



Dominance Based Crossover Operator for Evolutionary Multi-objective Algorithms

Olga Roudenko, Marc Schoenauer

► To cite this version:

Olga Roudenko, Marc Schoenauer. Dominance Based Crossover Operator for Evolutionary Multi-objective Algorithms. Parallel Problem Solving from Nature 2004, 2004, Birmingham, pp.812-821. inria-00000095

HAL Id: inria-00000095

<https://inria.hal.science/inria-00000095>

Submitted on 29 May 2005

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Dominance Based Crossover Operator for Evolutionary Multi-objective Algorithms DRAFT VERSION

Olga Rudenko, Marc Schoenauer

TAO Team, INRIA Futurs
LRI, bat. 490, Université Paris-Sud
91405 Orsay Cedex, France
Olga.Roudenko@lri.fr, Marc.Schoenauer@inria.fr

Abstract.

In spite of the recent quick growth of the Evolutionary Multi-objective Optimization (EMO) research field, there has been few trials to adapt the general variation operators to the particular context of the quest for the Pareto-optimal set. The only exceptions are some mating restrictions that take in account the distance between the potential mates – but contradictory conclusions have been reported. This paper introduces a particular mating restriction for Evolutionary Multi-objective Algorithms, based on the Pareto dominance relation: the partner of a non-dominated individual will be preferably chosen among the individuals of the population that it dominates. Coupled with the BLX crossover operator, two different ways of generating offspring are proposed. This recombination scheme is validated within the well-known NSGA-II framework on three bi-objective benchmark problems and one real-world bi-objective constrained optimization problem. An acceleration of the progress of the population toward the Pareto set is observed on all problems.

Introduction

The idea of *restricted mating* is not new in Evolutionary Computation: Goldberg [7] already suggested to forbid, or at least restrict, the crossover between too different individuals (i.e. that are too far apart for some distance on the genotypic space) – which makes sense for single-objective problems as soon as the population has started to accumulate on different fitness peaks, as recombining individuals from different peaks would often lead to lethal individuals. This idea has been transposed in the framework of Evolutionary Multi-objective Algorithms (EMAs) by Hajela and Lin [8], and by Fonseca and Fleming [6]. Nevertheless, Zitzler and Thiele [17] did not observe any improvement when mating similar individuals. On the other hand, Horn et al. [9] present an argument supporting mating of dissimilar individuals: in the multi-objective framework, because the population diversity is enforced, the information provided by very different solutions can be combined in such way that a new type of (good) compromises can hopefully be obtained. Nevertheless, Schaffer reported the absence

of the improvement when mating dissimilar individuals. To sum up, no clear conclusion can be drawn from existing experiments on the usefulness of restricted mating based on the (dis)similarity between mates.

On a real-world design problem, using a very specific representation, Wildman et Parks [13] have investigated different pairing strategies based on maximizing or minimizing different similarity measures. In particular, the similarity in the sense of the dominance rank has been considered, and enforcing the mating of the individuals from the elite archive with the individuals from the population, in an archive-based EMA, has been observed to be beneficial.

However, in all studies enumerated above, the efficiency of the proposed mating restrictions has been measured by the quality of the final non-dominated solutions, without addressing the issue of computational time. In this paper, we propose a restricted mating strategy whose main effect is to accelerate the progress of the population of an EMA toward the Pareto set. The main idea is fairly simple, and consists in using the Pareto dominance relation when choosing a mate for the best (non-dominated) individuals. Note that a more detailed presentation (in French) can be found in [10].

The paper is organized as follows. Next section briefly introduces evolutionary multi-objective optimization, and describes in more detail the NSGA-II algorithm, one of the best performing EMA to-date, that will be used in all experiments. Two slightly different implementations of the dominance-based crossover operator are then proposed in Section 2, based on BLX- α crossover, used throughout this study. Section 3 presents some experimental results witnessing the acceleration of the progress toward the Pareto set when using the proposed mating restrictions. Finally, Section 4 gives some guidelines for a more rigorous and complete validation of the proposed strategy, as well as for its possible refinements.

1 Evolutionary Multi-objective Optimization

Multi-objective optimization aims at simultaneously optimizing several contradictory objectives. For such kind of problems, there does not exist a single optimal solution, and compromises have to be made.

An element of the search space x is said to *Pareto-dominate* another element y if x is not worse than y with respect to all objectives, and is strictly better than y with respect to at least one objective. The set of all elements of the search space that are not Pareto-dominated by any other element is called the *Pareto set* of the multi-objective problem at hand: it represents the best possible compromises with respect to the contradictory objectives.

Solving a multi-objective problem amounts to choose one solution among those non-dominated solutions, and some decision arguments have to be given. Unlike classical optimization methods, that generally find one of the Pareto optimal solutions by making the initial optimization problem single-objective, EMAs are to-date the only algorithms that directly search for the whole Pareto

set, allowing decision makers to choose one of the Pareto solutions with more complete information.

1.1 Pareto-based Evolutionary Algorithms

In order to find a good approximation of the Pareto set (a uniform and well spread sampling of the non-dominated solutions, close to the actual Pareto set of the problem at hand), EMAs have to enforce some progress toward the Pareto set while, at the same time, preserving diversity between the non-dominated solutions.

Numerous evolutionary methods have been designed in the past years for the particular task of searching for the Pareto set (the interested reader will find a good summary in [2]). The best performing among them (NSGA-II [3], SPEA2 [16], PESA [1]) are directly based on the Pareto dominance relation, that actually ensures progressing toward the non-dominated set.

Among the diversity preserving techniques, some were transposed to EMAs from single-objective EAs (such as sharing, for instance), while others, like the crowding distance described in next subsection, are specific to the multi-objective framework.


Another recognized important feature of EMAs is elitism [15], directly related to the notion of the Pareto dominance in EMAs: the non-dominated individuals can be preserved either by maintaining an archive (SPEA2 and PESA) or by using a deterministic replacement procedure (NSGA-II).

1.2 NSGA-II

The NSGA-II algorithm has been proposed by Deb et al. in 2001 [3]. The progress toward the Pareto set is here favored by using a selection based on the *Pareto ranking*, that divides the population into a hierarchy of non-dominated subsets, as illustrated by figure 1(a). All non-dominated individuals of the population are first labeled as being of rank 1; they are then temporarily removed from the population, and the process is repeated: the non-dominated individuals of the remainder of the population are given rank 2, and so on, until the whole population is ranked.

NSGA-II diversity preserving technique is based on the *crowding distance* - one of the possible estimations of the density of the solutions belonging to the same non-dominated subset. The crowding distance of each individual i is computed as follows: the non-dominated subset to which the individual i belongs is ordered following each of the objectives; for each objective m , the distance $d_i^{(m)} = f_m(i+1) - f_m(i-1)$ between the surrounding neighbors of individual i according to objective m is computed (Fig. 1(b)); the sum over all objectives of these distances is the crowding distance of individual i .

The following comparison operator \succ is then used during the Darwinian stages (selection and replacement) of NSGA-II:



(a) Ensuring progress toward the Pareto set (b) Preserving diversity technique

Fig. 1. *NSGA-II comparison criteria*

$$\begin{aligned}
 x \succ y \text{ iff } & \text{rank}(x) < \text{rank}(y) \\
 & \text{or } \text{rank}(x) = \text{rank}(y) \\
 & \text{and } \text{crowding_dist}(x) > \text{crowding_dist}(y)
 \end{aligned}$$

NSGA-II selection is based on tournament: it chooses an individual for reproduction by uniformly drawing T individuals (generally, $T = 2$) from the population and returning the best of them with respect to the comparison operator \succ . NSGA-II replacement is deterministic: it consists in merging parents and offspring together, and choosing the N best individuals in that global population, again using comparison operator \succ . The algorithm NSGA-II is said to be elitist because the best (less crowded non-dominated) individuals are preserved from one generation to another.

2 Dominance-based crossover

The basic mechanism of the proposed mating restriction consists in allowing the mating of each of the best individuals only with an individual it dominates (if any), where *best individuals* means non-dominated individuals when applying NSGA-II, or members of the archive when using SPEA2 ou PESA.

The rationale behind using the dominance relation to form the couples for the crossover is the following. If x dominates y , then x is better than y for all objectives. Hence, the direction $y \rightarrow x$ is likely to improve all criteria simultaneously. Furthermore, a natural continuation of the same idea is to bias the distribution of the offspring toward the dominant parent, as better individuals are more likely to be found close to it. However, it is clear that success of this idea depends on the behavior of the objective functions in the region of the decision space where the mated individuals sit.

The resulting crossover, called Dominance-Based Crossover (DBX) will proceed as follows: a first mate is chosen using the usual selection procedure of the EMA at hand (e.g. tournament based on the \succ operator for NSGA-II). If the chosen individual is non-dominated and dominates some other individuals in the population, its mate is chosen among those. Otherwise, the mate is chosen us-

ing the usual selection procedure. In any case, crossover then proceeds with the chosen operator.

In this study, the well-known BLX- α crossover ($0 < \alpha < 1$), proposed by Eshelman and Schaffer [5] for real representations, has been used. Formally, given two parents $(x_i)_{i \in [1, n]}$ and $(y_i)_{i \in [1, n]}$, this operator produces an offspring by a random linear recombination of the parents as follows:

$$((x_i)_{i \in [1, n]}, (y_i)_{i \in [1, n]}) \longrightarrow (\phi_i x_i + (1 - \phi_i) y_i)_{i \in [1, n]}, \quad (1)$$

where $\phi_i = U[\alpha, 1 + \alpha]$. In our particular case, given a non-dominated individual $(x_i)_{i \in [1, n]}$ from the NSGA-II population, two possible strategies will be considered to generate the offspring:

1. **Symmetric DBX:** The mate $(y_i)_{i \in [1, n]}$ is chosen from the list of the individuals dominated by $(x_i)_{i \in [1, n]}$, if any, by tournament otherwise, and $\phi_i = U[-0.5, 1.5]$ in Equation (1).
2. **Biased DBX:** Similarly, the mate $(y_i)_{i \in [1, n]}$ is chosen from the list of the individuals dominated by $(x_i)_{i \in [1, n]}$, if any, but now $\phi_i = U[0.5, 1.5]$ in Equation (1), i.e. the offspring will be closer to the first parent $(x_i)_{i \in [1, n]}$.

3 Experimental results

3.1 Experimental conditions

This section presents some experimental results, on three standard benchmark problems [15] and on an industrial problem [12]. All experiments are run with population size 100 and tournament size 2. The two DBX crossovers are compared to the standard BLX-0.5, the crossover rate is set to 0.9. The uniform mutation operator is applied with rate 0.05. The algorithms run for at most 150 (resp. 250) generations for ZDT-benchmarks (resp. the industrial problem).

3.2 Bi-objective ZDT benchmarks

For each of ZDT1–ZDT3 test problems 31 NSGAII runs have been performed for biased DBX, symmetric DBX and standard BLX-0.5 operators starting with the same 31 initial populations. The non-dominated individuals over all 31 runs have been calculated at each 10th generation approximately until the moment when the whole population is non-dominated, that means that DBX crossover is not applied any longer. These snapshots corresponding to biased DBX and standard BLX-0.5 are shown in the figures 2, 3 and 4 for ZDT1, ZDT2 and ZDT3 respectively. On all three test problems, a small but steady acceleration of the progress toward the Pareto front is observed when using DBX crossover. On those problems, very similar results have been obtained with the symmetric and biased DBX operators (the snapshots corresponding to the symmetric DBX are not shown for the space reasons).




Fig. 2. ZDT1: the black bullets (for the biased DBX) and the gray crosses (for the standard BLX-0.5) represent the non-dominated individuals over 31 runs at generations 10, 20, 30 and 40.




Fig. 3. ZDT2: the black bullets (for the biased DBX) and the gray crosses (for the standard BLX-0.5) represent the non-dominated individuals over 31 runs at generations 20, 30, 40 and 50.




Fig. 4. ZDT3: the black bullets (for the biased DBX) and the gray crosses (for the standard BLX-0.5) represent the non-dominated individuals over 31 runs at generations 10, 20, 30 and 40.

3.3 Constrained bi-objective optimization

This industrial problem consists in optimizing the structural geometry, described by 13 continuous parameters, of the front crash member of the car in order to minimize its mass while maximizing the internal energy absorbed during a crash (two competitive objectives) under 8 constraints arising from the acoustic and static mechanical domains [12].

The constraints have been handled using the so-called *infeasibility objective approach* [14]: the aggregated sum of the scaled constraint violations (that can be viewed as a measure of distance separating each individual from the feasible region) was considered as an additional optimization criterion – the infeasibility objective. NSGA-II was hence applied to a three-objective optimization problem, with the difference that the infeasibility objective had a higher priority than both others in the \succ operator. In the other words, every individual located closer to the feasible region is considered better than any individual further away from it, regardless the corresponding values of the other objectives, the mass and the absorbed internal energy.

DBX operators can be employed in this context exactly like for unconstrained problems, using the modified \succ comparison. On every run, the restricted mating allowed faster approach of the feasible region, as well as faster feasibility of the whole population, as illustrated by Figure 5-a. Moreover, some differences with the results on the test problems ZDT have been steadily observed. First, biased DBX was more efficient than symmetric DBX. Then, and more importantly,

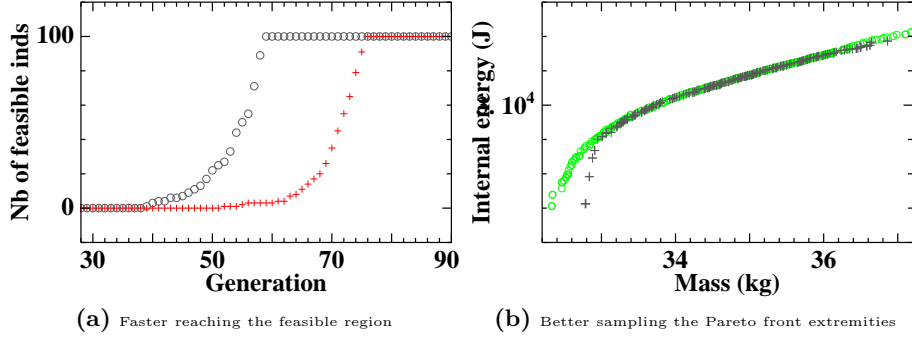


Fig. 5. *Effect of the dominance-based mating restriction in the presence of constraints: biased DBX (rings) VS standard BLX-0.5 (crosses)*

the use of mating restriction not only accelerated the search, it also provided solutions of better quality at the extremities of the Pareto front, as it can be seen on figure 5-b.

4 Discussion and future work

For all four problems considered in this study, the DBX operators (that only allow the mating of dominant individuals with individuals they dominate) have been shown to accelerate the progress of the populations toward the Pareto set. Moreover, for the optimization of the car front crash member, it also allowed finding solutions of better quality at the extremities of the Pareto set that could not be reached when using the usual recombination strategy.

The observed acceleration is relatively small, but is systematic. When observing NSGA-II dynamics (the typical NSGA-II evolution process is illustrated by Figure 6), we realize that in fact only a small number of DBX crossovers are actually applied. First, note that restricted mating is applied at most as many times as there are non-dominated individuals in the population. There are few non-dominated individuals at the beginning of evolution (Fig. 6-a,b), but each of them dominates a lot of other individuals (Fig. 6-b). As the population gets closer to the Pareto set, the number of non-dominated individuals rapidly increases (Fig. 6-a,c), but each of them dominates only a few individuals if any (Fig. 6-c). Finally, when the whole population is non-dominated (Fig. 6-d), DBX crossover cannot be applied because no individual actually dominates anyone in the population! The actual rate of application of the DBX operator is maximal about the moment when the half of the population gets the rank 1 (Fig. 6-c).

One possible improvement would be to increase the rate of application of DBX operators at the beginning of evolution by mating each non-dominated individual with not only one but several individuals it dominates.

Note that dynamic behavior of the populations described here above and, in particular, the disappearance of the dominated individuals is due to the the




Fig. 6. Evolution of the number of non-dominated individuals during a run of NSGA-II

replacement procedure in NSGA-II, and might not be so prominent with other EMAs. It is hence worth investigating the use of DBX restricted mating with other EMAs paradigms, such as SPEA2 or PESA, for instance.

Furthermore, there exist other crossover operators used in EAs (such as SBX [2], for example) that could be applied together with dominance based restricted mating, instead of BLX- α .

One more issue, that needs to be thoroughly investigated, is the efficiency of the DBX strategy when solving problems with more than two objectives. Indeed, in such situations, the “visual” analysis of the populations dynamics performed in the present work will not be possible anymore, and our plans include the use of so-called *running metrics* [4].

Indeed, the present study could use some performance metrics as well, instead of the graphical plots of the population dynamics in the objective space. However, in our opinion, such presentation is much clearer and makes things more explicit, hence it is more appropriate to the introduction of the new operator.

But probably the most important issue is to find a way to actually benefit from the reported acceleration in practice. Whereas its usefulness is obvious for costly applications, where the number of evaluations has to be limited, the question remains of how to detect that the algorithm is getting stuck, thus saving the computational time that restricted mating allows us to spare. An efficient stopping criterion is needed before any actual benefit can be gained from the observed acceleration. Such criterion has been proposed by the first author [10,11], but it was beyond the scope of this paper. Nevertheless, together

with the DBX operators described in this paper, it allowed an actual saving of about 8% of the computation time. However, further evaluation and refinement of that stopping criterion are still needed before definite strong conclusions can be drawn.

References

1. D. W. Corne, J. D. Knowles, and M.L. Oates. The Pareto Envelope-based Selection Algorithm for Multiobjective Optimization. In M. Schoenauer and al., editors, *Proc. PPSN 2000*, pages 839–848. LNCS 1917, Springer Verlag, 2000.
2. K. Deb. *Multi-Objective Optimization using Evolutionary Algorithms*. John Wiley, 2001.
3. K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A Fast Elitist Non-dominated Sorting Genetic Algorithm for Multi-objective Optimization: NSGA-II. In M. Schoenauer and al., editors, *Proc. PPSN 2000*, pages 849–858. LNCS 1917, Springer Verlag, 2000.
4. K. Deb and S. Jain. Running performance metrics for evolutionary multi-objective optimization. Technical Report 2002004, Indian Institute of Technology, Kanpur, May 2002.
5. L. Eshelman and J. D. Schaffer. Real-coded genetic algorithms and interval-schemata. In L. D. Whitley, editor, *Foundations of Genetic Algorithms 2*, pages 187–202, Morgan Kaufmann, 1993.
6. Carlos M. Fonseca and Peter J. Fleming. Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In *ICGA'93*, pages 416–423. Morgan Kaufmann, 1993.
7. D. E. Goldberg. *Genetic algorithms in search, optimization and machine learning*. Addison Wesley, 1989.
8. P. Hajela and C. Y. Lin. Genetic search strategies in multicriterion optimal design. *Structural Optimization*, 4:99–107, 1992.
9. J. Horn, S. N. Nafpliotis, and D. E. Goldberg. A niched pareto genetic algorithm for multiobjective optimization. In Z. Michalewicz et al., editors, *Proc. of ICEC'94*, pages 82–87. IEEE Press, 1994.
10. O. Rudenko. *Application des algorithmes évolutionnaires aux problèmes d'optimisation multi-objectif avec contraintes*. PhD thesis, École Polytechnique, 2004.
11. O. Rudenko and M. Schoenauer. A steady performance stopping criterion for pareto-based evolutionary algorithms. In *Proc. 6th Intl Conf. on Multi Objective Programming and Goal Programming*, 2004.
12. O. Rudenko, M. Schoenauer, T. Bosio, and R. Fontana. A multiobjective evolutionary algorithm for car front end design. In P. Collet et al., editors, *Artificial Evolution'2001*, pages 205–216, LNCS 2310, Springer Verlag, 2001.
13. A. Wildman and G. Parks. A comparative study of selective breeding strategies in a multiobjective genetic algorithm. In C. M. Fonseca et al., editors, *Proc. EMO'03*, pages 418–432. LNCS 2632, Springer Verlag, 2003.
14. J. Wright and H. Loosemore. An infeasibility objective for use in constrained pareto optimization. In E. Zitzler et al., editors, *Proc. EMO'01*, pages 256–268, LNCS 1993, Springer-Verlag, 2001.
15. E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, 8(2):125–148, 2000.

16. E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the strength pareto evolutionary algorithm. Technical Report 103, Computer Engineering and Networks Laboratory, ETH, Zurich, Switzerland, 2001.
17. E. Zitzler and L. Thiele. Multiobjective optimization using evolutionary algorithms - a comparative case study. In A.-E. Eiben, T. Bäck, M. Schoenauer, and H.-P. Schwefel, editors, *Proc. PPSN'98*, pages 292–301. LNCS 1498, Springer Verlag, 1998.