that the first ridge point sampled by the (1+1) EA has more than  $2k$  trailing ones.

The exponentially small probability given in Theorem 7 implies that at least an exponentially large number of restarts of the (1+1) EA are necessary to observe at least one run that identifies the global optimum in polynomial time. Thus, any restart strategy using the (1+1) EA as the base algorithm, whether the runs are in parallel or sequential, cannot solve RIDGEWITHBRANCHES in polynomial time. Hence, also a  $(\mu+1)$  EA using deterministic crowding as diversity mechanism fails w.o.p. We now present the main result of this subsection which determines that the  $(\mu+1)$  EA with inverse binary tournament selection identifies all the optima of RIDGEWITHBRANCHES efficiently, hence also the global one wherever it is designated to be.

**THEOREM 8.** *The probability that the  $(\mu+1)$  EA with polynomial population size  $\mu = \omega(n^3)$  and inverse binary tournament selection, optimises the function RIDGEWITHBRANCHES<sub>j</sub> for any  $j \in \{1, \dots, k-1\}$  in  $O(\mu^2 \sqrt{n})$  fitness function evaluations is at least  $1 - 2^{-\Omega(\sqrt{n})}$ .*

**PROOF.** We will initially follow some steps from the proof of Theorem 8 in [Jansen et al. 2005]. A randomly initialised bit-string satisfies  $x' \neq 0^{m'}$  and  $m''/3 < \text{ONEMAX}(x'') < 2m''/3$  with probability  $1 - 2^{-\Omega(\sqrt{n})}$ . The probability that all individuals in the initial population satisfy the same conditions has the same asymptotic order of  $1 - 2^{-\Omega(\sqrt{n})}$  due to the union bound. For any  $j \in \{1, \dots, k-2\}$ , let  $b_j$  to denote the  $j$ th branch point  $0^{m'} 0^{m''-j \cdot k} 1^{j \cdot k}$ . Given that the initial population satisfies the conditions and no bit-strings with higher fitness are found before, the expected time until the first branch point  $b_1 := 0^{m'} 0^{m''-k} 1^k$  is sampled for the first time is less than the sum of the expected  $O(\mu n \log n)$  evaluations for finding the  $0^n$  bit-string and expected  $O(\mu n^{3/2})$  fitness function evaluations for improving the function  $g$ ,  $2k = \Theta(\sqrt{n})$  times for each individual. Consequently, the probability that the actual time until  $b_1$  is sampled is  $O(\mu n^2 \log n)$  is  $1 - 2^{-\Omega(\sqrt{n})}$ . We will next bound the probability of sampling a solution  $y$  with  $\text{RIDGEWITHBRANCHES}(y) > \text{RIDGEWITHBRANCHES}(b_1)$  before sampling  $b_1$ . Considering levels of strings with a given number of zero-bits, due to symmetry, each bit-string on a level have the same probability of being sampled before the  $0^n$  bitstring, the first Ridge solution, is sampled. Between levels  $m''/3$  and  $m'' - 2k$ , there are  $2^{\Omega(k)}$  bit-strings in total and the probability of sampling a Ridge solution on these levels in the first  $O(\mu n^2 \log n)$  steps is at most  $2^{-\Omega(\sqrt{n})}$  by the union bound. Therefore, with probability  $1 - 2^{-\Omega(\sqrt{n})}$  our population will consist only of Ridge solutions with worse fitness than  $b_1$  at some point during optimisation.

Next, we will consider a phase where, for some fixed  $j \in \{0, \dots, k-2\}$ , all initial individuals in the population are Ridge solutions with more than  $j \cdot k$  and less than  $(j+1) \cdot k$  trailing ones. We will now bound the probability that the off-branch solutions take over the population before the Ridge solution  $0^{m'} 0^{m''-(j+1) \cdot k-1} 1^{(j+1) \cdot k+1}$  is sampled for the first time. Since this Ridge solution has higher fitness value than all off-branch solutions with smaller number of trailing ones, once created, this solution (which we will refer as a *safe* solution) cannot be removed by adding more off-branch solutions to the population.

It is exponentially unlikely that an off-branch solution of the form  $0^{m'} y 0^{m''-(j+2)k} 1^{(j+1) \cdot k}$  is created from a parent with less than  $j \cdot k$  trailing ones. In particular, this event requires at least  $\Omega(\sqrt{n})$  specific bits to be flipped which occurs with probability at most  $n^{-\Omega(\sqrt{n})}$ . We will condition on this event never happening and absorb this probability into our failure probability of  $2^{-\Omega(\sqrt{n})}$  in our claim. Thus, we can now categorise the newly created off-branch solutions into two types

according to their parents: (1) whose parents are Ridge solutions with at least  $j \cdot k$  trailing ones and (2) whose parents are also off-branch solutions with  $(j + 1) \cdot k$  trailing ones. From a Ridge solution, the conditional probability of creating a safe solution given that either an off-branch solution or a safe solution has been created is  $\Theta(1/\sqrt{n})$  since only  $\Theta(\sqrt{n})$  bits contribute to the additive  $|x|$  term in the definition of  $g$ . Thus, with probability  $1 - 2^{-\Omega(\sqrt{n})}$  we obtain a safe solution before obtaining  $n/2$  type-1 off-branch solutions.

Now, we turn to the off-branch solutions whose parents are also off-branch solutions (type-2). Given that there are  $k$  off-branch individuals in the population and no safe solutions, the probability of selecting an off-branch solution as parent is at most  $(k/\mu)^2$  since the off-branch solution have higher fitness value than the rest of the population and inverse tournament selection returns the candidate with the lowest fitness.

Pessimistically assuming that the off-branch sub-population initially has  $n/2$  individuals mutated from Ridge solutions, we will bound the time until the total number of off-branch solution increases from  $n/2$  to  $n$ , *i.e.*, we create  $n/2$  type-2 individuals. We will pessimistically assume that the probability of picking an off-branch individual is  $(n/\mu)^2$ , thus, using Chernoff bounds on geometric variables [Doerr 2011], it takes at least,  $\frac{n \cdot (\mu/n)^2}{3} = \Omega(\mu^2/n)$  iterations with probability  $1 - 2^{-\Omega(\mu^2/n)}$  before there are  $n$  off-branch individuals in the population. On the other hand in  $O(\mu n^{3/2})$  iterations, the whole population consists of individuals with at least as fit as the branch point  $b_{j+1}$  since at most  $O(\sqrt{n})$  Ridge points has to be traversed for each solution, where the improvement probability is exactly  $1/n$ . W.o.p. this process takes less than  $O(\mu n^2) \subset o(\mu^2/n)$  considering  $\mu = \omega(n^3)$ . Since we have already established that only  $n$  individuals can be on the off-branch, at least a  $1 - o(1)$  fraction of the population is at the branch point  $b_{j+1}$ . Thus, a branch point is selected as parent with constant probability and it takes at most  $O(n^2)$  iterations before a safe solution is created with probability at least  $1 - 2^{-\Omega(\sqrt{n})}$ .

Once the safe solution is created we apply the same line of argument in reverse to now show that off-branch individuals optimise the ONEMAX prefix before the safe solution takes over the population. However, this line of argumentation is relatively easier since the number of off-branch individuals are initially asymptotically larger than the number of safe solutions. Given that the population consists only of off-branch and safe solutions, in  $O(\mu n^2 \log n)$  iterations the local optima  $0^{m'} 1^k 0^{m'' - (j+2) \cdot k} 1^{(j+1) \cdot k}$  is found w.o.p. since a  $1 - o(1)$  fraction of the population is optimising the ONEMAX landscape. Increasing the number of safe solutions from  $\sqrt{n}/2$  to  $\sqrt{n}$  takes at least  $\frac{\sqrt{n} \cdot (\mu/\sqrt{n})^2}{2}$  iterations w.o.p. due to Chernoff bounds on geometric variables. Since the takeover time is asymptotically larger than the ONEMAX optimisation time, we can conclude that at every branch point the local optima is sampled and the safe solution is added to the population with probability at least  $1 - 2^{-\Omega(\sqrt{n})}$ . After the local optima is found, the off-branch solution can no longer improve and the progress is stagnated until the safe solution takes over the population in at most  $O(\mu^2)$  iterations. Since this take over has to happen at all  $\sqrt{n}/2$  branch points, the total takeover time is at most  $O(\mu^2 \sqrt{n})$  with probability  $1 - 2^{-\Omega(\sqrt{n})}$  and this term dominates the runtime.  $\square$

## 6.2 A Function Class with Two Slopes

In the previous subsection we considered a function with many different slopes and only one global optimum. We showed that inverse binary selection makes the  $(\mu+1)$  EA efficient for the function by allowing it to explore all the slopes one at a time, while all other considered evolutionary algorithms require exponential time w.o.p. In this subsection we study a function class with two different slopes that have to be explored in parallel. This function class was originally introduced by Witt to illustrate problem characteristics where large parent populations are essential [Witt 2006]. More

precisely, he defined a function called  $f$  where the (1+1) EA w.o.p. reaches a local optimum which needs exponential time to escape from, while the classical  $(\mu+1)$  EA with uniform parent selection and population size  $\mu \geq n/\ln(en)$  finds the global optimum in expected polynomial time. Overall, according to which of the two optima is the global one, parent populations may be either essential or detrimental w.o.p. (i.e., restarts do not change the outcome). Our aim in this subsection is to show that by using inverse tournament selection, the  $(\mu+1)$  EA can efficiently identify both optima with high probability either in a single run (large tournaments) or via restarts (smaller tournaments). We first prove that the function with the original global optimum can also be optimized efficiently by the  $(\mu+1)$  EA with inverse parent selection (here inverse binary tournament selection) and a sufficiently large population. Afterwards we will show experimentally that both optima can be identified already for not too large population and tournament sizes i.e., the algorithm is preferable to the (1+1) EA and the  $(\mu+1)$  EA with uniform selection.

For simplicity, we do not use the original definition  $f$  from [Witt 2006] here, but a version that still allows for an exponential performance difference using considerably shorter proofs. We remark that the simplified version does not allow a proof of polynomial time in expectation but still a polynomial time bound that occurs w.o.p. i.e., efficient optimisation w.o.p. Although we think that the  $(\mu+1)$  EA with inverse binary tournament selection and appropriate  $\mu$  indeed has expected polynomial runtime on the original function  $f$ , we do not find the additional proof overhead that would be necessary for such a result appropriate.

The function we study is adapted from [Witt 2006] and is a weighted sum of two independent optimisation problems defined on disjoint subspaces of  $\{0, 1\}^n$ . On the last  $\ell := n^{1/3}$  bits, the well-known LEADINGONES problem has to be solved, more precisely for  $x = (x_1, \dots, x_n)$  we define  $\text{LSO}(x) := \sum_{i=1}^{\ell} \prod_{j=1}^i x_{n-\ell+j}$  as the number of *leading suffix ones* value. On the first  $m := n - \ell$  bits, a ONEMAX problem has to be solved, hence we define  $\text{PO}(x) := \sum_{i=1}^m x_i$  as the number of *prefix ones*. The function is then defined by

$$\text{TWOGRADIENTS}(x) := \begin{cases} n^2 \text{LSO}(x) + \text{PO}(x) & \text{if } \text{PO}(x) \leq 2m/3 \\ n^2 \ell - m - 1 + \text{PO}(x) & \text{otherwise.} \end{cases}$$

The function TWOGRADIENTS has a global optimum at  $\text{LSO}(x) = \ell$  and  $\text{PO}(x) = 2m/3$ . To reach this, it is typically required that the number of leading suffix ones is maximized before the number of prefix ones grows beyond  $2m/3$ . Any search point where  $\text{PO}(x) > 2m/3$  is better than any search point with non-optimal LSO-value. Hence, if number of prefix ones is increased faster than the leading suffix ones and a PO-value of  $2m/3$  is exceeded before optimising the suffix then the function leads to a local optimum at  $\text{PO}(x) = m$ . This local optimum is a trap since at least  $m/3$  prefix ones have to be flipped to reach the global optimum from there.

We now show (similarly to the original result from [Witt 2006]) that the (1+1) EA w.o.p. fails to optimize TWOGRADIENTS efficiently.

**THEOREM 9.** *With probability at least  $1 - 2^{-\Omega(n^{1/3})}$ , the (1+1) EA needs time  $2^{\Omega(n \log n)}$  to optimize TWOGRADIENTS.*

**PROOF.** The main idea is to show that w.o.p. the (1+1) EA reaches a so-called bad search point where  $\text{PO}(x) \geq 3m/4$  and  $\text{LSO}(x) < \ell$ . From all bad search points, it is necessary to flip at least  $m/12 = \Omega(n)$  bits simultaneously to reach the global optimum. The probability of this is at most  $1/(m/12)! = 2^{-\Omega(n \log n)}$  and according to a union bound, with probability  $1 - 2^{-\Omega(n \log n)}$  this does not happen in  $2^{cn \log n}$  iterations if  $c$  is chosen as a sufficiently small constant.

To show that a bad search point is reached w.o.p., we study the growth of the PO-value carefully. First of all, by Chernoff bounds with probability  $1 - 2^{-\Omega(n)}$ , it holds  $\text{PO}(x_0) \geq m/3$  for the initial

search point  $x_0$ . Moreover,  $\Pr(\text{LSO}(x_0) \geq \ell/4) = 2^{-\ell/4}$  since this event requires to initialize the first  $\ell/4$  suffix to all ones.

We assume  $\text{PO}(x_0) \geq m/3$  and  $\text{LSO}(x) \leq \ell/4$  to happen, which altogether happens with probability  $1 - 2^{-\Omega(n)} - 2^{-\Omega(\ell)} = 1 - 2^{-\Omega(n^{1/3})}$ . Next, we note from the definition of `TWOGRADIENTS` that the PO-value can only decrease in steps increasing the LSO-value. Finally, we note that the number of leading suffix ones can increase by more than one when the leftmost zero in the suffix flips. More precisely, it increases by the number of so-called free-riders, the bits being random after the leftmost suffix zero. It is well known that the number of free-riders is stochastically bounded by a geometric distribution with parameter  $1/2$ .

We claim that it is sufficient to observe the following set of four events in a phase of length  $s := 3em$  to create a bad search point, pessimistically assuming that no bad search is created before:

- There are at least  $m/2$  PO-increasing steps,
- there are at most  $\ell/4$  LSO-increasing steps,
- the total number of prefix bits flipped in at most  $\ell/4$  LSO-increasing steps is at most  $\ell/2$ ,
- the total number of free-riders in at most  $\ell/4$  LSO-increasing steps is at most  $\ell/3$ .

If all these events together happen, the number of prefix ones increases to at least  $m/3 + m/2 - \ell/2 \geq 3m/4$  and the number of leading suffix ones to at most  $\ell/4 + \ell/4 + \ell/3 < \ell$ , which corresponds to a bad search point. Hence, we only have to verify that each of the four events happens with probability at least  $1 - 2^{-\Omega(n^{1/3})}$ , from which the theorem follows via a union bound.

The first event happens with probability  $1 - 2^{-\Omega(m)}$  according to Chernoff bounds since each of the  $s = 3em$  steps flips one of the least  $m/4$  zero-bits with probability at least  $(m/4)(1/n)(1 - 1/n)^{n-1} \geq 1/(5e)$  for  $n$  large enough. The second event occurs with probability  $1 - 2^{-\Omega(\ell)}$  by Chernoff bounds since the probability of an LSO-increasing step is  $1/n$ , resulting in an expected number of at most  $3e$  such steps. The third event is also shown to occur with probability  $1 - 2^{-\Omega(\ell)}$  via Chernoff bounds since each of the  $m$  prefix bits is flipped independently with probability  $1/n$ . Finally, using Chernoff bounds for the geometric distribution, also the probability of the fourth event is bounded by  $1 - 2^{-\Omega(\ell)}$ .  $\square$

We now present the positive result, showing that the  $(\mu+1)$  EA both with uniform and inverse binary tournament selection optimize `TWOGRADIENTS` efficiently for sufficiently large  $\mu$ .

**THEOREM 10.** *Let  $\mu = \sqrt{n}$ . Then the  $(\mu+1)$  EA both with uniform and with inverse binary tournament parent selection optimize `TWOGRADIENTS` in time  $O(n^{2.15})$  with probability  $1 - 2^{-\Omega(n^{0.15})}$ .*

**PROOF.** The proof strategy is again inspired by the corresponding result for the function  $f$  in [Witt 2006]. We consider a phase of length  $s := 22\ell n^{1.15}$  and show that w.o.p. the number of leading suffix ones rises to its maximum  $\ell$  before an individual with at least  $2m/3$  prefix ones is created. From there the global optimum is reachable within  $O(n^{2.15})$  generations w.o.p.

In more detail, we first consider the growth of the potential  $L$ , defined as the maximum number of leading suffix ones within the population. During this analysis, we assume that no individual with more than  $2m/3$  prefix ones is created within the phase. As we will prove below, this assumption is valid for  $s$  generations with probability  $1 - 2^{-\Omega(n^{0.983})}$ . Given a current potential  $L$ , we call an individual top if it has  $L$  or more leading suffix ones. We consider the grow of the number of top individuals, noting that this number is non-decreasing since only individuals of lowest fitness are removed and the factor  $n^2$  in the definition of `TWOGRADIENTS` gives higher fitness to individuals with more leading suffix ones regardless of the prefix ones. Let  $T$  be the number of top individuals. The probability of increasing this number by choosing and cloning such an individual is at least  $\frac{T}{\mu}(1 - 1/n)^n \geq \frac{T}{2e\mu}$  for uniform parent selection and at least  $\frac{T^2}{\mu^2}(1 - 1/n)^n \geq \frac{T^2}{2e\mu^2}$  for inverse binary

tournament selection since it suffices to choose two top individuals for the tournament. Hence, by standard argument for takeover times as in [Witt 2006], the expected time to reach  $\mu$  top individuals is at most

$$2e\mu \sum_{T=1}^{\mu} 1/T \leq 2e\mu(\ln \mu + 1)$$

for uniform and

$$2e\mu^2 \sum_{T=1}^{\mu} 1/T^2 \leq 2e\mu^2 \frac{\pi^2}{6} \leq 9\mu^2$$

for inverse binary tournament selection. We assume that  $n$ , and thereby  $\mu$ , is large enough so that  $2e\mu(\ln \mu + 1) \leq 9\mu^2$  holds. Given  $\mu$  top individuals and pessimistically assuming the  $L$ -value has not increased in the meantime, the probability of increasing the  $L$ -value is at least  $1/n$  since only the leftmost suffix zero has to flip. Hence, altogether, the expected time to increase the  $L$ -value is at most  $9\mu^2 + n = 10n$  by our choice of  $\mu$  and the expected time to reach  $\ell$  suffix ones in all individuals is at most  $10\ell n + 9\mu^2 \leq 11\ell n$  (still assuming no individual with at least  $2m/3$  prefix ones). By applying Markov's inequality and repeating this argumentation for  $n^{0.15}$  subphases of length  $20\ell n$ , we obtain that the probability of reaching  $\ell$  suffix ones for all individuals within  $22\ell n^{1.15} = O(n^{1.484})$  generations is at least  $1 - 2^{-\Omega(n^{0.15})}$ . From this point on the algorithm proceeds as if it was optimising ONEMAX on the  $m$  prefix bits only. Adjusting the preceding analysis (using  $m$  instead of  $\ell$  improvements, each pessimistically happening with probability at least  $1/(2en)$ ), the remaining time to reach the global optimum is easily bounded by  $O(mn^{1.15}) = O(n^{2.15})$  with probability  $1 - 2^{-\Omega(n^{0.15})}$ .

It remains to prove that within  $s$  generations no individual with at least  $2m/3$  prefix ones is created w.o.p. For this purpose, we use a straightforward generalization of the family tree argument from [Witt 2006], pessimistically assuming that each vertex in the tree at time  $t$  twice has the chance of  $1/\mu$  to receive a new child. This caters for the inverse binary tournament selection and increases the bound on the depth from Lemma 2 in [Witt 2006] only by a factor of 2; hence the probability that the depth  $D(t)$  at time  $t$  is  $6t/\mu$  or bigger is  $2^{-\Omega(t/\mu)}$ . We use this statement with  $t := s$ , implying that the probability of the depth being at least  $d^* := 22\ell n^{1.15}/\mu = 22n^{0.983}$  is  $2^{-\Omega(n^{0.983})}$ . Hence, in order to reach  $2m/3$  prefix ones within  $s$  steps, either the tree must have depth greater than  $d^*$ , or it must contain a so-called bad path of length at most  $d^*$  along which the number of prefix ones is increased by at most  $m/12$ . Here we assume that all initial individuals have at most  $7m/12$  prefix ones, which happens with probability  $1 - 2^{-\Omega(n)}$  by Chernoff bounds and a union bound over  $\sqrt{n}$  individuals.

The probability of increasing the number of ones by at least  $m/12 = \Omega(n)$  on a fixed path of length  $d^* = o(n)$  in the family tree is  $2^{-\Omega(n)}$  by Chernoff bounds since such a path is an independent sequence of mutations without selection. Only if this increase happens, the path is bad. By considering all possible paths in a union bound, the probability of a bad path emerging in  $s$  generations is at most

$$\begin{aligned} \sum_{k=1}^{d^*} \binom{s}{k} \left(\frac{2}{\mu}\right)^k 2^{-\Omega(n)} &= \sum_{k=1}^{d^*} \binom{d^* \mu}{k} \left(\frac{2}{\mu}\right)^k 2^{-\Omega(n)} \\ &\leq \sum_{k=1}^{d^*} \left(\frac{ed^* \mu}{k}\right)^k \left(\frac{2}{\mu}\right)^k 2^{-\Omega(n)} = \sum_{k=1}^{d^*} \left(\frac{2ed^*}{k}\right)^k 2^{-\Omega(n)} \\ &\leq d^* \left(\frac{2ed^*}{d^*}\right)^{d^*} 2^{-\Omega(n)} \leq 2^{O(d^*)} 2^{-\Omega(n)} = 2^{-\Omega(n)}, \end{aligned}$$

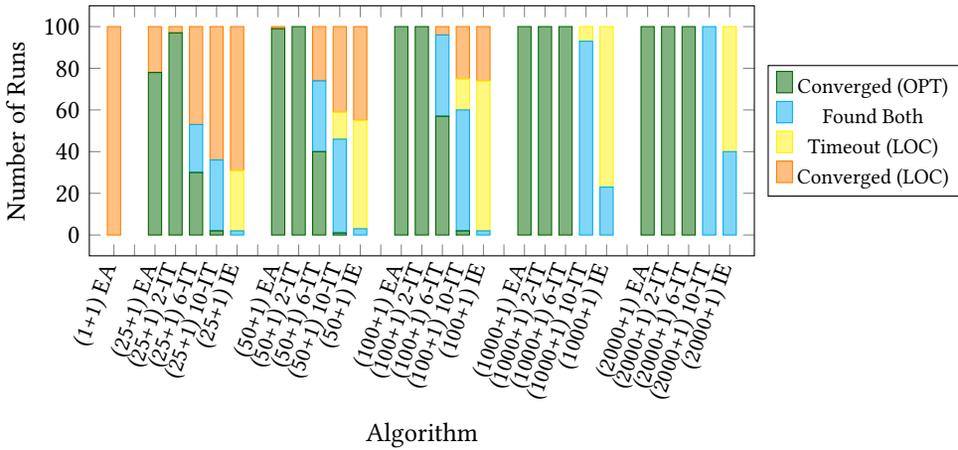


Fig. 6. Outcomes of 100 runs of algorithms using various selection operators and population sizes on TwoGradients with  $n = 1000$  (and  $\ell = 10$ ).

where the first inequality used the well-known estimate  $\binom{n}{k} \leq (en/k)^k$ , the second one that the terms of the sum are monotone increasing within the considered range of  $k$  and the final estimation that  $d^* = o(n)$ . Summing up all failure probabilities in a union bound completes the proof.  $\square$

We now present some experimental results to provide evidence that the inverse selection operators allow to identify both optima with reasonably high probability, while uniform selection can only identify the local optimum for small population sizes (including the (1+1) EA) and only the global optimum for large enough population sizes.

To illustrate the behaviour of the considered selection mechanisms, we have performed 100 independent runs of each mechanism and population size on TwoGradients with  $n = 1000$  (and hence  $\ell = 10$ ).

We define four possible outcomes for the runs: converging on the global optimum (OPT) without having constructed the local optimum (LOC), converging on the local optimum without having constructed the global optimum, constructing both optima before converging, or exceeding the computational budget<sup>6</sup> without converging (a possibility for the Inverse Elitist and Inverse Tournament selection mechanisms).

Figure 6 shows the proportion of the runs achieving specific outcomes. As expected, the (1+1) EA fails to find the global optimum in all 100 runs, while increasing the population size allows both uniform and inverse tournament mechanisms to find the global optima. Additionally, larger inverse tournament sizes enable the algorithm to construct both optima. For large enough population and tournament sizes (see eg.,  $\mu = 2n$  and  $s = 10$ ) both optima are identified in each run. For smaller population sizes either of the optima can be identified with reasonable probability (see eg.,  $\mu = 50$  and  $s = 6$ ), hence a restart strategy is effective.

<sup>6</sup>For practical reasons, we consider the global optimum to be unreachable by mutation of the local optimum, and sample waiting times for inverse tournament selection to select a viable ancestor when the majority of the population is composed of individuals on the local optimum. A timeout is reported when these waiting times cause the total optimisation time to exceed  $2^{63} - 1$ , or when the inverse elitist mechanism has constructed the local optimum while individuals with higher fitness on the branch leading to the global optimum still exist in the population (as these can neither be removed nor chosen for reproduction).

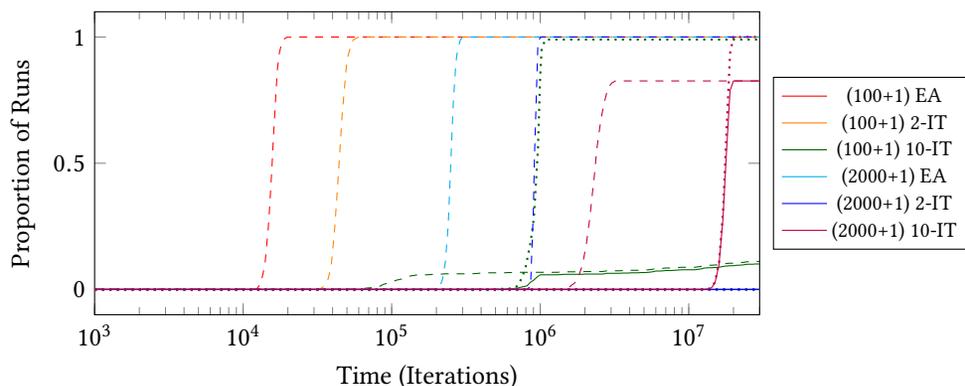


Fig. 7. Proportion of 1000 runs which find both the local and global optima (solid), at least the global optimum (dashed), or at least the local optimum (dotted) within a certain number of iterations on TwoGRADIENTS with  $n = 1000$  (and  $\ell = 10$ ) for various selection operators and population sizes.

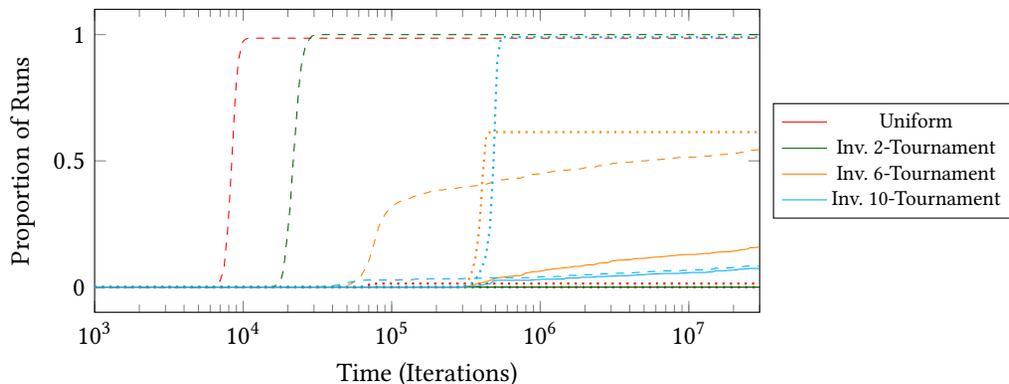


Fig. 8. Proportion of 1000 runs which find both the local and global optima (solid), at least the global optimum (dashed), or at least the local optimum (dotted) within a certain number of iterations on TwoGRADIENTS with  $n = 1000$  (and  $\ell = 10$ ) for various selection operators with  $\mu = 50$ .

Figures 7 and 8 give an estimate of the probability that a particular type of optimum (or both optima) have been constructed after a certain number of fitness iterations. With sufficiently large population and tournament sizes, the inverse tournament selection mechanism is shown to be able to find both the local and global optima with reasonable probability. This indicates that it could beneficially be combined with a restart strategy to reliably find both optima, while our theoretical results show that the  $(1+1)$  EA would require an exponential number of restarts to reliably find the global optimum. Additionally, our experimental results show that the  $(\mu + 1)$  EA does not find the local optimum when  $\mu$  is large. Hence, if the fitness values of the optima were exchanged, the  $(\mu + 1)$  EA would not have found the global optimum in any of the runs.

## 7 NP-COMPLETE PROBLEMS

In the previous section we have shown good explorative performance of the  $(\mu+1)$  EA with lower selective pressure than uniform on various multimodal pseudo-Boolean benchmark functions with different characteristics. In this section we verify the global exploration effectiveness of lowering

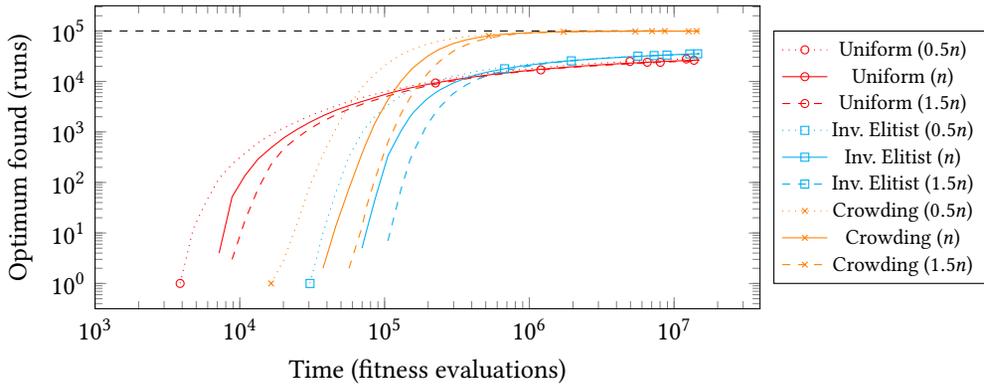


Fig. 9. Number of runs finding a satisfying solution for MAXSAT (combined over 100 independent runs on each of the 1000 instances in the SATLIB uf100 problem class) as a function of the elapsed number of fitness evaluations using various parent selection mechanisms and population sizes.

the selective pressure for two classical NP-Complete combinatorial optimisation problems, MAXSAT and the multidimensional knapsack problem (MKP).

## 7.1 MaxSat

In the MAXSAT problem, a set of clauses over  $n$  Boolean variables is given, and the objective is to find a truth-value assignment which maximises the number of satisfied clauses.

To evaluate the performance of Inverse Elitist and Inverse Tournament selection mechanisms, we use the uniform Random-3-SAT instances from SATLIB, which “tend to be particularly hard for both systematic SAT solvers and stochastic local search algorithms” [Hoos and Stützle 2000]. Solutions are represented as an  $n$ -bit string, and interpreted as truth-value assignments by setting  $x_i$  to true if and only if bit  $i$  is set to 1. The fitness value of a solution is simply the number of clauses it satisfies.

In particular, we focus on the instances in the uf100 and uf250 classes (composed of 1000 and 100 instances, respectively), which are satisfiable uniform Random-3-SAT instances with  $n = 100$  and  $n = 250$  variables respectively. We perform 100 independent runs of each algorithm on each instance in the problem class, and record the number of iterations it takes to produce a solution which satisfies all clauses.

Figure 9 compares the number of successful runs (i.e. those that find a satisfying truth-value assignment) of the Uniform and Inverse Elitist selection mechanisms, as well as Deterministic Crowding, for varying population sizes as a function of the available fitness evaluation budget on the uf100 problem class.

Uniform selection is the quickest mechanism to produce satisfying solutions for some instances. However, as the number of fitness evaluations increases, it achieves fewer successful runs than both Inverse Elitist selection and Deterministic Crowding. Additionally, while smaller populations (i.e.  $\mu = n/2$ ) generally produce satisfying solutions faster, all three considered population sizes seem to converge to similar values for each algorithm as the fitness evaluation budget increases. This suggests that there are many initial search points from which the a global optimum can be reached within the allocated number of fitness evaluations, hence an explanation of why the parallel (1+1) EAs due to deterministic crowding perform best.

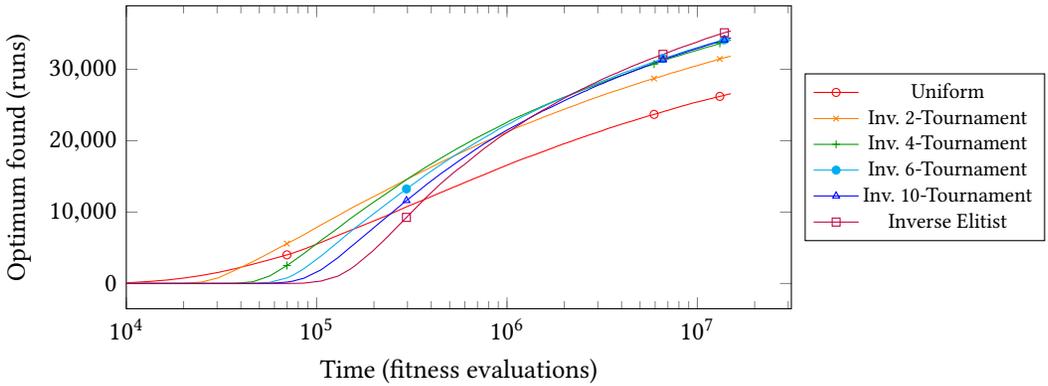


Fig. 10. Number of runs finding a satisfying solution for MAXSAT (combined over 100 independent runs on each of the 1000 instances in the SATLIB uf100 problem class) as a function of the elapsed number of fitness evaluations using various parent selection mechanisms and  $\mu = n = 100$ .

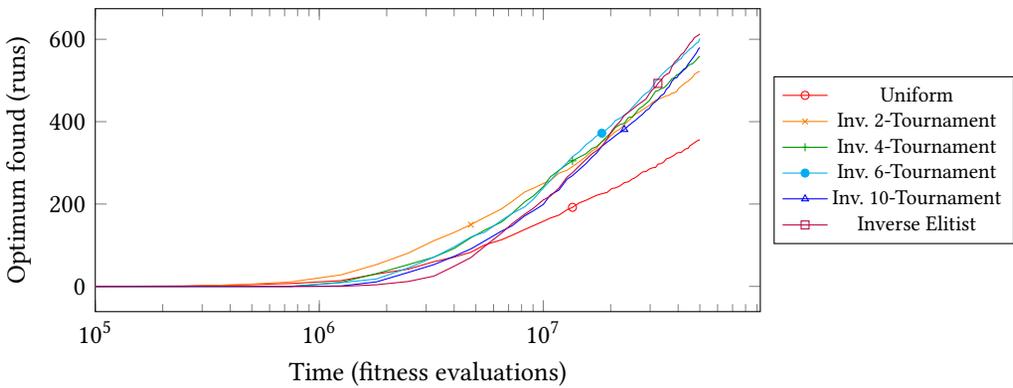


Fig. 11. Number of runs finding a satisfying solution for MAXSAT (combined over 100 independent runs on each of the 100 instances in the SATLIB uf250 problem class) as a function of the elapsed number of fitness evaluations using various parent selection mechanisms and  $\mu = n = 250$ .

Figures 10 and 11 respectively compare the performance of the uniform, inverse tournament, and inverse elitist selection mechanisms with  $\mu = n$  on the uf100 and uf250 problem classes (i.e. instances with  $n = 100$  and  $n = 250$  respectively). With a sufficient fitness evaluation budget, the inverse tournament and inverse elitist selection mechanisms outperform uniform selection, being able to find the satisfying assignments more frequently. As expected, large inverse tournaments behave similarly to the deterministic Inverse Elitist selection mechanism (being slower to find the first satisfying assignments, but eventually solving more instances), while smaller inverse tournaments behave similarly to the Uniform selection mechanism (quicker to find solutions on easier instances, but less successful in the longer runs). Notably, there exists a range of fitness evaluation budgets where the inverse tournament mechanisms outperform both the Uniform and Inverse Elitist mechanisms. This suggests that on these problem instances, the inverse selection mechanisms are able to preserve a greater amount of solution diversity during the optimisation process at the cost of a reduced hillclimbing speed. This increased diversity enables the global

optimum to be found more reliably compared to the uniform selection mechanism, which may converge on an initially-promising branch leading to a local optimum that is difficult to escape from sooner in the optimisation process.

To confirm that the inverse selection mechanisms perform better than Uniform selection on the MAXSAT instances, and that the difference observed in the experiments is statistically significant, we apply the Wilcoxon signed-rank test [Wilcoxon 1945], and use the Holm-Bonferroni method [Holm 1979] to account for multiple comparisons with a significance level  $\alpha = 0.01$ . We consider two algorithm performance metrics: the fitness of the best solution constructed within the allocated budget ( $B = 15,000,000$  and  $B = 50,000,000$  fitness evaluations for the *uf100* and *uf250* instance classes respectively), and the time taken to find a satisfying variable assignment (using the budget as an upper bound if no such solution is found during a run). When performing the comparisons between the performance of two algorithms on a problem class (i.e. *uf100* and *uf250*), runs on all instances within the problem class are used, and runs on the same instance are paired; thus each comparison on the *uf100* problem class uses the results of  $100 \times 1000 = 10^5$  runs of each algorithm, and each comparison on the *uf250* problem class uses the results of  $100 \times 100 = 10^4$  runs of each algorithm. For both problem classes, and all three considered population sizes ( $\mu \in \{n/2, n, 3n/2\}$ ), the performed two-sided statistical tests confirm that there was a significant difference in both best solution fitness and average runtime between the Uniform selection mechanism and other selection mechanisms considered, while one-sided statistical tests confirm that the inverse tournament selection mechanisms outperform the Uniform selection mechanism (i.e. find better solutions after at most  $B$  iterations, and find the optimum solutions faster within a budget of  $B$  iterations).

## 7.2 Multidimensional Knapsack

In the Multidimensional Knapsack problem (MKP), the objective is to select a subset of the  $n$  available items which maximises the value of the selected items, subject to constraints in  $m$  dimensions. More formally, given capacities  $b_i$  ( $1 \leq i \leq m$ ), item rewards  $p_j$  ( $1 \leq j \leq n$ ), and weights  $r_{ij}$ , the goal is to choose  $x_j \in \{0, 1\}$  to maximise  $\sum_{j=1}^n p_j x_j$  subject to  $\sum_{j=1}^n r_{ij} x_j \leq b_i$  for all  $i = 1, \dots, m$ .

We represent solutions as an  $n$ -bit string, with bit  $j$  determining whether item  $j$  is included in the selection or not. To guide the EAs toward valid solutions (i.e. those that satisfy the capacity constraints), we define the fitness of the solutions as follows: let  $W = 1 + \sum_{j=1}^n p_j$ , then

$$f(x) = \left( \sum_{j=1}^n x_j p_j \right) + W \cdot \sum_{i=1}^m \min \left( 0, b_i - \sum_{j=1}^n r_{ij} x_j \right).$$

The fitness of solutions which satisfy the constraints is the sum of the rewards  $p_j$  of the selected items. If any constraint is violated, the fitness will be negative (as the rightmost sum will produce a negative value), and will improve if an item which contributes to a violating constraint is removed.

To evaluate the performance of the selection mechanisms on MKP, we use the *mknep* instances from the OR-Library [Beasley 1990], originally constructed by Chu and Beasley [Chu and Beasley 1998] to evaluate the performance of a genetic algorithm on the problem. This set contains 270 instances in total, and consists of 10 instances for each combination of  $m \in \{5, 10, 30\}$ ,  $n \in \{100, 250, 500\}$  and constraint tightness ratios  $\alpha \in \{0.25, 0.5, 0.75\}$ . As the instances were randomly generated, the fitness value of the global optimum is not known. Hence, instead of considering the time required to reach the global optimum, we record the fitness value of a solution found after a certain number of fitness evaluations, and normalise these values by dividing by the fitness value of the best known solution for the instance (combining the best fitness values listed in OR-Library with those found during the runs).

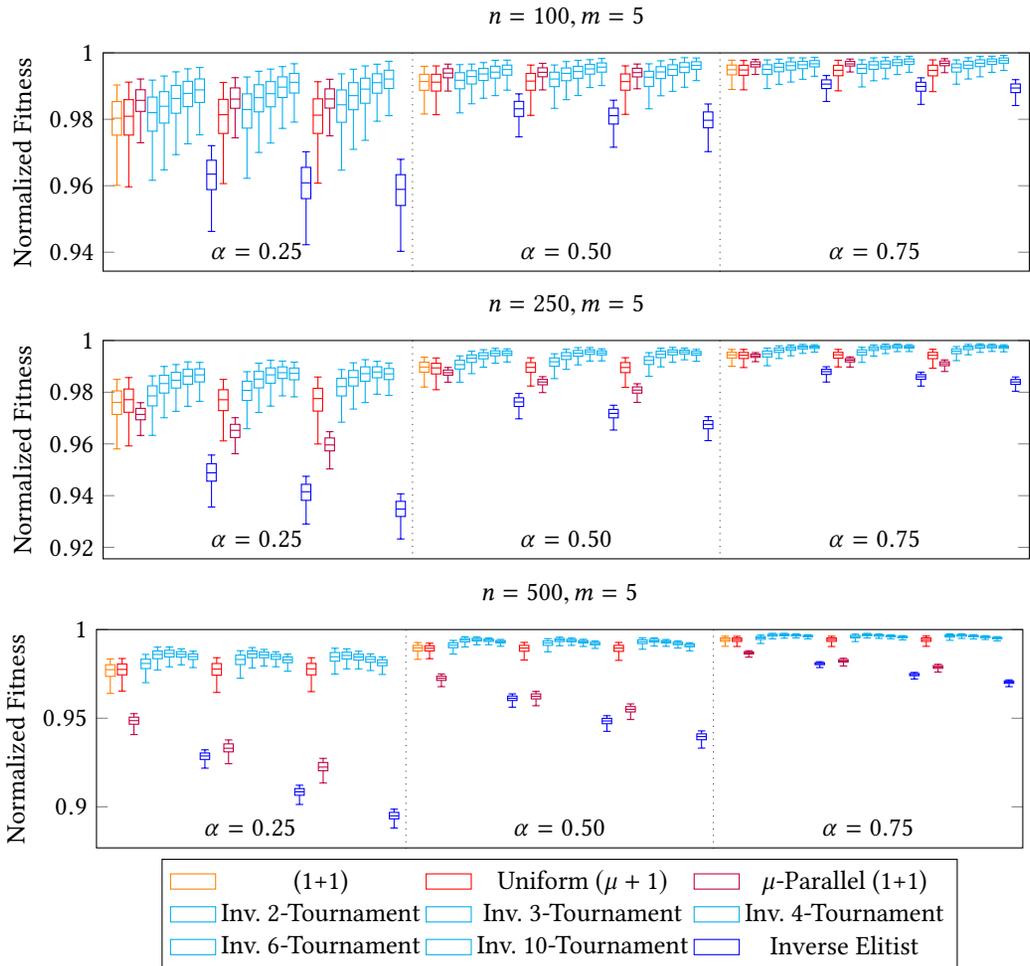


Fig. 12. Solution quality produced by a  $(\mu + 1)$  EA after 30,000,000 iterations, using various selection mechanisms and choices of  $\mu$  on 5-dimensional MKP instances from OR-LIB with varying constraint tightness  $\alpha$ . Each block repeats the population-employing algorithms using  $\mu = n/2$ ,  $\mu = n$ , and  $\mu = 3n/2$  in that order.

Figures 12, 13, 14 show the results aggregated over 100 independent runs for each algorithm and problem instance, giving all algorithms a budget of 30,000,000 fitness evaluations. We note that the larger Inverse Tournament sizes, such as the Inverse-10-Tournament tend to produce solutions of higher fitness than the standard  $(1 + 1)$  EA, the standard  $(\mu + 1)$  EA, and  $\mu$  parallel runs of the standard  $(1 + 1)$  EA (each given  $30,000,000/\mu$  fitness evaluations). The fully Inverse Elitist selection mechanism does not perform as well. This may be explained by the mechanism requiring a larger mutation to occur in order to escape from any local optimum it constructs before any higher-fitness individuals can be considered.

To evaluate whether the observed differences in solution fitness after 30,000,000 iterations are significant, and to confirm that the Inverse Tournament selection mechanisms outperform Uniform selection and  $\mu$  Parallel runs mechanisms, we apply the same Wilcoxon signed-rank test [Wilcoxon 1945], and Holm-Bonferroni method [Holm 1979] as in the previous section for MAXSAT again using

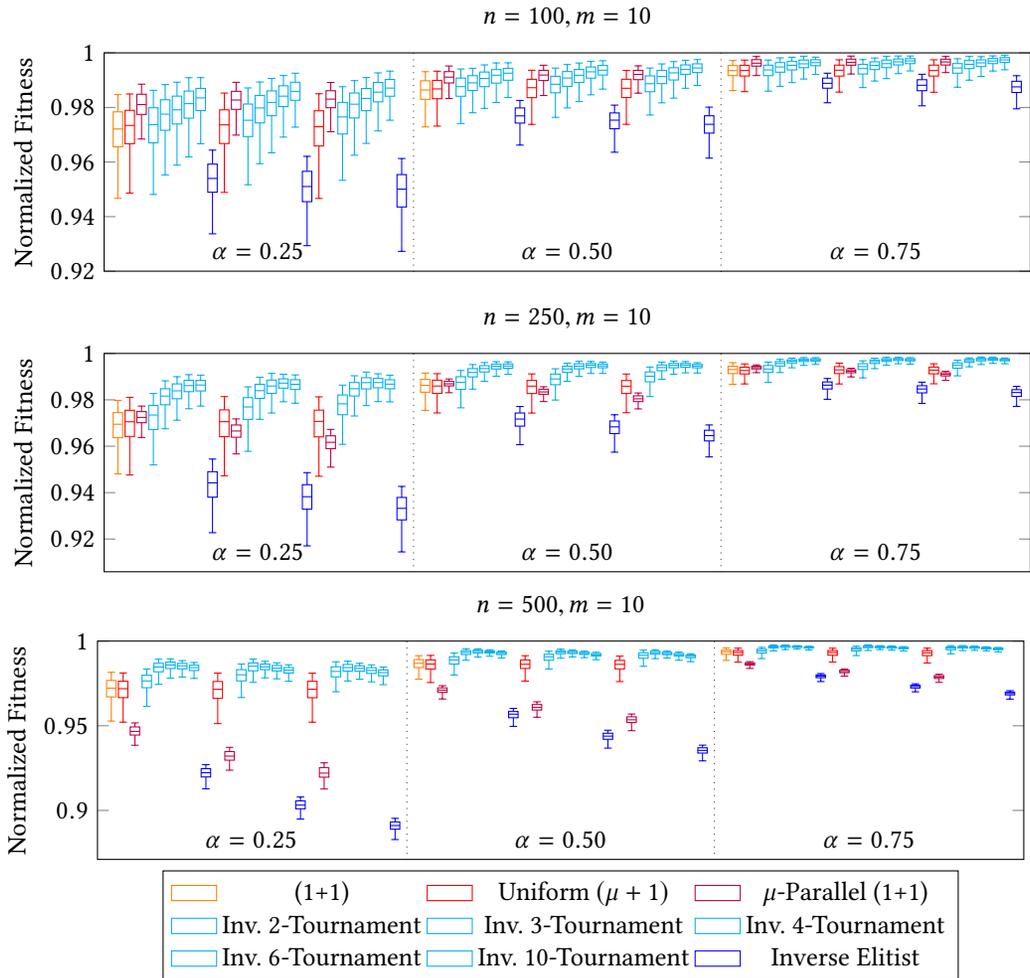


Fig. 13. Solution quality produced by a  $(\mu + 1)$  EA after 30,000,000 iterations, using various selection mechanisms and choices of  $\mu$  on 10-dimensional MKP instances from OR-LIB with varying constraint tightness  $\alpha$ . Each block repeats the population-employing algorithms using  $\mu = n/2$ ,  $\mu = n$ , and  $\mu = 3n/2$  in that order.

a significance level of  $\alpha = 0.01$ . When performing the comparisons between the performance of two algorithms, runs on all available instances are used, and runs on the same instance are paired; thus, each comparison is based on the results of  $100 \times 270 = 27,000$  runs of each algorithm. The performed two-sided tests confirm that for all considered population sizes ( $\mu \in \{n/2, n, 3n/2\}$ ) there is a significant difference between the solution fitness produced when using Uniform selection and all other algorithms, as well as between the  $\mu$  parallel runs of a  $(1 + 1)$  EA and all other algorithms, while the performed one-sided tests confirm that the Inverse Tournament mechanisms produce solutions with higher fitness (compared to Uniform selection and  $\mu$  parallel  $(1 + 1)$  runs) at the end of the  $B$ -iteration budget.

For these problem instances and fitness evaluation budget, the poor performance of the  $\mu$  parallel runs of the  $(1 + 1)$  EA, particularly as  $\mu$  increases, may be explained by the existence of many basins of attraction toward local optima that are difficult to escape from. As the  $\mu (1 + 1)$  EAs do

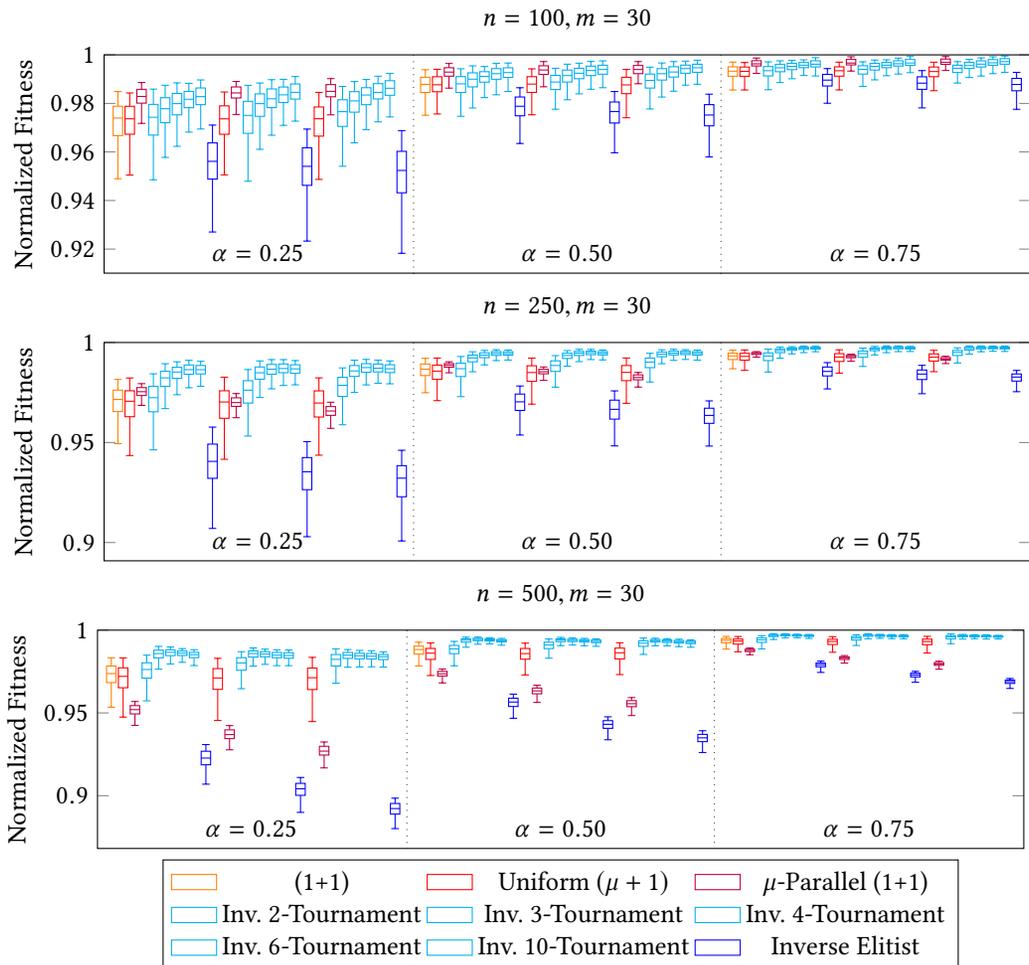


Fig. 14. Solution quality produced by a  $(\mu + 1)$  EA after 30,000,000 iterations, using various selection mechanisms and choices of  $\mu$  on 30-dimensional MKP instances from OR-LIB with varying constraint tightness  $\alpha$ . Each block repeats the population-employing algorithms using  $\mu = n/2$ ,  $\mu = n$ , and  $\mu = 3n/2$  in that order.

not interact during the optimisation process, once a  $(1 + 1)$  EA has reached a local optimum, its remaining computational budget is spent waiting for a mutation that escapes the local optimum. If a large mutation is required to escape, waiting for it to occur may be wasteful when compared to using fitness evaluations to explore the regions of the search space around higher-fitness solutions found by the other search trajectories.

The Inverse Tournament selection mechanisms are able to remove local optima from the population by selecting a better-fitness individual for reproduction, re-allocating the remaining fitness evaluations to explore a higher-fitness region of the search space, enabling the mechanism to more gracefully cope with increasing population sizes. The size of the inverse tournament balances the ease of such re-allocations with the diversity-preserving effects of inverse selection: small tournaments are more likely to converge the population on a single trajectory sooner, while larger

tournaments may use more fitness evaluations to explore the region around a local optimum before ejecting it from the population.

## 8 DISCUSSION

In the previous sections we have shown several advantages of using lower selective pressure than uniform for standard steady-state EAs i.e., that only use mutation. In this section we present a discussion of how previous literature has argued that uniform parent selection induces too high selective pressure also in steady-state GAs that use crossover as well as mutation.

Striving for a good balance between exploration and exploitation has arguably been the hottest topic in evolutionary computation since its conception. Already in his seminal book, John Holland introduced and analysed two forms of reproductive schemes to serve as a basis for the field of genetic algorithms [Holland 1992]. One was the canonical generational model, where  $\mu$  individuals were selected according to fitness proportional selection and the offspring replaced the entire population, independent of their quality. The other scheme used *overlapping generations* where only one new offspring was created in each generation: an individual was selected for reproduction with fitness proportional selection and an individual was removed uniformly at random from the parent population to make space for the new offspring. Hence, individuals would have an expected lifetime of  $\mu$  generations rather than only one generation as in generational models. The idea was that if the average fitness of the population didn't change much over the lifetime of an individual  $i$ , then the expected number of offspring of that individual would be the same as in the generational model (i.e.,  $f(i)/\bar{f}$ , which was the basis for Holland's schema theorem).

In practice, the average fitness of the population changes over the lifetime of individuals in the overlapping model. An empirical study by Kenneth De Jong showed how the use of overlapping generations introduced genetic drift (i.e., faster takeover times) which increased in severity with the decrease in size of the offspring population (i.e.,  $\lambda = 1$  leads to the most severe genetic drift) [De Jong and Sarma 1993]. As a result, the generational model was preferred by researchers and became the standard as well as the backbone of the canonical genetic algorithm [Goldberg 1989]. Since fitness proportional selection was applied to select parents, fitness-scaling functions had to be used to distinguish between individuals and avoid premature convergence once all the population had similar fitness values.

In 1985 James Baker reported experiments where the selection probability of individuals was allocated according to their rank in the population rather than according to their fitness relative to the population's average fitness. Rank-based selection not only resolved the scaling problems due to fitness proportional selection (which had previously also been avoided by Brindle's Tournament selection [Goldberg and Deb 1991]) but, more importantly, allowed for an easy way to control the selective pressure and keep it constant throughout the optimisation process [Whitley 1989]. Rank selection allows to explicitly set the expected number of offspring of each individual as a function of their rank in the population. Baker's intention was to slow down convergence to achieve more accurate optimisation. With the introduction of the Genitor (GENetic ImplemenTOR) algorithm, the first of what Syswerda termed "steady-state" GAs [Syswerda 1991], the opportunity to use steady-state schemes to gain control over the conflict between exploration and exploitation was also envisaged. Whitley states that "It can be argued that there are only two primary factors in genetic search: population diversity and selective pressure. These two factors are inversely related [...] As selective pressure is increased, the search focuses on the top individuals in the population, but because of this "exploitation" genetic diversity is lost. Reducing the selective pressure (or using larger population sizes) increases "exploration" because more genotypes and thus more schemata are involved in the search [...] [Hence] selective pressure and population diversity should

be controlled as directly as possible [...] Selective pressure can be *simply* controlled by allocating reproductive trials according to rank” [Whitley 1989].

However, the envisaged balance between exploration and exploitation was never achieved with the Genitor algorithm or other steady-state GAs because the same parent selection operators as used in generational schemes were adopted (typically rank selection with a bias of at least 1.5 in Genitor and fitness-proportional selection by Syswerda) while at the same time low quality individuals were deleted for replacement (the worst in Genitor, exponential inverse ranking in Syswerda’s version). Naturally, these choices led to considerably greater genetic drift than that witnessed by De Jong (when, in his overlapping GA, he deleted random individuals) as pointed out by an excellent theoretical analysis of the takeover times of various selection operators by Goldberg and Deb [Goldberg and Deb 1991]. In this paper we have shown how the steady-state scheme can indeed be used to largely reduce the exploration/exploitation conflict. While decreasing the selective pressure to uniform or even lower is indeed impractical (if not impossible) in generational schemes [Corus et al. 2018a], it can be easily implemented in steady-state algorithms.

In particular, the artificially introduced *elitism* in steady-state GAs *protects* the best individuals so long as better ones are not identified. Hence, there is no hurry to exploit the best individuals before they disappear: as long as they really are good, they will not be removed from the population. This simple insight sheds light on how the conflict between exploration and exploitation may be largely reduced with steady-state schemes: there is no need for a perfect balance between exploration and exploitation. Exploration may proceed for as long as desired because the best individuals will always be waiting to be exploited as long as better solutions are not identified first. In particular, steady-state algorithms allow for the selective pressure to be *reduced* to low levels that are impossible to achieve with the generational model. This is indeed the distinguishing power of steady-state EAs and what makes overlapping generations special. In this paper we have provided theoretical and empirical evidence of the advantages that can be gained by exploiting this insight. To the best of our knowledge no previous studies of reducing the parent selective pressure below uniform have been undertaken. We have recently discovered that the insight presented in this paper to avoid premature convergence (i.e., preventing the best individuals from reproducing) was briefly mentioned by Bäck and Hoffmeister in 1991 and termed “left extinctive selection” [Bäck and Hoffmeister 1991]. However, they deemed that the idea might not have any practical relevance and, to the best of our knowledge, was abandoned (the idea is indeed impractical in the context of the  $(\mu, \lambda)$  generational GAs they were using in this particular work: a steady-state scheme is required). On the other hand, studies on replacement strategies to reduce genetic drift have been undertaken which led to smaller takeover times than in standard steady-state GAs [Branke et al. 1999; Smith 2007].

## 9 CONCLUSION

In this paper we have analysed the effects of decreasing the parent selection pressure of a standard steady-state EA well below the commonly used uniform selection and have shown the benefits of this decrease without requiring any further algorithmic sophistication. In particular, we first proved that the  $(\mu+1)$  EA and  $(\mu+1)$  RLS evolving populations of arbitrary polynomial size via uniform parent selection, or via operators with higher selective pressure, fail to identify both optima of the TwoMAX benchmark function with high probability. On the other hand, we proved that the algorithms are efficient at locating both optima for reasonable population sizes if Inverse Elitist selection is used instead. An experimental analysis explains why Inverse Tournament selection with small tournament sizes is even more powerful for global optimisation: for the tested problem sizes it has higher success probabilities for TwoMAX and is capable of escaping local optima of

TRUNCATEDTWO MAX<sub>k</sub>. In particular, in order to escape from local optima some probability of not selecting the worst individuals is crucial.

Afterwards we analysed the impressive performance of the low selective pressure algorithms for more sophisticated multi-modal problems generalised from previously used benchmark functions from the literature which we named RIDGEWITHBRANCHES and TWOGRADIENTS. The two functions have different characteristics. The former consists of a series of local optima that should be identified one at a time, while the latter consists of two gradients with different steepnesses which should be optimised in parallel. While the Inverse Tournament ( $\mu+1$ ) EA is efficient for all the problems, we are not aware of any other algorithm, with or without diversity, that can efficiently optimise all the three function classes TWO MAX, RIDGEWITHBRANCHES and TWOGRADIENTS. In fact we are not aware of any other algorithm that is effective for RIDGEWITHBRANCHES alone. The latter is an exemplary benchmark function to highlight the explorative capabilities of steady-state EAs with low selective pressure: the algorithm explores one branch at a time, identifying one local optimum after another, and only abandoning each one if it does not turn out to be the global optimum.

The last section of the paper presented an experimental analysis of the algorithms for the classical MAXSAT and Multidimensional Knapsack problems. For both problems the inverse selection operators outperform uniform selection. However, for the MaxSat instances single trajectory algorithms with restarts outperform all the population based EAs, while for the knapsack instances the opposite is true.

We conclude with a final comment regarding the optimisation time of the inverse selection algorithms. We have presented very effective algorithms for multi-modal optimisation with minimal sophistication (i.e., no additional parameters compared to basic EAs). For all the problems considered in this paper, we see no reason to use a selection operator with higher pressure than inverse tournaments of constant size. For difficult multi-modal problems with many optima, extremely fast global optimisation cannot reasonably be expected. For easier unimodal problems, the selective pressure may be increased by simply decreasing the population size (i.e., if the problem is easy, then large populations are not necessary). For instance, consider the simple OneMax problem. The expected runtime of the inverse tournament ( $\mu+1$ ) EA is  $O(\mu n \ln n)$  which is the same asymptotic expected runtime of the algorithm using uniform selection and deterministic crowding as diversity mechanism [Friedrich et al. 2009]. The expected runtime can be reduced to the best achievable via standard bit mutation-only (i.e.,  $O(n \ln n)$ ) simply by reducing the population size to a constant (i.e., effectively increasing the selective pressure). Furthermore, we believe that the standard ( $\mu+1$ ) GA (i.e., with crossover) coupled with inverse tournament selection with constant tournament and population sizes is faster than any standard bit mutation-only evolutionary algorithm for OneMax. This result may be proven by simply re-calculating the takeover times in the analysis of [Corus and Oliveto 2020]. In particular, also the result that the algorithm is faster than any unary unbiased search heuristic, if identical copies of parents are not re-evaluated, holds. Indeed, future work should explore theoretically and experimentally the effects of lower selective pressure on more sophisticated algorithms using crossover, more advanced mutation operators [Corus et al. 2018b, 2019; Doerr et al. 2017], and problem-tailored ones [Kautz and Selman 1993, 1996].

### Acknowledgments

The authors would like to thank the members of the EC-Theory group in the UK for initial discussions regarding the analysis of steady-state EAs with large population sizes for TWO MAX. A special thank you also goes to Christian Gießen for his collaboration in the preliminary versions of this paper. Part of this work was funded by the EPSRC under grant agreement N. EP/M004252/1.

On Steady-State Evolutionary Algorithms and Selective Pressure: Why Inverse Rank-Based Allocation of Reproductive Trials is Best 37

## REFERENCES

- T. Bäck and F. Hoffmeister. 1991. A Survey of Evolution Strategies. In *Proceedings of the Fourth International Conference on Genetic Algorithms*, R.K Belew and L.B. Booker (Eds.). Morgan Kaufmann, 92–99.
- T. Bäck, F. Hoffmeister, and HP. Schwefel. 1991. A Survey of Evolution Strategies. In *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann, 2–9.
- J. E. Beasley. 1990. OR-Library: distributing test problems by electronic mail. *Journal of the operational research society* 41, 11 (1990), 1069–1072. <http://people.brunel.ac.uk/~mastjjb/jeb/info.html>.
- B. Bollobas. 1985. *Random Graphs*. Academic Press, London - New York.
- Jürgen Branke, Massimo Cutaia, and Heinrich Dold. 1999. Reducing Genetic Drift in Steady State Evolutionary Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 1999)*. Morgan Kaufmann, 68–74.
- P.C. Chu and J.E. Beasley. 1998. A Genetic Algorithm for the Multidimensional Knapsack Problem. *Journal of Heuristics* 4, 1 (1998), 63–86.
- D. Corus, D.-C. Dang, A. V. Eremeev, and P. K. Lehre. 2018a. Level-Based Analysis of Genetic Algorithms and Other Search Processes. *IEEE Transactions on Evolutionary Computation* 22, 5 (2018), 707–719.
- D. Corus and P. S. Oliveto. 2018. Standard Steady State Genetic Algorithms Can Hillclimb Faster than Mutation-only Evolutionary Algorithms. *IEEE Transactions on Evolutionary Computation* 22, 5 (2018), 720–732.
- D. Corus and P. S. Oliveto. 2020. On the benefits of populations for the exploitation speed of standard steady-state genetic algorithms. *Algorithmica* (2020), -. <https://doi.org/10.1007/s00453-020-00743-1> In press (Online).
- D. Corus, P. S. Oliveto, and D. Yazdani. 2018b. Fast Artificial Immune Systems. In *Proceedings of the International Conference on Parallel Problem Solving from Nature (PPSN XV) (Lecture Notes in Computer Science (LNCS))*, Vol. 11102. Springer, 67–78.
- D. Corus, P. S. Oliveto, and D. Yazdani. 2019. Artificial Immune Systems Can Find Arbitrarily Good Approximations for the NP-hard Number Partitioning Problem. *Artificial Intelligence* 247 (2019), 180–196.
- D.-C. Dang, T. Friedrich, T. Kötzing, M. S. Krejca, P. K. Lehre, P. S. Oliveto, D. Sudholt, and A. M. Sutton. 2018. Escaping Local Optima Using Crossover with Emergent Diversity. *IEEE Transactions on Evolutionary Computation* 22, 3 (2018), 484–497.
- C. Darwin. 1859. *On the origin of the species*. John Murray.
- C. Darwin. 1868. *The Variation of Animals and Plants under Domestication*. John Murray.
- K. De Jong and J. Sarma. 1993. Generation Gaps Revisited. In *Proceedings of Foundations of Genetic Algorithms II (FOGA 1993)*, 19–28.
- A. de Perthuis de Laillevault, B. Doerr, and C. Doerr. 2015. Money for Nothing: Speeding Up Evolutionary Algorithms Through Better Initialization. In *Proceedings of the Genetic and Evolutionary Computation Conference, (GECCO 2015)*. ACM Press, 815–822.
- B. Doerr. 2011. Analyzing Randomized Search Heuristics: Tools from Probability Theory. In *Theory of Randomized Search Heuristics: Foundations and Recent Developments*, B. Doerr and A. Auger (Eds.). World Scientific, 1–20.
- B. Doerr, H. Ph. Le, R. Makhmara, and T. D. Nguyen. 2017. Fast Genetic Algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2017)*, 777–784.
- T. Friedrich, P. S. Oliveto, D. Sudholt, and C. Witt. 2009. Analysis of Diversity-Preserving Mechanisms for Global Exploration. *Evolutionary Computation* 17, 4 (2009), 455–476.
- D. E. Goldberg. 1989. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman.
- D. E. Goldberg and K. Deb. 1991. A comparative analysis of selection schemes used in genetic algorithms. In *Proceedings of Foundations of Genetic Algorithms I (FOGA 1991)*, 69–93.
- J. H. Holland. 1992. *Adaptation in natural and artificial systems: an introductory analysis with applications to biology, control, and artificial intelligence*. MIT press.
- S. Holm. 1979. A Simple Sequentially Rejective Multiple Test Procedure. *Scandinavian Journal of Statistics* 6, 2 (1979), 65–70.
- H. H. Hoos and T. Stützle. 2000. SATLIB: An Online Resource for Research on SAT. *SAT 2000 (2000)*, 283–292.
- T. Jansen, K. A. De Jong, and I. Wegener. 2005. On the Choice of the Offspring Population Size in Evolutionary Algorithms. *Evolutionary Computation* 13, 4 (2005), 413–440.
- H. Kautz and B. Selman. 1993. Domain-Independent Extension to GSAT: Solving Large Structured Satisfiability Problems. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI'93)*, 290–295.
- H. Kautz and B. Selman. 1996. Pushing the envelope: planning, propositional logic, and stochastic search. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI'96)*, 1194–1201.
- P. K. Lehre. 2010. Negative drift in populations. In *Proceedings of the International Conference on Parallel Problem Solving from Nature (PPSN XI) (Lecture Notes in Computer Science (LNCS))*, Vol. 6328, 244–253.
- J. Lengler. 2018. A General Dichotomy of Evolutionary Algorithms on Monotone Functions. In *Proceedings of the International Conference on Parallel Problem Solving from Nature (PPSN XV) (Lecture Notes in Computer Science (LNCS))*, Vol. 11101. Springer, 3–15.

- P. S. Oliveto, D. Sudholt, and C. Zarges. 2018. On the benefits and risks of using fitness sharing for multimodal optimisation. *Theoretical Computer Science* 773 (2018), 53–70.
- E. Covantes Osuna and D. Sudholt. 2017. Analysis of the clearing diversity-preserving mechanism. In *Proceedings of Foundations of Genetic Algorithms (FOGA 2017)*. ACM Press, 55–63.
- M. Preuss. 2015. *Multimodal Optimization by Means of Evolutionary Algorithms*. Springer.
- J. Smith. 2007. On Replacement Strategies in Steady State Evolutionary Algorithms. *Evolutionary Computation* (2007), 29–59.
- H. Spencer. 1864. *The principles of biology*. Williams and Norgate.
- D. Sudholt. 2019. The Benefits of Population Diversity in Evolutionary Algorithms: A Survey of Rigorous Runtime Analyses. In *Theory of Randomized Search Heuristics in Discrete Search Spaces*, B. Doerr and F. Neumann (Eds.). Springer, 359–404.
- A. Sutton. 2018. Crossover can simulate bounded tree search on a fixed-parameter tractable optimization problem. In *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2018)*. ACM Press, 1531–1538.
- G. Syswerda. 1989. Uniform crossover in genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms*. 2–9.
- G. Syswerda. 1991. A Study of Reproduction in Generational and Steady-State Genetic Algorithms. In *Proceedings of Foundations of Genetic Algorithms I (FOGA 1991)*. 94–101.
- D. Whitley. 1989. The Genitor algorithm and selection pressure: why rank-based allocation of reproductive trials is best. In *Proceedings of the Third International Conference on Genetic Algorithms*. 116–121.
- F. Wilcoxon. 1945. Individual Comparisons by Ranking Methods. *Biometrics Bulletin* 1, 6 (1945), 80–83.
- C. Witt. 2006. Runtime Analysis of the  $(\mu + 1)$  EA on Simple Pseudo-Boolean Functions. *Evolutionary Computation* 14, 1 (2006), 65–86.
- C. Witt. 2014. Fitness levels with tail bounds for the analysis of randomized search heuristics. *Inform. Process. Lett.* 114, 1-2 (2014), 38–41.