Adaptive Mutation Control in Panmictic And Spatially Distributed Multi-objective Evolutionary Algorithms

Marco Laumanns[1], Günter Rudolph[2], and Hans-Paul Schwefel[2]

[1] ETH Zürich, Institut TIK, CH–8092 Zürich, Switzerland
[2] Universität Dortmund, Fachbereich Informatik, D–44221 Dortmund, Germany

Abstract. This paper addresses the problem of controlling mutation strength in multi-objective evolutionary algorithms. Adaptive parameter control is one major issue in the field of evolutionary computation, and several methods have been proposed and applied successfully for single objective optimization problems. In this study we examine whether these results carry over to the multi-objective case and what kind of modifications must be taken to meet the difficulties and pitfalls of conflicting objectives.

1 Introduction

Evolutionary algorithms have shown to be a useful auxiliary tool for approximating the Pareto set of multi-objective optimization problems. Much effort has been taken to cope with the evaluation and selection of solutions on the basis of partially ordered objective spaces. Research focusing on the role of the variation operators in evolutionary multi-objective optimization has remained comparatively rare. Though most algorithms simply apply standard operators from the single-objective case, their behavior in the presence of multiple objectives may differ substantially, particularly concerning the adaptation of mutation intensities [1].

When large search spaces are to be explored, adaptive variation operators are mandatory to achieve both, a satisfactory rate of progress towards the optimum and a high precision of solutions. In this study we explore the behavior of a standard self-adaptive evolution strategy (SA-ES) with Pareto-based selection and a spatially distributed predator-prey EA [2] for different multi-objective test functions (see Table 1).

2 Test scenario

We examine an algorithm's ability to produce a sequence of solutions that converges to the Pareto set. Therefore the average population distance to the set of efficient solutions (in decision variable space) is computed and compared to the average expected mutation step size. The search space is $\mathbb{R}^n$ with $n = 100$ and an individual is coded – as usual in evolution strategies – as an n-dimensional vector of floating point numbers for the decision variables. Additional strategy parameters $\sigma_i \in \mathbb{R}^+, i \in \{1, \dots, n_\sigma\}$ represent the standard deviation ('step size') of the normal-distributed random vectors used for mutation. Discrete recombination is used for the decision variables and intermediate recombination is used for the mutation step sizes.

Test functions	$F1 : \mathbb{R}^n \mapsto \mathbb{R}^2$	$F2 : \mathbb{R}^n \mapsto \mathbb{R}^2$
	$\sum_{i=1}^{n} x_i^2$ $(x_1 - 2)^2 + \sum_{i=2}^{n} x_i^2$	$\sum_{i=1}^{n} (\lvert x_i \rvert^{0.8} + \sin^3(x_i))$ $\sum_{i=1}^{n-1} (-10e^{-0.2\sqrt{x_i^2 + x_{i+1}^2}})$.

Table 1. Test functions $F1$ (multi-sphere model) and F2 (multi-modal)

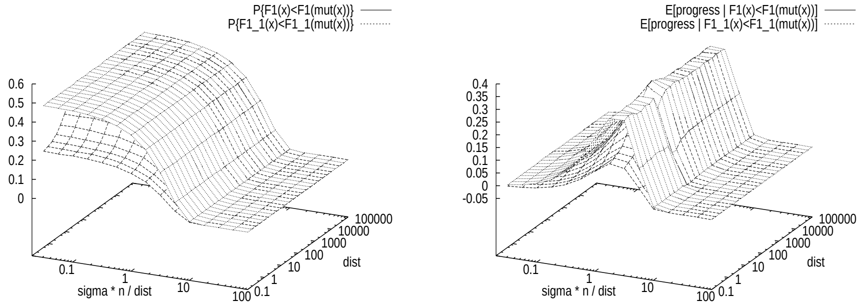

Fig. 1. Estimated success probabilities (left) and expected normalized progress (right) for mutations (step size σ) of individuals with distance *dist* to the optimum on F1.

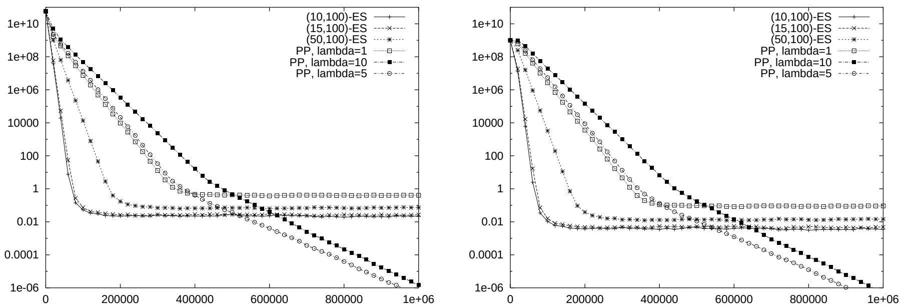

Fig. 2. Median population distance (left) and expected step length (right) over time (number of function evaluations) on F1

3 Behavior on the multi-sphere model (F1)

It is known from the single-objective case that standard self-adaptive evolution strategies exhibit linear order convergence on the sphere model (here: first component of $F1$). For the multi-sphere model, however, this property seems to be valid only individuals which are 'sufficiently far away' from the Pareto-optimal set. After a period of geometrically decreasing population distance to the Pareto set, the solutions suddenly start to oscillate around a small but fixed final distance (Fig. 2). An explanation can be deduced from Fig. 1: Not only the success rate decreases, but also the normalized average progress for successful mutations. Thus, it is increasingly difficult to proceed closer to the Pareto set, if successes are judged via the dominance relation.

A remedy might be found in the single-objective selection scheme of the predator-prey-EA. In the original version exactly $\lambda = 1$ offspring was produced for each individual deleted by a predator, but this does not improve convergence (Fig. 2). When $\lambda > 1$ offspring are produced, and the best one is selected for survival, further convergence can be achieved. However, for increasing λ the rate of convergence decreases.

4 Behavior on a multi-modal function (F2)

It has been shown recently that self-adaptive mutations may not guarantee convergence to the global optimum in multi-modal fitness landscapes [3]. Simulations on the first component of F2 as a multi-modal single objective function revealed that it is very difficult to reach the basin of attraction of the global optimal solution. On average, the (15/2,100)-ES was able to adjust 90 % of the decision variables; applying the weaker binary tournament selection instead lead to 95 % adjustment. These problems also hinder convergence of the (15,100)-ES in the multi-objective case.

The predator-prey EA, however, seems to overcome this problem of premature convergence in both cases, if $\lambda > 0$ is chosen. Again, large λ values mean more individuals to evaluate and thus shows slower convergence rates in respect to the number of objective function evaluations.

5 Conclusions

In this study we examined the problem of controlling mutation strength in multi-objective evolutionary algorithms. The results show that the standard self-adaptive evolution strategies have difficulties to converge to the Pareto set due to the low success probability of Pareto based selection. Alternative selection methods like the predator-prey approach discussed here are a possibility to overcome this limitation. A disadvantage of this method, however, is the lower rate of convergence due to the less efficient single-criterion selection. It may be speculated whether a hybrid method using Pareto-based selection for a fast but rough localization of the Pareto set and single-criterion selection for an accurate approximation would be appropriate.

Another possible way to improve the convergence properties of self-adaptive MOEAs might be a combination of low selection pressure (to diminish the danger of premature convergence) and elitism (to prevent possible divergence caused by increasing mutation step sizes and low selection pressure). Also the topic of diversity needs a thorough investigation in connection with convergence in self-adaptive MOEAs. Finally, it should be emphasized that these results are entirely based on numerical experiments and empirical evaluation. Theoretical analysis could fundamentally improve the understanding of the effects of these heuristic methods.

References

1. Günter Rudolph. On a multi-objective evolutionary algorithm and its convergence to the pareto set. In *IEEE International Conference on Evolutionary Computation (ICEC'98)*, pages 511–516, Piscataway, NJ, 1998. IEEE.
2. Marco Laumanns, Günter Rudolph, and Hans-Paul Schwefel. A spatial predator-prey approach to multi-objective optimization: A preliminary study. In Agoston E. Eiben, Thomas Bäck, Marc Schoenauer, and Hans-Paul Schwefel, editors, *Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V)*, pages 241–249, Berlin, Germany, 1998. Springer.
3. Günter Rudolph. Self-adaptive mutations may lead to premature convergence. Technical Report CI-73/99, Department of Computer Science/XI, University of Dortmund, 1999.