Runtime Analysis of a Simple Multi-Objective Evolutionary Algorithm

Oliver Giel*

Fachbereich Informatik, Lehrstuhl 2, Universität Dortmund,
Dortmund, Germany
oliver.giel@cs.uni-dortmund.de

Abstract. Practical knowledge on the design and application of multi-objective evolutionary algorithms (MOEAs) is available but well-founded theoretical analyses of the runtime are rare. Laumanns, Thiele, Zitzler, Welzel and Deb (2002) have started such an analysis for two simple mutation-based algorithms (SEMO and FEMO) for combinatorial optimization problems. These algorithms search locally in the neighborhood of their current population by selecting an individual and flipping one randomly chosen bit. Due to their local search operator they cannot escape from local optima, and, therefore, they have no finite expected runtime in general.

We investigate the runtime of a variant of SEMO whose mutation operator flips each bit independently. It is proven that its expected runtime is $O(n^n)$ for all objective functions $f : \{0,1\}^n \to \mathbb{R}^m$, i. e., independently of the number of objectives m. There are bicriteria problems among the hardest problems for this algorithm. Moreover, for each d between 2 and n, a bicriteria problem with expected runtime $\Theta(n^d)$ is presented. This shows that bicriteria problems cover the full range of potential runtimes of this variant of SEMO. For the problem LOTZ (leading ones trailing zeroes), the runtime does not increase substantially if we use the global search operator. Finally, we consider the problem MOCO (multi-objective counting ones). We show that the conjectured bound $O(n^2 \log n)$ on the expected runtime is wrong for both variants of SEMO. In fact, MOCO is almost a worst case example for SEMO if we consider the expected runtime; however, the runtime is $O(n^2 \log n)$ with high probability.

Keywords. Runtime analysis, multi-objective evolutionary algorithms

Knowledge on the design and application of multi-objective evolutionary algorithms (MOEAs) is immense and has increased considerably in recent years. But theoretical analyses of their runtime are still rare. The works [1,2,3] and [4] focus

* Supported by the Deutsche Forschungsgemeinschaft (DFG) as part of the Collaborative Research Center "Computational Intelligence" (SFB 531).

Dagstuhl Seminar Proceedings 04461
Practical Approaches to Multi-Objective Optimization
http://drops.dagstuhl.de/opus/volltexte/2005/271

on the limit behavior, i. e., under what conditions an algorithm can find a set of optimal solutions when time goes to infinity. A common approach to learn how EAs work is to analyze basic EAs for simple objective functions f. The authors of [5] started such an analysis for two simple (and closely related) mutation based MOEAs. Their base algorithm SEMO (simple evolutionary multi-objective optimizer) can be outlined like this:

$P := \{x\}$ where x is chosen uniformly from $\{0,1\}^n$.
Loop:
  Choose $x \in P$ uniformly.
  $x' :=$ mutation(x).
  If no element in P weakly dominates x',
    add x' to P and remove all individuals dominated by x'.

The idea of this algorithm is to keep a population P of points that do not (weakly) dominate each others. In each step, an individual x is selected from P for mutation. If no point in P dominates the offspring x', then x' is added to P and all points dominated by x' are removed from P. For a finite decision space, the hope is that the population will contain for each point y of the Pareto front one point of its pre-image $f^{-1}(y)$ after some time. Of course, this is not guaranteed to happen if the mutation operator produces only points in the neighborhood of the current population.

In the following, we restrict to the case where the decision space is $\{0,1\}^n$ and consider objective functions $f\colon \{0,1\}^n \to \mathbb{R}^m$. (The results hold for any partially orded set in place of $\mathbb{R}^m$ but they apply only to the case where the search space is $\{0,1\}^n$. It does not seem possible to generalize the results or proof techniques to real-parameter problems.) The authors of [5] have started to analyze SEMO with the mutation operator that flips a uniformly chosen bit. In this text, we call their algorithm the *local* SEMO. The local SEMO has a finite expected runtime only for certain objective functions $f\colon \{0,1\}^n \to \mathbb{R}^m$. By "runtime" we denote the random number of objective function evaluations (essentially the number of iterations of the loop) until the population contains some pre-image point for every point of the Pareto front.

We analyze SEMO with the mutation operator that flips each bit independently with probability $1/n$. We call this variant the *global* SEMO. We first investigate the worst case and show that the global SEMO has an expected runtime of $O(n^n)$ for all objective functions $f\colon \{0,1\}^n \to \mathbb{R}^m$. Note that this bound is independent of the number of objectives m. Moreover, the bound matches the worst-case runtime of typical EAs working in the scenario of single-objective optimization. We also show that a lower bound of $\Omega(n^n)$ can easily be obtained from a bicriteria problem. Thus, considering expected runtimes, no problem with many objectives is essentially harder for the global SEMO than certain bicriteria problems. We broaden this result by showing that bicriteria problems can have any of the following expected runtimes: $\Theta(n^2), \Theta(n^3), \ldots, \Theta(n^n)$ and also $O(n \log n)$ is possible. This is obtained from a problem which is basically a discrete version of Schaffer's function $x \mapsto (x^2, (x-2)^2)$. This result shows that bicriteria problems can have almost any runtime in the full range of potential runtimes.

We then revisit two multi-objective problems that have been studied in the literature before. In [5] the bicriteria problem LOTZ (leading ones trailing zeroes) is investigated. The two conflicting objectives of a bitstring are its number of leading ones and its number of trailing zeroes. The set of Pareto optimal solutions implies all bitstrings of the form $1^i 0^{n-i}$. For the local SEMO, the authors of [5] have proved that the expected runtime is $\Theta(n^3)$ for LOTZ. Flipping only one bit in each step assures certain properties that simplify the analysis. These properties do not carry over to the global SEMO. We prove that the expected runtime of the global SEMO is also $O(n^3)$ and that this bound is not exceeded with overwhelming probability. Thus, independent bitflips do not increase the runtime substantially. Finally, we study the problem MOCO (multi-objective counting ones). Letting $||x|| \in \{0, \ldots, n\}$ denote the number of 1-bits of a bitstring x, $\varphi(x) := 2\pi ||x||/n$ defines an angle in the range $[0, 2\pi]$. Then, $\text{MOCO}(x) := (\cos\varphi(x), \sin\varphi(x))$. Pareto optimal solutions either have n 1-bits or at most $n/4$ 1-bits. Here, the conjecture was that both the local and the global SEMO have an expected runtime of $\Theta(n^2 \log n)$ [6]. Thierens' conjecture was based on calculations with assumptions and on experiments that suggest this expected runtime. We show that the local SEMO has no finite expected runtime for this function and that the expected runtime of the global SEMO is large, namely $n^{\Omega(n)}$. Thus, the expected runtime of SEMO does not reflect the behavior observed in practice. We prove the following revised conjecture. For the local and global SEMO, the runtime is $\Theta(n^2 \log n)$ with probability $1 - o(1)$.

The conference version of this work [7] is available and also a more comprehensive technical report [8] including the results on MOCO. A journal version of this work has been submitted.

References

1. Rudolph, G.: Evolutionary search for minimal elements in partially ordered finite sets. In: Proc. of the 7th Annual Conf. on Evolutionary Programming. (1998) 345–353
2. Rudolph, G.: On a multi-objective evolutionary algorithm and its convergence to the Pareto set. In: Proc. of the 5th IEEE Conf. on Evolutionary Computation. (1998) 511–516
3. Rudolph, G.: Evolutionary search under partially ordered fitness sets. In: Proc. of the Internat. NAISO Congress on Information Science Innovations (ISI 2001). (2001) 818–822
4. Rudolph, G., Agapie, A.: Convergence properties of some multi-objective evolutionary algorithms. In: Proc. of the 2000 Congress on Evolutionary Computation (CEC 2000). (2000) 1010–1016
5. Laumanns, M., Thiele, L., Zitzler, E., Welzl, E., Deb, K.: Running time analysis of multi-objective evolutionary algorithms on a simple discrete optimization problem. In: Proc. of the 7th Internat. Conf. on Parallel Problem Solving From Nature (PPSN VII). LNCS 2439 (2002) 44–53

6. Thierens, D.: Convergence time analysis for the multi-objective counting ones problem. In: Proc. of the 2nd Internat. Conf. on Evolutionary Multi-Criterion Optimization (EMO 2003). LNCS 2632 (2003) 355–364
7. Giel, O.: Expected runtimes of a simple multi-objective evolutionary algorithm. In: Proceedings of the 2003 Congress on Evolutionary Computation (CEC 2003). Volume 3. (2003) 1918–1925
8. Giel, O.: Runtime analyses for a simple multi-objective evolutionary algorithm. Technical Report CI-155/03, Universität Dortmund (2003)