



This is a repository copy of *Simple hyper-heuristics control the neighbourhood size of randomised local search optimally for LeadingOnes.*

White Rose Research Online URL for this paper:
<https://eprints.whiterose.ac.uk/146353/>

Version: Accepted Version

Article:

Lissovai, A., Oliveto, P. and Warwicker, J.A. (2020) Simple hyper-heuristics control the neighbourhood size of randomised local search optimally for LeadingOnes. *Evolutionary Computation*, 28 (3). pp. 437-461. ISSN 1063-6560

https://doi.org/10.1162/evco_a_00258

© 2019 Massachusetts Institute of Technology Press. This is an author-produced version of a paper subsequently published in *Evolutionary Computation*. Uploaded in accordance with the publisher's self-archiving policy.

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 including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Simple Hyper-heuristics Control the Neighbourhood Size of Randomised Local Search Optimally for LeadingOnes*

Andrei Lissovoi a.lissovoi@sheffield.ac.uk
Department of Computer Science, University of Sheffield, UK

Pietro S. Oliveto p.oliveto@sheffield.ac.uk
Department of Computer Science, University of Sheffield, UK

John Alasdair Warwicker j.warwicker@sheffield.ac.uk
Department of Computer Science, University of Sheffield, UK

Abstract

Selection hyper-heuristics (HHs) are randomised search methodologies which choose and execute heuristics during the optimisation process from a set of low-level heuristics. A machine learning mechanism is generally used to decide which low-level heuristic should be applied in each decision step. In this paper we analyse whether sophisticated learning mechanisms are always necessary for HHs to perform well. To this end we consider the most simple HHs from the literature and rigorously analyse their performance for the LEADINGONES benchmark function. Our analysis shows that the standard Simple Random, Permutation, Greedy and Random Gradient HHs show no signs of learning. While the former HHs do not attempt to learn from the past performance of low-level heuristics, the idea behind the Random Gradient HH is to continue to exploit the currently selected heuristic as long as it is successful. Hence, it is embedded with a reinforcement learning mechanism with the shortest possible memory. However, the probability that a promising heuristic is successful in the next step is relatively low when perturbing a reasonable solution to a combinatorial optimisation problem. We generalise the ‘simple’ Random Gradient HH so success can be measured over a fixed period of time τ , instead of a single iteration. For LEADINGONES we prove that the *Generalised Random Gradient (GRG)* HH can learn to adapt the neighbourhood size of Randomised Local Search to optimality during the run. As a result, we prove it has the best possible performance achievable with the low-level heuristics (Randomised Local Search with different neighbourhood sizes), up to lower order terms. We also prove that the performance of the HH improves as the number of low-level local search heuristics to choose from increases. In particular, with access to k low-level local search heuristics, it outperforms the best-possible algorithm using any subset of the k heuristics. Finally, we show that the advantages of GRG over Randomised Local Search and Evolutionary Algorithms using standard bit mutation increase if the anytime performance is considered (i.e., the performance gap is larger if approximate solutions are sought rather than exact ones). Experimental analyses confirm these results for different problem sizes (up to $n = 10^8$) and shed some light on the best choices for the parameter τ in various situations.

1 Introduction

Many successful applications of randomised search heuristics to real-world optimisation problems have been reported. Despite these successes, it is still difficult to decide which particular search heuristic is a good choice for the problem at hand and what parameter settings should be used. In particular, while it is well understood that each heuristic will be efficient on some classes of problems and inefficient on others (Wolpert and Macready, 1997), very little guidance is available explaining how to choose an algorithm for a given problem. The high level idea behind the field of hyper-heuristics (HHs) is to overcome this difficulty

*An extended abstract of this manuscript has appeared at the 2017 Genetic and Evolutionary Computation Conference (GECCO 2017) (Lissovoi et al., 2017).

by evolving the search heuristic for the problem rather than choosing one in advance (or applying several arbitrary ones until a satisfactory solution is found). The overall goal is to automate the design and the tuning of the algorithm and parameters for the optimisation problem, hence to achieve a more generally applicable system.

HHs are usually classified into two categories: *generation* HHs and *selection* HHs (Burke et al., 2013). The former typically work offline and generate new heuristics from components of existing heuristics, while the latter typically work online and repeatedly select from a set of low-level heuristics which one to apply in the next state of the optimisation process. The low-level heuristics can be further classified as either *construction heuristics*, which build a solution incrementally, or *perturbation heuristics*, which start with a complete solution and try to iteratively improve it.

In this paper, selection-perturbation HHs will be considered. Selection HHs belong to the wide field of *algorithm selection from algorithm portfolios* (Bezerra et al., 2018). A noteworthy example of such a system is the SATZilla approach for the SAT problem (Xu et al., 2008). We refer to (Kotthoff, 2016) for an overview of approaches to algorithm selection for combinatorial optimisation problems. It should be noted that HHs may also be used for other purposes. For example, a HH may be used as a parameter control method if it is asked to select at each state between different heuristics that only differ by one parameter value, rather than between different general purpose algorithms. See (Doerr and Doerr, 2018) for a comprehensive review of parameter control techniques. In general though, while the aim of parameter control mechanisms is to essentially choose between different parameter values for a given algorithm, selection HHs normally work at a higher level, since their aim is to learn which low-level heuristics (i.e., algorithms) work best at different states of the optimisation process i.e., very different algorithms may perform better in different areas of the search space.

Selection HHs consist of two separate components: (1) a *heuristic selection method*, often referred to as the *learning mechanism*, to decide which heuristic to be applied in the next step of the optimisation process, and (2) a *move acceptance operator* to decide whether the newly produced search points should be accepted.

Move acceptance operators are classified into *deterministic* ones, which make the same decisions independent of the state of the optimisation process, and *non-deterministic* ones, which might make different decisions for the same solutions at different times.

The majority of heuristic selection methods in the literature apply machine learning techniques that generate scores for each heuristic based on their past performance. A commonly used method for the purpose is reinforcement learning (Cowling et al., 2001; Nareyek, 2004; Burke et al., 2003). However, also considerably simpler heuristic selection methods have been successfully applied in the literature. These include, among others: *Simple Random* heuristic selection which just chooses heuristics uniformly at random; *Random Permutation* which generates a random ordering of the low-level heuristics and applies them in that order; *Greedy* heuristic selection which applies all the heuristics to the current search point and returns the best found solution; *Random Gradient* which chooses a random heuristic and continues to apply it as long as it is successful, and *Random Permutation Gradient* which is a combination of the Permutation and Random Gradient methodologies (Cowling et al., 2001, 2002; Berberoğlu and Uyar, 2011).

Apart from the Greedy one, an obvious advantage of these simple methods is that they execute very quickly. On the other hand, it is unclear how well they perform in terms of how effective they are at identifying (near) optimal solutions in a short time. While the Simple Random, Random Permutation and Greedy mechanisms do not attempt to learn anything from the past performance of the applied heuristics, Random Gradient and Permutation Random Gradient can still be considered as intelligent selection mechanisms because they embed a reinforcement learning mechanism, albeit with the shortest memory length possible. Experimental work has suggested that such mechanisms may be useful for highly rugged search landscapes (Burke et al., 2013). In this paper we will rigorously evaluate the performance of simple heuristic selection methods for a unimodal benchmark problem and focus most of our attention on the Random Gradient method that uses ‘minimal’ intelligence.

Numerous successes of selection HHs for NP-hard optimisation problems have been reported, including scheduling (Cowling et al., 2001, 2002; Gibbs et al., 2010), timetabling (Özcan et al., 2010), vehicle routing (Asta and Özcan, 2014), cutting and packing (López-Camacho et al., 2014). However, their theoretical understanding is very limited. Some insights into the behaviour of selection HHs have been achieved via landscape analyses (Maden et al., 2009; Ochoa et al., 2009a,b). Concerning their performance, the most sophisticated HH that has been analysed is the one considered by Doerr et al. (2016), where the different

neighbourhood sizes k of the Randomised Local Search (RLS $_k$) algorithm were implicitly used as low-level heuristics and a reinforcement learning mechanism was applied as the heuristic selection method. The authors proved that this HH can track the best local search neighbourhood size (i.e., the best fitness-dependent number of distinct bits to be flipped) for ONEMAX, while they showed experimentally that it outperforms RLS $_1$ and the (1+1) Evolutionary Algorithm ((1+1) EA) for LEADINGONES. The few other available runtime analyses consider the less sophisticated heuristic selection methodologies which are the focus of this paper.

Lehre and Özcan (2013) presented the first runtime analysis of selection-perturbation HHs. They considered the Simple Random HH (Cowling et al., 2001) that at each step randomly chooses between a 1-bit flip and a 2-bit flip operator as low-level heuristics, and presented an example benchmark function class, called GAPPATH, where it is necessary to use more than one of the low-level heuristics to optimise the problem. Similar example functions have also been constructed by He et al. (2012).

A comparative time-complexity analysis of selection HHs has been presented by Alanazi and Lehre (2014). They considered several of the common simple selection mechanisms, namely Simple Random, Permutation, Random Gradient and Greedy (Cowling et al., 2001, 2002) and analysed their performance for the standard LEADINGONES benchmark function when using a low-level heuristic set consisting of a 1-bit flip and a 2-bit flip operator (i.e., the same set previously considered by Lehre and Özcan (2013)). Their runtime analyses show that the four simple selection mechanisms have the same asymptotic expected runtime, while an experimental evaluation indicates that the runtimes are indeed equivalent already for small problem dimensions. Recently, additive Reinforcement Learning selection was also shown (under mild assumptions) to often have asymptotically equivalent performance to Simple Random selection, including for the same problem setting (i.e., LEADINGONES selecting between 1-bit flip and 2-bit flip operators) (Alanazi and Lehre, 2016). In particular, the results indicate that selection mechanisms such as additive Reinforcement Learning and Random Gradient do not learn to exploit the more successful low-level heuristics and end up having the same performance as Simple Random selection.

The main idea behind Random Gradient is to continue to exploit the currently selected heuristic as long as it is successful. Unlike construction heuristics, where iterating a greedy move on a currently successful heuristic may work for several consecutive construction steps, the probability that a promising heuristic is successful in the next step is relatively low when perturbing a reasonable solution to a combinatorial optimisation problem.

To this end, in this paper we propose to generalise the Random Gradient selection-perturbation mechanism such that success can be measured over a fixed period of time τ , which we call the learning period, instead of doing so after each iteration. We refer to the generalised HH as Generalised Random Gradient (GRG). We use the LEADINGONES benchmark function and the 1-bit flip and a 2-bit flip operators as low-level heuristics to show that with this simple modification Random Gradient can be surprisingly fast, even if no sophisticated machine learning mechanism is used to select from the set of low-level heuristics. We first derive the exact leading constants in the expected runtimes of the simple mechanisms for LEADINGONES, thus proving that they all have expected runtime $\frac{\ln(3)}{2}n^2 + o(n^2) \approx 0.54931n^2 + o(n^2)$, confirming what was implied by the experimental analysis of Alanazi and Lehre (2014). This result indicates that all the simple mechanisms essentially choose operators at random in each iteration. Thus they have worse performance than the single operator that always flips one bit (i.e., RLS $_1$), which takes $\frac{1}{2}n^2$ expected iterations to optimise LEADINGONES. We then provide upper bounds on the expected runtime of GRG for the same function. We rigorously prove that the generalised HH has a better expected running time than RLS $_1$ for appropriately chosen values for the parameter τ . Furthermore, we prove that GRG can achieve an expected runtime of $\frac{1+\ln(2)}{4}n^2 + o(n^2) \approx 0.42329n^2 + o(n^2)$ for LEADINGONES, when τ satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{2} - \varepsilon)n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{2}$. This is the best possible expected runtime for an unbiased (1+1) black box algorithm using only the same two mutation operators. We refer to such performance as *optimal* for the HH as no better expected runtime, up to lower order terms, may be achieved with any combination of the available low-level heuristics.

Afterwards we turn our attention to low-level heuristic sets that contain arbitrarily many operators (i.e., $\{1, \dots, k\}$ -bit flips, $k = \Theta(1)$) as commonly used in practical applications. We first show that including more operators is detrimental to the performance of the simple mechanisms for LEADINGONES. Then, we prove that the performance of GRG improves as the number of available operators to choose from increases. In particular, when choosing from k operators as low-level heuristics, GRG is in expectation faster than the best possible performance achievable by any algorithm using any combination of $m < k$ of the operators,

Algorithm 1 Simple Selection Hyper-heuristic (Cowling et al., 2001, 2002; Alanazi and Lehre, 2014)

```
1: Choose  $s \in S$  uniformly at random
2: while stopping conditions not satisfied do
3:   Choose  $h \in H$  according to the learning mechanism
4:    $s' \leftarrow h(s)$ 
5:   if  $f(s') > f(s)$  then
6:      $s \leftarrow s'$ 
```

for any $k = \Theta(1)$.

We conclude the paper with an experimental analysis of the HH for increasing problem sizes, up to $n = 10^8$. The experiments confirm that GRG outperforms its low-level heuristics. For two operators, it is shown how proper choices for the parameter τ can lead to the near optimal performance already for these problem sizes. When the HH is allowed to choose between more than two operators, the experiments show that, already for the considered problem sizes, the performance of GRG improves with more operators for appropriate choices of the learning period τ .

The paper is structured as follows. In Section 2, we formally introduce the HH framework together with the simple selection HHs and GRG. In Section 3 we analyse the simple and generalised HHs using two low-level heuristics for the LEADINGONES benchmark function. In Section 4 we present the results for the HHs that choose between an arbitrary number of low-level heuristics. Section 5 presents the experimental analysis. In the conclusion we present a discussion and some avenues for future work.

Compared to its conference version (Lissovoi et al., 2017) this manuscript has been considerably extended. The results of Section 3 have been generalised to hold for any values of τ less than $(\frac{1}{2} - \varepsilon) n \ln n$, for some constant $0 < \varepsilon < \frac{1}{2}$. In general, the results have been considerably strengthened. In particular, compared to the extended abstract, in the present manuscript we prove that GRG optimises LEADINGONES in the best possible expected runtime achievable, up to lower order terms. Section 4 is a completely new addition, while a more comprehensive set of experiments is included in Section 5.

2 Preliminaries

2.1 The Hyper-heuristic Framework and the Heuristic Selection Methods

Let S be a finite search space, H a set of low-level heuristics and $f : S \rightarrow \mathbb{R}^+$ a fitness function. Algorithm 1 shows the pseudocode representation for a simple selection HH as used in previous experimental and theoretical work (Cowling et al., 2001, 2002; Alanazi and Lehre, 2014). Throughout the paper, the term runtime refers to the number of fitness evaluations used by Algorithm 1 until it finds the optimum. Note that the strict inequality in line 5 of Algorithm 1 is consistent with previous HH literature (Alanazi and Lehre, 2014), and using ' \geq ' instead would make no difference for the benchmark problem we consider in this work (i.e., LEADINGONES).

Different HHs are obtained from the framework described in Algorithm 1 according to which heuristic selection method is applied in Step 3. While sophisticated machine learning mechanisms are usually used, the following *learning mechanisms* have also been commonly considered in the literature to solve combinatorial optimisation problems (Cowling et al., 2001, 2002):

- **Simple Random**, which selects a low-level heuristic $h \in H$ independently with probability p_h in each iteration. Each heuristic is usually selected uniformly at random i.e., $p_h = \frac{1}{|H|}$;
- **Permutation**, which generates a random ordering of the low-level heuristics in H and returns them in that sequence when called by the HH;
- **Greedy**, which applies all the low-level heuristics in H in parallel and returns the best found solution;
- **Random Gradient**, which randomly selects a low-level heuristic $h \in H$, and keeps using it as long as it obtains improvements.

Algorithm 2 Generalised Random Gradient Hyper-heuristic

```
1: Choose  $x \in S$  uniformly at random
2: while stopping conditions not satisfied do
3:   Choose  $h \in H$  uniformly at random
4:    $c_t \leftarrow 0$ 
5:   while  $c_t < \tau$  do
6:      $c_t \leftarrow c_t + 1; x' \leftarrow h(x)$ 
7:     if  $f(x') > f(x)$  then
8:        $c_t \leftarrow 0; x \leftarrow x'$ 
```

Alanazi and Lehre (2014) derived upper and lower bounds on the expected runtime of Algorithm 1 for the LEADINGONES benchmark function when using these four simple heuristic selection mechanisms equipped with $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$. The first low-level heuristic flips one bit chosen uniformly at random while the second one chooses two bits to flip with replacement (i.e., the same bit may flip twice).

The runtime analysis performed by Alanazi and Lehre (2014) only provided bounds of the same asymptotic order on the expected runtime for all four mechanisms. An experimental analysis they carried out suggested, however, that all mechanisms have the same performance as just choosing the low-level heuristics at random for LEADINGONES. We will prove this by deriving the exact expected runtimes of the simple mechanisms in Section 3. Differently from the other three mechanisms, the idea behind Random Gradient is to try to learn from the past performance of the heuristics. However, by making a heuristic selection decision in every iteration, the mechanism does not have enough time to appreciate whether the selected heuristic is a good choice for the current optimisation state.

In this paper we generalise the Random Gradient mechanism to allow a longer period of time to appreciate whether a low-level heuristic is successful or not, before deciding whether to select a different low-level heuristic. Our aim is to maintain the intrinsic ideas of the simple learning mechanism while generalising it sufficiently to allow for learning to take place. The pseudocode of the resulting HH is given in Algorithm 2, while the proposed learning mechanism works as follows:

Generalised Random Gradient (GRG): A low-level heuristic is chosen uniformly at random (*Decision Stage*) and run for a learning period of fixed time τ . If an improvement is found before the end of the period, then a new period of time τ is immediately initialised (i.e., c_t in Algorithm 2 is set to 0 immediately) (*Exploitation Stage*). If the chosen operator fails to provide an improvement in τ iterations, a new operator is chosen at random.

Our aim is to prove that the simple GRG HH runs in the best possible expected runtime achievable for LEADINGONES using mutation operators with different neighbourhood sizes (i.e., $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ mutation operators) as low-level heuristics. We will first derive the best possible expected runtime achievable by applying the k operators in any order, and then prove that GRG matches it up to lower order terms.

Throughout the paper, when discussing the $m\text{BITFLIP}$ operator, we refer to the mutation operator which flips $m = \Theta(1)$ bits in the bit-string with replacement; that is, it is possible to flip and re-flip the same bit within the same mutation step. Since this has been used in previous literature on the topic (Lehre and Özcan, 2013; Alanazi and Lehre, 2014, 2016), we naturally continue with this choice. We will also prove that the presented results hold also if operators that flip m bits without replacement are used (i.e., operators that select a new bit-string with Hamming distance m from the original bit-string). In particular, we will show that any performance differences are limited to lower order terms in the expected runtimes. The latter operators are well known RLS algorithms with neighbourhood size m (i.e., RLS_m).

2.2 The Pseudo-Boolean Benchmark Function LEADINGONES

The LEADINGONES (LO) benchmark function counts the number of consecutive one-bits in a bit string before the first zero-bit:

$$\text{LEADINGONES}(x) := \sum_{i=1}^n \prod_{j=1}^i x_j.$$

The unrestricted black box complexity¹ of LEADINGONES is $O(n \log \log n)$ (Afshani et al., 2013) and there exist randomised search heuristics with expected runtimes of $o(n^2)$. Recently some Estimation of Distribution Algorithms (EDAs) have been presented which have surprisingly good performance for the problem. Doerr and Krejca (2018) introduced a modified compact Genetic Algorithm (cGA) called sig-cGA that, rather than updating the frequency vector in every generation, does so only once it notices a significance in its history of samples (i.e., once a significant number of 1s or 0s are sampled in a certain bit position, as opposed to a uniform binomial distribution of samples). They proved that the algorithm optimises LEADINGONES, ONEMAX and BINVAL in $O(n \log n)$ expected function evaluations. Two other algorithms have been proven to optimise LEADINGONES in the same asymptotic expected time: a stable compact Genetic Algorithm (scGA) that biases updates that favour frequencies that move towards $\frac{1}{2}$ (Friedrich et al., 2016) and a Convex Search Algorithm (CSA) using binary uniform convex hull recombination (for sufficiently large populations and with an appropriate restart strategy) (Moraglio and Sudholt, 2017). While the latter two algorithms have exceptional performance for LEADINGONES, their runtime is very poor for ONEMAX, respectively providing runtimes at least exponential and super-polynomial in the problem size with high probability (Doerr and Krejca, 2018).

The unbiased black box complexity of LEADINGONES is $\Theta(n^2)$ (Lehre and Witt, 2012). If biased mutation operators are allowed but truncation selection is imposed then no asymptotic improvement may be achieved over the unbiased black box complexity. Indeed, Doerr and Lengler (2018) recently proved that the best possible asymptotic performance of any (1+1) elitist black-box algorithm for LEADINGONES is $\Omega(n^2)$. This bound is matched by the performance of simple well studied heuristics. RLS has an expected runtime of $0.5n^2$ fitness function evaluations (Buzdalov and Buzdalova, 2015) and it is well known that the standard (1+1) EA (with mutation rate $\frac{1}{n}$) has an expected runtime of $\frac{\epsilon-1}{2}n^2 - o(n^2) \approx 0.85914n^2 - o(n^2)$ (Böttcher et al., 2010). Böttcher et al. (2010) also showed that the best static mutation rate for the (1+1) EA is $\approx \frac{1.5936}{n}$, which improves the expected runtime to $\frac{\epsilon}{4}n^2 \pm O(n) \approx 0.77201n^2 \pm O(n)$. Furthermore, they showed that the (1+1) EA with an appropriately chosen dynamic mutation rate (i.e., $\frac{1}{\text{LO}(x)+1}$) can outperform any static choice, giving an expected runtime of approximately $0.68n^2 \pm o(n^2)$. Note that the (1+1) EA is slowed down by the approximately 37% of the iterations in which no bits are flipped (Jansen and Zarges, 2011). The best possible (1+1) EA variant that uses standard bit mutation (discounting iterations in which no bits flip) with the best possible mutation rate at each step has an expected runtime of approximately $0.404n^2 \pm o(n^2)$ (Doerr et al., 2019). This is still worse than the best expected performance achievable by an unbiased (1+1) black box algorithm, approximately $0.388n^2 \pm o(n^2)$ (Doerr and Wagner, 2018; Doerr, 2018). In this paper we will prove that GRG equipped with a sufficient number of low-level heuristics can match this optimal expected runtime.

2.3 Mathematical Techniques

We now introduce some important tools which will be used in the analyses.

Theorem 1 Wald’s Equation (Wald, 1944; Doerr and Künnemann, 2013) Let T be a random variable with bounded expectation and let X_1, X_2, \dots be non-negative random variables with $E(X_i | T \geq i) \leq C$. Then

$$E\left(\sum_{i=1}^T X_i\right) \leq E(T) \cdot C.$$

Theorem 2 Additive Drift Theorem (He and Yao, 2001; Oliveto and Yao, 2011) Let $\{X_t\}_{t \geq 0}$ be a Markov process over a finite set of states S , and $g : S \rightarrow \mathbb{R}_0^+$ a function that assigns a non-negative real number to every state. Let the time to reach the optimum be $T := \min\{t \geq 0 \mid g(X_t) = 0\}$. If there exists $\delta > 0$ such that at any time step $t \geq 0$ and at any state $X_t \geq 0$ the following condition holds:

$$E(g(X_t) - g(X_{t+1}) \mid g(X_t) > 0) \geq \delta,$$

¹That is, the expected number of fitness function evaluations performed by the best-possible black-box algorithm until the optimum is found, i.e., there are no restrictions on which queries to the black box (i.e., which search points to be evaluated) the algorithm is allowed to make.

then

$$E(T \mid X_0) \leq \frac{g(X_0)}{\delta}$$

and

$$E(T) \leq \frac{E(g(X_0))}{\delta}.$$

3 Hyper-Heuristics Are Faster Than Their Low-level Heuristics

In this section we show that GRG can outperform its constituent low-level heuristics even when the latter are efficient for the problem at hand. To achieve this, we equip the HH with the minimal set of low level heuristics: $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$. We first introduce a lower bound on the expected runtime of all algorithms which use only the 1BITFLIP and 2BITFLIP operators for LEADINGONES. Afterwards we will show that GRG matches this expected runtime up to lower order terms.

To achieve the first result we will rely on the following theorem.

Theorem 3 (Böttcher et al. (2010); Buzdalov and Buzdalova (2015); Doerr (2018)) *Consider the run of an unbiased (1+1) black box algorithm optimizing the LEADINGONES function, and let T be the first time the optimum is generated. Given random initialisation:*

$$E(T) = \frac{1}{2} \sum_{i=0}^{n-1} A_i,$$

where A_i is the expected time needed to find an improvement given a solution with fitness i .

Theorem 3 holds for any unbiased (1+1) black box algorithm for LEADINGONES (Doerr, 2018). Since we analyse unbiased (1+1) black box algorithms, each differing by their mutation operator, we can apply Theorem 3 in our analysis.

A_i is the expected time to find an improvement when the current solution x consists of i leading 1-bits. Hence, the theorem implies that the overall expected runtime is minimised if the heuristic $h \in H$ with lowest expected improvement time A_i is applied whenever the current solution has i leading 1-bits, for all $0 \leq i < n$. Naturally, the heuristic with lowest expected time A_i is the one with the highest *success* probability (i.e., the probability of producing a fitness improvement).

The 1BITFLIP operator has a success probability of $P(\text{Imp}_1 \mid \text{LO}(x) = i) = \frac{1}{n}$. The success probability of the 2BITFLIP operator is $P(\text{Imp}_2 \mid \text{LO}(x) = i) = \frac{1}{n} \cdot \frac{n-i-1}{n} \cdot 2 = \frac{2n-2i-2}{n^2}$. Hence, for $i \geq \frac{n}{2}$, the 1BITFLIP operator is more effective (i.e., has a higher probability of success and hence a lower expected waiting time for each improvement A_i), while 2BITFLIP is preferable before. Note that both are equally effective when $\text{LO}(x) = \frac{n}{2} - 1$. The expected runtime of an algorithm that applies these operators in such a way is a lower bound on the expected runtime of all unbiased (1+1) black box algorithms using only the same two operators.

Theorem 4 *The best-possible expected runtime of any unbiased (1+1) black box algorithm using only the 1BITFLIP and 2BITFLIP operators for LEADINGONES is $\frac{1+\ln(2)}{4}n^2 + O(n) \approx 0.42329n^2 + O(n)$.*

Proof: We partition the analysis into two phases: Phase 1, where the 2BITFLIP operator is used for the first half of the search, and Phase 2, where the 1BITFLIP operator is used for the second half of the search. We consider the expected runtimes of each phase separately (respectively, $E(T_1)$ and $E(T_2)$) and sum them to find the total expected runtime.

To find the expected runtime for Phase 1 we can consider every point x such that $\text{LO}(x) = i \geq \frac{n}{2}$ to be optimal and we calculate the expected runtime for the 2BITFLIP operator to find any such point (i.e., $A_i = 0$ for $i \geq \frac{n}{2}$). Since the probability of improvement of the 2BITFLIP operator is $P(\text{Imp}_2 \mid \text{LO}(x) = i) =$

$\frac{2n-2i-2}{n^2}$, the expected time to make an improvement for $0 \leq i < \frac{n}{2}$ is $A_i = \frac{n^2}{2n-2i-2}$. Hence, for Phase 1,

$$\begin{aligned} E(T_1) &= \frac{1}{2} \sum_{i=0}^{n-1} A_i = \frac{1}{2} \sum_{i=0}^{\frac{n}{2}-1} \frac{n^2}{2n-2i-2} = \frac{1}{2} \sum_{k=\frac{n}{2}}^{n-1} \frac{n^2}{2k} \\ &= \frac{n^2}{4} \sum_{k=\frac{n}{2}}^{n-1} \frac{1}{k} = \frac{n^2}{4} \left(H_{n-1} - H_{\frac{n}{2}-1} \right), \end{aligned}$$

where H_x is the x_{th} Harmonic number, which can be bounded by $\ln(n+1) + \gamma - \frac{1}{n+1} \leq H_n \leq \ln(n+1) + \gamma - \frac{1}{2(n+1)}$ and γ is the Euler-Mascheroni constant ($\gamma \approx 0.57722\dots$). Hence, we can bound the term $H_{n-1} - H_{\frac{n}{2}-1}$ from below by $\ln(2)$, and from above by $\ln(2) + \frac{3}{2n}$, which gives the following bounds on $E(T_1)$:

$$\ln(2) \cdot \frac{n^2}{4} \leq E(T_1) \leq \left(\ln(2) + \frac{3}{2n} \right) \cdot \frac{n^2}{4}.$$

Since the upper bound on $E(T_1)$ is at most $\frac{3n}{8} = O(n)$ greater than the lower bound of $\ln(2) \cdot \frac{n^2}{4}$, we have:

$$\frac{\ln(2)}{4} n^2 \leq E(T_1) \leq \frac{\ln(2)}{4} n^2 + \frac{3n}{8}.$$

The expected runtime for Phase 2 can be calculated using Theorem 3 with $A_i = n$ for $i \geq \frac{n}{2}$. Similarly to Phase 1, we consider $A_i = 0$ for $i < \frac{n}{2}$.

$$E(T_2) = \frac{1}{2} \sum_{i=0}^{n-1} A_i = \frac{1}{2} \sum_{i=\frac{n}{2}}^{n-1} n = \frac{1}{2} \cdot n \cdot \frac{n}{2} = \frac{n^2}{4}.$$

Hence, the total expected runtime is:

$$\frac{1 + \ln(2)}{4} n^2 \leq E(T_1) + E(T_2) \leq \frac{1 + \ln(2)}{4} n^2 + \frac{3n}{8}.$$

□

Thus, we have proven a theoretical lower bound of $\frac{1+\ln(2)}{4} n^2 + O(n) \approx 0.42329n^2 + O(n)$ for the expected runtime of any unbiased (1+1) black box algorithm using the 1BITFLIP and 2BITFLIP operators for LEADINGONES, which is better in the leading constant than the expected runtime of RLS (i.e., $0.5n^2$).

The rest of the section is structured as follows. In Subsection 3.1 we show that the simple HH mechanisms all perform equivalently to Simple Random for LEADINGONES, up to lower order terms. In Subsection 3.2 we show that GRG is much faster and matches the optimal expected runtime derived in Theorem 4, up to lower order terms.

3.1 Simple Mechanisms

In this section we show that the standard simple heuristic selection mechanisms (i.e., Simple Random, Permutation, Greedy and Random Gradient) all have the same expected runtime for LEADINGONES, up to lower order terms. Note that, from the experiments performed by Alanazi and Lehre (2014), we know that these lower order terms do not have any visible impact on the average runtime already for small problem sizes.

The following theorem derives the expected runtime for the Simple Random mechanism. The subsequent corollary extends the result to the other mechanisms.

Theorem 5 *Let p_1 be the probability of choosing the 1BITFLIP mutation operator, and $1 - p_1$ the probability of choosing the 2BITFLIP mutation operator. Then, the expected runtime of the Simple Random mechanism using $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$, with $p_1 \in (0, 1)$, for LEADINGONES is $\frac{1}{4(1-p_1)} \ln\left(\frac{2-p_1}{p_1}\right) n^2 + o(n^2)$. If $p_1 = 0$ the expected runtime is infinite. If $p_1 = 1$, the expected runtime is $0.5n^2$.*

Proof: If $p_1 = 0$, only the 2BITFLIP operator is used. There is a non-zero probability of reaching the point $x = 1^{n-1}0$, which cannot be improved by using the 2BITFLIP operator. Hence, by the law of total probability, the expected runtime is infinite.

If $p_1 = 1$, only the 1BITFLIP operator is used, resulting in exactly RLS, which has expected runtime $0.5n^2$ (Buzdalov and Buzdalova, 2015).

If $p_1 \in (0, 1)$, the expected runtime is $\frac{1}{2} \sum_{i=1}^n A_i$, where A_i is the expected time needed to find an improvement given a solution with fitness i (Theorem 3). In each iteration, the 1BITFLIP operator is chosen with probability p_1 and leads to a fitness improvement with probability $P(\text{Imp}_1 \mid \text{LO}(x) = i) = \frac{1}{n}$; the 2BITFLIP operator is chosen with probability $1 - p_1$ and leads to a fitness improvement with probability $P(\text{Imp}_2 \mid \text{LO}(x) = i) = \frac{2n-2i-2}{n^2}$. Hence, $(A_i)^{-1} = p_1 \cdot \frac{1}{n} + (1 - p_1) \cdot \frac{2n-2i-2}{n^2}$, and

$$A_i = \frac{1}{\frac{p_1}{n} + \frac{(1-p_1)(2n-2i-2)}{n^2}} = \frac{n^2}{2(1-p_1)(n-i-1) + np_1}.$$

Hence, by Theorem 3, the expected runtime is

$$\frac{1}{2} \sum_{i=0}^{n-1} \frac{n^2}{2(1-p_1)(n-i-1) + np_1} = \frac{n^2}{2} \sum_{k=1}^n \frac{1}{(2p_1-2)k + (2-p_1)n} \quad (1)$$

$$= \frac{n^2}{2(2p_1-2)} \left(\ln \left(\frac{1 + \frac{(2-p_1)}{2p_1-2}}{\frac{(2-p_1)}{2p_1-2}} \right) + o(1) \right) = \frac{n^2}{4(1-p_1)} \ln \left(\frac{2-p_1}{p_1} \right) + o(n^2), \quad (2)$$

where Equation 1 becomes Equation 2 via the following simplification:

$$\begin{aligned} \sum_{k=1}^n \frac{1}{a \cdot k + b \cdot n} &= \frac{1}{a} \cdot \sum_{k=1}^n \frac{1}{k + \frac{b}{a} \cdot n} = \frac{1}{a} \left(\sum_{k=1}^{(1+\frac{b}{a})n} \frac{1}{k} - \sum_{k=1}^{\frac{b}{a}n} \frac{1}{k} \right) \\ &= \frac{1}{a} \left(\ln \left(\frac{1 + \frac{b}{a}}{\frac{b}{a}} \right) + o(1) \right) \end{aligned}$$

with $a = 2p_1 - 2$ and $b = 2 - p_1$. □

When $p_1 = 0.5$ (i.e., there is an equal chance of choosing each operator in each iteration), the standard Simple Random mechanism has an expected improvement time of $A_i = \frac{2n^2}{3n-2i-2}$, and an expected runtime of $\frac{\ln(3)}{2}n^2 + o(n^2) \approx 0.54931n^2 + o(n^2)$. The expected runtime improves with increasing p_1 , hence the optimal choice is $p_1 = 1$ (i.e., RLS).

Corollary 6 *The expected runtime of the Permutation, Greedy and Random Gradient mechanisms using $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$ for LEADINGONES is $\frac{\ln(3)}{2}n^2 + o(n^2) \approx 0.54931n^2 + o(n^2)$.*

Proof: Given a bit-string x with $\text{LO}(x) = i$, let p_i be the probability that at least one improvement is constructed within two fitness function evaluations. For the Permutation and Greedy mechanisms, we have $p_i = \frac{1}{n} + (1 - \frac{1}{n}) \cdot \frac{2n-2i-2}{n^2}$ as either the 1BITFLIP operator can succeed in one iteration, or it can fail and then the 2BITFLIP operator can succeed in the next iteration. To derive an upper bound on the expected optimisation time, we note that the difference between $\frac{2}{p_i}$ (i.e., the expected waiting time for an improvement to be constructed in terms of fitness evaluations) and the A_i waiting times for the standard Simple Random mechanism with $p_1 = 0.5$ (i.e., $A_i = \frac{2n^2}{3n-2i-2}$) is at most constant, and thus the difference between the expected runtimes of these mechanisms and Simple Random is limited to lower order terms. We note that with probability at most $\frac{2}{n^2}$, both mutations (the two mutations performed in parallel by the Greedy mechanism, or the two mutations performed sequentially by the Permutation mechanism) considered are improvements. As this occurs at most a constant number of times in expectation (i.e., $O(n^2) \cdot \frac{2}{n^2} = O(1)$), and the maximum expected waiting time for any improving step is $O(n)$, the lower bound differs from the upper bound by at most an $O(1) \cdot O(n) = O(n)$ term. Thus, the expected runtime of these mechanisms is also $\frac{\ln(3)}{2} \cdot n^2 + o(n^2)$.

For the Random Gradient mechanism, we note that the probability that an operator, when repeated following a success, is successful again is at most $\frac{2}{n}$ (an upper bound on the success probability of the 2BITFLIP operator). Since there are at most n improvements to be made throughout the search space, the expected number of repetitions which produce an improvement is at most 2. If the chosen operator is not successful, the Random Gradient mechanism behaves identically to the Simple Random mechanism. Its expected runtime is therefore at least the expected runtime of the Simple Random mechanism less an $O(n)$ term, and at most the expected runtime of the Simple Random mechanism plus n (as there are at most n iterations where the mechanisms differ in operator selection). Thus, its expected runtime is also $\frac{\ln(3)}{2} \cdot n^2 + o(n^2)$. \square

We point out that the lower bound for the Random Gradient HH contradicts the lower bound of $\frac{n^2}{9}(4 + 3 \ln(\frac{10}{3})) \approx 0.846n^2 + o(n^2)$ found by Alanazi and Lehre (2014). However, their bound results from a small mistake in their proof. They should have found a lower bound of $\frac{n^2}{9}(3 \ln(\frac{10}{3})) + o(n^2) \approx 0.401n^2 + o(n^2)$, which agrees with our derived expected runtime of $\frac{\ln 3}{2}n^2 + o(n^2) \approx 0.549n^2 + o(n^2)$.

3.2 Generalised Random Gradient

In this subsection we present a rigorous theoretical analysis of GRG for LEADINGONES. The main result of this subsection is that GRG is able to match, up to lower order terms, the best-possible performance of any algorithm using the 1BITFLIP and 2BITFLIP operators for LEADINGONES, as presented in Theorem 4. We state the main result in Corollary 7 now, with the proof at the end of the subsection.

Corollary 7 [Of Theorem 8] *The expected runtime of the Generalised Random Gradient hyper-heuristic using $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$ for LEADINGONES, with τ that satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{2} - \varepsilon) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{2}$, is at most $\frac{1+\ln(2)}{4}n^2 + o(n^2) \approx 0.42329n^2 + o(n^2)$.*

This improves significantly upon the result of Lissovoi et al. (2017). While the previous result showed that the HH outperformed its low level heuristics, Corollary 7 also shows that the expected runtime matches the optimal one up to lower order terms. Our previous work only considered setting $\tau = cn$ for constant c , a learning period for which optimal expected runtime could not be proven.

Before proving the corollary, we first present the necessary prerequisite results. The following theorem is very general as it provides an upper bound on the expected runtime of GRG for any value of τ smaller than $\frac{1}{2}n \ln n$. In particular, Theorem 8 allows us to identify values of τ for which the expected runtime of the HH is the optimal one achievable using the two operators. This is the result highlighted in Corollary 7 and depicted in Figure 1.

Our proof technique partitions the search space into w stages, each representing an equal range of fitness values. The expected runtime of the algorithm does not depend on w . However, w does affect the upper bound on the runtime we obtain. The greater the number of stages, the tighter the upper bound the theorem statement provides as long as $w = o\left(\frac{n}{\exp(\tau/n)}\right)$.

Theorem 8 *The expected runtime of the Generalised Random Gradient hyper-heuristic using $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$ for LEADINGONES, with $\tau \leq (\frac{1}{2} - \varepsilon) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{2}$, is at most*

$$\frac{n^2}{2} \left(\sum_{j=1}^w \frac{\frac{2\tau}{n} + e^{\frac{\tau}{n}} M_1(1) + e^{2\frac{\tau}{n}(1-\frac{j}{w})} M_2(j, w)}{(e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2)w} \right) + o(n^2)$$

where

$$M_2(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \text{ or } \frac{1}{2(1-\frac{j}{w})} > \frac{\tau}{n} \\ \frac{1}{2(1-\frac{j}{w})} & \text{otherwise} \end{cases}$$

and $M_1(x) := \min\{\frac{\tau}{n}, x\}$, with $w \in \omega\left(\frac{\tau}{n}\right) \cap o\left(\frac{n}{\exp(\tau/n)}\right)$.

Figure 1 presents the theoretical upper bounds from Theorem 8 for a range of linear τ values. For $\tau = 5n$, GRG already outperforms RLS, giving an expected runtime of $0.46493n^2 + o(n^2)$. For $\tau = 30n$, the




Figure 1: The leading constants in the theoretical upper bounds on the expected number of fitness function evaluations required by the Generalised Random Gradient hyper-heuristic to find the LEADINGONES optimum ($w = 100,000$). 2_{Opt} is the leading constant in the best expected runtime achievable using the 1BITFLIP and 2BITFLIP operators in any order.

performance improves to $0.42385n^2 + o(n^2)$, matching the best possible performance from Theorem 4 up to 3 decimal places. We have seen that GRG is able to exactly match this optimal performance, up to lower order terms, for $\tau = \omega(n)$ in Corollary 7.

Proof: [Of Theorem 8]

For the purpose of this proof, we partition the optimisation process into w stages based on the value of the LEADINGONES fitness function: during stage j , the LEADINGONES value of the current solution is at least $\frac{(j-1)n}{w}$ and less than $\frac{jn}{w}$. After all the w stages have been completed, the global optimum with a LEADINGONES value of n will have been found.

We will provide an upper bound on the expectation of the runtime T of GRG for LEADINGONES by summing up the expected values of T_j (the expected time between the first random operator choice in each stage of the optimisation process and the moment when the necessary fitness value to enter the following stage is reached). More precisely, our analysis of $E(T_j)$ requires the stage to start with a random choice of mutation operator, we bound $E(T) \leq \sum_{j=1}^w (E(T_j) + E(S_{j+1}))$ where S_{j+1} is a random variable denoting the expected number of iterations between the first solution in stage $j+1$ being constructed and the first random operator choice occurring in that stage. We will later show that with proper parameter choices, the contribution of the S_{j+1} terms can be bounded by $o(n^2)$, and therefore they do not affect the leading constant in the overall bound.

Let us now consider T_j . Recall that a mutation operator is selected uniformly at random and is allowed to run until it fails to produce an improvement within τ consecutive iterations. Let N_j be a random variable denoting the number of random operator choices the HH performs during stage j , and let $X_{j,1}, \dots, X_{j,N_j}$ be the number of iterations the HH runs each chosen operator for. We note that $T_j = \sum_{k=1}^{N_j} X_{j,k}$ is a sum of a number of non-negative variables, and we will later show that $E(N_j)$ is bounded; thus, $E(T_j)$ can be bounded by applying Wald's equation (Theorem 1):

$$E(T_j) = E\left(\sum_{k=1}^{N_j} X_{j,k}\right) \leq E(N_j)E(X_j), \quad (3)$$

where $E(X_j)$ denotes an upper bound on all $E(X_{j,k})$ in stage j .

To bound $E(N_j)$, the expected number of times the random operator selection is performed during stage j , we will provide a lower bound on the expected number of improvements found following the operator selection and apply the Additive Drift Theorem (Theorem 2) to find the expected number of random operator selections that occur before a sufficient number of improvements for entering the next stage have been found.

Let F_1 and F_2 denote the events that the 1BITFLIP and 2BITFLIP operators fail to find an improvement during τ iterations. For the 1BITFLIP operator (with probability of fitness improvement in one iteration $\frac{1}{n}$), this event occurs with probability $P(F_1) = (1 - \frac{1}{n})^\tau$ throughout the process, which is within $(e^{-\frac{\tau}{n}} - \frac{1}{n}, e^{-\frac{\tau}{n}}]$ (since $(1 - \frac{1}{n})^n \leq \frac{1}{e} \leq (1 - \frac{1}{n})^{n-1}$). For the 2BITFLIP operator (with probability of fitness improvement in one iteration $\frac{2n-2i-2}{n^2}$), recall that during stage j , the ancestor individual has at most $i = \frac{jn}{w} - 1$ and at least $i = \frac{(j-1)n}{w}$ 1-bits. Thus:

$$P(F_2) \leq \left(1 - 2 \cdot \frac{1}{n} \cdot \frac{n - (\frac{jn}{w} - 1) - 1}{n}\right)^\tau \leq e^{-2\frac{\tau}{n}(1 - \frac{j}{w})},$$

and

$$\begin{aligned} P(F_2) &\geq \left(1 - 2 \cdot \frac{1}{n} \cdot \frac{n - \frac{(j-1)n}{w} - 1}{n}\right)^\tau \\ &> \left(1 - \frac{2 \cdot (1 - \frac{j-1}{w})}{n}\right)^\tau \geq e^{-\frac{2\tau}{n}(1 - \frac{j-1}{w})} - \frac{1}{n}. \end{aligned}$$

A geometric distribution with parameter $p = P(F_1)$ (or $p = P(F_2)$) can be used to model the number of improvements that the 1BITFLIP (or 2BITFLIP) operator finds prior to failing to find an improvement for τ iterations. The expectation of this distribution is $\frac{1-p}{p} = \frac{1}{p} - 1$. We can model these events with a geometric random variable since GRG immediately restarts a period of τ iterations with the chosen heuristic when a successful mutation occurs (i.e., it sets $c_t = 0$ in Algorithm 2). Hence, there are no periods of τ consecutive unsuccessful mutations apart from the period when the operator finally fails. Combined over both operators, the expected number of improvements D_j produced following a single random operator selection during stage j is:

$$E(D_j) = \frac{1}{2} \left(\frac{1}{P(F_1)} - 1\right) + \frac{1}{2} \left(\frac{1}{P(F_2)} - 1\right) \geq \frac{e^{\frac{\tau}{n}}}{2} + \frac{e^{\frac{2\tau}{n}(1 - \frac{j}{w})}}{2} - 1,$$

where we inserted the upper bounds on $P(F_1)$ and $P(F_2)$ to obtain the inequality. We will use this expectation as the drift on the progress of the randomly chosen operator in the Additive Drift Theorem to provide an upper bound on $E(N_j)$. Recall that each stage requires advancing through at most $\frac{n}{w}$ fitness values. Since bits beyond the leading 1-bit prefix and the first 0-bit remain uniformly distributed, at most $\frac{n}{2w}$ improvements by mutation are required in expectation. If each step of a random process in expectation contributes $E(D_j)$ improvements by mutation, then the expected number of steps required to complete stage j is at most

$$E(N_j) \leq \frac{\frac{n}{2w}}{E(D_j)} \leq \frac{n}{\left(e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1 - \frac{j}{w})} - 2\right) w}, \quad (4)$$

by the Additive Drift Theorem.

We apply Wald's equation to bound $E(X_j)$, the expected number of iterations before a selected mutation operator fails to produce an improvement for τ iterations: let S be the number of improvements found by the operator before it fails, and W_1, \dots, W_S be the number of iterations it took to find each of those improvements; then, once selected, the 1BITFLIP operator runs for

$$E(X_j \mid \text{1BITFLIP}) = \tau + \sum_{k=1}^S E(W_k) = \tau + E(S)E(W_1 \mid S \geq 1, \text{1BITFLIP}),$$

where τ accounts for the iterations immediately before failure and the sum counts the expected number of iterations preceding each constructed improvement. Recall that $E(S) = \frac{1}{P(F_1)} - 1$ by the properties of the geometric distribution and observe that $E(W_1 \mid S \geq 1) = E(W_1 \mid W_1 \leq \tau) \leq \min\{\tau, E(W_1)\}$. Using a

waiting time argument (i.e., considering the expectation of a geometric random variable) gives $E(W_1) = n$, and with a similar argument for the 2BITFLIP operator, $E(W_2) \leq \frac{n^2}{2(n - (\frac{j}{w} - 1))} = \frac{n}{2(1 - \frac{j}{w})}$. Hence:

$$\begin{aligned} E(X_j | 1\text{BITFLIP}) &\leq \tau + \left(\frac{1}{P(F_1)} - 1 \right) \min\{\tau, n\} \\ E(X_j | 2\text{BITFLIP}) &\leq \tau + \left(\frac{1}{P(F_2)} - 1 \right) \min \left\{ \tau, \frac{n}{2(1 - \frac{j}{w})} \right\}. \end{aligned}$$

Combining these conditional expectations with lower bounds on $P(F_1)$ and $P(F_2)$ yields

$$\begin{aligned} E(X_j) &\leq \frac{1}{2} \cdot E(X_j | 1\text{BITFLIP}) + \frac{1}{2} \cdot E(X_j | 2\text{BITFLIP}) \tag{5} \\ &\leq \tau + \frac{\min\{\tau, n\}}{2(e^{-\frac{\tau}{n}} - \frac{1}{n})} + \frac{\min \left\{ \tau, \frac{n}{2(1 - \frac{j}{w})} \right\}}{2 \left(e^{-\frac{2\tau}{n}(1 - \frac{j-1}{w})} - \frac{1}{n} \right)} \\ &\leq \frac{n}{2} \left(\frac{2\tau}{n} + \frac{\min\{\frac{\tau}{n}, 1\}}{e^{-\frac{\tau}{n}} - \frac{1}{n}} + \frac{\min \left\{ \frac{\tau}{n}, \frac{1}{2(1 - \frac{j}{w})} \right\}}{e^{-\frac{2\tau}{n}(1 - \frac{j-1}{w})} - \frac{1}{n}} \right) \\ &\leq \frac{n}{2} \left(\frac{2\tau}{n} + e^{\frac{\tau}{n}} \min \left\{ \frac{\tau}{n}, 1 \right\} + e^{\frac{2\tau}{n}(1 - \frac{j}{w})} \min \left\{ \frac{\tau}{n}, \frac{1}{2(1 - \frac{j}{w})} \right\} \right) + O(1), \end{aligned}$$

using $e^{\frac{2\tau}{n}(1 - \frac{j-1}{w})} = e^{\frac{2\tau}{wn}} \cdot e^{\frac{2\tau}{n}(1 - \frac{j}{w})} = (1 + o(1)) \cdot e^{\frac{2\tau}{n}(1 - \frac{j}{w})}$, since $w = \omega\left(\frac{\tau}{n}\right)$, and $\frac{a}{b - \frac{c}{n}} = \frac{a}{b} + \frac{ac}{b^2n - bc} = \frac{a}{b} + O\left(\frac{1}{n}\right)$ to limit the contributions of the $-\frac{1}{n}$ terms in the denominators to lower order terms.

If $j = w$, then $\min \left\{ \frac{\tau}{n}, \frac{1}{2(1 - \frac{j}{w})} \right\}$ is undefined. However, we know that the chosen operator would only be applied for a maximum of τ steps. We thus refer to the functions $M_1(x) = \min\{\frac{\tau}{n}, x\}$ and $M_2(j, w)$ from now on, where

$$M_2(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \text{ or } \frac{1}{2(1 - \frac{j}{w})} > \frac{\tau}{n} \\ \frac{1}{2(1 - \frac{j}{w})} & \text{otherwise.} \end{cases}$$

Finally, we can bound the overall expected runtime of the HH. Recall that S_j denotes the number of iterations the HH spends in stage j while using the random operator chosen during stage $j - 1$, and since $E(S_j) \leq \max(E(X_j | 1\text{BITFLIP}), E(X_j | 2\text{BITFLIP})) < 2E(X_j) = O(n \cdot \exp(\frac{2\tau}{n}))$ and $w = o\left(\frac{n}{\exp(\frac{2\tau}{n})}\right)$, it follows that $\sum_{j=1}^w E(S_{j+1}) = o(n^2)$.

Substituting the bounds provided in Eq 4 and Eq. 5 into Eq. 3 yields the theorem statement:

$$\begin{aligned} E(T) &\leq \sum_{j=1}^w (E(T_j) + E(S_{j+1})) \leq \left(\sum_{j=1}^w E(N_j)E(X_j) \right) + o(n^2) \\ &\leq o(n^2) + \frac{n^2}{2} \times \sum_{j=1}^w \frac{\frac{2\tau}{n} + e^{\frac{\tau}{n}} M_1(1) + e^{\frac{2\tau}{n}(1 - \frac{j}{w})} M_2(j, w)}{\left(e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1 - \frac{j}{w})} - 2 \right) w}. \end{aligned}$$

□

We can now prove the main result of this section (i.e., Corollary 7), which states that the expected runtime of GRG using $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$ for LEADINGONES with τ that satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{2} - \varepsilon) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{2}$, is at most $\frac{1 + \ln(2)}{4} n^2 + o(n^2) \approx 042329n^2 + o(n^2)$.

Proof: [Of Corollary 7]

Consider first the terms $M_1(1)$ and $M_2(j, w)$. Since $\tau = \omega(n)$, we have $M_1(1) = 1$. For $M_2(j, w)$ note that for the first half of the search, $M_2(j, w) = \frac{1}{2(1-\frac{j}{w})}$, and in the second half of the search, $M_2(j, w) = O(w) = n^{o(1)}$ (excluding the case when $j = w$ in which $M_2(j, w) = \frac{\tau}{n} = n^{o(1)}$). Note also that $\frac{\tau}{n} = \omega(1)$. Hence, we can simplify the sum from Theorem 8 as follows:

$$\begin{aligned}
E(T) &\leq \frac{n^2}{2} \left(\sum_{j=1}^w \frac{\frac{2\tau}{n} + e^{\frac{\tau}{n}} M_1(1) + e^{\frac{2\tau}{n}(1-\frac{j}{w})} M_2(j, w)}{(e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2)w} \right) + o(n^2) \\
&= \frac{n^2}{2} \left(\sum_{j=1}^w \frac{\frac{2\tau}{n} + e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1-\frac{j}{w})} M_2(j, w)}{(e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2)w} \right) + o(n^2) \\
&= \frac{n^2}{2} \left(\sum_{j=1}^w \frac{1}{w} \left(\frac{\frac{2\tau}{n} + 2}{e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2} + \frac{e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1-\frac{j}{w})} M_2(j, w) - 2}{e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2} \right) \right) \\
&\quad + o(n^2) \\
&= \frac{n^2}{2} \left(\sum_{j=1}^w \frac{1}{w} \left(\frac{e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1-\frac{j}{w})} M_2(j, w) - 2}{e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2} \right) \right) + o(n^2).
\end{aligned}$$

The exponential terms will asymptotically dominate the summation term (since $\frac{\tau}{n} = \omega(1)$). When $j > \frac{w}{2}$, we note that $\exp(\frac{\tau}{n}) > \exp(\frac{2\tau}{n}(1-\frac{j}{w}))$ and this term will dominate the other (vice-versa for $j < \frac{w}{2}$). At $j = \frac{w}{2}$, the two values are equal (with the term in the summand equal to 1). Hence, we can split the sum into three parts:

$$\begin{aligned}
E(T) &\leq \frac{n^2}{2} \left(\sum_{j=1}^w \frac{1}{w} \left(\frac{e^{\frac{\tau}{n}} + e^{\frac{2\tau}{n}(1-\frac{j}{w})} M_2(j, w) - 2}{e^{\frac{\tau}{n}} + e^{2\frac{\tau}{n}(1-\frac{j}{w})} - 2} \right) \right) + o(n^2) \\
&= \frac{n^2}{2} \left(\sum_{j=1}^{w/2-1} \frac{1}{w} \left(\frac{e^{\frac{2\tau}{n}(1-\frac{j}{w})} \cdot \frac{1}{2(1-\frac{j}{w})}}{e^{\frac{2\tau}{n}(1-\frac{j}{w})}} \right) \right) + \frac{n^2}{2w} + \frac{n^2}{2} \left(\sum_{j=w/2+1}^w \frac{1}{w} \left(\frac{e^{\frac{\tau}{n}}}{e^{\frac{\tau}{n}}} \right) \right) \\
&\quad + o(n^2) \\
&= \frac{n^2}{4} \left(\sum_{j=1}^{w/2} \frac{1}{w-j} \right) + \frac{n^2}{4} + o(n^2) = \left(\frac{1 + \ln(2)}{4} \right) n^2 + o(n^2).
\end{aligned}$$

Note that at $j = w$, the respective ‘dominating terms’ are $\exp(\frac{\tau}{n})$ and $\exp(\frac{2\tau}{n}(1-\frac{w}{w})) = 1$, and $M(j, w) = \frac{\tau}{n} < \frac{1}{2} \ln(n)$. Hence, $\exp(\frac{\tau}{n})$ will dominate. □

4 Increasing the Choice of Low-level Heuristics Leads to Improved Performance

Runtime analyses of HHs are easier if the algorithms can only choose between two operators. However, in realistic contexts, HHs have to choose between many more operators.

In this section, we consider the previously analysed GRG HH, yet extend the set H of low-level heuristics to be of size $k \geq 2$ i.e., $|H| = k = \Theta(1)$. We extend upon the previous analysis by considering the k operators as k different mutation operators, each flipping between 1 and k bits with replacement uniformly at random. This more accurately represents HH approaches employed for real-world problems.

As before, we consider the LEADINGONES benchmark function. We will rigorously show that the performance of the simple mechanisms deteriorates as the number ($|H| = k = \Theta(1)$) of operators increases. In

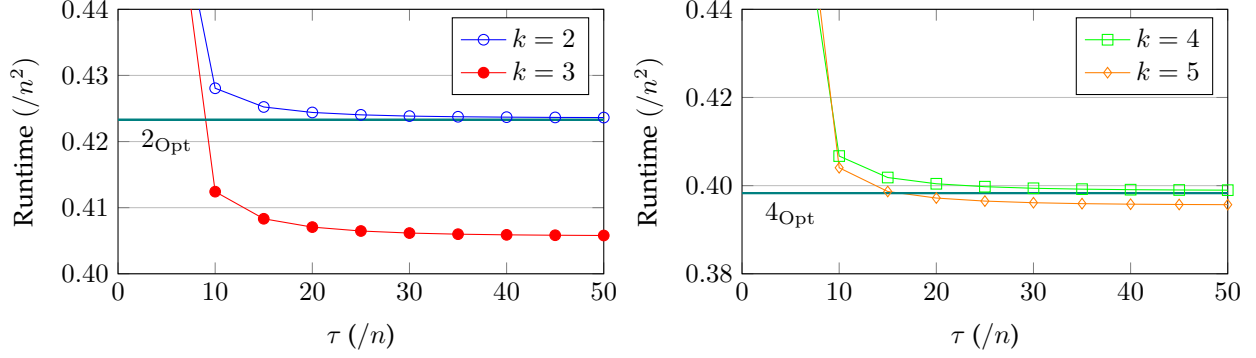


Figure 2: A comparison of the optimal expected runtimes of the Generalised Random Gradient hyper-heuristic with k -operators against the leading constant in the theoretical upper bound of the expected runtime of the Generalised Random Gradient hyper-heuristic with $k + 1$ -operators; for $k = 2$ and $k = 4$. 2_{Opt} ($\approx 0.42329n^2 + o(n^2)$) and 4_{Opt} ($\approx 0.39830n^2 + o(n^2)$) are the best possible runtimes achievable by the GRG hyper-heuristic with access to 2 and 4 low-level heuristics respectively, from Corollary 16.

addition, we prove decreasing upper bounds on the expected runtime of GRG as the number of operators increases.

The main result of this section is the following theorem, which states that GRG that chooses between k stochastic mutation operators has better expected performance than any algorithm, including the best-possible, using less than k stochastic mutation operators for LEADINGONES. We present the statement of Theorem 9 now, and the proof at the end of this section.

Theorem 9 *The expected runtime for LEADINGONES of the Generalised Random Gradient hyper-heuristic using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$, with τ that satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{k} - \varepsilon) n \ln(n)$ for some constant $0 < \varepsilon < \frac{1}{k}$, is smaller than the best-possible expected runtime of any unbiased (1+1) black box algorithm using any strict subset of $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$.*

Theorem 9 highlights the power of HHs as general-purpose problem solvers. The inclusion of more heuristics to the set of low-level heuristics is implied to be preferable, showcasing the impressive learning capabilities of even simple HHs. Figure 2 highlights the meaning of Theorem 9 for $k = 3$ and $k = 5$. The expected runtime of GRG with three operators is better than the best possible expected runtime achievable using the first two operators in any combination. The figure also highlights that with five operators the expected runtime is better than any one achievable using fewer operators.

We now present the necessary prerequisite analysis to obtain the result. We first derive the best-possible expected runtime achievable by any unbiased (1+1) black box algorithm using $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ mutation operators for LEADINGONES. Before we prove this, we introduce the following two helpful lemmata (which hold for all problem sizes $n \geq 1$).

Lemma 10 *Given a bit-string x with $LO(x) = i$, the probability of improvement of a mutation operator that flips $m = \Theta(1)$ bits with replacement ($m\text{BITFLIP}$) is*

$$P(\text{Imp}_m \mid LO(x) = i) = m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} + O\left(\frac{1}{n^2}\right).$$

Proof: In the context of the LEADINGONES function, the way for the $m\text{BITFLIP}$ operator to make an improvement is to flip the first 0-bit after the leading 1-bits, whilst keeping the leading 1-bits in the prefix unchanged. This occurs with probability $m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}$.

It is also possible to make fitness improvements while one or more bits of the prefix (i.e., within the first $i + 1$ bits) are flipped, each an even number of times. This can occur by flipping and re-flipping between 1 and $\lfloor \frac{m-1}{2} \rfloor$ bits. Firstly, the first 0-bit must be flipped into a 1-bit, which occurs with probability at most $\frac{m}{n}$. Then, an even number of the remaining bitflips can be used to flip bits in the prefix, which need to be

reflipped to maintain the prefix. The remaining flips will occur within the suffix $n - i - 1$ bits and these events can occur in any order. Hence, the probability of producing an improvement while flipping and reflipping an even number of bits (> 0) in the prefix is at most:

$$\begin{aligned} & \frac{m}{n} \cdot \sum_{j=1}^{\lfloor \frac{m-1}{2} \rfloor} \left[\binom{m-1}{2j} \cdot \frac{1}{i+1} \cdot \left(\frac{i+1}{n-1} \right)^{2j} \cdot \left(\frac{n-i-1}{n-1} \right)^{m-2j-1} \right] \\ &= \frac{m}{n} \cdot \binom{m-1}{2} \cdot \frac{1}{i+1} \cdot \left(\frac{i+1}{n-1} \right)^2 \cdot \left(\frac{n-i-1}{n-1} \right)^{m-3} + O\left(\frac{1}{n^2}\right) = O\left(\frac{1}{n^2}\right), \end{aligned}$$

given $m = \Theta(1)$. Hence, the probability of improvement of the m BITFLIP operator for LEADINGONES is

$$P(\text{Imp}_m \mid \text{LO}(x) = i) = m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n} \right)^{m-1} + O\left(\frac{1}{n^2}\right).$$

□

Lemma 11 *Given a bit-string x with $\text{LO}(x) = i$, consider two mutation operators flipping $a = \Theta(1)$ and $b = \Theta(1)$ bits with replacement respectively (a BITFLIP and b BITFLIP), with $b > a$. Then (excluding $i = n - 1$) b BITFLIP has a higher probability of success than a BITFLIP (i.e., $P(\text{Imp}_b \mid \text{LO}(x) = i) > P(\text{Imp}_a \mid \text{LO}(x) = i)$) when*

$$i < n \left(1 - \left(\frac{b}{a} \right)^{\frac{1}{a-b}} \right) - 1.$$

Lemma 11 follows from Lemma 10. In particular, an operator that flips m bits (i.e., m BITFLIP) has a higher probability of success than an operator that flips $m - 1$ bits (i.e., $m - 1$ BITFLIP) or less when $i < \frac{n}{m} - 1$. It is worth noting that only an operator which flips an odd number of bits can make progress when $\text{LO}(x) = n - 1$, and the 1BITFLIP operator has the best success probability at this point (i.e., $P(\text{Imp}_1 \mid \text{LO}(x) = i) = \frac{1}{n}$). We additionally note through a simple calculation that the 1BITFLIP operator has the highest probability of success in the second half of the search space (i.e., when $\frac{n}{2} \leq i < n$).

In Theorem 12 we present the best-possible expected runtime for any unbiased (1+1) black box algorithm using $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ as mutation operators for LEADINGONES. Similar results have recently been presented by Doerr (2018) and Doerr and Wagner (2018) for mutation operators that flip bits without replacement. The expected runtimes are the same up to lower order terms.

Theorem 12 *The best-possible expected runtime of any unbiased (1+1) black box algorithm using $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$ for LEADINGONES is*

$$E(T_{k,\text{Opt}}) = \frac{1}{2} \left(\sum_{i=0}^{\frac{n}{k}-1} \frac{1}{k \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n} \right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{i=\frac{n}{m+1}}^{\frac{n}{m}-1} \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n} \right)^{m-1}} \right) \pm o(n^2).$$

Proof: We refer to an operator as optimal if it yields the highest probability of success based on the current $\text{LO}(x) = i$ value. From Lemma 11 the m BITFLIP operator is optimal when $\frac{n}{m+1} \leq i \leq \frac{n}{m} - 1$ (unless $m = k$, in which case k BITFLIP is optimal for $0 \leq i \leq \frac{n}{k} - 1$).

The best-possible expected runtime is achieved by an algorithm that always uses the optimal operator. Such an algorithm will apply the 1BITFLIP operator while $\text{LO}(x) \geq \frac{n}{2}$ and the other operators will have $P(\text{Imp}_m \mid \text{LO}(x) = i) = \Theta\left(\frac{1}{n}\right)$ in their optimal regions. Hence, any $O\left(\frac{1}{n^2}\right)$ terms from the $P(\text{Imp}_m \mid \text{LO}(x) = i)$ terms (for $2 \leq m \leq k$) are asymptotically insignificant and their contributions can be grouped into a lower order $o(n^2)$ term in the final runtime.

Hence, by standard waiting time arguments, coupled with Theorem 3, we have that the best-possible expected runtime when using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ is

$$E(T_{k,\text{Opt}}) = \frac{1}{2} \left(\sum_{i=0}^{\frac{n}{k}-1} \frac{1}{k \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n} \right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{i=\frac{n}{m+1}}^{\frac{n}{m}-1} \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n} \right)^{m-1}} \right) \pm o(n^2).$$

□

In particular, taking limits as $n \rightarrow \infty$, we have $E(T_{1,\text{Opt}}) = \frac{1}{2}n^2$, $E(T_{2,\text{Opt}}) = \frac{1+\ln(2)}{4}n^2 \approx 0.42329n^2$, $E(T_{3,\text{Opt}}) = \left(\frac{1}{3} + \frac{\ln(2)}{2} - \frac{\ln(3)}{4}\right)n^2 \approx 0.40525n^2$, $E(T_{5,\text{Opt}}) = \left(\frac{3721}{11520} + \frac{\ln(2)}{2} - \frac{\ln(3)}{4}\right)n^2 \approx 0.39492n^2$. A closed form result for $E(T_{k,\text{Opt}})$ is difficult to find as is the limit for the best-possible expected runtime as $k \rightarrow \infty$. A numerical analysis by Doerr and Wagner (2018) suggests an expected runtime of $E(T_{\infty,\text{Opt}}) \approx 0.388n^2 \pm o(n^2)$.

To emphasise the practical importance of our result, we first prove that the performance of simple mechanisms considered in the literature (Cowling et al., 2001, 2002; Alanazi and Lehre, 2014) worsens with the increase of the low-level heuristic set size.

4.1 ‘Simple’ Mechanisms

We will now see how the simple learning mechanisms (Simple Random, Permutation, Greedy and Random Gradient) perform when having to choose between k operators (i.e., $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$) for LEADINGONES. We will show that incorporating more operators is detrimental to the performance of the simple mechanisms for LEADINGONES.

We start again by stating the expected runtime of the Simple Random mechanism, and use this as a basis for the other three mechanisms. Recall that the standard Simple Random mechanism chooses each operator uniformly at random in each iteration (i.e., with probability $\frac{1}{k}$ when using k operators).

Theorem 13 *The expected runtime of the Simple Random mechanism using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$ for LEADINGONES is*

$$\frac{k}{2} \cdot \sum_{i=0}^{n-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} - o(n^2).$$

Note that the expected runtime of the Simple Random mechanism increases with k . Hence, incorporating more operators is detrimental to its performance. In particular, the expected runtimes when using 1, 2 and 3 operators are respectively (up to lower order terms) $\frac{1}{2}n^2$, $\frac{\ln(3)}{2}n^2 \approx 0.54931n^2$ and $\left(\frac{3\sqrt{2}}{4} \arctan\left(\frac{\sqrt{2}}{2}\right)\right)n^2 \approx 0.65281n^2$.

Proof: [Of Theorem 13]

Recall from Lemma 10 that the probability of an improvement when flipping m bits with replacement is

$$P(\text{Imp}_m \mid \text{LO}(x) = i) = m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} + O\left(\frac{1}{n^2}\right).$$

The Simple Random mechanism will choose an operator at random in each iteration. Since $|H| = k = \Theta(1)$, the probability that the randomly chosen operator finds an improvement in one iteration is

$$\begin{aligned} \frac{1}{k} \cdot \sum_{m=1}^k P(\text{Imp}_m \mid \text{LO}(x) = i) &= \frac{1}{k} \cdot \sum_{m=1}^k \left(m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} + O\left(\frac{1}{n^2}\right) \right) \\ &= \frac{1}{kn} + \frac{1}{k} \cdot \sum_{m=2}^k \left(m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} + O\left(\frac{1}{n^2}\right) \right) \\ &= \Omega\left(\frac{1}{n}\right). \end{aligned}$$

Hence, using the same nomenclature as in Theorem 3, the expected time for an improvement is $A_i = \frac{1}{\frac{1}{k} \cdot \sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} - o(n)$, where the lower order $O\left(\frac{1}{n^2}\right)$ terms from $P(\text{Imp}_m \mid \text{LO}(x) = i)$ are collected in the lower order $-o(n)$ term. Thus, we can sum up the expectations as stated in Theorem 3 to prove the

theorem statement:

$$\begin{aligned} \frac{1}{2} \sum_{i=0}^{n-1} A_i &= \frac{1}{2} \cdot \sum_{i=0}^{n-1} \frac{1}{\frac{1}{k} \cdot \sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} - o(n^2) \\ &= \frac{k}{2} \cdot \sum_{i=0}^{n-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} - o(n^2). \end{aligned}$$

□

We now prove the same deteriorating performance for the other simple HHs.

Corollary 14 *The expected runtime of the Permutation, Greedy and Random Gradient mechanisms using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$ for LEADINGONES is*

$$\frac{k}{2} \cdot \sum_{i=0}^{n-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \pm o(n^2).$$

Proof: Given a bit-string x with $\text{LO}(x) = i$, let p_i be the probability that at least one fitness improvement is constructed within k fitness function evaluations. For the Greedy and Permutation mechanisms we have,

$$\begin{aligned} p_i &= P(\text{Imp}_1 + (1 - P(\text{Imp}_1)) \cdot P(\text{Imp}_2) + \dots + \left(\prod_{j=1}^{k-1} (1 - P(\text{Imp}_j)) \right) \cdot P(\text{Imp}_k)) \\ &= \sum_{m=1}^k \left(\left(\prod_{j=1}^{m-1} \left(1 - j \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{j-1} \right) \right) \cdot m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} \right) \\ &= \sum_{m=1}^k \left((1 - o(1)) \cdot m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} \right) = \sum_{m=1}^k ((1 - o(1)) \cdot P(\text{Imp}_m)), \end{aligned}$$

where we denoted $P(\text{Imp}_m \mid \text{LO}(x) = i)$ as $P(\text{Imp}_m)$ for brevity.

To move from the second equation to the third, we used that $P(\text{Imp}_m) = O\left(\frac{1}{n}\right)$ for all $1 \leq m \leq k = \Theta(1)$ and that $\prod_{j=1}^{m-1} (1 - P(\text{Imp}_j)) = (1 - O\left(\frac{1}{n}\right))^m = 1 - o(1)$. To derive an upper bound on the expected optimisation time we note that the difference between $\frac{k}{p_i}$ (i.e., the expected waiting time for an improvement to be constructed), and the A_i expected waiting times used to prove Theorem 13 is at most constant (for $k = \Theta(1)$). Thus, the difference between the expected runtimes of these mechanisms and Simple Random is limited to lower order terms. With probability at most $\frac{k(k-1)}{n^2}$ at least two mutations of the k considered ones are improvements (two parallel successes of the Greedy mechanism, or two mutations performed sequentially by the Permutation mechanism). Since this occurs at most a sublinear number of times in expectation, and the maximum expected waiting time for any improving step is $O(n)$, the lower bound differs from the upper bound by at most an $o(n^2)$ term. Thus, the expected runtime of these mechanisms is also $\frac{k}{2} \sum_{i=0}^{n-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \pm o(n^2)$.

For the Random Gradient mechanism, we note that the probability that an operator that is applied following a success is successful again is at most $\frac{k}{n}$ (an upper bound on the success probability of the $k\text{BITFLIP}$ operator by Lemma 10). Since there are at most n improvements to be made throughout the search space, the expected number of repetitions which produce an improvement is at most $k = \Theta(1)$. If the chosen operator is not successful, the Random Gradient mechanism behaves identically to the Simple Random mechanism. Its expected runtime is therefore at least the expected runtime of the Simple Random mechanism less an $o(n^2)$ term, and at most the expected runtime of the Simple Random mechanism plus n because there are at most n iterations where the mechanisms differ in operator selection. Thus, its expected runtime is also $\frac{k}{2} \sum_{i=0}^{n-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \pm o(n^2)$. □




Figure 3: The leading constants in the theoretical upper bounds on the expected number of fitness function evaluations required by the k -operator Generalised Random Gradient hyper-heuristic to find the LEADINGONES optimum (we have used a value of $w = 100,000$ in Theorem 15).

4.2 The Generalised Random Gradient Hyper-heuristic Has the Best Possible Performance Achievable

In this subsection, we present a rigorous theoretical analysis of the expected runtime of GRG using $k = \Theta(1)$ operators for LEADINGONES. The following general theorem provides an upper bound on the expected runtime of GRG using $k = \Theta(1)$ low-level stochastic mutation heuristics of different neighbourhood size for any value of τ smaller than $\frac{1}{k}n \ln n$. The theorem allows us to identify values of τ for which the expected runtime of GRG is the optimal expected runtime that may be achieved by using k operators. This result will be highlighted in Corollary 16 for values of $\tau = \omega(n)$. The main result of this section has been presented in Theorem 9, which shows that increasing the number of operators (i.e., $|H|$) that GRG has access to leads to faster expected runtimes and, in particular, expected runtimes that are strictly smaller than that of any unbiased (1+1) black box algorithm using any strict subset of the same set of operators.

Similarly to Theorem 8, the proof technique partitions the search space into w stages, each representing an equal range of fitness values. As before, the expected runtime of the algorithm does not depend on w . However, w does affect the upper bound we obtain. The greater the number of stages, the tighter the upper bound the theorem provides as long as $w = o\left(\frac{n}{\exp(k\tau/n)}\right)$.

Theorem 15 *The expected runtime of the Generalised Random Gradient hyper-heuristic using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$, for LEADINGONES, with $\tau \leq \left(\frac{1}{k} - \varepsilon\right) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{k}$, is at most*

$$\frac{n^2}{2} \cdot \sum_{j=1}^w \left(\frac{\left(k \cdot \frac{\tau}{n} + \left[\sum_{m=1}^k e^{m \frac{\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}} \cdot M_m(j, w) \right] \right)}{w \cdot \left(\left[\sum_{m=1}^k e^{m \frac{\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}} \right] - k \right)} \right) \pm o(n^2)$$

where for $m \geq 2$,

$$M_m(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \text{ or } \frac{1}{m(1-j/w)^{m-1}} > \frac{\tau}{n} \\ \frac{1}{m(1-j/w)^{m-1}} & \text{otherwise} \end{cases}$$

and $M_1(j, w) = \min\left\{\frac{\tau}{n}, 1\right\}$, with $w \in \omega\left(\frac{\tau}{n}\right) \cap o\left(\frac{n}{\exp(k\tau/n)}\right)$.

Theorem 15 presents upper bounds on the expected runtime of GRG, with access to an arbitrary number $k = \Theta(1)$ of low-level stochastic mutation heuristics of different neighbourhood sizes for LEADINGONES using any learning period $\tau \leq \left(\frac{1}{k} - \varepsilon\right) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{k}$. The bound on the expected runtime improves as τ increases up to the limit provided by the theorem. Figure 3 shows the relationship

k	$\tau = 5n$	$\tau = 50n$	$\tau = 100n$	$\tau = \frac{1}{10}n \ln(n)$
2	0.46493	0.42363	0.42329	0.42329
3	0.46802	0.40579	0.40525	0.40525
4	0.48102	0.39897	0.39830	0.39830
5	0.49630	0.39568	0.39492	0.39492
11	3.090×10^{23}	8785.8	0.38987	0.38987
18	1.886×10^{44}	5.363×10^{24}	1034.8	0.38899

Table 1: The upper bounds in the leading constants found from various parameter combinations of the number of operators k and the learning period τ in Theorem 15 with $w = 100,000$. The values reported in the fourth column are the best possible for each k . As the number of operators increases, larger values of τ are required for good performance ($\tau = \omega(n)$ is necessary to achieve optimal performance). Note that while Theorem 15 allows us to show optimal leading constants for $\tau = \omega(n)$, the upper bounds it provides for large values of k and small values of τ are far from tight due to the generality of the theorem.

between the duration of the learning period and theoretical upper bounds for different k -operator variants of GRG from Theorem 15. The upper bounds found by GRG with more operators are better than the ones with fewer operators, as implied by Corollary 12. In particular, the upper bounds for GRG with k -operators outperforms the best possible expected runtime for any algorithm that has fewer than k operators, as implied by Theorem 9. We depicted this result explicitly for $k = 3$ and $k = 5$ in Figure 2. When $\tau = \omega(n)$ and $\tau \leq (\frac{1}{k} - \varepsilon)$ for some constant $0 < \varepsilon < \frac{1}{k}$, GRG is able to find the best possible runtime achievable with the low-level heuristics, up to lower order terms, as will be presented in Corollary 16.

Doerr and Wagner (2018) calculated that the best expected runtime for any unbiased (1+1) black box algorithm using such mutation operators for LEADINGONES is (up to lower order terms) $\approx 0.388n^2$. Corollary 16 states that GRG can match this theoretical performance limit up to one decimal place with 4 low-level heuristics, up to two decimal places with 11 low-level heuristics and up to three decimal places with 18 low-level heuristics. In Table 1 we present some of the most interesting parameter combinations of k and τ .

Proof: [Of Theorem 15]

We follow the proof idea for GRG with two operators (Theorem 8). That is, we partition the optimisation process into w stages based on the value of the LEADINGONES fitness function: during stage j , the LEADINGONES value of the current solution is at least $\frac{(j-1)n}{w}$ and less than $\frac{jn}{w}$. After all w stages have been completed, the global optimum with a LEADINGONES value of n will have been found.

We derive an upper bound on the expectation of the runtime T by summing up the expected values of T_j , the expected time between the first random operator choice in each stage of the optimisation process and the moment when the necessary fitness value to enter the following stage is reached (that is, the time before the first random operator choice occurs in the next stage). As our analysis of $E(T_j)$ requires the stage to start with a random choice of the mutation operator, we bound $E(T) \leq \sum_{j=1}^w (E(T_j) + E(S_{j+1}))$ where S_{j+1} denotes the expected number of iterations between the first solution in stage $j + 1$ that is constructed and the first random operator choice occurring in that stage. We will later show that with proper parameter choices, the contribution of the S_{j+1} terms can be bounded by $o(n^2)$, therefore they do not affect the leading constant in the overall bound.

Let us now consider T_j , the number of iterations the HH spends in stage j . Recall that a mutation operator is selected uniformly at random and is allowed to run until it fails to produce an improvement within τ consecutive iterations. Let N_j be a random variable denoting the number of random operator choices that the HH performs during stage j , and let $X_{j,1}, \dots, X_{j,N_j}$ be the number of iterations where the HH runs each chosen operator for. We note that $T_j = \sum_{k=1}^{N_j} X_{j,k}$ is a sum of a number of non-negative variables, and we will later show that $E(N_j)$ is bounded. Thus, $E(T_j)$ can be bounded by applying Wald's

equation (Theorem 1):

$$E(T_j) = E\left(\sum_{k=1}^{N_j} X_{j,k}\right) \leq E(N_j)E(X_j), \quad (6)$$

where $E(X_j)$ denotes an upper bound on all $E(X_{j,k})$ in stage j .

To bound $E(N_j)$, the expected number of times the random operator selection is performed during stage j , we will provide a lower bound on the expected number of improvements found following the operator selection and apply the Additive Drift Theorem (Theorem 2) to find the expected number of random operator selections that occur before a sufficient number of improvements for entering the next stage have been found.

We will again denote $P(\text{Imp}_m \mid \text{LO}(x) = i)$ as $P(\text{Imp}_m)$ throughout this proof, for brevity. Recall from Lemma 10 that the improvement probability for a $m\text{BITFLIP}$ operator at $\text{LO}(x) = i$ (with $m = \Theta(1)$) is

$$P(\text{Imp}_m) = m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} + O\left(\frac{1}{n^2}\right).$$

Since either the $m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}$ terms will asymptotically dominate the lower order term or the $P(\text{Imp}_1) = \frac{1}{n}$ term will asymptotically dominate (due to the presence of the 1BITFLIP operator within the set H), the lower order $O\left(\frac{1}{n^2}\right)$ terms will be insignificant. In particular, we can combine the contributions of the $O\left(\frac{1}{n^2}\right)$ terms into an $o(n^2)$ term in the final bound on the expected runtime and omit them from the $P(\text{Imp}_m)$ terms in the calculations.

Let F_m denote the event that the $m\text{BITFLIP}$ operator fails to find an improvement during τ iterations. This will occur with probability $P(F_m) = (1 - P(\text{Imp}_m))^\tau$. Within stage j , the ancestor individual has at most $i = \frac{jn}{w} - 1$ and at least $i = \frac{(j-1)n}{w}$ leading 1-bits. Hence,

$$\begin{aligned} P(F_m) &\leq \left(1 - m \cdot \frac{1}{n} \left(\frac{n - \left(\frac{jn}{w} - 1\right) - 1}{n}\right)^{m-1}\right)^\tau \leq \left(1 - m \cdot \frac{1}{n} \left(1 - \frac{j}{w}\right)^{m-1}\right)^\tau \\ &\leq e^{-\frac{m\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}}, \end{aligned}$$

and

$$\begin{aligned} P(F_m) &\geq \left(1 - m \cdot \frac{1}{n} \left(\frac{n - \left(\frac{(j-1)n}{w}\right) - 1}{n}\right)^{m-1}\right)^\tau \geq \left(1 - m \cdot \frac{1}{n} \left(1 - \frac{j-1}{w}\right)^{m-1}\right)^\tau \\ &\geq e^{-\frac{m\tau}{n} \left(1 - \frac{j-1}{w}\right)^{m-1}} - \frac{1}{n}. \end{aligned}$$

A geometric distribution with parameter $p = P(F_m)$ can be used to model the number of improvements that the $m\text{BITFLIP}$ operator finds prior to failing to find an improvement for τ iterations. The expectation of this distribution is $\frac{1-p}{p} = \frac{1}{p} - 1$. We can model these events with a geometric random variable since GRG immediately restarts a period of τ iterations with the chosen heuristic when a successful mutation occurs (i.e., it sets $c_t = 0$ in Algorithm 2). Hence, there are no periods of τ consecutive unsuccessful mutations apart from the period when the operator finally fails. Combined over all k operators, the expected number of improvements D_j produced following a single random operator selection during stage j is greater than:

$$E(D_j) \geq \sum_{m=1}^k \frac{1}{k} \left(\frac{1}{P(F_m)} - 1\right) \geq \frac{1}{k} \left(\sum_{m=1}^k e^{\frac{m\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}}\right) - 1,$$

where the last inequality is obtained by inserting the upper bounds for $P(F_m)$. We use this expectation as the drift in the Additive Drift Theorem to provide an upper bound for $E(N_j)$. Recall that for each stage at most $\frac{n}{w}$ fitness values have to be crossed. Since bits beyond the leading 1-bits prefix and the first zero bit remain uniformly distributed, in expectation $\frac{n}{2w}$ improvements by mutation are required. If each step of

a random process contributes $E(D_j)$ expected improvements by mutation, the expected number of steps required to complete stage j is at most

$$E(N_j) \leq \frac{\frac{n}{2w}}{E(D_j)} \leq \frac{k}{2} \cdot \frac{n}{w} \cdot \frac{1}{\left(\sum_{m=1}^k e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}}\right) - k}, \quad (7)$$

by the Additive Drift Theorem.

We apply Wald's equation to bound $E(X_j)$, the expected number of iterations before a selected mutation operator fails to produce an improvement for τ iterations: let S_m be the number of improvements found by the operator before it fails, and $W_{m,1}, W_{m,2}, \dots, W_{m,S_m}$ be the number of iterations it took to find each of those improvements. Then, once selected, the m BITFLIP operator runs for

$$E(X_j | m\text{BITFLIP}) = \tau + \sum_{s=1}^{S_m} E(W_{m,s}) = E(S_m)E(W_{m,1} | S_m \geq 1),$$

where τ accounts for the iterations immediately before the failure and the sum counts the expected number of iterations preceding each constructed improvement. Recall that $E(S_m) = \frac{1}{P(F_m)} - 1$ by the properties of the geometric distribution and observe that $E(W_{m,1} | S_m \geq 1) = E(W_{m,1} | W_{m,1} \leq \tau) \leq \min\{\tau, E(W_{m,1})\}$. Using a waiting time argument (i.e., considering the expectation of a geometric random variable with success probability $p = P(\text{Imp}_m)$) for $E(W_{m,1})$, we get,

$$\begin{aligned} E(X_j | m\text{BITFLIP}) &\leq \tau + \left(\frac{1}{P(F_m)} - 1\right) \cdot \min\left\{\tau, \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n - (\frac{jn}{w} - 1) - 1}{n}\right)^{m-1}}\right\} \\ &= \tau + \left(\frac{1}{e^{-\frac{m\tau}{n}(1-\frac{j-1}{w})^{m-1}} - \frac{1}{n}} - 1\right) \cdot \min\left\{\tau, \frac{n}{m(1-\frac{j}{w})^{m-1}}\right\} \\ &\leq \tau + e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \cdot \min\left\{\tau, \frac{n}{m(1-\frac{j}{w})^{m-1}}\right\} + O(1), \end{aligned}$$

with $e^{\frac{m\tau}{n}(1-\frac{j-1}{w})^{m-1}} = (1 + o(1)) \cdot e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}}$ for $m = \Theta(1)$ and $w = \omega\left(\frac{\tau}{n}\right)$. Combining these conditional expectations yields,

$$\begin{aligned} E(X_j) &\leq \sum_{m=1}^k \frac{1}{k} \cdot E(X_j | m\text{BITFLIP}) \\ &\leq \tau + \sum_{m=1}^k \left[\frac{1}{k} \cdot e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \cdot \min\left\{\tau, \frac{n}{m(1-\frac{j}{w})^{m-1}}\right\} \right] + O(k) \\ &\leq n \left(\frac{\tau}{n} + \sum_{m=1}^k \left[\frac{1}{k} \cdot e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \cdot \min\left\{\frac{\tau}{n}, \frac{1}{m(1-\frac{j}{w})^{m-1}}\right\} \right] \right) + O(k) \\ &\leq \frac{n}{k} \left(\frac{k\tau}{n} + \sum_{m=1}^k \left[e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \cdot \min\left\{\frac{\tau}{n}, \frac{1}{m(1-\frac{j}{w})^{m-1}}\right\} \right] \right) + O(k). \end{aligned} \quad (8)$$

At $j = w$, we have that $\min\left\{\frac{\tau}{n}, \frac{1}{m(1-\frac{j}{w})^{m-1}}\right\}$ is undefined. However, we know that the chosen operator would only be applied for a maximum of τ steps. Hence, we define $M_1(j, w) = \min\left\{\frac{\tau}{n}, 1\right\}$ and for $m \geq 2$,

$$M_m(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \text{ or } \frac{1}{m(1-\frac{j}{w})^{m-1}} > \frac{\tau}{n} \\ \frac{1}{m(1-\frac{j}{w})^{m-1}} & \text{otherwise.} \end{cases}$$

Finally, we return to bounding the overall expected runtime of the HH. Recall that S_j denotes the number of iterations that the algorithm spends in stage j while using the random operator chosen during stage $j-1$ and since $E(S_j) \leq \max(E(X_j \mid m\text{BITFLIP})_{m=1,\dots,k} < kE(X_j) = O(n \cdot \exp(\frac{k\tau}{n}))$ and $w = o(\frac{n}{\exp(\frac{k\tau}{n})})$, it follows that $\sum_{j=1}^w E(S_{j+1}) = o(n^2)$.

Substituting the bounds provided in Eq. 7 and Eq. 8 into Eq. 6 yields the theorem statement:

$$\begin{aligned} E(T) &\leq \sum_{j=1}^w (E(T_j) + E(S_{j+1})) \leq \left(\sum_{j=1}^w E(N_j)E(X_j) \right) + o(n^2) \\ &\leq o(n^2) + \frac{n^2}{2} \times \sum_{j=1}^w \frac{1}{w} \cdot \frac{\frac{k\tau}{n} + \sum_{m=1}^k \left[e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \cdot M_m(j, w) \right]}{\left(\sum_{m=1}^k e^{\frac{m\tau}{n}(1-\frac{j}{w})^{m-1}} \right) - k}. \end{aligned}$$

□

We now show that for appropriate values of the parameter τ , the k -operator GRG runs in the best possible expected runtime achievable for a mechanism using k -operators, up to lower order terms. The following corollary provides an upper bound on the expected runtime for sufficiently large learning periods (i.e., $\tau = \omega(n)$). The bound matches the best expected runtime achievable proven in Theorem 12.

Corollary 16 *The expected runtime of the Generalised Random Gradient hyper-heuristic using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$ for LEADINGONES, with τ that satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{k} - \varepsilon) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{k}$, is at most*

$$\frac{1}{2} \left(\sum_{i=0}^{\frac{n}{k}-1} \frac{1}{k \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{i=\frac{n}{m+1}}^{\frac{n}{m}-1} \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \right) \pm o(n^2).$$

Proof: This proof follows a similar structure to the proof of Corollary 7, i.e., we aim to simplify the sum from Theorem 15.

Theorem 15 states that,

$$E(T) \leq \frac{n^2}{2} \cdot \sum_{j=1}^w \left(\frac{\left(k \cdot \frac{\tau}{n} + \left[\sum_{m=1}^k e^{m \frac{\tau}{n} (1-\frac{j}{w})^{m-1}} \cdot M_m(j, w) \right] \right)}{w \cdot \left(\left[\sum_{m=1}^k e^{m \frac{\tau}{n} (1-\frac{j}{w})^{m-1}} \right] - k \right)} \right) \pm o(n^2).$$

Since $\tau = \omega(n)$, we have that $\frac{\tau}{n} = \omega(1)$. Hence, we can simplify the $M_m(j, w)$ terms. Recall that we have $M_1(j, w) = \min\{\frac{\tau}{n}, 1\} = 1$, and for $m \geq 2$,

$$M_m(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \text{ or } \frac{1}{m(1-\frac{j}{w})^{m-1}} > \frac{\tau}{n} \\ \frac{1}{m(1-\frac{j}{w})^{m-1}} & \text{otherwise.} \end{cases}$$

Since $\frac{\tau}{n} = \omega(1)$, we have that $\frac{\tau}{n} > \frac{1}{m(1-\frac{j}{w})^{m-1}}$. Hence, we can simplify $M_m(j, w)$ as follows:

$$M_m(j, w) := \begin{cases} \frac{\tau}{n} & \text{if } j = w \\ \frac{1}{m(1-\frac{j}{w})^{m-1}} & \text{otherwise.} \end{cases}$$

In particular, for $j \neq w$, we have that $M_m(j, w) = \Theta(1)$, for $m \geq 1$.

We can further simplify the sum from Theorem 15 by noting that, since $\frac{\tau}{n} = \omega(1)$, the exponential term in each summand will be asymptotically dominant. In particular, the $\frac{k\tau}{n}$ term in the numerator and the $-k$ term in the denominator will be asymptotically dominated. Hence, we can relegate these terms into the $o(n^2)$ lower order term.

With these modifications we can simplify the sum as follows,

$$\begin{aligned} E(T) &\leq \frac{n^2}{2} \cdot \sum_{j=1}^w \left(\frac{\left(k \cdot \frac{\tau}{n} + \left[\sum_{m=1}^k e^{m \frac{\tau}{n} (1 - \frac{j}{w})^{m-1}} \cdot M_m(j, w) \right] \right)}{w \cdot \left(\left[\sum_{m=1}^k e^{m \frac{\tau}{n} (1 - \frac{j}{w})^{m-1}} \right] - k \right)} \right) \pm o(n^2) \\ &\leq \frac{n^2}{2} \cdot \sum_{j=1}^w \left(\frac{\left(\sum_{m=1}^k e^{m \frac{\tau}{n} (1 - \frac{j}{w})^{m-1}} \cdot M_m(j, w) \right)}{w \cdot \left(\left[\sum_{m=1}^k e^{m \frac{\tau}{n} (1 - \frac{j}{w})^{m-1}} \right] \right)} \right) \pm o(n^2). \end{aligned}$$

To further simplify the sum, we consider when each of the k exponential summand terms will be asymptotically dominant in the numerator. Since $\frac{\tau}{n} = \omega(1)$, whichever exponential term has the largest exponent will asymptotically dominate the other terms. In particular, the m_{th} term will dominate when $\frac{m\tau}{n} (1 - \frac{j}{w})^{m-1}$ is the largest amongst $1 \leq m \leq k$. Comparing the terms at m and $m + 1$, we see that the m_{th} term dominates the $(m + 1)_{th}$ term when $j \geq \frac{w}{m+1}$. Continuing these calculations, we get that the m_{th} term will dominate all the others (i.e., it will have the largest exponent) when $\frac{w}{m+1} \leq j \leq \frac{w}{m} - 1$ for $m < k$, and the k_{th} term will dominate when $j \leq \frac{w}{k} - 1$. If a term is asymptotically dominated, we can relegate it into the $o(n^2)$ lower order term. Hence, the only terms remaining in the numerator in the period when the m_{th} operator dominates will be the m_{th} exponential term multiplied by $M_m(j, w)$. The denominator will similarly only contain the m_{th} exponential term, multiplied by w . The sum will hence simplify to a sum of $M_m(j, w)$ terms:

$$\begin{aligned} E(T) &\leq \frac{n^2}{2} \cdot \left(\sum_{j=1}^{\frac{w}{k}} M_k(j, w) + \sum_{m=1}^{k-1} \sum_{j=\frac{w}{m+1}+1}^{\frac{w}{m}} M_m(j, w) \right) \pm o(n^2). \\ &= \frac{n^2}{2} \cdot \left(\sum_{j=1}^{\frac{w}{k}} \frac{1}{k \left(1 - \frac{j}{w}\right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{j=\frac{w}{m+1}+1}^{\frac{w}{m}} \frac{1}{m \left(1 - \frac{j}{w}\right)^{m-1}} \right) \pm o(n^2). \\ &= \frac{n}{2} \cdot \left(\sum_{j=1}^{\frac{w}{k}} \frac{1}{k \cdot \frac{1}{n} \cdot \left(1 - \frac{j}{w}\right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{j=\frac{w}{m+1}+1}^{\frac{w}{m}} \frac{1}{m \cdot \frac{1}{n} \cdot \left(1 - \frac{j}{w}\right)^{m-1}} \right) \pm o(n^2). \end{aligned}$$

Recall that stage j refers to when $\frac{(j-1)n}{w} \leq i \leq \frac{jn}{w} - 1$ where i is the LO value of the current solution. In particular, if we substitute $i \geq \frac{(j-1)n}{w}$, (or $j \leq \frac{iw}{n} + 1$) into each summand, we get,

$$\frac{1}{m \cdot \frac{1}{n} \cdot \left(1 - \frac{j}{w}\right)^{m-1}} \leq \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}}.$$

Furthermore, using the upper and lower bounds on i to adjust the bounds on the sums (which will cancel out the multiplicative n term) gives the final result:

$$\begin{aligned} E(T) &\leq \frac{n}{2} \cdot \left(\sum_{j=1}^{\frac{w}{k}} \frac{1}{k \cdot \frac{1}{n} \cdot \left(1 - \frac{j}{w}\right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{j=\frac{w}{m+1}+1}^{\frac{w}{m}} \frac{1}{m \cdot \frac{1}{n} \cdot \left(1 - \frac{j}{w}\right)^{m-1}} \right) \pm o(n^2) \\ &\leq \frac{1}{2} \cdot \left(\sum_{i=0}^{\frac{n}{k}-1} \frac{1}{k \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{k-1}} + \sum_{m=1}^{k-1} \sum_{i=\frac{n}{m+1}}^{\frac{n}{m}-1} \frac{1}{m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \right) \pm o(n^2). \end{aligned}$$

□

We can now prove the main result of this section, i.e., Theorem 9, which states that the expected runtime of GRG using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ and $k = \Theta(1)$ with τ that satisfies both $\tau = \omega(n)$ and $\tau \leq (\frac{1}{k} - \varepsilon) n \ln(n)$, for some constant $0 < \varepsilon < \frac{1}{k}$, for LEADINGONES is smaller than the best-possible expected runtime for any unbiased (1+1) black box algorithm using any strict subset of $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$.

Proof: [Of Theorem 9]

Corollary 16 shows that GRG with k mutation operators can match the best-possible performance of any unbiased (1+1) black box algorithm with k mutation operators for LEADINGONES from Theorem 12, up to lower order terms. The results from Corollary 16 and Theorem 12 imply that the only difference in the best case expected runtime for GRG with k operators (i.e., using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$) and the best-possible expected runtime for an algorithm with $k-1$ operators (i.e., using $\{1\text{BITFLIP}, \dots, (k-1)\text{BITFLIP}\}$) occurs when $0 \leq i \leq \frac{n}{k} - 1$. In this region, the best case expected runtime of the k -operator GRG matches the expected runtime of an algorithm using only the $k\text{BITFLIP}$ operator, up to lower order terms, while the expected runtime of the best-possible $k-1$ -operator algorithm matches the expected runtime of an algorithm using only the $(k-1)\text{BITFLIP}$ operator, up to lower order terms.

We know from Lemma 11 that the expected runtime to optimise this area when applying the $k\text{BITFLIP}$ operator will be smaller than the expected runtime when applying the $(k-1)\text{BITFLIP}$ operator, and thus using the $k\text{BITFLIP}$ operator will be faster. Since in the rest of the search space, the best case k -operator GRG and the best-possible $k-1$ -operator algorithm will have the same expected performance (up to lower order terms), the best case expected runtime of GRG with k operators will be smaller than the best-possible expected runtime of any algorithm with $k-1$ operators.

Reiterating the argument for the best-possible performance of an algorithm with $(1, \dots, m)\text{BITFLIP}$ operators ($m < k-1$) allows us to conclude that the best case expected runtime of the k -operator GRG is faster than the best-possible performance of any algorithm with $m < k$ operators. In particular, GRG with $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ is faster than the best-possible algorithm using any strict subset of $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$. Lemma 11 combined with a similar argument as used above with $k-1$ operators implies that GRG will outperform the best-possible algorithm equipped with a subset of the k operators in the area of the search space where the latter algorithm is missing any of the $\{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ operators (i.e., if the best-possible algorithm is missing the $m\text{BITFLIP}$ operator, GRG will outperform it in the region where the LEADINGONES fitness value satisfies $\frac{n}{m+1} \leq \text{LO}(x) \leq \frac{n}{m} - 1$), hence its expected runtime is smaller. \square

4.3 Anytime Performance

In this subsection, we present the expected fixed target running times of the algorithms considered thus far. Fixed target running times have been previously presented by Carvalho Pinto and Doerr (2017). In order to achieve the runtimes presented in Corollary 17 (and depicted in Figure 4) we have adapted Theorem 13 and Theorem 15 to sum up to a fixed target $\text{LO}(x) = X \leq n$, while keeping the problem size as n (respectively, w) within the summands.

Corollary 17 *Let $X \leq n$. The expected time needed to reach for the first time a search point x of LEADINGONES-value at least X is at most:*

- $\frac{k}{2} \cdot \sum_{i=0}^{X-1} \frac{1}{\sum_{m=1}^k m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1}} \pm o(n^2)$ for the Simple Random mechanism with k operators;
- $\frac{n^2}{2} \cdot \left(\sum_{j=1}^{\lceil X \cdot \frac{w}{n} \rceil} \frac{\left(k \cdot \frac{\tau}{n} + \left[\sum_{m=1}^k e^{m \frac{\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}} \cdot M_m(j, w) \right] \right)}{w \cdot \left(\left[\sum_{m=1}^k e^{m \frac{\tau}{n} \left(1 - \frac{j}{w}\right)^{m-1}} \right]_{-k} \right)} \right) + o(n^2)$ for GRG using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$, with $\tau \leq \left(\frac{1}{k} - \varepsilon\right) n \ln(n)$ for some constant $0 < \varepsilon < \frac{1}{k}$, where $M_m(j, w)$ are as defined in Theorem 15.

Figure 4 shows the expected fixed target running times using the set $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ for the Simple Random mechanism with $k = 1, 2, 3$ and GRG with $\tau = 10n$ and $k = 2, 3, 4$ for LEADINGONES. We have used $w = 1,000$ in Theorem 15 to achieve the GRG plots.

We see that all the GRG variants outperform the respective Simple Random variants for any fixed target independent of how many operators are used by GRG. All GRG variants are close to matching the anytime performance of the optimal GRG and the greater the number of operators the HH has access to, the better the anytime performance. For example, GRG with 4 operators outperforms all other algorithms for every fixed target $0 < X \leq n$.




Figure 4: Expected fixed target running times for LEADINGONES, with dimension $n = 10,000$. The values shown in the legend for Simple Random and GRG variants with $\tau = 10n$ are the normalised (by n^2) expected optimisation times.

Interestingly, the Simple Random mechanisms with more operators outperform those with less operators for smaller targets. For example, for a fixed target of $\text{LO}(x) = \frac{n}{2} = 5,000$, the expected first hitting time of Simple Random with 2 operators is $\approx 18.9\%$ smaller than that of RLS (i.e., Simple Random with 1 operator, Simple₁), while Simple Random with 3 operators is $\approx 26.0\%$ faster than RLS.

4.4 Extension of the Results to Standard Randomised Local Search Heuristics

In this subsection, we will extend the results presented so far in Section 4 such that they also hold for hyper-heuristics that select from the set $H = \{\text{RLS}_1, \dots, \text{RLS}_k\}$. In particular, we will show that the main result of this section (i.e., Theorem 9) also holds for RLS algorithms with different neighbourhood sizes (that flip distinct bits without replacement). For this purpose, we will show that the difference in the improvement probabilities of the RLS_m algorithms and the $m\text{BITFLIP}$ operators (from Lemma 10) are limited to lower order $O(\frac{1}{n^2})$ terms. A simple argument will then prove that the remaining results of this section also hold for RLS_m operators.

Lemma 18 *Given a bit-string x with $\text{LO}(x) = i \leq n - m$, the probability of improvement of an RLS_m algorithm that flips $m = \Theta(1)$ bits without replacement is*

$$P(\text{Imp}_{\text{RLS}_m} \mid \text{LO}(x) = i) = m \cdot \frac{1}{n} \cdot \left(\frac{n-i-1}{n}\right)^{m-1} \pm O\left(\frac{1}{n^2}\right).$$

Proof: To prove the statement we will show that the differences between the $P(\text{Imp}_m \mid \text{LO}(x) = i)$ from Lemma 10 and $P(\text{Imp}_{\text{RLS}_m} \mid \text{LO}(x) = i)$ for each $m = \Theta(1)$ are limited to lower order $O(\frac{1}{n^2})$ terms.

Doerr and Wagner (2018) showed that the improvement probability of RLS_m at the state $\text{LO}(x) = i$ is

$$P(\text{Imp}_{\text{RLS}_m} \mid \text{LO}(x) = i) = \frac{\binom{n-i-1}{m-1}}{\binom{n}{m}}.$$

Note that the improvement probability of RLS_m for $\text{LO}(x) = i > n - m$ is 0.

A simple calculation shows that the largest difference between the improvement probability of RLS_m and the improvement probability of the $m\text{BITFLIP}$ mutation operator (from Lemma 10) occurs when $\text{LO}(x) =$

$i = 0$. We now calculate this difference (dropping the state dependence for brevity):

$$\begin{aligned}
P(\text{Imp}_{\text{RLS}_m}) - P(\text{Imp}_m) &\leq \frac{\binom{n-1}{m-1}}{\binom{n}{m}} - \left(m \cdot \frac{1}{n} \cdot \left(\frac{n-1}{n} \right)^{m-1} + O\left(\frac{1}{n^2}\right) \right) \\
&= \frac{m}{n} - \left(\frac{m}{n} \cdot \left(\frac{n-1}{n} \right)^{m-1} + O\left(\frac{1}{n^2}\right) \right) \\
&= \frac{m}{n} \left(1 - \left(1 - \frac{1}{n} \right)^{m-1} \right) - O\left(\frac{1}{n^2}\right) \\
&= \frac{m}{n} \left(1 - \left(1 - \frac{m-1}{n} \pm o\left(\frac{1}{n}\right) \right) \right) - O\left(\frac{1}{n^2}\right) \\
&= \frac{m(m-1)}{n^2} \pm O\left(\frac{1}{n^2}\right) = O\left(\frac{1}{n^2}\right),
\end{aligned}$$

since $m = \Theta(1)$. Hence, the difference in the improvement probabilities of the two mutation operators is limited to lower order $O\left(\frac{1}{n^2}\right)$ terms and the lemma statement holds. \square

Through a simple application of Lemma 18 (rather than applying Lemma 10), the results for the simple mechanisms presented in Subsection 4.1 (i.e., Theorem 13 and Corollary 14), and the general results for GRG (i.e., Theorem 15 and Corollary 16), also hold for the hyper-heuristics using the heuristic set $H = \{\text{RLS}_1, \dots, \text{RLS}_k\}$, up to lower order $\pm o(n^2)$ terms.

We can also extend Theorem 12 and Theorem 9 to algorithms using the heuristic set $\{\text{RLS}_1, \dots, \text{RLS}_k\}$. We know from Theorem 12 that the $m\text{BITFLIP}$ operator is optimal (i.e., it has the highest probability of improvement amongst all BITFLIP operators) during the time when $\frac{n}{m+1} \leq i \leq \frac{n}{m} - 1$. Doerr (2018) showed that RLS_m is optimal (i.e., has the highest probability of improvement among all RLS operators) when $\frac{n-m}{m+1} \leq i \leq \frac{n}{m} - \frac{m-1}{m}$. The differences in the expected runtime of the best-possible algorithms using the two heuristic sets would therefore amount to the expected time spent improving from (at most) $m+1 = \Theta(1)$ different LEADINGONES fitness values. Since the expected time to improve from each of these fitness values is at most $\Theta(n)$, the difference between the expected runtime of the best-possible algorithm using $\{\text{RLS}_1, \dots, \text{RLS}_k\}$ operators and that using $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ operators, for LEADINGONES is therefore limited to lower order $o(n^2)$ terms. Hence, Theorem 12 also holds for the heuristic set $\{\text{RLS}_1, \dots, \text{RLS}_k\}$, up to lower order $\pm o(n^2)$ terms. With similar arguments, Theorem 9 also holds for sets of RLS operators.

5 Complementary Experimental Analysis

In the previous sections we proved that GRG performs efficiently for the LEADINGONES benchmark function for large enough problem sizes n . In this section we present some experimental results to shed light on its performance for different problem sizes up to $n = 10^8$. All parameter combinations have been simulated 10,000 times.

In order to efficiently handle larger problem dimensions experimentally, we do not simulate each individual mutation performed by the HH, but rather sample the waiting times for a fitness-improving mutation to occur using a geometric distribution (with the success probability p depending on the current operator and LEADINGONES value of the current solution). Specifically, suppose T_{im} is a random variable denoting the number of mutations required to get one fitness-improving mutation. Since T_{im} counts the number of independent trials each with probability p of success up to and including the first successful trial, $P(T_{\text{im}} \leq k) = 1 - (1-p)^k$ by the properties of the geometric distribution. Given access to a uniform $(0, 1)$ -distributed random variable U , T_{im} can be sampled by computing $\left\lceil \frac{\log(1-U)}{\log(1-p)} \right\rceil$.

5.1 Two Low-level Heuristics ($k = 2$)

We first consider GRG using two operators only (i.e., $H = \{1\text{BITFLIP}, 2\text{BITFLIP}\}$) and look at the impact on the average runtime of the parameter τ and of the problem size n .




Figure 5: Average number of fitness function evaluations required by the hyper-heuristics with $k = 2$ operators to find the LEADINGONES optimum in relation to the duration of the learning period for $n = 10,000$ (solid), $n = 50,000$ (dashed).




Figure 6: Average number of fitness function evaluations required for the Generalised Random Gradient hyper-heuristic with $k = 2$ operators to find the LEADINGONES optimum as the problem size n increases.

Figure 5 shows how the average runtimes of GRG for LEADINGONES vary with the duration of the learning period τ for problem sizes $n = 10,000$ and $n = 50,000$. The performance of GRG clearly depends on the choice of τ . It is worth noting that as the problem size increases, for $\tau \approx 0.55n \ln(n)$, the runtime seems to be approaching the optimal performance proven in Corollary 7 (i.e., $\frac{1+\ln(2)}{4}n^2 \approx 0.42329n^2$). For well chosen τ values, the HH beats the expected runtime of RLS and also the experimental runtime for the recently presented reinforcement learning HH that chooses between different neighbourhood sizes for RLS_k (Doerr et al., 2016). They reported an average runtime of $0.450n^2$ for the parameter choices they used. As τ increases past $0.6n^2$, we see a detriment in the performance of the HH. It is worth noting, however, that for $n = 50,000$, it is required that $\tau > 1.5n \ln(n) = 811,483$ to be worse than the RLS expected runtime of $0.5n^2$, indicating that the parameter is robust.

Figure 6 shows the effects of increasing the problem size n for a variety of fixed values of τ . We can see that an increased problem size leads to faster average runtimes. In particular, the HH requires a problem size of at least 200 before it outperforms RLS. The performance difference between the τ values decreases with increased n , indicating that further increasing n would lead to similar, optimal performance for a large range of values, as implied by Corollary 7. For $n = 10^8$, the runtime for $\tau = 0.6n \ln n$ is $\approx 0.42716n^2$, only slightly deviated from the optimal value of $\approx 0.42329n^2$.




Figure 7: Experimental fixed target running times for LEADINGONES, with dimension $n = 10,000$. The values shown in the legend for Simple Random and GRG variants with $\tau = 10n$ are the normalised (by n^2) average optimisation times.

5.2 More Than Two Low-level Heuristics ($k \geq 2$)

We now consider GRG using k operators, $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$, and look at the impact of the parameter τ .

Figure 7 shows the average fixed target running times using the set $H = \{1\text{BITFLIP}, \dots, k\text{BITFLIP}\}$ for the Simple Random mechanism with $k = 1, 2, 3$ operators, and GRG with $\tau = 10n$ and $k = 2, 3, 4$ for LEADINGONES. We have used a problem size of $n = 10,000$. The conclusions drawn from the experiments match the theoretical results shown in Figure 4. We see that although the final average runtime of the Simple Random mechanisms with more operators increases, they are faster initially. In particular, while $\text{LO}(x) \leq \frac{n}{2} = 5,000$, Simple Random with 3 operators is preferable to the variant with 2 or 1 (i.e., RLS). This Simple Random variant even outperforms GRG (with our arbitrarily chosen learning period $\tau = 10n$) for fixed targets smaller than $\approx 7,000$. For GRG we see, just like in Figure 4, that GRG with more operators is more effective than GRG with fewer operators for any fixed target. However, the difference is smaller than the theoretical results suggest, implying that while more operators are preferable, GRG with fewer operators is still effective for these problem sizes i.e., as the problem size increases so does the difference in performance in favour of larger sets H .

Figure 8 shows the average runtimes for GRG when it chooses between 2 to 5 mutation operators with different neighbourhood sizes for a LEADINGONES problem size $n = 100,000$. We see that incorporating more operators can be beneficial to the performance of GRG. Whilst the 2-operator HH achieves a best average performance of $\approx 0.43742n^2$, the 5-operator HH achieves a best average performance of $\approx 0.42906n^2$. It is, however, important to set τ appropriately to achieve the best performance. For $0 \leq 0.35n \ln(n) \leq \tau$ and $\tau \geq 1.5n \ln(n)$, the 2-operator HH outperforms the others. The results imply that the more operators are incorporated in the HH, the shorter the range of τ values for which it performs the best in comparison with the HH with fewer operators. We know from Corollary 16 that for a sufficiently large problem size, GRG with 5-operators will tend towards the theoretical optimal expected runtime of $\approx 0.39492n^2$. Furthermore, we know from Theorem 9 that GRG equipped with more operators will be faster.

Recently Doerr and Wagner (2018) analysed a (1+1) EA for LEADINGONES, where the mutation rate is updated on-the-fly. The number of bits to be flipped is sampled from a binomial distribution $B(n, p)$ and, differently from standard bit mutation, they do not allow a 0BITFLIP to occur by resampling until a non-zero bitflip occurs². Then a multiplicative comparison-based update rule, similar to the 1/5th rule from

²Note that while not allowing 0-bit flips is a reasonable choice for an extremely simplified EA such as the (1+1) EA, it is doubtful that it is a good idea in general. For instance, we conjecture that the expected runtime of the $(\mu + 1)$ EA (Witt, 2006) and $(\mu + 1)$ GA (Corus and Oliveto, 2018, 2019) would deteriorate for ONEMAX and LEADINGONES, while non-elitist EAs and GAs would require




Figure 8: Average number of fitness function evaluations required by the Generalised Random Gradient hyper-heuristic with k operators to find the LEADINGONES optimum, $n = 100,000$.

combinatorial optimisation (Kern et al., 2004) is applied to update the parameter p (i.e., the mutation rate). A successful mutation increases the mutation rate by a multiplicative factor $A > 1$, while an unsuccessful mutation multiplies the mutation rate by a multiplicative factor $b < 1$.

An experimental analysis is performed to tentatively identify the leading constants in the expected runtime for the best combinations of A and b . The best leading constant that has been identified is 0.4063 (using the best parameter configuration for which an average of at least 1000 runs has been reported; i.e., $A = 1.2$ and $b = 0.85$). While it is unclear whether their identified value indeed scales with the problem size (they only test problem sizes up to $n = 1500$), such a leading constant is worse than the $\approx 0.388n^2$ achieved by GRG. However, we note that the best possible expected runtime that such an EA that learns to automatically adapt the neighbourhood size to optimality can achieve is $\approx 0.404n^2$ (Doerr et al., 2019), due to the random sampling of the number of bits to flip, i.e., the best adaptive RLS is faster than the best adaptive EA.

6 Conclusion

Our foundational understanding of the performance of hyper-heuristics (HHs) is improving. Algorithm selection from algorithm portfolio systems and selection HHs generally use sophisticated machine learning algorithms to identify online (during the run) which low-level mechanisms have better performance during different stages of the optimisation process. Recently it has been proven that a reinforcement learning mechanism, which assigns different scores to each low-level heuristic according to how well it performs, allows a simple HH to run in the best expected runtime achievable with RLS_k low-level heuristics for the ONEMAX benchmark problem, up to lower order terms (Doerr et al., 2016).

In this paper we considered whether sophisticated learning mechanisms are always necessary for a HH to effectively learn to apply the right low-level heuristic.

We considered four of the most simple learning mechanisms commonly applied in the literature to combinatorial optimisation problems, namely, Simple Random, Permutation, Greedy and Random Gradient and showed that they all have the same performance for the LEADINGONES benchmark function when equipped with RLS_k low-level heuristics. While the former three mechanisms do not attempt to learn from the past performance of the low-level heuristics, the idea behind Random Gradient is to continue applying a heuristic so long as it is successful. We argued that looking at the performance of a heuristic after one single application is insufficient to appreciate whether it is a good choice or not. To this end, we gener-

exponential time to optimise any function with a unique optimum (Corus et al., 2018; Oliveto and Witt, 2014, 2015). An alternative implementation that allows 0-bit flips to occur, but simply avoids evaluating offspring that are identical to their parents, would solve the issue while at the same time producing the results reported by Doerr and Wagner (2018). Such an idea was suggested by Carvalho Pinto and Doerr (2017).

alised the existing Random Gradient learning mechanism to allow success to be measured over a longer period of time and called such time the *learning period*. We showed that the Generalised Random Gradient (GRG) HH can learn to adapt the neighbourhood size $k = \Theta(1)$ of RLS_k optimally during the run for the LEADINGONES benchmark function. As a byproduct, we proved that, up to lower order terms, GRG has the best possible runtime achievable by any algorithm that uses the same low level heuristics. In particular, it is faster than well-known unary unbiased evolutionary and local-search algorithms, including the (1+1) Evolutionary Algorithm ((1+1) EA) and Randomised Local Search (RLS). Furthermore, we also showed that for targets smaller than n (i.e., anytime performance), the advantages of GRG over RLS and EAs using standard bit mutation are even greater (i.e., they are even better as approximation algorithms for the problem).

To apply the generalised HH, a value for the learning period τ is required. Although our results indicate that τ is a fairly robust parameter (i.e., for $n = 10,000$, GRG achieved faster experimental runtimes than that of the (1+1) EA for all tested values of τ between 1 and 10^6 , and faster than RLS for all tested values of τ between 28,000 and 120,000), setting it appropriately will lead to optimal performance. Clearly τ must be large enough to have at least a constant expected number of successes within τ steps, if the HH has to learn about the operator performance. Naturally, setting too large values of τ may lead to large runtimes since switching operators requires $\Omega(\tau)$ steps.

We have also rigorously shown that the performance of the simple mechanisms deteriorates as the choice of operators with different neighbourhood sizes increases, while the performance of GRG improves with a larger choice, as desired for practical applications. In particular, GRG is able to outperform in expectation any unbiased (1+1) black box algorithm with access to strictly smaller sets of operators.

Recently, Doerr et al. (2018) have equipped the GRG HH with an adaptive update rule to automatically adapt the parameter τ throughout the run (i.e., the learning period can change its duration during the optimisation process). They proved that the HH is able to achieve the same optimal performance, up to lower order terms, for the LEADINGONES benchmark function. This so called Adaptive Random Gradient HH, equipped with two operators, experimentally outperforms the best setting of GRG, confirming that τ should not be fixed throughout the optimisation process.

Several further directions can be explored in future work. Firstly, the performance of GRG on a broader class of problems, including classical ones from combinatorial optimisation, should be rigorously studied. Secondly, more sophisticated HHs that have shown superior performance in practical applications should be analysed, such as the machine learning approaches which keep track of the historical performance of each low-level heuristic and use this information to decide which one to be applied next. In particular, it should be highlighted when more sophisticated learning mechanisms are required and the reasons behind the requirement. Thirdly, the understanding of HHs that switch between elitist and non-elitist low-level heuristics for multimodal functions should be improved (Lissovoi et al., 2019). Finally, considering more sophisticated low-level heuristics (e.g. with different population sizes) will bring a greater understanding of the general performance of selection HHs and their wider application in real-world optimisation. Another area where a foundational theoretical understanding is lacking is that of parameter tuning (Hall et al., 2019).

Acknowledgements This work was supported by EPSRC under grant EP/M004252/1.

References

- Afshani, P., Agrawal, M., Doerr, B., Doerr, C., Larsen, K. G., and Mehlhorn, K. (2013). The query complexity of finding a hidden permutation. In Brodnik, A., López-Ortiz, A., Raman, V., and Viola, A., editors, *Space-Efficient Data Structures, Streams, and Algorithms: Papers in Honor of J. Ian Munro on the Occasion of His 66th Birthday*, pages 1–11. Springer.
- Alanazi, F. and Lehre, P. K. (2014). Runtime analysis of selection hyper-heuristics with classical learning mechanisms. In *IEEE Congress on Evolutionary Computation, CEC '14*, pages 2515–2523. IEEE.
- Alanazi, F. and Lehre, P. K. (2016). Limits to learning in reinforcement learning hyper-heuristics. In *Evolutionary Computation in Combinatorial Optimization, EvoCOP '16*, pages 170–185. Springer.
- Asta, S. and Özcan, E. (2014). An apprenticeship learning hyper-heuristic for vehicle routing in HyFlex. In *IEEE Symposium on Evolving and Autonomous Learning Systems, EALS '14*, pages 65–72. IEEE.

- Berberoglu, A. and Uyar, A. Ş. (2011). Experimental comparison of selection hyper-heuristics for the short-term electrical power generation scheduling problem. In *Applications of Evolutionary Computation, EvoApplications '11*, pages 444–453. Springer.
- Bezerra, L. C. T., López-Ibáñez, M., and Stützle, T. (2018). A large-scale experimental evaluation of high-performing multi- and many-objective evolutionary algorithms. *Evolutionary Computation*, 26(4):621–656.
- Böttcher, S., Doerr, B., and Neumann, F. (2010). Optimal fixed and adaptive mutation rates for the leadingones problem. In *Parallel Problem Solving from Nature, PPSN '10*, pages 1–10. Springer.
- Burke, E., Kendall, G., and Soubeiga, E. (2003). A tabu-search hyperheuristic for timetabling and rostering. *Journal of Heuristics*, 9(6):451–470.
- Burke, E. K., Gendreau, M., Hyde, M., Kendall, G., Ochoa, G., Özcan, E., and Qu, R. (2013). Hyper-heuristics: A survey of the state of the art. *Journal of the Operational Research Society*, 64(12):1695–1724.
- Buzdalov, M. and Buzdalova, A. (2015). Can onemax help optimizing leadingones using the EA+RL method? In *IEEE Congress on Evolutionary Computation, CEC '15*, pages 1762–1768. IEEE.
- Carvalho Pinto, E. and Doerr, C. (2017). Towards a more practice-aware runtime analysis. In *Conference on Artificial Evolution, EA '17*, pages 298–305. Springer.
- Corus, D., Dang, D., Eremeev, A. V., and Lehre, P. K. (2018). Level-based analysis of genetic algorithms and other search processes. *IEEE Transactions on Evolutionary Computation*, 22(5):707–719. To appear at GECCO '19.
- Corus, D. and Oliveto, P. S. (2018). Standard steady state genetic algorithms can hillclimb faster than mutation-only evolutionary algorithms. *IEEE Transactions on Evolutionary Computation*, 22(5):720–732.
- Corus, D. and Oliveto, P. S. (2019). On the benefits of populations on the exploitation speed of standard steady-state genetic algorithms. *CoRR*, abs/1903.10976.
- Cowling, P., Kendall, G., and Soubeiga, E. (2001). A hyperheuristic approach to scheduling a sales summit. In *Practice and Theory of Automated Timetabling, PATAT '01*, pages 176–190. Springer.
- Cowling, P., Kendall, G., and Soubeiga, E. (2002). Hyperheuristics: A tool for rapid prototyping in scheduling and optimisation. In *Applications of Evolutionary Computing, EvoWorkshops '02*, pages 1–10. Springer.
- Doerr, B. (2018). Better runtime guarantees via stochastic domination. *CoRR*, abs/1801.04487.
- Doerr, B. and Doerr, C. (2018). Theory of parameter control for discrete black-box optimization: Provable performance gains through dynamic parameter choices. *CoRR*, abs/1804.05650.
- Doerr, B., Doerr, C., and Lengler, J. (2019). Self-adjusting mutation rates with provably optimal success rules. *CoRR*, abs/1902.02588.
- Doerr, B., Doerr, C., and Yang, J. (2016). k-bit mutation with self-adjusting k outperforms standard bit mutation. In *Parallel Problem Solving from Nature, PPSN '16*, pages 824–834. Springer.
- Doerr, B. and Krejca, M. S. (2018). Significance-based estimation-of-distribution algorithms. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '18*, pages 1483–1490. ACM.
- Doerr, B. and Künnemann, M. (2013). How the $(1+\lambda)$ evolutionary algorithm optimizes linear functions. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '13*, pages 1589–1596. ACM.
- Doerr, B., Lissovoi, A., Oliveto, P. S., and Warwicker, J. A. (2018). On the runtime analysis of selection hyper-heuristics with adaptive learning periods. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '18*, pages 1015–1022. ACM.
- Doerr, C. and Lengler, J. (2018). The $(1+1)$ elitist black-box complexity of leadingones. *Algorithmica*, 80(5):1579–1603.
- Doerr, C. and Wagner, M. (2018). Simple on-the-fly parameter selection mechanisms for two classical discrete black-box optimization benchmark problems. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '18*, pages 943–950. ACM.

- Friedrich, T., Kötzing, T., and Krejca, M. S. (2016). EDAs cannot be balanced and stable. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '16*, pages 1139–1146. ACM.
- Gibbs, J., Kendall, G., and Özcan, E. (2010). Scheduling English football fixtures over the holiday period using hyper-heuristics. In *Parallel Problem Solving from Nature, PPSN '10*, pages 496–505. Springer.
- Hall, G. T., Oliveto, P. S., and Sudholt, D. (2019). On the impact of the cutoff time on the performance of algorithm configurators. *CoRR*, abs/1904.06230. To appear at GECCO '19.
- He, J., He, F., and Dong, H. (2012). Pure strategy or mixed strategy? In *Evolutionary Computation in Combinatorial Optimization, EvoCOP '12*, pages 218–229. Springer.
- He, J. and Yao, X. (2001). Drift analysis and average time complexity of evolutionary algorithms. *Artificial Intelligence*, 127(1):57–85.
- Jansen, T. and Zarges, C. (2011). Analysis of evolutionary algorithms: From computational complexity analysis to algorithm engineering. In *Proceedings of the Workshop on Foundations of Genetic Algorithms, FOGA '11*, pages 1–14. ACM.
- Kern, S., Müller, S. D., Hansen, N., Büche, D., Ocenasek, J., and Koumoutsakos, P. (2004). Learning probability distributions in continuous evolutionary algorithms – a comparative review. *Natural Computing*, 3(1):77–112.
- Kotthoff, L. (2016). *Algorithm Selection for Combinatorial Search Problems: A Survey*, pages 149–190. Springer.
- Lehre, P. K. and Özcan, E. (2013). A runtime analysis of simple hyper-heuristics: To mix or not to mix operators. In *Proceedings of the Workshop on Foundations of Genetic Algorithms, FOGA '13*, pages 97–104. ACM.
- Lehre, P. K. and Witt, C. (2012). Black-box search by unbiased variation. *Algorithmica*, 64(4):623–642.
- Lissovoi, A., Oliveto, P. S., and Warwicker, J. A. (2017). On the runtime analysis of generalised selection hyper-heuristics for pseudo-boolean optimisation. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '17*, pages 849–856. ACM.
- Lissovoi, A., Oliveto, P. S., and Warwicker, J. A. (2019). On the time complexity of algorithm selection hyper-heuristics for multimodal optimisation. In *AAAI Conference on Artificial Intelligence, AAAI '19*. To Appear.
- López-Camacho, E., Terashima-Marin, H., Ross, P., and Ochoa, G. (2014). A unified hyper-heuristic framework for solving bin packing problems. *Expert Systems with Applications*, 41(15):6876–6889.
- Maden, İ., Uyar, S., and Özcan, E. (2009). Landscape analysis of simple perturbative hyperheuristics. In *Conference on Soft Computing, Mendel '09*, pages 16–22. Mendel.
- Moraglio, A. and Sudholt, D. (2017). Principled design and runtime analysis of abstract convex evolutionary search. *Evolutionary Computation*, 25(2):205–236.
- Nareyek, A. (2004). Choosing search heuristics by non-stationary reinforcement learning. In *Metaheuristics: Computer Decision-Making*, pages 523–544. Springer.
- Ochoa, G., Qu, R., and Burke, E. K. (2009a). Analyzing the landscape of a graph based hyper-heuristic for timetabling problems. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO '09*, pages 341–348. ACM.
- Ochoa, G., Vázquez-Rodríguez, J. A., Petrovic, S., and Burke, E. (2009b). Dispatching rules for production scheduling: a hyper-heuristic landscape analysis. In *IEEE Congress on Evolutionary Computation, CEC '09*, pages 1873–1880. IEEE.
- Oliveto, P. S. and Witt, C. (2014). On the runtime analysis of the simple genetic algorithm. *Theoretical Computer Science*, 545:2 – 19.
- Oliveto, P. S. and Witt, C. (2015). Improved time complexity analysis of the simple genetic algorithm. *Theoretical Computer Science*, 605:21 – 41.
- Oliveto, P. S. and Yao, X. (2011). Runtime analysis of evolutionary algorithms for discrete optimization. In Auger, A. and Doerr, B., editors, *Theory of Randomized Search Heuristics: Foundations and Recent Developments*, pages 21–52. World Scientific.

- Özcan, E., Misir, M., Ochoa, G., and Burke, E. K. (2010). A reinforcement learning - great-deluge hyper-heuristic for examination timetabling. *Applied Metaheuristic Computing*, 1(1):39–59.
- Wald, A. (1944). On cumulative sums of random variables. *The Annals of Mathematical Statistics*, 15(3):283–296.
- Witt, C. (2006). Runtime analysis of the $(\mu + 1)$ EA on simple pseudo-boolean functions. *Evolutionary Computation*, 14(1):65–86.
- Wolpert, D. H. and Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE Transactions on Evolutionary Computation*, 1(1):67–82.
- Xu, L., Hutter, F., Hoos, H. H., and Leyton-Brown, K. (2008). Satzilla: Portfolio-based algorithm selection for sat. *Journal of Artificial Intelligence Research*, 32(1):565–606.