Comparison of Multiobjective Evolutionary Algorithms: Empirical Results

Eckart Zitzler[1], Kalyanmoy Deb[2], and Lothar Thiele[1]

[1]Computer Engineering and Networks Laboratory (TIK)
Department of Electrical Engineering
Swiss Federal Institute of Technology (ETH) Zurich
CH-8092 Zürich, Switzerland
E-mail: {zitzler,thiele}@tik.ee.ethz.ch

[2]Kanpur Genetic Algorithms Laboratory (KanGAL)
Department of Mechanical Engineering
Indian Institute of Technology Kanpur
Kanpur, PIN 208 016, India
E-mail: deb@iitk.ac.in

TIK-Report No. 70
Institut für Technische Informatik und Kommunikationsnetze, ETH Zürich
Gloriastrasse 35, ETH-Zentrum, CH–8092 Zürich, Switzerland

February 12, 1999

Abstract

In this paper, we provide a systematic comparison of various evolutionary approaches to multiobjective optimization using six carefully chosen test functions. Each test function involves a particular feature that is known to cause difficulty in the evolutionary optimization process, mainly in converging to the Pareto-optimal front (e.g., multimodality and deception). By investigating these different problem features separately, it is possible to identify the kind of problems to which a certain technique is well suited or not. However, in contrast to what was suspected beforehand, the experimental results indicate a hierarchy of the algorithms under consideration. Furthermore, the emerging effects give evidence that the suggested test functions provide sufficient complexity to compare multiobjective optimizers. Finally, elitism is shown to be an important factor for improving evolutionary multiobjective search.

1 Motivation

Evolutionary algorithms (EAs) have become established as *the* method at hand to explore the Pareto-optimal front in multiobjective optimization problems. This is not only because there are hardly any alternatives in the field of multiobjective optimization; due to their inherent parallelism and their capability to exploit similarities of solutions by crossover, they are able to capture several Pareto-optimal solutions in a single optimization run. The numerous applications and the rapidly growing interest in the area of multiobjective EAs take this fact into account.

After the first pioneering studies on evolutionary multiobjective optimization appeared in the mid-eighties (Schaffer 1984; Schaffer 1985; Fourman 1985), a couple of different EA implementations were proposed in the years 1991–1994 (Kursawe 1991; Hajela and Lin 1992; Fonseca and Fleming 1993; Horn, Nafpliotis, and Goldberg 1994; Srinivas and Deb 1994). Later, these approaches (and variations of them) were successfully applied to various multiobjective optimization problems (Ishibuchi and Murata 1996; Cunha, Oliviera, and Covas 1997; Valenzuela-Rendón and Uresti-Charre 1997; Fonseca and Fleming 1998b; Parks and Miller 1998). In recent years, some researchers have investigated particular topics of evolutionary multiobjective search, such as convergence to the Pareto-optimal front (Veldhuizen and Lamont 1998; Rudolph 1998), niching (Obayashi, Takahashi, and Takeguchi 1998), and elitism (Parks and Miller 1998; Obayashi, Takahashi, and Takeguchi 1998), while others have concentrated on developing new evolutionary techniques (Laumanns, Rudolph, and Schwefel 1998; Zitzler and Thiele 1998a).

In spite of this variety, there is a lack of studies which compare the performance and different aspects of the several approaches. Consequently, the question arises, which implementations are suited to which sort of problem and what are the specific advantages and drawbacks, respectively, of different techniques.

First steps in this direction have been made in both theory and practice. On the theoretical side, Fonseca and Fleming (1995b) discussed the influence of different fitness assignment strategies on the selection process. On the practical side, Zitzler and Thiele (1998b) used a NP-hard 0/1 knapsack problem to compare several multiobjective EAs.

In this paper, we provide a systematic comparison of six multiobjective EAs, including a random search strategy as well as a single-objective EA using objective aggregation. The basis of this empirical study is formed by a set of well-defined, domain-independent test functions which allow the investigation of independent problem features. We thereby draw upon results presented in (Deb 1998), where problem features that may make convergence of EAs to the Pareto-optimal front difficult are identified and, furthermore, methods of constructing appropriate test functions are suggested. The functions considered here cover the range of convexity, non-convexity, discrete Pareto fronts, multimodality, deception, and biased search spaces. Hence, we are able to systematically compare the approaches based on different kinds of difficulty and to determine more exactly where certain techniques are advantageous or have trouble. In this context, we also examine further factors such as population size and elitism.

The paper is structured as follows: Section 2 introduces key concepts of multiobjective optimization and defines the terminology used in this paper mathematically. We then give a brief overview of the multiobjective EAs under consideration with special emphasis on the differences between them. The test functions, their construction and their choice, are the subject of Section 4, which is followed by a discussion about performance metrics to assess the quality of trade-off fronts. Afterwards, we present the experimental results in Section 6 and investigate further aspects like elitism (Section 7) and population size (Section 8) separately. A discussion of the results as well as future perspectives is given in the last chapter.

2 Definitions

Traditionally, optimization problems involving multiple, conflicting objectives are approached by aggregating the objectives into a scalar function and solving the resulting single-objective optimization problem. In contrast, in this study we are concerned with finding a set of optimal trade-offs, the so-called Pareto-optimal front. In the following, we formalize this well-known concept and also define the difference between local and global Pareto-optimal fronts.

A multiobjective search space is partially ordered in the sense that two arbitrary solutions are related to each other in two possible ways: either one dominates the other or neither dominates.

Definition 1 *Let us consider, without loss of generality, a multiobjective minimization problem with m decision variables (parameters) and n objectives:*

$$\begin{array}{ll} \textit{Minimize} & \mathbf{y} = f(\mathbf{x}) = (f_1(\mathbf{x}), \ldots, f_n(\mathbf{x})) \\ \textit{where} & \mathbf{x} = (x_1, \ldots, x_m) \in X \\ & \mathbf{y} = (y_1, \ldots, y_n) \in Y \end{array} \tag{1}$$

and where $\mathbf{x}$ *is called* decision vector*,* X parameter space*,* $\mathbf{y}$ objective vector*, and* Y objective space*. A decision vector* $\mathbf{a} \in X$ *is said to* dominate *a decision vector* $\mathbf{b} \in X$ *(also written as* $\mathbf{a} \succ \mathbf{b}$*) if and only if*

$$\begin{array}{lll} \forall i \in \{1, \ldots, n\}: & f_i(\mathbf{a}) \leq f_i(\mathbf{b}) & \wedge \\ \exists j \in \{1, \ldots, n\}: & f_j(\mathbf{a}) < f_j(\mathbf{b}) & \end{array} \tag{2}$$

Additionally, in this study we say $\mathbf{a}$ covers $\mathbf{b}$ ($\mathbf{a} \succeq \mathbf{b}$) *if and only if* $\mathbf{a} \succ \mathbf{b}$ *or* $f(\mathbf{a}) = f(\mathbf{b})$*.*

Based on the above relation, we can define nondominated and Pareto-optimal solutions:

Definition 2 *Let* $\mathbf{a} \in X$ *be an arbitrary decision vector.*

1. *The decision vector* $\mathbf{a}$ *is said to be* nondominated *regarding a set* $X' \subseteq X$ *if and only if*

$$\nexists \mathbf{a}' \in X' : \mathbf{a}' \succ \mathbf{a} \tag{3}$$

If it is clear within the context which set X' *is meant, we simply leave it out.*

2. *The decision vector* $\mathbf{a}$ *is* Pareto-optimal *if and only if* $\mathbf{a}$ *is nondominated regarding* X*.*

Pareto-optimal decision vectors are not dominated by any other decision vectors in the search space and represent, in our terminology, globally optimal solutions. However, analogous to single-objective optimization problems there may also be local optima which constitute a nondominated set within a certain neighborhood. This corresponds to the concepts of global and local Pareto-optimal sets introduced by Deb (1998):

Definition 3 *Consider a set of decision vectors* $X' \subseteq X$*.*

1. *The set* X' *is denoted as a* local Pareto-optimal set *if and only if*

$$\forall \mathbf{a}' \in X' : \nexists \mathbf{a} \in X : \mathbf{a} \succ \mathbf{a}' \wedge ||\mathbf{a} - \mathbf{a}'|| < \epsilon \tag{4}$$

where $|| \cdot ||$ *is a corresponding distance metric and* $\epsilon > 0$*.*

2. *The set* X' *is called a* global Pareto-optimal set *if and only if*

$$\forall \mathbf{a}' \in X' : \nexists \mathbf{a} \in X : \mathbf{a} \succ \mathbf{a}' \tag{5}$$

Note that a global Pareto-optimal set does not necessarily contain all Pareto-optimal solutions. If we refer to the entirety of the Pareto-optimal solutions, we simply write Pareto-optimal set or front.

3 Evolutionary Multiobjective Optimization

Two major problems must be addressed when an evolutionary algorithm is applied to multiobjective optimization:

1. How to accomplish fitness assignment and selection, respectively, in order to guide the search towards the Pareto-optimal set.
2. How to maintain a diverse population in order to prevent premature convergence and achieve a well distributed trade-off front.

Concerning the first issue, one can distinguish between non-Pareto and Pareto-based approaches. The last makes direct use of the dominance relation from Definition 1; Goldberg (1989) was the first to suggest a Pareto-based fitness assignment strategy. The second issue plays a central role in multimodal search in general. Fitness sharing (Goldberg and Richardson 1987) is most commonly used; other techniques, e.g., crowding and its derivates, are only rarely incorporated in multiobjective EAs.

Among the two non-Pareto approaches considered in this study is the Vector Evaluated Genetic Algorithm (Schaffer 1984; Schaffer 1985). Here, selection is done for each objective separately, filling equally sized proportions of the mating pool. The mating pool is shuffled before crossover and mutation are performed. No explicit niching mechanism is used in this implementation.

Hajela and Lin (1992) proposed another non-Pareto algorithm which is based on weighted-sum aggregation, where an individual is assessed by summing up the weighted objective values. Since the corresponding linear combinations are directly encoded in the chromosome, each individual may be evaluated regarding a different weight combination. Moreover, diversity among the linear combinations is achieved by fitness sharing, enabling the parallel evolution of families of solutions.

Furthermore, three Pareto-based EAs are investigated which belong to the first implementations using Pareto-dominance for fitness assignment. Additionally, a recent Pareto-based approach is included in the comparison.

Fonseca and Fleming (1993) proposed a ranking procedure, where an individual's rank equals the number of population members by which it is dominated. Reproduction probabilities are determined by means of exponential ranking, and afterwards the raw fitness values are averaged and shared among individuals having identical ranks. Finally, stochastic universal sampling (SUS) is used to fill the mating pool. The basic concept has been extended meanwhile by, e.g., adaptive fitness sharing and continuous introduction of random immigrants (Fonseca and Fleming 1995a; Fonseca and Fleming 1998a), which is, however, not regarded in this study.

Another concept which combines binary tournament selection and Pareto-dominance is the Niched Pareto Genetic Algorithm presented in (Horn and Nafpliotis 1993; Horn, Nafpliotis, and Goldberg 1994). Two competing individuals and a comparison set of t_{dom} individuals are picked at random from the population. If one of the competing individuals is dominated by any member of the set and the other is not, then the latter is chosen as the winner of the tournament. Otherwise, the result of the tournament is decided by fitness sharing.

In their Nondominated Sorting Genetic Algorithm, Srinivas and Deb (1994) introduced the concept of nondomination level. The idea behind it is figuratively speaking to peel off the different trade-off fronts in the population step by step. The first nondominated front extracted belongs to nondomination level one, the nondominated set of the remaining population members constitutes the next level, and so forth. Accordingly, the fitness of an individual relates to its nondomination level, where level one has highest reproduction probability. In addition, fitness is shared within each level of nondomination. However, as opposed to the other Pareto-based algorithms under consideration, Srinivas' and Deb's implementation uses fitness sharing on the parameter space instead of the objective space.

The fourth of the Pareto-based strategies, the Strength Pareto Evolutionary Algorithm (Zitzler and Thiele 1998a), incorporates ideas from studies on immune systems and coevolution by maintaining a second, external population. This set, which contains the nondominated solutions found so far, is continuously updated and, in case its size exceeds a given maximum, reduced by clustering. Fitness is assigned to individuals in both population and external set interdependently. The fitness of a member of the external set is antiproportional to the number of individuals in the population which it dominates. Vice versa, the fitness of a population member is antiproportional to the sum of the fitness values of those external solutions by which it is dominated. Finally, individuals from the population as well as the external set take part in the selection process.

For a thorough discussion of different evolutionary approaches to multiobjective optimization, the interested reader is referred to (Fonseca and Fleming 1995b; Tamaki, Kita, and Kobayashi 1996; Horn 1997).

4 Test Functions for Multiobjective Optimizers

Deb (1998) has identified several features which may cause difficulties for multiobjective EAs in i) converging to the Pareto-optimal front and ii) maintaining diversity within the population. Concerning the first issue, multimodality, deception, and isolated optima are well-known problem areas in single-objective evolutionary optimization. The second issue is important in order to achieve a well distributed nondominated front. However, certain characteristics of the Pareto-optimal front may prevent an EA from finding diverse Pareto-optimal solutions: convexity or non-convexity, discreteness, and non-uniformity. For each of the six problem features mentioned a corresponding test function is constructed following the guidelines in (Deb 1998). We thereby restrict ourselves to only two objectives, in order to investigate the simplest case first. In our opinion, two-dimensional problems already reflect essential aspects of multiobjective optimization. Moreover, we do not consider maximization or mixed minimization/maximization problems.

Each of the test functions defined below is structured in the same manner and consists itself of three functions f_1, g, h (Deb 1998, p.15):

$$\begin{array}{lrcl} \text{Minimize} & \mathcal{T}(\mathbf{x}) & = & (f_1(x_1), f_2(\mathbf{x})) \\ \text{subject to} & f_2(\mathbf{x}) & = & g(x_2, \ldots, x_m) h(f_1(x_1), g(x_2, \ldots, x_m)) \\ \text{where} & \mathbf{x} & = & (x_1, \ldots, x_m) \end{array} \tag{6}$$

The function f_1 is a function of the first decision variable only, g is a function of the remaining $m-1$ variables, and h takes the function values of f_1 and g. The test functions differ in these three functions as well as in the number of variables m and in the values the variables may take.

Definition 4 *We introduce six test functions $\mathcal{T}_1, \ldots, \mathcal{T}_6$ that follow the scheme given in Equation 6:*

- *The test function $\mathcal{T}_1$ has a convex Pareto-optimal front; the solutions are uniformly distributed in the search space:*

$$\begin{array}{lcl} f_1(x_1) & = & x_1 \\ g(x_2, \ldots, x_m) & = & 1 + 9 \cdot \sum_{i=2}^{m} x_i \\ h(f_1, g) & = & 1 - \sqrt{f_1/g} \end{array} \tag{7}$$

  where $m = 30$ and $x_i \in [0, 1]$. The Pareto-optimal front is formed with $g(\mathbf{x}) = 1$.

- *The test function $\mathcal{T}_2$ is the non-convex counterpart to $\mathcal{T}_1$:*

$$\begin{array}{lcl} f_1(x_1) & = & x_1 \\ g(x_2, \ldots, x_m) & = & 1 + 9 \cdot \sum_{i=2}^{m} x_i \\ h(f_1, g) & = & 1 - (f_1/g)^2 \end{array} \tag{8}$$

where $m = 30$ *and* $x_i \in [0,1]$. *The Pareto-optimal front is formed with* $g(\mathbf{x}) = 1$.

- *The test function* $\mathcal{T}_3$ *represents the discreteness features; its Pareto-optimal front consists of several non-contiguous convex parts, the search space is unbiased:*

$$\begin{array}{lcl} f_1(x_1) & = & x_1 \\ g(x_2,\ldots,x_m) & = & 1 + 9 \cdot \sum_{i=2}^{m} x_i \\ h(f_1,g) & = & 1 - \sqrt{f_1/g} - (f_1/g)\sin(10\pi f_1) \end{array} \tag{9}$$

where $m = 30$ *and* $x_i \in [0,1]$. *The Pareto-optimal front is formed with* $g(\mathbf{x}) = 1$. *The introduction of the sine function in* h *causes discontinuity in the Pareto-optimal front. However, there is no discontinuity in the search space.*

- *The test function* $\mathcal{T}_4$ *contains* 21^9 *local Pareto-optimal fronts and therefore tests for the EA's ability to deal with multimodality:*

$$\begin{array}{lcl} f_1(x_1) & = & x_1 \\ g(x_2,\ldots,x_m) & = & 1 + 10(m-1) + \sum_{i=2}^{m}(x_i^2 - 10\cos(4\pi x_i)) \\ h(f_1,g) & = & 1 - \sqrt{f_1/g} \end{array} \tag{10}$$

where $m = 10$, $x_1 \in [0,1]$ *and* $x_2,\ldots,x_m \in [-5,5]$. *The global Pareto-optimal front is formed with* $g(\mathbf{x}) = 1$ *, the best local Pareto-optimal front with* $g(\mathbf{x}) = 1.25$. *Note that not all local Pareto-optimal fronts are distinguishable in the objective space.*

- *The test function* $\mathcal{T}_5$ *describes a deceptive problem and distinguishes itself from the other test functions in that* x_i *represents a binary string:*

$$\begin{array}{lcl} f_1(x_1) & = & 1 + u(x_1) \\ g(x_2,\ldots,x_m) & = & \sum_{i=2}^{m} v(u(x_i)) \\ h(f_1,g) & = & 1/f_1 \end{array} \tag{11}$$

where $u(x_i)$ *gives the number of ones in the bit vector* x_i *(unitation),*

$$v(u(x_i)) = \left\{ \begin{array}{ll} 2 + u(x_i) & \textit{if}\;\; u(x_i) < 5 \\ 1 & \textit{if}\;\; u(x_i) = 5 \end{array} \right\}.$$

and $m = 11$, $x_1 \in \{0,1\}^{30}$ *and* $x_2,\ldots,x_m \in \{0,1\}^5$. *The true Pareto-optimal front is formed with* $g(\mathbf{x}) = 10$, *while the best deceptive Pareto-optimal front includes all solutions for which* $g(\mathbf{x}) = 11$. *The global Pareto-optimal front as well as the local ones are convex.*

- *The test function* $\mathcal{T}_6$ *includes two difficulties caused by the non-uniformity of the search space: firstly, the Pareto-optimal solutions are non-uniformly distributed along the global Pareto front (the front is biased for solutions for which* $f_1(\mathbf{x})$ *is near one); secondly, the density of the solutions is least near the Pareto-optimal front and highest away from the front:*

$$\begin{array}{lcl} f_1(x_1) & = & 1 - \exp(-4x_1)\sin^6(6\pi x_1) \\ g(x_2,\ldots,x_m) & = & 1 + 9 \cdot (\sum_{i=2}^{m} x_i)^{0.25} \\ h(f_1,g) & = & 1 - (f_1/g)^2 \end{array} \tag{12}$$

where $m = 10$, $x_i \in [0,1]$. *The Pareto-optimal front is formed with* $g(\mathbf{x}) = 1$ *and is convex.*

We will discuss each function in more detail in Section 6, where the corresponding Pareto-optimal fronts are visualized as well.

5 Metrics of Performance

Comparing different optimization techniques experimentally always involves the notion of performance. In the case of multiobjective optimization, the definition of quality is substantially more complex than for single-objective optimization problems, because the optimization goal itself consists of multiple objectives:

- The distance of the resulting nondominated set to the Pareto-optimal front should be minimized.
- A good (in most cases uniform) distribution of the solutions found is desirable. The assessment of this criterion might be based on a certain distance metric.
- The extent of the obtained nondominated front should be maximized, i.e., for each objective a wide range of values should be covered by the nondominated solutions.

In the literature, some attempts can be found to formalize the above definition (or parts of it) by means of quantitative metrics. Performance assessment by means of weighted-sum aggregation was introduced by Esbensen and Kuh (1996). Thereby, a set X' of decision vectors is evaluated regarding a given linear combination by determining the minimum weighted-sum of all corresponding objective vectors of X'. Based on this concept, a sample of linear combinations is chosen at random (with respect to a certain probability distribution) and the minimum weighted-sums for all linear combinations are summed up and averaged. The resulting value is taken as a measure of quality. A drawback of this metric is that only the "worst" solution determines the quality value per linear combination. Although several weight combinations are used, non-convex regions of the trade-off surface contribute to the quality more than convex parts and may, as a consequence, dominate the performance assessment. Finally, the distribution as well as the extent of the nondominated front is not considered.

Another interesting way of performance assessment was proposed by Fonseca and Fleming (1996). Given a set $X' \subseteq X$ of nondominated points, a boundary function divides the search space into two regions: the points not dominated by or equal to members of X' and the points covered by X'. They call this particular function, which can also be seen as the locus of the family of tightest goal vectors known to be attainable, the attainment surface. Taking multiple optimization runs into account, a method is described to compute an average attainment surface by using auxiliary straight lines and sampling their intersections with the attainment surfaces obtained. As a result, the samples represented by the average attainment surface can be assessed relatively by means of statistical tests and therefore allow comparison of the performance of two or more multiobjective optimizers. Drawbacks of this approach are that i) the characteristics of the distribution are falsified by the construction of the average attainment surface and ii) it remains unclear how the quality difference can be expressed, i.e., how much better algorithm A is than algorithm B. However, Fonseca and Fleming describe ways of meaningful statistical interpretation in contrast to the other studies considered here, and furthermore, their methodology seems to be well suited for visualization of the outcomes of several runs.

In the context of investigations on convergence to the Pareto-optimal front, some authors (Rudolph 1998; Veldhuizen and Lamont 1998) have considered the distance of a given set to the Pareto-optimal set in the same way as the function $\mathcal{M}_1$ defined below. The distribution was not taken into account, because the focus was not on this matter. However, in comparative studies distance alone is not sufficient for performance evaluation, since extremely differently distributed fronts may have the same distance to the Pareto-optimal front.

Two complementary metrics of performance were presented in (Zitzler and Thiele 1998a; Zitzler and Thiele 1998b). On the one hand, the size of the dominated area in the objective space is taken under consideration; on the other hand, a pair of nondominated sets is compared by calculating the fraction of each set that is covered by the other set.

The area combines all three criteria (distance, distribution, and extent) into one, and therefore sets differing in more than one criterion may not be distinguished. Also, points located in convex regions may be overrated. The second metric is in some way similar to the comparison methodology proposed in (Fonseca and Fleming 1996). It can be used to show that the outcomes of an algorithm dominate the outcomes of another algorithm, although, it does not tell how much better it is. We give its definition here, because it is used in the remainder of this paper.

Definition 5 *Let $X', X'' \subseteq X$ be two sets of decision vectors. The function $\mathcal{C}$ maps the ordered pair (X',X") to the interval [0,1]:*

$$\mathcal{C}(X', X'') := \frac{|\{\mathbf{a''} \in X''; \exists\, \mathbf{a'} \in X' : \mathbf{a'} \succeq \mathbf{a''}\}|}{|X''|} \tag{13}$$

The value $\mathcal{C}(X', X'') = 1$ means that all points in X'' are dominated by or equal to points in X'. The opposite, $\mathcal{C}(X', X'') = 0$, represents the situation when none of the points in X'' are covered by the set X'. Note that both $\mathcal{C}(X', X'')$ and $\mathcal{C}(X'', X')$ have to be considered, since $\mathcal{C}(X', X'')$ is not necessarily equal to $\mathcal{C}(X'', X')$ (e.g., if X' dominates X'' then $\mathcal{C}(X', X'') = 1$ and $\mathcal{C}(X'', X') = 0$).

In summary, it may be said that performance metrics are hard to define and probably it will not be possible to define a single metric which allows for all criteria in a meaningful way. Along with that problem, the statistical interpretation associated with a performance comparison is costly and still needs to be answered, since multiple significance tests are involved and thus tools from analysis of variance may be required.

In this study, we have chosen a visual presentation of the results together with the application of the metric from Definition 5. The reason for this is that we would like to investigate i) whether test functions can adequately test specific aspects of each multi-objective algorithm and ii) whether any visual hierarchy of the chosen algorithms exists. However, for a deeper investigation of some of the algorithms (which is the subject of future work), we suggest the following metrics which allow assessment of each of the criteria listed at the beginning of this section separately.

Definition 6 *Given a set of pairwise nondominating decision vectors $X' \subseteq X$, a neighborhood parameter $\sigma > 0$ (to be chosen appropriately), and a distance metric $||\cdot||$. We introduce three functions to assess the quality of X' regarding the parameter space:*

1. *The function $\mathcal{M}_1$ gives the average distance to the Pareto-optimal set $\overline{X} \subseteq X$:*

$$\mathcal{M}_1(X') := \frac{1}{|X'|} \sum_{\mathbf{a'} \in X'} \min\{||\mathbf{a'} - \overline{\mathbf{a}}||; \overline{\mathbf{a}} \in \overline{X}\} \tag{14}$$

2. *The function $\mathcal{M}_2$ takes the distribution in combination with the number of nondominated solutions found into account:*

$$\mathcal{M}_2(X') := \frac{1}{|X' - 1|} \sum_{\mathbf{a'} \in X'} |\{\mathbf{b'} \in X'; ||\mathbf{a'} - \mathbf{b'}|| < \sigma\}| \tag{15}$$

3. *The function $\mathcal{M}_3$ considers the extent of the front described by X':*

$$\mathcal{M}_3(X') := \sqrt{\sum_{i=1}^{m} \max\{||a'_i - b'_i||; \mathbf{a'}, \mathbf{b'} \in X'\}} \tag{16}$$

Analogous, we define three metrics $\mathcal{M}_1^*$, $\mathcal{M}_2^*$, and $\mathcal{M}_3^*$ on the objective space. Let $Y', \overline{Y} \subseteq Y$ be the sets of objective vectors that correspond to X' and $\overline{X}$ respectively, and $\sigma^* > 0$

and $||\cdot||^*$ *as before:*

$$\mathcal{M}_1^*(Y') := \frac{1}{|Y'|} \sum_{\mathbf{p'} \in Y'} \min\{||\mathbf{p'} - \overline{\mathbf{p}}||^*; \overline{\mathbf{p}} \in \overline{Y}\} \tag{17}$$

$$\mathcal{M}_2^*(Y') := \frac{1}{|Y' - 1|} \sum_{\mathbf{p'} \in Y'} |\{\mathbf{q'} \in Y'; ||\mathbf{p'} - \mathbf{q'}||^* < \sigma^*\}| \tag{18}$$

$$\mathcal{M}_3^*(Y') := \sqrt{\sum_{i=1}^{n} \max\{||p'_i - q'_i||^*; \mathbf{p'}, \mathbf{q'} \in Y'\}} \tag{19}$$

While $\mathcal{M}_1$ and $\mathcal{M}_1^*$ are intuitive, $\mathcal{M}_2$ and $\mathcal{M}_3$ (respectively $\mathcal{M}_2^*$ and $\mathcal{M}_3^*$) need further explanation. The distribution metrics give a value within the interval $[0, 1]$; it reflects the average fraction of members of X' (Y') which lie outside the neighborhood of an arbitrary solution in X' (Y'). Obviously, the higher the value the better the distribution for an appropriate neighborhood parameter (e.g., $\mathcal{M}_2^*(X') = 1$ means that for each objective vector there is no other objective vector within σ^*-distance to it). The functions $\mathcal{M}_3$ and $\mathcal{M}_3^*$ use the maximum extent in each dimension to estimate the range to which the fronts spreads out. In the case of two objectives, this equals the distance of the two outer solutions.

6 Comparison of Different Evolutionary Approaches

6.1 Methodology

We compare eight algorithms on the six proposed test functions:

1. RAND: A random search algorithm.
2. FFGA: Fonseca's and Fleming's multiobjective EA (1993).
3. NPGA: The Niched Pareto Genetic Algorithm (Horn, Nafpliotis, and Goldberg 1994).
4. HLGA: Hajela's and Lin's weighted-sum based approach (1992).
5. VEGA: Vector Evaluated Genetic Algorithm (Schaffer 1985).
6. NSGA: The Nondominated Sorting Genetic Algorithm (Srinivas and Deb 1994).
7. SOEA: A single-objective evolutionary algorithm using weighted-sum aggregation.
8. SPEA: The Strength Pareto Evolutionary Algorithm (Zitzler and Thiele 1998a).

The multiobjective EAs as well as RAND were executed 30 times on each test problem, where the population was monitored for nondominated solutions and the resulting nondominated set was taken as the outcome of one optimization run. Here, RAND serves as an additional point of reference and randomly generates a certain number of individuals per generation, according to the rate of crossover and mutation (but neither crossover and mutation nor selection are performed). Hence, the number of fitness evaluations was the same as for the EAs. In contrast, 100 simulation runs were considered in case of SOEA, each run optimizing towards another randomly chosen linear combination of the objectives. The nondominated solutions among all solutions generated in the 100 runs form the trade-off front achieved by SOEA on a particular test function.

Independent of the algorithm and the test function, each simulation run was carried out using the following parameters:

Number of generations	:	250
Population size	:	100
Crossover rate	:	0.8
Mutation rate	:	0.01
Niching parameter σ_{share}	:	0.48862
Domination pressure t_{dom}	:	10

The niching parameter was calculated using the guidelines given in (Deb and Goldberg 1989) assuming the formation of ten independent niches. Since NSGA uses genotypic fitness sharing on $\mathcal{T}_5$, a different value $\sigma_{\text{share}} = 34$ was chosen for this particular case. Concerning NPGA, the recommended value for $t_{\text{dom}} = 10\%$ of the population size was taken (Horn and Nafpliotis 1993). Furthermore, for reasons of fairness SPEA ran with a population size of 80 where the external nondominated set was restricted to 20.

Regarding the implementations of the algorithms, one chromosome was used to encode the m parameters of the corresponding test problem. Each parameter is represented by 30 bits; the parameters $x_2, \ldots, x_m$ only comprise 5 bits for the deceptive function $\mathcal{T}_5$. Moreover, all approaches except FFGA were realized using binary tournament selection with replacement, in order to avoid effects cause by different selection schemes. Furthermore, since fitness sharing may produce chaotic behavior in combination with tournament selection, a slightly modified method is incorporated here, named *continuously updated sharing* (Oei, Goldberg, and Chang 1991). In FFGA, the originally proposed stochastic universal sampling is employed, because fitness assignment is closely related to this particular selection algorithm.

6.2 Simulation Results

In Figures 1 to 6, the nondominated fronts achieved by the different algorithms are visualized. Per algorithm and test function, the outcomes of the first five runs were unified, and then the dominated solutions were removed from the union set; the remaining points are plotted in the figures. Also shown are the Pareto-optimal fronts (lower curves) as well as additional reference curves (upper curves). The latter curves allow a more precise evaluation of the obtained trade-off fronts and were calculated by adding $0.1 \cdot |\max\{f_2(\mathbf{x})\} - \min\{f_2(\mathbf{x})\}|$ to the f_2 values of the Pareto-optimal points. The space between Pareto-optimal and reference fronts represents about 10% of the corresponding objective space. However, the curve resulting for the deceptive function $\mathcal{T}_5$ is not appropriate for our purposes, since it lies above the fronts produced by the random search algorithm. Instead, we consider all solutions with $g(\mathbf{x}) = 2 \cdot 10$, i.e., for which the parameters are set to the deceptive attractors.

In addition to the graphical presentation, the different algorithms were assessed in pairs using the $\mathcal{C}$ metric from Definition 5. For an ordered algorithm pair (A_1, A_2), there is a sample of 30 $\mathcal{C}$ values according to the 30 runs performed. Each value is computed on the basis of the nondominated sets achieved by A_1 and A_2 with the same initial population. Here, *box plots* (Chambers, Cleveland, Kleiner, and Tukey 1983) are used to visualize the distribution of these samples (Figure 7). A box plot consists of a box summarizing 50% of the data. The upper and lower ends of the box are the upper and lower quartiles, while a thick line within the box encodes the median. Dashed appendages summarize the spread and shape of the distribution.[1] Furthermore, the shortcut REFS in Figure 7 stands for reference set and represents for each test function a set of 100 equidistant points which are uniformly distributed on the corresponding reference curve.

Generally, the simulation results prove that all multiobjective EAs do better than the random search algorithm. However, the box plots reveal that HLGA, NPGA, and FFGA do not always cover the randomly created trade-off front completely. Furthermore, it

[1]Note that outside values are not plotted in Figure 7 in order to prevent overloading of the presentation.

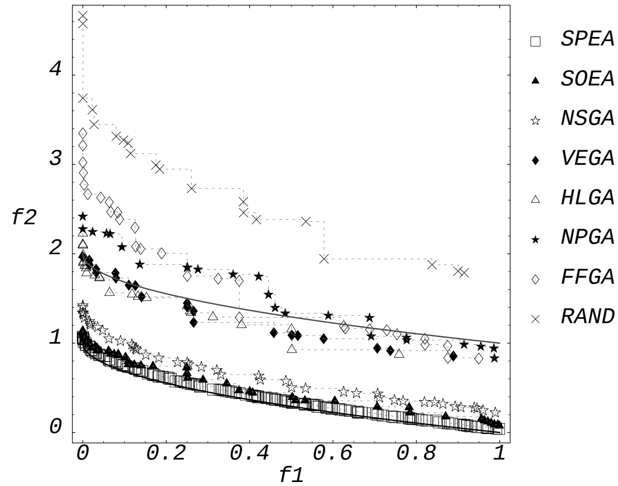

Figure 1: Test function $\mathcal{T}_1$ (convex)

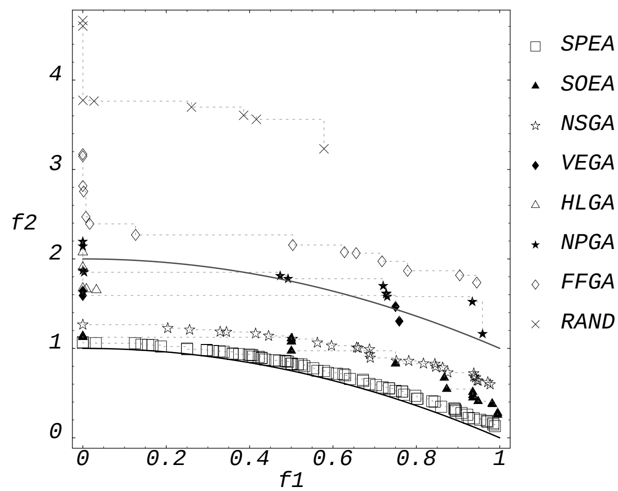

Figure 2: Test function $\mathcal{T}_2$ (non-convex)

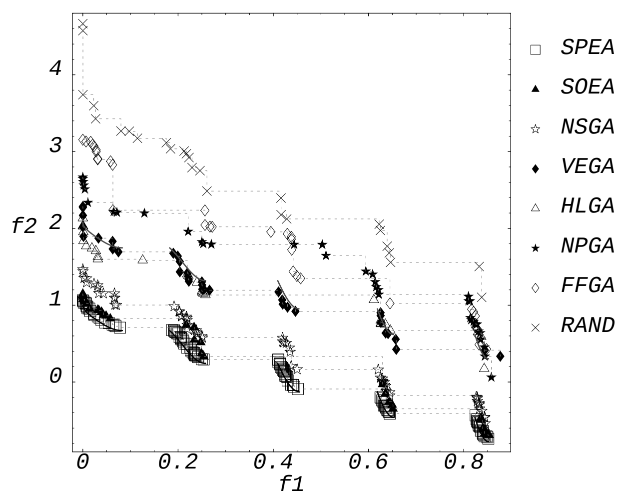

Figure 3: Test function $\mathcal{T}_3$ (discrete)

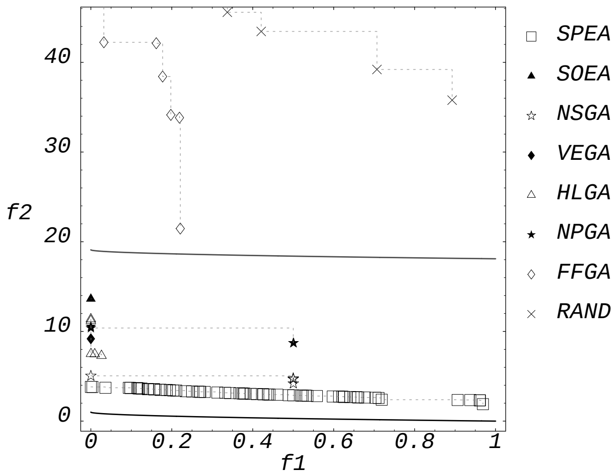

Figure 4: Test function $\mathcal{T}_4$ (multimodal)

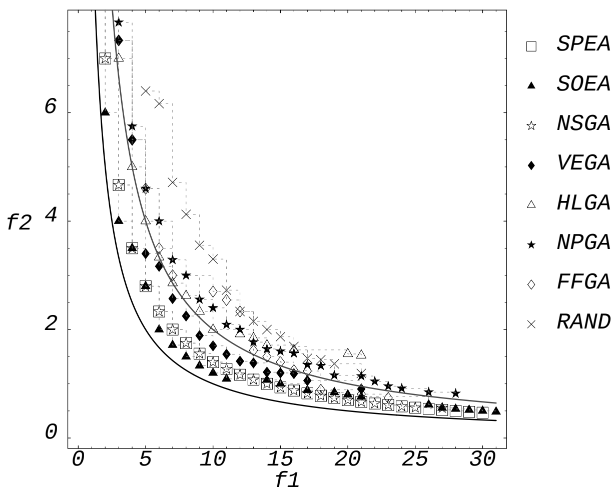

Figure 5: Test function $\mathcal{T}_5$ (deceptive)

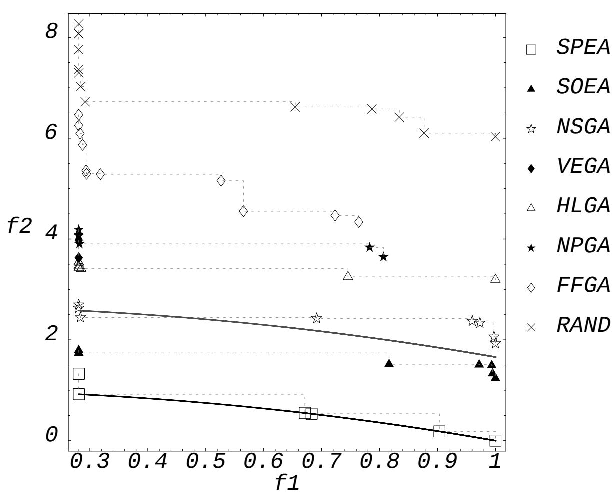

Figure 6: Test function $\mathcal{T}_6$ (non-uniform)

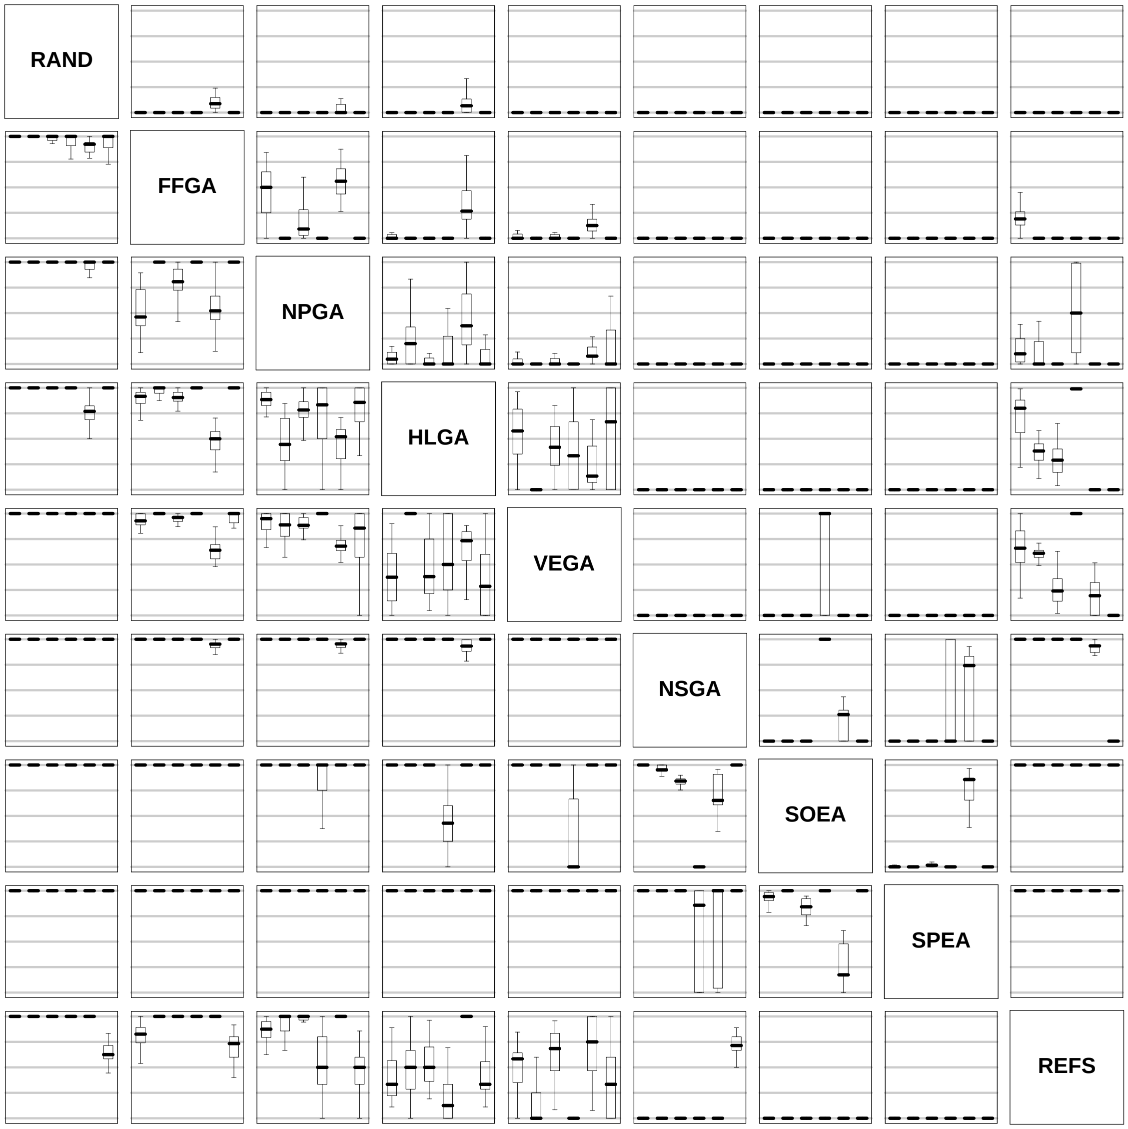

Figure 7: Box plots based on the $\mathcal{C}$ metric. Each rectangle contains six box plots representing the distribution of the $\mathcal{C}$ values for a certain ordered pair of algorithms; the leftmost box plot relates to $\mathcal{T}_1$, the rightmost to $\mathcal{T}_6$. The scale is 0 at the bottom and 1 at the top per rectangle. Furthermore, each rectangle refers to algorithm A associated with the corresponding row and algorithm B associated with the corresponding column and gives the fraction of B covered by A ($\mathcal{C}(A, B)$).

can be observed that NSGA clearly outperforms the other non-elitist multiobjective EAs regarding both distance to the Pareto-optimal front and distribution of the nondominated solutions. This confirms the results presented in (Zitzler and Thiele 1998b). Furthermore, it is remarkable that VEGA performs well compared to NPGA and FFGA, although some serious drawbacks of this approach are known (Fonseca and Fleming 1995b). The reason for this might be that we consider the off-line performance here in contrast to other studies which examine the on-line performance (Horn and Nafpliotis 1993; Srinivas and Deb 1994). Finally, the best performance is provided by SPEA, which makes explicit use of the concept of elitism. Apart from $\mathcal{T}_5$, it even outperforms SOEA, in spite of substantially lower computational effort and although SOEA uses an elitist strategy as well. This observation leads to the question of whether elitism would increase the performance of the other multiobjective EAs. We will investigate this matter in the next section.

Considering the different problem features separately, convexity seems to cause the least amount of difficulty for the multiobjective EAs. All algorithms evolved reasonably distributed fronts, although there was a difference in the distance to the Pareto-optimal set. On the non-convex test function $\mathcal{T}_2$, however, HLGA, VEGA, and SOEA have difficulties finding intermediate solutions, as linear combinations of the objectives tend to prefer solutions strong in at least one objective (Fonseca and Fleming 1995b, p.4). Pareto-based algorithms have advantages here, but only NSGA and SPEA evolved a sufficient number of nondominated solutions. In the case of $\mathcal{T}_3$ (discreteness), HLGA and VEGA are superior to both FFGA and NPGA. While the fronts achieved by the former cover about 25% of the reference set on average, the latter come up with 0% coverage. Among the considered test functions, $\mathcal{T}_4$ and $\mathcal{T}_5$ seem to be the hardest problems, since none of the algorithms was able to evolve a global Pareto-optimal set. The results on the multimodal problem indicate that elitism is helpful here; SPEA is the only algorithm which found a widely distributed front. Remarkable is also that NSGA and VEGA outperform SOEA on $\mathcal{T}_4$. Again, the comparison with the reference set reveals, that HLGA and VEGA (100% coverage) surpass NPGA (50% coverage) and FFGA (0% coverage). Concerning the deceptive function, SOEA is best, followed by SPEA and NSGA. Among the remaining EAs, VEGA appears to be preferable here, covering about 20% of the reference set, while the others cover 0% in all runs. Finally, it can be observed that the biased search space together with the non-uniform represented Pareto-optimal front ($\mathcal{T}_6$) makes it difficult for the EAs to evolve a well-distributed nondominated set. This also affects the distance to the global optimum, as even the fronts produced by NSGA do not cover the points in the reference set.

7 Elitism in Multiobjective Search

As opposed to single-objective optimization, where the best solution is always copied into the next population, the incorporation of elitism in multiobjective EAs is substantially more complex. Instead of one best solution, we have here an elite set whose size can be considerable compared to the population. This fact involves two questions which must be answered in this context:

- **Population $\Longrightarrow$ Elite Set:**
  Which solutions are kept for how long in the elite set?
- **Elite Set $\Longrightarrow$ Population:**
  When and how are which members of the elite set re-inserted into the population?

Often used is the concept of maintaining an external set of solutions which are nondominated among all individuals generated so far. In each generation, a certain percentage of the population is filled up or replaced by members of the external set—these members are either selected at random (Ishibuchi and Murata 1996) or according to other criteria, such

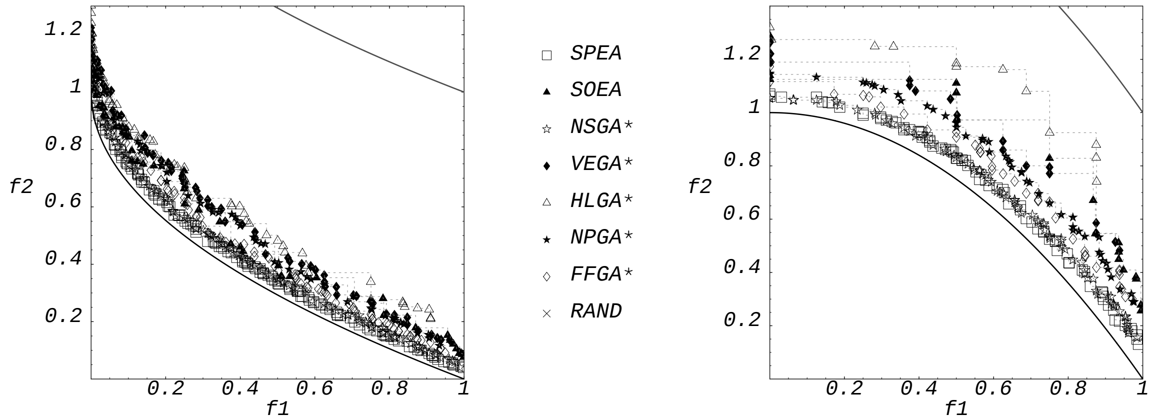

Figure 8: Results on the test functions $\mathcal{T}_1$ (left) and $\mathcal{T}_2$ (right) using elitism.

as the period that an individual has stayed in the set (Parks and Miller 1998). Another promising way of elitism provides the so-called $(\lambda + \mu)$ selection mainly used in the area of evolutionary strategies (Bäck 1996), where parents and offspring compete against each other. Rudolph (1998) examines a simplified version of a multiobjective EA originally presented in (Kursawe 1991) which is based on (1+1) selection.

In this study, the elitism mechanism proposed in (Zitzler and Thiele 1998a) was generalized and implemented in FFGA, NPGA, HLGA, VEGA, and NSGA as follows: Let P denote the current population of size N and $\overline{P}$ a second, external population which keeps the nondominated solutions found so far; the size of $\overline{P}$ is restricted to $\overline{N}$.

Step 1: Generate the initial population P and set $\overline{P} = \emptyset$.

Step 2: Set $P' = P + \overline{P}$ (multi-set union) and perform fitness assignment on the extended population P' of size $N' = N + \overline{N}$.

Step 3: Update external population by copying all nondominated members of P to $\overline{P}$ and afterwards removing double or dominated individuals from $\overline{P}$.

Step 4: If $|\overline{P}| > \overline{N}$ then calculate reduced nondominated set $\overline{P_r}$ of size $\overline{N}$ by clustering and set $\overline{P} = \overline{P_r}$.

Step 5: Select N individuals out of the N' individuals in P' and perform crossover and mutation to create the next population P''.

Step 6: Substitute P by P'' and go to Step 2 if the maximum number of generations is not reached.

The elitism variants of the algorithms are marked by an asterisk in order to distinguish them from the techniques originally proposed by the corresponding authors. Note that the clustering procedure in Step 4 requires a distance metric. In case of NSGA*, the phenotypic distance on the parameter space was considered, while the other algorithms used the phenotypic distance on the objective space.

The results for $\mathcal{T}_1$ and $\mathcal{T}_2$ are shown in Figure 8.[2] Obviously, elitism is helpful on these two functions, although the visual presentation has to be interpreted with care as only five runs are considered. For instance, NSGA* and SPEA seem to perform equally well here using those particular parameter settings. Moreover, the figures indicate that

[2]The experiments were performed as described in Section 6; however, N was set to 80 and $\overline{N}$ to 20, similar to SPEA.

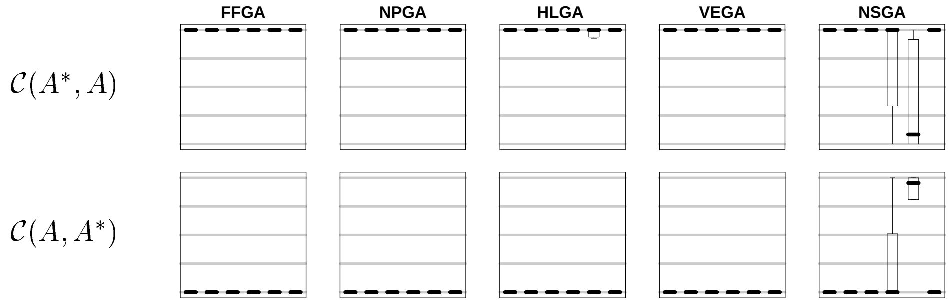

Figure 9: Box plots comparing each non-elitism algorithm A with its elitism-variant A^*.

elitism can even help multiobjective EAs to surpass the performance of a weighted-sum single-objective EA in spite of significantly lower computational effort. However, both test functions and the metric used are not sufficient here to also compare the elitist variants with each other. Testing different elitist strategies and different elitist multiobjective EAs on more difficult test functions will be the subject of future work.

Nevertheless, we have compared each algorithm with its elitist variant based on the $\mathcal{C}$ metric. As can be seen in Figure 9, elitism appears to be an important factor to improve evolutionary multiobjective optimization. Only in one case (NSGA on the deceptive problem) was the performance of the elitist variant worse than the non-elitist version. Investigation of this matter will also be an important part of an elitism study.

8 Influence of the Population Size

On two test functions ($\mathcal{T}_4$ and $\mathcal{T}_5$), none of the algorithms under consideration was able to find a global Pareto-optimal set regarding the chosen parameters. Therefore, several runs were performed in order to investigate the influence of the population size as well as the maximum number of generations converging towards the Pareto-optimal front.

In Figure 10, the outcomes of multiple NSGA runs are visualized. On the deceptive test function $\mathcal{T}_4$, NSGA found a subset of the globally optimal solutions using a population size of 1000. In contrast, $\mathcal{T}_5$ seems to be a difficult test problem, since even a population size of 10000 was not sufficient to converge to the optimal trade-off front after 250 generations. This did also not change when the maximum number of generations was increased substantially (10000). In the later case, the resulting front was (using a population size of 500) almost identical to the one achieved by NSGA* running 1000 generations. However, the incorporation of elitism finally enabled NSGA to find a global Pareto-optimal set after 10000 generations.

To sum up, one may say that the choice of the population size strongly influences the EA's capability to converge towards the Pareto-optimal front. Obviously, small populations do not provide enough diversity among the individuals. Increasing the population size, however, does not automatically yield an increase in performance, as can be observed with the multimodal function. The same holds for the number of generations to be simulated. Elitism, on the other hand, seems to be an appropriate technique to prevent premature convergence. Even after 1000 generations, better solutions, and finally Pareto-optimal solutions, evolved with $\mathcal{T}_4$.

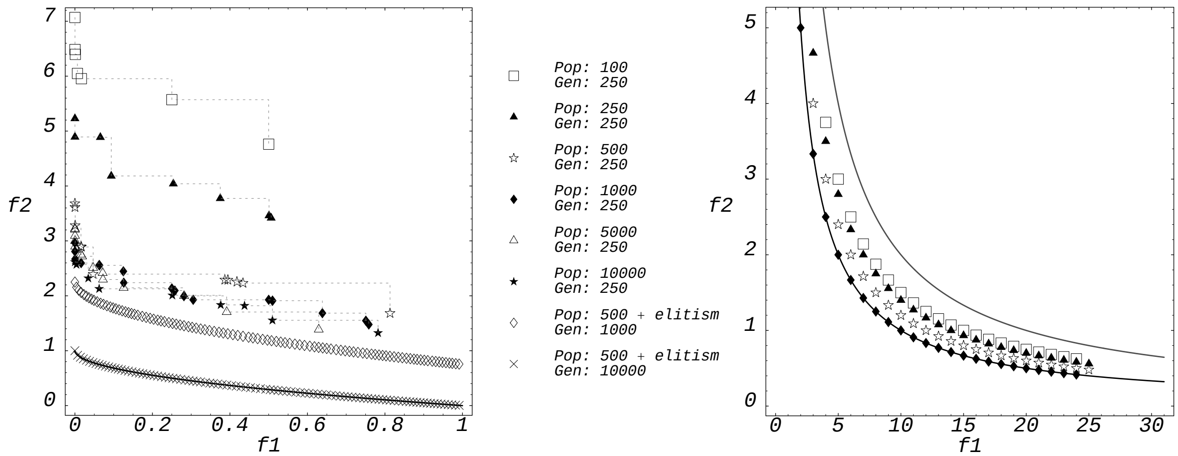

Figure 10: Comparison of different population sizes on the test functions $\mathcal{T}_4$ (left) and $\mathcal{T}_5$ (right) using NSGA. On $\mathcal{T}_4$, two runs with elitism were performed for 1000 and 10000 generations.

9 Conclusions

We have carried out a systematic comparison of several multiobjective EAs on six different test functions. Major results are:

- The suggested test functions provide sufficient complexity to compare different multiobjective optimizers. Multimodality and deception seem to cause the most difficulty for evolutionary approaches. However, non-convexity is also a problem feature which mainly weighted-sum based algorithms appear to have problems with.
- A clear hierarchy of algorithms emerges regarding the distance to the Pareto-optimal front in descending order of merit:
  1. SPEA (Zitzler and Thiele 1998a).
  2. NSGA (Srinivas and Deb 1994).
  3. VEGA (Schaffer 1985).
  4. HLGA (Hajela and Lin 1992)
  5. NPGA (Horn, Nafpliotis, and Goldberg 1994).
  6. FFGA (Fonseca and Fleming 1993).

  While there is a clear performance gap between SPEA and NSGA as well as between NSGA and the remaining algorithms, the fronts achieved by VEGA, HLGA, NPGA, and FFGA are rather close together. However, the results indicate that VEGA might be slightly superior to the other three EAs, while NPGA achieves fronts closer to the global optimum as FFGA. Moreover, it seems that VEGA and HLGA have difficulties evolving well-distributed trade-off fronts on the non-convex function.
- Elitism is an important factor in evolutionary multiobjective optimization. On the one hand, this statement is supported by the fact that SPEA i) clearly outperforms all algorithms on five of the six test functions and ii) is the only method among the ones under consideration which incorporates elitism as a central part of the algorithm. On the other hand, the performance of the other algorithms improved significantly when SPEA's elitist strategy was included (cf. Figure 9). Preliminary results indicate that NSGA with elitism equals the performance of SPEA.

This study forms a good basis to combine promising aspects of different algorithms into a new approach that shows good performance on all test problems. The experimental results suggest that such an algorithm may be constructed, where probably the nondominated sorting classification as well as elitism play a major role. Several issues must be addressed, ranging from the question of how elitism is implemented most effectively to the problem of whether distance metrics should operate on the parameter space or the objective space. In this context, the suggested performance metrics could be useful to compare techniques quantitatively, overcoming the limitations of the $\mathcal{C}$ metric used here.

Acknowledgments

Kalyanmoy Deb acknowledges support from Alexander von Humboldt Foundation, Germany, during the course of this study. He also wishes to thank Lothar Thiele and Eckart Zitzler for inviting him to ETH, Zurich, where a plan of this work was conceived.

References

Bäck, T. (1996). *Evolutionary Algorithms in Theory and Practice.* New York, etc.: Oxford University Press.

Chambers, J. M., W. S. Cleveland, B. Kleiner, and P. A. Tukey (1983). *Graphical Methods for Data Analysis.* Pacific Grove, California: Wadsworth & Brooks/Cole Publishing Company.

Cunha, A. G., P. Oliviera, and J. Covas (1997). Use of genetic algorithms in multicriteria optimization to solve industrial problems. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, San Francisco, California, pp. 682–688. Morgan Kaufmann.

Deb, K. (1998). Multi-objective genetic algorithms: Problem difficulties and construction of test functions. Technical Report No. CI-49/98, Department of Computer Science/XI, University of Dortmund.

Deb, K. and D. E. Goldberg (1989). An investigation of niche and species formation in genetic function optimization. In J. D. Schaffer (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*, pp. 42–50. Morgan Kaufmann.

Esbensen, H. and E. S. Kuh (1996). Design space exploration using the genetic algorithm. In *IEEE International Symposium on Circuits and Systems (ISCAS'96)*, pp. 500–503.

Fonseca, C. M. and P. J. Fleming (1993). Genetic algorithms for multiobjective optimization: Formulation, discussion and generalization. In S. Forrest (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms*, San Mateo, California, pp. 416–423. Morgan Kaufmann.

Fonseca, C. M. and P. J. Fleming (1995a). Multiobjective genetic algorithms made easy: Selection, sharing and mating restrictions. In *First International Conference on Genetic Algorithms in Engineering Systems: Innovations and Applications (GALESIA 95)*, pp. 45–52. The Institution of Electrical Engineers.

Fonseca, C. M. and P. J. Fleming (1995b). An overview of evolutionary algorithms in multiobjective optimization. *Evolutionary Computation 3*(1), 1–16.

Fonseca, C. M. and P. J. Fleming (1996). On the performance assessment and comparison of stochastic multiobjective optimizers. In *Fourth International Conference on Parallel Problem Solving from Nature (PPSN-IV)*, pp. 584–593.

Fonseca, C. M. and P. J. Fleming (1998a, January). Multiobjective optimization and multiple constraint handling with evolutionary algorithms—part i: A unified formulation. *IEEE Transactions on Systems, Man, and Cybernetics 28*(1), 26–37.

Fonseca, C. M. and P. J. Fleming (1998b, January). Multiobjective optimization and multiple constraint handling with evolutionary algorithms—part ii: Application example. *IEEE Transactions on Systems, Man, and Cybernetics 28*(1), 38–47.

Fourman, M. P. (1985). Compaction of symbolic layout using genetic algorithms. In J. J. Grefenstette (Ed.), *Proceedings of an International Conference on Genetic Algorithms and Their Applications*, pp. 141–153.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning*. Reading, Massachusetts: Addison-Wesley.

Goldberg, D. E. and J. Richardson (1987). Genetic algorithms with sharing for multimodal function optimization. In *Genetic Algorithms and their Applications: Proceedings of the Second International Conference on Genetic Algorithms*, Hillsdale, NJ, pp. 41–49. Lawrence Erlbaum.

Hajela, P. and C.-Y. Lin (1992). Genetic search strategies in multicriterion optimal design. *Structural Optimization 4*, 99–107.

Horn, J. (1997). *Handbook of Evolutionary Computation*, Chapter F1.9 Multicriteria decision making. Bristol (UK): Institute of Physics Publishing.

Horn, J. and N. Nafpliotis (1993, July). Multiobjective optimization using the niched pareto genetic algorithm. IlliGAL Report 93005, Illinois Genetic Algorithms Laboratory, University of Illinois, Urbana, Champaign.

Horn, J., N. Nafpliotis, and D. E. Goldberg (1994). A niched pareto genetic algorithm for multiobjective optimization. In *Proceedings of the First IEEE Conference on Evolutionary Computation, IEEE World Congress on Computational Computation*, Volume 1, Piscataway, NJ, pp. 82–87. IEEE Service Center.

Ishibuchi, H. and T. Murata (1996). Multi-objective genetic local search algorithm. In *Proceedings of 1996 IEEE International Conference on Evolutionary Computation (ICEC'96)*, pp. 119–124.

Kursawe, F. (1991). A variant of evolution strategies for vector optimization. In *Parallel Problem Solving from Nature — Proc. 1st Workshop PPSN*, Berlin, pp. 193–197. Springer.

Laumanns, M., G. Rudolph, and H.-P. Schwefel (1998). A spatial predator-prey approach to multi-objective optimization: A preliminary study. In *Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V)*, pp. 241–249.

Obayashi, S., S. Takahashi, and Y. Takeguchi (1998). Niching and elitist models for mogas. In *Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V)*, pp. 260–269.

Oei, C. K., D. E. Goldberg, and S.-J. Chang (1991, December). Tournament selection, niching, and the preservation of diversity. IlliGAL Report 91011, University of Illinois at Urbana-Champaign, Urbana, IL 61801.

Parks, G. T. and I. Miller (1998). Selective breeding in a multiobjective genetic algorithm. In *Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V)*, pp. 250–259.

Rudolph, G. (1998). On a multi-objective evolutionary algorithm and its convergence to the pareto set. In *IEEE International Conference on Evolutionary Computation (ICEC'98)*, pp. 511–516. IEEE.

Schaffer, J. D. (1984). *Multiple Objective Optimization with Vector Evaluated Genetic Algorithms*. Ph. D. thesis, Vanderbilt University. Unpublished.

Schaffer, J. D. (1985). Multiple objective optimization with vector evaluated genetic algorithms. In J. J. Grefenstette (Ed.), *Proceedings of an International Conference on Genetic Algorithms and Their Applications*, pp. 93–100.

Srinivas, N. and K. Deb (1994). Multiobjective optimization using nondominated sorting in genetic algorithms. *Evolutionary Computation 2*(3), 221–248.

Tamaki, H., H. Kita, and S. Kobayashi (1996). Multi-objective optimization by genetic algorithms: A review. In *Proceedings of 1996 IEEE International Conference on Evolutionary Computation (ICEC'96)*, pp. 517–522.

Valenzuela-Rendón, M. and E. Uresti-Charre (1997). A non-generational genetic algorithm for multiobjective optimization. In *Proceedings of the Seventh International Conference on Genetic Algorithms*, San Francisco, California, pp. 658–665. Morgan Kaufmann.

Veldhuizen, D. A. V. and G. B. Lamont (1998). Evolutionary computation and convergence to a pareto front. In *Genetic Programming 1998 Conference*.

Zitzler, E. and L. Thiele (1998a, May). An evolutionary algorithm for multiobjective optimization: The strength pareto approach. Technical Report 43, Computer Engineering and Networks Laboratory (TIK), Swiss Federal Institute of Technology (ETH) Zurich, Gloriastrasse 35, CH-8092 Zurich, Switzerland. Submitted to IEEE Transactions on Evolutionary Computation.

Zitzler, E. and L. Thiele (1998b). Multiobjective optimization using evolutionary algorithms — a comparative case study. In *Fifth International Conference on Parallel Problem Solving from Nature (PPSN-V)*, pp. 292–301.