

SPEA2: IMPROVING THE STRENGTH PARETO EVOLUTIONARY ALGORITHM FOR MULTI-OBJECTIVE OPTIMIZATION

Eckart Zitzler*, Marco Laumanns, and Lothar Thiele

Computer Engineering and Networks Laboratory (TIK)

Department of Information Technology and Electrical Engineering

Swiss Federal Institute of Technology (ETH) Zurich

Gloriastr. 35, CH-8092 Zurich, Switzerland

e-mail: {zitzler,laumanns,thiele}@tik.ee.ethz.ch

web page: <http://www.tik.ee.ethz.ch/aroma>

Abstract. The Strength Pareto Evolutionary Algorithm (SPEA)¹ is a relatively recent technique for finding or approximating the Pareto-optimal set for multi-objective optimization problems. In different studies^{1,2} SPEA has shown very good performance in comparison to other multiobjective evolutionary algorithms, and therefore it has been a point of reference in various recent investigations.³ Furthermore, it has been used in different applications.⁴ In this paper, an improved version, namely SPEA2, is proposed, which incorporates in contrast to its predecessor a fine-grained fitness assignment strategy, a density estimation technique, and an enhanced archive truncation method. The comparison of SPEA2 with SPEA and two other modern elitist methods, PESA and NSGA-II, on different test problems yields promising results.

Key words: multiobjective optimization, evolutionary algorithms, diversity preservation, elitism

1 INTRODUCTION

After the first studies on evolutionary multiobjective optimization (EMO) in the mid-1980s, a number of Pareto-based techniques were proposed in 1993 and 1994, which demonstrated the capability of EMO algorithms to approximate the set of optimal trade-offs in a single optimization run. These approaches did not incorporate elitism explicitly, but a few years later the importance of this concept in multiobjective search was recognized and supported experimentally.² A couple of elitist multiobjective evolutionary algorithms were presented at this time, among others SPEA.¹ SPEA, an acronym for Strength Pareto Evolutionary Algorithm, was one of the first techniques that were extensively compared to several existing evolution-based methods.^{1,2} As it clearly outperformed the (non-elitist) alternative approaches under consideration, it has been used as a point of reference by vari-

ous researchers.³ Meanwhile further progress has been made and recently proposed methods, for instance NSGA-II⁵ and PESA,³ were shown to outperform SPEA on certain test problems. Furthermore, new insights into the behavior of EMO algorithms improved our knowledge about the basic principles and the main factors of success in EMO.^{6,7}

In this paper, SPEA2 is presented for which we tried to incorporate most recent results and to eliminate the potential weaknesses of its predecessor (see the corresponding technical report for further details⁸). In particular, the main differences of SPEA2 in comparison to SPEA are: i) an improved fitness assignment scheme which takes for each individual into account how many individuals it dominates and it is dominated by, ii) a nearest neighbor density estimation technique, and iii) a new archive truncation methods that guarantees the preservation of boundary solutions.

2 THE SPEA2 ALGORITHM

The main loop of SPEA2 is as follows:

Algorithm 1

- Input: N (population size)
 \bar{N} (archive size)
 T (maximum number of generations)
- Output: \mathbf{A} (nondominated set)
- Step 1: *Initialization:* Generate an initial population \mathbf{P}_0 and create the empty archive (external set) $\bar{\mathbf{P}}_0 = \emptyset$. Set $t = 0$.
- Step 2: *Fitness assignment:* Calculate fitness values of individuals in \mathbf{P}_t and $\bar{\mathbf{P}}_t$ (cf. Section 2.1).
- Step 3: *Environmental selection:* Copy all nondominated individuals in \mathbf{P}_t and $\bar{\mathbf{P}}_t$ to $\bar{\mathbf{P}}_{t+1}$. If size of $\bar{\mathbf{P}}_{t+1}$ exceeds \bar{N} then reduce $\bar{\mathbf{P}}_{t+1}$ by means of the truncation operator, otherwise if size of $\bar{\mathbf{P}}_{t+1}$ is less than \bar{N} then fill $\bar{\mathbf{P}}_{t+1}$ with dominated individuals in \mathbf{P}_t and $\bar{\mathbf{P}}_t$ (cf. Section 2.2).
- Step 4: *Termination:* If $t \geq T$ or another stopping criterion is satisfied then set \mathbf{A} to the set of decision vectors represented by the nondominated individuals in $\bar{\mathbf{P}}_{t+1}$. Stop.
- Step 5: *Mating selection:* Perform binary tournament selection with replacement on $\bar{\mathbf{P}}_{t+1}$ in order to fill the mating pool.
- Step 6: *Variation:* Apply recombination and mutation operators to the mating pool and set \mathbf{P}_{t+1} to the resulting population. Increment generation counter ($t = t + 1$) and go to Step 2.

In contrast to SPEA, SPEA2 uses a fine-grained fitness assignment strategy which incorporates density information as will be described in Section 2.1. Furthermore, the archive size is fixed, i.e., whenever the number of nondominated individuals is less than the predefined archive size, the archive is filled up by dominated individuals; with SPEA, the archive size may vary over time. In addition, the clustering technique, which is invoked when the nondominated front exceeds the archive limit, has been replaced by an alternative truncation method which has similar features but does not lose boundary points. Details on the environmental selection procedure will be given in Section 2.2. Finally, another difference to SPEA is that only members of the archive participate in the mating selection process.

2.1 Fitness Assignment

In Step 2 of the SPEA2 main loop, each individual \mathbf{i} in the archive $\overline{\mathbf{P}}_t$ and the population \mathbf{P}_t is assigned a strength value $S(\mathbf{i})$, representing the number of solutions it dominates $S(\mathbf{i}) = |\{\mathbf{j} \mid \mathbf{j} \in \mathbf{P}_t + \overline{\mathbf{P}}_t \wedge \mathbf{i} \succ \mathbf{j}\}|$ where $|\cdot|$ denotes the cardinality of a set, $+$ stands for multiset union and the symbol \succ corresponds to the Pareto dominance relation. On the basis of the S values, the raw fitness $R(\mathbf{i})$ of an individual \mathbf{i} is calculated $R(\mathbf{i}) = \sum_{\mathbf{j} \in \mathbf{P}_t + \overline{\mathbf{P}}_t, \mathbf{j} \succ \mathbf{i}} S(\mathbf{j})$. That is the raw fitness is determined by the strengths of its dominators in both archive and population, as opposed to SPEA where only archive members are considered in this context. Please note that fitness is to be minimized here.

In addition, density information is incorporated to discriminate between individuals having identical raw fitness values. The density estimation technique used in SPEA2 is an adaptation of the k -th nearest neighbor method,⁹ where the density at any point is a (decreasing) function of the distance to the k -th nearest data point. Here, we simply take the inverse of the distance to the k -th nearest neighbor as the density estimate. To be more precise, for each individual \mathbf{i} the distances (in objective space) to all individuals \mathbf{j} in archive and population are calculated and stored in a list. After sorting the list in increasing order, the k -th element gives the distance sought, denoted as σ_i^k . As a common setting, we use k equal to the square root of the sample size,⁹ thus, $k = \sqrt{N + \overline{N}}$. Afterwards, the density $D(\mathbf{i})$ corresponding to \mathbf{i} is defined by $D(\mathbf{i}) = \frac{1}{\sigma_i^k + 2}$. In the denominator, two is added to ensure that its value is greater than zero and that $D(\mathbf{i}) < 1$. Finally, adding $D(\mathbf{i})$ to the raw fitness value $R(\mathbf{i})$ of an individual \mathbf{i} yields its fitness $F(\mathbf{i})$ $F(\mathbf{i}) = R(\mathbf{i}) + D(\mathbf{i})$.

2.2 Environmental Selection

The archive update operation performed in Step 3 of Algorithm 1 works as follows. First all nondominated individuals, i.e., those which have a fitness lower than one, are copied to the archive of the next generation $\overline{\mathbf{P}}_{t+1} = \{\mathbf{i} \mid \mathbf{i} \in \mathbf{P}_t + \overline{\mathbf{P}}_t \wedge F(\mathbf{i}) < 1\}$. If the nondominated front fits exactly into the archive ($|\overline{\mathbf{P}}_{t+1}| = \overline{N}$) the environmental selection step is completed. Otherwise, there can be two situations: Either the archive is too small ($|\overline{\mathbf{P}}_{t+1}| < \overline{N}$) or too large ($|\overline{\mathbf{P}}_{t+1}| > \overline{N}$). In the first case, the best $\overline{N} - |\overline{\mathbf{P}}_{t+1}|$ dominated individuals in the previous archive and population are copied to the new archive. This can be implemented by sorting the multiset $\mathbf{P}_t + \overline{\mathbf{P}}_t$ according to the fitness values and copy the first $\overline{N} - |\overline{\mathbf{P}}_{t+1}|$ individuals \mathbf{i} with $F(\mathbf{i}) \geq 1$ from the resulting ordered list to $\overline{\mathbf{P}}_{t+1}$. In the second case, when the size of the current nondominated (multi)set exceeds \overline{N} , an archive truncation procedure is invoked which iteratively removes individuals from $\overline{\mathbf{P}}_{t+1}$ until $|\overline{\mathbf{P}}_{t+1}| = \overline{N}$. Here, at each iteration that individual \mathbf{i} is chosen for removal for which $\mathbf{i} \leq_d \mathbf{j}$ for all $\mathbf{j} \in \overline{\mathbf{P}}_{t+1}$ with

$$\mathbf{i} \leq_d \mathbf{j} \quad :\Leftrightarrow \quad \forall 0 < k < |\overline{\mathbf{P}}_{t+1}| : \sigma_i^k = \sigma_j^k \quad \vee \\ \exists 0 < k < |\overline{\mathbf{P}}_{t+1}| : \left[\left(\forall 0 < l < k : \sigma_i^l = \sigma_j^l \right) \wedge \sigma_i^k < \sigma_j^k \right]$$

where σ_i^k denotes the distance of \mathbf{i} to its k -th nearest neighbor in $\overline{\mathbf{P}}_{t+1}$. In other words, the individual which has the minimum distance to another individual is chosen at each stage; if there are several individuals with minimum distance the tie is broken by considering the second smallest distances and so forth.

Name	Type	Domain	Objective functions
SPH- $m^{10,11}$	min	$[-10^3, 10^3]^n$ $n = 100$	$f_j(x) = \sum_{1 \leq i \leq n, i \neq j} (x_i)^2 + (x_j - 1)^2$ $1 \leq j \leq m, \quad m = 2, 3$
ZDT6 ²	min	$[0, 1]^n$ $n = 100$	$f_1(x) = 1 - \exp(-4x_1) \sin^6(6\pi x_1)$ $f_2(x) = g(x) [1 - (f_1(x)/g(x))^2]$ $g(x) = 1 + 9 \cdot ((\sum_{i=2}^n x_i)/(n-1))^{0.25}$
QV ¹²	min	$[-5, 5]^n$ $n = 100$	$f_1(x) = (\frac{1}{n} \sum_{i=1}^n (x_i^2 - 10 \cos(2\pi x_i) + 10))^{\frac{1}{4}}$ $f_2(x) = (\frac{1}{n} \sum_{i=1}^n ((x_i - 1.5)^2 - 10 \cos(2\pi(x_i - 1.5)) + 10))^{\frac{1}{4}}$
KUR ¹³	min	$[-10^3, 10^3]^n$ $n = 100$	$f_1(x) = \sum_{i=1}^n (x_i ^{0.8} + 5 \cdot \sin^3(x_i) + 3.5828)$ $f_2(x) = \sum_{i=1}^{n-1} (1 - e^{-0.2\sqrt{x_i^2 + x_{i+1}^2}})$
KP-750- m^1	max	$\{0, 1\}^n$ $n = 750$	$f_j(x) = \sum_{i=1}^n x_i \cdot p_{i,j}$ s.t. $g_j(x) = \sum_{i=1}^n x_i \cdot w_{i,j} \leq W_j, \quad 1 \leq j \leq m, \quad m = 2, 3, 4$ $p_{i,j}$ (profit values), $w_{i,j}$ (weight values) randomly chosen

Table 1: Test problems used in this study. The objective functions are given by $f_j, 1 \leq j \leq m$, where m denotes the number of objectives and n the number of decision variables.

3 COMPARATIVE CASE STUDY

3.1 Experimental Design

The behavior of SPEA2 is compared to SPEA, NSGA-II and PESA on a number of test functions. The algorithms are implemented according to their description in the literature. As the main feature under concern is the fitness assignment and the selection processes, our implementations only differ in these respects, where the other operators (recombination, mutation, sampling) remain identical. For each algorithm we used identical population and archive sizes. The test functions are summarized in Tab. 1, where both combinatorial and continuous problems were chosen.

For each algorithm and each problem, 30 runs with different random seeds have been carried out. For the quality or performance measure we apply a volume-based approach according to (Zitzler and Thiele, 1999)¹ with slight modifications. Here, a reference volume between the origin and an utopia point – defined by the profit sums of all items in each objective – is taken into account. The aim is to minimize the fraction of that space, which is not dominated by any of the final archive members. We consider this as the most appropriate scalar indicator since it combines both the distance of solutions (towards some utopian trade-off surface) and the spread of solutions. For each run, we measure the (normalized) size of the nondominated objective space over time, which leads to a sample of 30 values for each time step in each experiment.

3.2 Results and Discussion

In all test cases, SPEA2 shows to constitute a significant improvement over its predecessor SPEA as it reaches better results on all considered problems. As the box plots in Fig. 1 show, the performance differences increase with the number of objectives. With four objectives, a clear distinction between NSGA-II and SPEA2 on the one hand and PESA and SPEA on the other hand can be made: Both PESA and SPEA, which do not guarantee the extreme solutions to be kept in the archive, appear to stagnate without having reached a well spread distribution of solutions. SPEA2 and NSGA-II seem to behave very similar on the different problems. In some cases NSGA-II reaches a broader spread and hence a better value of the performance

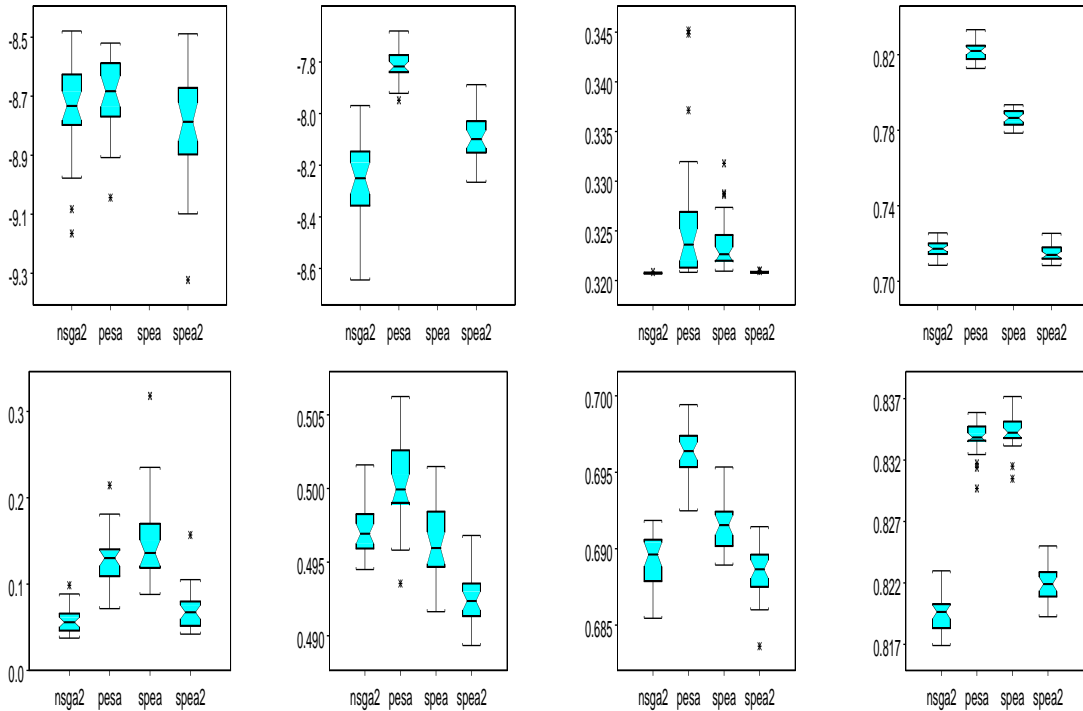


Figure 1: Box plots showing the distribution of the performance values at the end of the 30 runs. Upper row: SPH-2, SPH-3, ZDT6, and QV. Lower row: KUR, KP-750-2, KP-750-3, and KP-750-4.

measure, while SPEA2 provides a better distribution of points, especially when the number of objectives increases. PESA, however, tends to have difficulties to keep the outer solutions on certain test functions.

It is very instructive, however, to see how the performance develops over time, i.e. with the number of function evaluations. For many problems, PESA appears to be converging quicker at the beginning, which is probably due to its higher elitism intensity. Both NSGA-II and SPEA2, which also allow dominated individuals to maintain a minimum archive size, seem to make use of this increased diversity in the later stage of the run where they attain a broader distribution and hence better performance values.

4 CONCLUSIONS

In this study we have presented SPEA2, an improved elitist multi-objective evolutionary algorithm that employs an enhanced fitness assignment strategy compared to its predecessor SPEA as well as new techniques for archive truncation and density-based selection. Extensive numerical comparisons of SPEA2 with SPEA and with PESA and NSGA-II, two other recently proposed algorithms, have been carried out on various continuous and combinatorial test problems.

The key results of the comparison are:

- SPEA2 performs better than its predecessor SPEA on all problems.
- PESA has fastest convergence, probably due to its higher elitism intensity, but has difficulties on some problems because it does not always keep the boundary solutions.

- SPEA2 and NSGA-II show the best performance overall.
- In higher dimensional objective spaces, SPEA2 seems to have advantages over PESA and NSGA-II.

REFERENCES

- [1] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computation*, **3**(4), 257–271, (1999).
- [2] E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, **8**(2), 173–195, (2000).
- [3] D. W. Corne, J. D. Knowles, and M. J. Oates. The pareto envelope-based selection algorithm for multiobjective optimisation. In Marc Schoenauer et al., editor, *Parallel Problem Solving from Nature – PPSN VI*, pages 839–848, Berlin. Springer, (2000).
- [4] Michael Lahanas, Natasa Milickovic, Dimos Baltas, and Nikolaos Zamboglou. Application of multiobjective evolutionary algorithms for dose optimization problems in brachytherapy. In E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, and D. Corne, editors, *Proceedings of the First International Conference on Evolutionary Multi-Criterion Optimization (EMO 2001)*, volume 1993 of *Lecture Notes in Computer Science*, pages 574–587, Berlin. Springer-Verlag, (2001).
- [5] K. Deb, S. Agrawal, A. Pratap, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In Marc Schoenauer et al., editor, *Parallel Problem Solving from Nature – PPSN VI*, pages 849–858, Berlin. Springer, (2000).
- [6] Marco Laumanns, Eckart Zitzler, and Lothar Thiele. A unified model for multi-objective evolutionary algorithms with elitism. In *Congress on Evolutionary Computation (CEC-2000)*, pages 46–53, Piscataway, NJ. IEEE, (2000).
- [7] Marco Laumanns, Eckart Zitzler, and Lothar Thiele. On the effects of archiving, elitism, and density based selection in evolutionary multi-objective optimization. In E. Zitzler, K. Deb, L. Thiele, C. A. Coello Coello, and D. Corne, editors, *Proceedings of the First International Conference on Evolutionary Multi-Criterion Optimization (EMO 2001)*, volume 1993 of *Lecture Notes in Computer Science*, pages 181–196, Berlin. Springer-Verlag, (2001).
- [8] Eckart Zitzler, Marco Laumanns, and Lothar Thiele. SPEA2: Improving the strength pareto evolutionary algorithm. Technical Report 103, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich, Gloriastrasse 35, CH-8092 Zurich, Switzerland, May 2001.
- [9] B. W. Silverman. *Density estimation for statistics and data analysis*. Chapman and Hall, London, (1986).
- [10] J. David Schaffer. Multiple objective optimization with vector evaluated genetic algorithms. In John J. Grefenstette, editor, *Proceedings of an International Conference on Genetic Algorithms and Their Applications*, pages 93–100, Pittsburgh, PA. (1985). sponsored by Texas Instruments and U.S. Navy Center for Applied Research in Artificial Intelligence (NCARAI).
- [11] Marco Laumanns, Günther Rudolph, and Hans-Paul Schwefel. Mutation control and convergence in evolutionary multi-objective optimization. In *Proceedings of the 7th International Mendel Conference on Soft Computing (MENDEL 2001)*, Brno, Czech Republic, June 2001.
- [12] D. Quagliarella and A. Vicini. Coupling genetic algorithms and gradient based optimization techniques. In D. Quagliarella, J. Périaux, C. Poloni, and G. Winter, editors, *Genetic Algorithms and Evolution Strategy in Engineering and Computer Science — Recent advances and industrial applications*, pages 289–309. Wiley, Chichester, (1997).
- [13] Frank Kursawe. A variant of evolution strategies for vector optimization. In H.-P. Schwefel and R. Männer, editors, *Parallel Problem Solving from Nature*, pages 193–197, Berlin. Springer, (1991).