

SMS-EMOA: Multiobjective selection based on dominated hypervolume

Nicola Beume^{a,*}, Boris Naujoks^a, Michael Emmerich^b

^a University of Dortmund, Chair of Algorithm Engineering, 44221 Dortmund, Germany

^b University of Leiden, Leiden Institute for Advanced Computer Science, 2333 CA Leiden, Netherlands

Available online 27 September 2006

Abstract

The hypervolume measure (or \mathcal{S} metric) is a frequently applied quality measure for comparing the results of evolutionary multiobjective optimisation algorithms (EMOA). The new idea is to aim explicitly for the maximisation of the dominated hypervolume within the optimisation process. A steady-state EMOA is proposed that features a selection operator based on the hypervolume measure combined with the concept of non-dominated sorting. The algorithm's population evolves to a well-distributed set of solutions, thereby focussing on interesting regions of the Pareto front. The performance of the devised \mathcal{S} metric selection EMOA (SMS-EMOA) is compared to state-of-the-art methods on two- and three-objective benchmark suites as well as on aeronautical real-world applications.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Evolutionary computations; Evolutionary multiple objective optimisation; Performance assessment; Hypervolume measure; OR in aerodynamic industries

1. Introduction

Dealing with preferences can be a difficult task in multiobjective optimisation, especially if the knowledge about the structure of the set of optimal solutions is incomplete. By adopting the a posteriori approach, the decision maker postpones the final decision on preferences between conflicting objectives until a set of compromise solutions is known. Such sets can be obtained by Pareto optimisation [1,2].

In Pareto optimisation, the purpose is to optimise a problem of m objective functions $\mathbf{f} = (f_1, \dots, f_m)$ with $f_i : \mathbb{S} \rightarrow \mathbb{R}, i \in \{1, \dots, m\}$, and \mathbb{S} denoting a search (or: decision) space. Here, only minimisation is considered. In the examined examples, the decision space is a subset of \mathbb{R}^n , though the proposed algorithms can easily be generalised to other spaces. Since objectives can be conflicting, instead of searching for a single best solution, we seek for a set of good compromise solutions. Among solution vectors $\mathbf{y}, \mathbf{y}' \in \mathbb{R}^m$ a partial order is defined by the *dominance relation* “ \prec ”:

* Corresponding author.

E-mail address: nicola.beume@cs.uni-dortmund.de (N. Beume).

$$\mathbf{y} \prec \mathbf{y}' \quad (\mathbf{y} \text{ dominates } \mathbf{y}') : \iff \forall i \in \{1, \dots, m\} : y_i \leq y'_i \quad \text{and} \quad \exists j \in \{1, \dots, m\} : y_j < y'_j. \quad (1)$$

For $\mathbf{x}, \mathbf{x}' \in \mathbb{S}$, we write $\mathbf{x} \prec \mathbf{x}' : \iff \mathbf{f}(\mathbf{x}) \prec \mathbf{f}(\mathbf{x}')$. The minimal elements of the dominance relation are called *Pareto-optimal* and the set of Pareto-optimal decision vectors is called *Pareto(-optimal) set*: $\{\mathbf{x} \in \mathbb{S} \mid \nexists \mathbf{x}' \in \mathbb{S} : \mathbf{x}' \prec \mathbf{x}\}$. The corresponding image under \mathbf{f} in the solution space \mathbb{R}^m is called *Pareto front*. A point in the search space is called *non-dominated* within a set $A \subseteq \mathbb{S}$ if there is no point in A dominating it.

In practice, it is rarely possible to compute the true Pareto set with any algorithmic method, because of the black-box character of many objective functions and/or the complexity of the search space. Instead, algorithms are intended to yield approximations of the Pareto front and thereby the Pareto set.

Among Pareto optimisation techniques, EMOA established themselves in a leading position concerning research interests and progress. EMOA are applicable to a wide range of problems and their population based approach assures an inherent set of best qualified solutions. Therefore, they are well-suited methods, whenever analytical methods fail to obtain the precise Pareto set.

A variety of quality indicators is used in order to measure the quality of Pareto front approximations. Among them, the *hypervolume measure* or *\mathcal{S} metric* is of outstanding importance. It is a quality indicator that rewards the convergence towards the Pareto front as well as the representative distribution of points along the front. The hypervolume measure was originally proposed by Zitzler and Thiele [3], who called it the *size of dominated space* [4]. Let A denote the Lebesgue measure, then the \mathcal{S} metric is defined as

$$\mathcal{S}(B, \mathbf{y}_{\text{ref}}) := A\left(\bigcup_{\mathbf{y} \in B} \{\mathbf{y}' \mid \mathbf{y} \prec \mathbf{y}' \prec \mathbf{y}_{\text{ref}}\}\right), \quad B \subseteq \mathbb{R}^m. \quad (2)$$

Here, $\mathbf{y}_{\text{ref}} \in \mathbb{R}^m$ denotes a reference point that should be dominated by all Pareto-optimal solutions.

Fleischer [5] proved that, given a finite search space and a reference point, maximisation of the hypervolume measure is equivalent to finding the Pareto set. Therefore, a set of bounded size which obtains the maximal \mathcal{S} metric value that is possible for the given size, only consists of Pareto-optimal solutions. It has been empirically observed that for a fixed number of points, the maximisation of the hypervolume metric yields subsets of the Pareto front which are well-distributed (cf. [6,7]).

The absolute value of the \mathcal{S} metric depends on the choice of the reference point. This dependence can influence the ranking of Pareto front approximations (cf. [8]). Despite this disadvantage, the \mathcal{S} metric was used in several comparative studies of EMOA, e.g. [4,9,10].

A rather nice characteristic of the hypervolume measure is, that the ranking of the solutions is invariant to linear scaling of the objective functions. Furthermore, the reference point only influences the relative contribution of the points containing worst values (boundary points). There are further theoretical qualities of the \mathcal{S} metric which make it a rather fair measure, and preferable to other performance metrics [11].

Although it seems straightforward to consider the quality indicator directly in the selection operator, with a few exceptions [12,13,7], this option has hardly been considered for the design of EMOA, yet. This article presents an approach proposed recently by the authors [7], namely the selection of solution candidates according to their contribution to the dominated hypervolume. It deepens the discussion by emphasising on problems with three objectives.

The article is structured as follows. In Section 2, we provide a comprehensive description of the SMS-EMOA and relate it to similar state-of-the-art approaches. A detailed discussion on the algorithm's behaviour and qualities is given, including implementation as well as theoretical aspects. Besides, an efficient modification of the SMS-EMOA is presented which has an improved runtime and in some cases even higher quality than the basic SMS-EMOA. In addition, efficient algorithms for computing hypervolume contributions of non-dominated solutions are discussed, with a focus on the two- and three-dimensional case. In Section 3, we compare the performance of the SMS-EMOA to conventional EMOA on the most common benchmark problems for evolutionary multiobjective optimisation. The studies of design optimisation tasks from aeronautic industry, featuring two and three objectives, confirm the applicability of the new approach to real-world problems (Section 4). Only two- and three-dimensional problems are covered, since these cases are most frequently addressed in real-world applications. Finally, Section 5 outlines the main conclusions of this work.

2. The \mathcal{S} metric selection (SMS-) EMOA

The SMS-EMOA has been designed to cover a maximal hypervolume with a finite number of points. Furthermore, we wanted to diminish the problem of choosing the right reference point and aimed at a simple algorithm that can be parallelised in an asynchronous manner and extended by knowledge integrating techniques, such as approximate function evaluations. The SMS-EMOA combines ideas borrowed from other EMOA, like the well-established NSGA-II [14] and archiving strategies presented by Knowles et al. [6,8]. It is a steady-state algorithm founded on two pillars: (1) non-dominated sorting is used as a ranking criterion and (2) the hypervolume is applied as selection criterion to discard that individual, which contributes the least hypervolume to the worst-ranked front.

2.1. Details of the SMS-EMOA

The basic algorithm is described in Algorithm 1. Starting with an initial population of μ individuals, a new individual is generated by means of randomised variation operators.¹ The new individual will become a member of the next population, if replacing another individual leads to a higher quality of the population with respect to the \mathcal{S} metric.

Algorithm 1. SMS-EMOA

```

1:  $P_0 \leftarrow \text{init}()$  /* Initialise random population of  $\mu$  individuals */
2:  $t \leftarrow 0$ 
3: repeat
4:    $q_{t+1} \leftarrow \text{generate}(P_t)$  /* generate offspring by variation */
5:    $P_{t+1} \leftarrow \text{Reduce}(P_t \cup \{q_{t+1}\})$  /* select  $\mu$  best individuals */
6:    $t \leftarrow t + 1$ 
7: until termination condition fulfilled

```

The procedure `Reduce` described in Algorithm 2 selects the μ individuals of the subsequent population. The algorithm `fast-nondominated-sort` used in NSGA-II [14] is applied to partition the population in v sets $\mathcal{R}_1, \dots, \mathcal{R}_v$ according to the non-dominated sorting defined by Goldberg [15]. The subsets are called fronts and are provided with an index representing a hierarchical order (the level of domination) whereas the solutions within each front are mutually non-dominated. The first subset contains all non-dominated solutions of the original set Q . The second front consists of individuals that are non-dominated in the set $(Q \setminus \mathcal{R}_1)$, e.g. each member of \mathcal{R}_2 is dominated by at least one member of \mathcal{R}_1 . More general, the i th front consists of individuals that are non-dominated if the individuals of the fronts j with $j < i$ were removed from Q .

Algorithm 2. Reduce(Q)

```

1:  $\{\mathcal{R}_1, \dots, \mathcal{R}_v\} \leftarrow \text{fast-nondominated-sort}(Q)$  /* all  $v$  fronts of  $Q$  */
2:  $r \leftarrow \text{argmin}_{s \in \mathcal{R}_v} [\Delta_{\mathcal{S}}(s, \mathcal{R}_v)]$  /*  $s \in \mathcal{R}_v$  with lowest  $\Delta_{\mathcal{S}}(s, \mathcal{R}_v)$  */
3: return  $(Q \setminus \{r\})$  /* eliminate detected element */

```

Afterwards, one individual is discarded from the worst ranked front. Whenever this front comprises $|\mathcal{R}_v| > 1$ individuals, the individual $s \in \mathcal{R}_v$ is eliminated that minimises

$$\Delta_{\mathcal{S}}(s, \mathcal{R}_v) := \mathcal{S}(\mathcal{R}_v) - \mathcal{S}(\mathcal{R}_v \setminus \{s\}). \quad (3)$$

¹ If not stated otherwise, we employed the variation operators used by Deb et al. for their ϵ -MOEA algorithm [10]. These are the simulated binary crossover (SBX) and a polynomial mutation operator, described in detail in [1]. We used the implementation available on the KanGAL home page <http://www.iitk.ac.in/kangal/>.

The value of $\Delta_{\mathcal{S}}(s, \mathcal{R}_v)$ can be interpreted as the exclusive contribution of s to the \mathcal{S} metric value of its appropriate front. By definition of $\Delta_{\mathcal{S}}(s, \mathcal{R}_v)$, an individual, which dominates another is always kept and no non-dominated individual is replaced by a dominated one. This measure keeps those individuals which maximise the population's \mathcal{S} metric value, which implies that the covered hypervolume of a population cannot decrease by application of the Reduce operator. Thus, for Algorithm 1 the following invariant holds:

$$\mathcal{S}(P_t) \leq \mathcal{S}(P_{t+1}). \quad (4)$$

2.1.1. Steady-state selection

Due to the high computational effort of the hypervolume calculation (see Section 2.3), a steady state selection scheme is used. Since only one individual is created, only one has to be deleted from the population in each generation. Thus, the selection operator has to compute at most $\mu + 1$ values of the \mathcal{S} metric (exactly $\mu + 1$ values in case all solutions are non-dominated). These are the values of the subsets of the worst ranked front, in which one point of the front is left out, respectively. A $(\mu + \lambda)$ selection scheme would require the calculation of $\binom{\mu + \lambda}{\mu}$ possible \mathcal{S} metric values to identify an optimally composed population, maximising the \mathcal{S} metric net value. Besides, a steady-state scheme can easily be parallelised by means of asynchronously distributed function evaluations.

2.1.2. Population size

In contrast to other strategies that store non-dominated individuals in an archive, the SMS-EMOA keeps a population of non-dominated and dominated individuals at constant size. Keeping only non-dominated individuals, might lead to small or even single-membered populations and thus to a crucial loss of diversity. Modifications of SMS-EMOA with dynamic population size have been considered in [16]. These variants are interesting whenever an acceleration of the evolution process in its beginning is desired. To avoid the loss of diversity, we recommend to define a lower bound for the population size.

2.1.3. Handling of boundary solutions

In a two-dimensional solution space, the reference point \mathbf{y}_{ref} of the \mathcal{S} metric is only needed to calculate the hypervolume of the two extremal points that consist of one best and one worst objective value, respectively. We decided to omit \mathbf{y}_{ref} and always keep these extremal solutions.

In the case of more objective functions, there can be far more points, whose hypervolume contributions depend on the choice of the reference point. These exactly are the boundary points which contain at least one worst objective value. It is not intended to keep all boundary solutions, thus contributions are calculated for them as well. We used a dynamic handling of the reference point, where \mathbf{y}_{ref} is recalculated in each generation. It is chosen as the vector of the currently worst objective values increased by 1.0. Thus, the contribution of a point that contains a worst objective value depends on the remaining objectives as the difference to the reference point behaves neutral in the product. Thus, good non-extremal solutions might outperform extremal ones concerning the hypervolume. First observations advised, that the dynamic handling of the reference point is preferable to a static one.

2.2. Selection variants of SMS-EMOA

A competitive study of several selection modi within the SMS-EMOA framework was presented in [16]. The concepts all feature the contributing hypervolume in combination with non-dominated sorting, the common dominance criterion, or the number of dominating points. The latter will be described in more detail here.

The *number of dominating points* $d(s, P(t))$ is the number of points from set $P(t)$ that dominate point s , formally

$$d(s, P(t)) := |\{y \in P(t) | y \prec s\}|. \quad (5)$$

A modified Reduce procedure is given in Algorithm 3. The number of dominating points $d(s, P(t))$ is used as selection criterion in case of dominated solutions and the contributing hypervolume $\Delta_{\mathcal{F}}$ is applied whenever all individuals of $P(t)$ are non-dominated. If the population consists of more than one front, the individual with the highest $d(s, P(t))$ value among the solutions of the worst ranked front is discarded. Otherwise, the $d(s, P(t))$ values of all individuals equal zero and the $\Delta_{\mathcal{F}}$ selection is used instead.

Algorithm 3. Reduce(Q)

1: $\{\mathcal{R}_1, \dots, \mathcal{R}_v\} \leftarrow \text{nondominated-sort}(Q)$	/* all v fronts of Q */
2: if $v > 1$ then	
3: $r \leftarrow \operatorname{argmax}_{s \in \mathcal{R}_v} [d(s, Q)]$	/* $s \in \mathcal{R}_v$ with highest $d(s, Q)$ */
4: else	
5: $r \leftarrow \operatorname{argmin}_{s \in \mathcal{R}_1} [\Delta_{\mathcal{F}}(s, \mathcal{R}_1)]$	/* $s \in \mathcal{R}_1$ with lowest $\Delta_{\mathcal{F}}(s, \mathcal{R}_1)$ */
6: end if	
7: return $(Q \setminus \{r\})$	/* eliminate detected element */

One motivation for developing such a measure is the smaller runtime complexity compared to the hypervolume measure. Moreover, the motivation was to achieve a different ranking of dominated solutions in order to emphasise on sparsely filled regions of the solution space. Here, individuals are kept for proceeding generations to fill gaps in the Pareto front approximation. The contributing hypervolume is applied to distribute solutions well on the front they are lying on. Though the final purpose is to distribute solutions well on the first non-dominated front (\mathcal{R}_1), this is not an end in itself on the other fronts. The $d(s, P(t))$ measure favours solutions located in those areas, where the better fronts are sparsely populated. The idea is that emanating offspring solutions may rise in rank to better fronts and fill those vacancies. In areas where the non-dominated front is densely populated, it is of no use to keep dominated individuals.

The measure is visualised by means of an example in Fig. 1. The contributing hypervolume measure favours y_8 over y_9 but the $d(s, P(t))$ measure favours y_9 since there is only one dominating point. Solution y_9 is obviously much more interesting because of the big vacancy between y_5 and y_6 on the first front.

The number of dominating points can be calculated efficiently: the selection candidates are compared with all solutions to check the dominance. This is possible in time $O(m\mu^2)$, with m denoting the number of objectives and μ the population size. The expected runtime complexity is lower than the one of the basic SMS-EMOA. The worst case complexity of the SMS-EMOA is the same, because the population may consist of non-dominated solutions only. In this case, the $d(s, P(t))$ selection scheme is never applied. The average case complexity is assumed to be better.




Fig. 1. The hypervolume dominated by front $\mathcal{R}_2 = \{y_7, \dots, y_{10}\}$ is light-coloured and the contributions $\Delta_{\mathcal{F}}(y_8, \mathcal{R}_2)$ and $\Delta_{\mathcal{F}}(y_9, \mathcal{R}_2)$ are visualised by the attached dark rectangles. The hatched areas correspond to the regions containing points that dominate y_8 or y_9 , respectively. Here, $d(y_8, P(t)) = 4$ and $d(y_9, P(t)) = 1$, and $\Delta_{\mathcal{F}}(y_8, \mathcal{R}_2) > \Delta_{\mathcal{F}}(y_9, \mathcal{R}_2)$.

2.3. Calculation of contributing hypervolume $\Delta_{\mathcal{S}}$

The best known algorithms calculate the hypervolume in a runtime, that is polynomial in the number of points, but grows exponentially with the number of objectives. An overview of such algorithms is given next. Additionally, dedicated algorithms to directly calculate all values of $\Delta_{\mathcal{S}}$ at once in case of two and three objectives are presented. These algorithms are significantly faster than repetitive calls of procedures to compute the hypervolume as a whole.

In case of two objectives, we take the points of the worst-ranked non-dominated front and sort them in ascending order according to the values of the first objective function f_1 . We get a sequence that is additionally sorted in descending order concerning the f_2 values, because the points are mutually non-dominated. Given a sorted front $\mathcal{R}_v = \{s_1, \dots, s_{|\mathcal{R}_v|}\}$, $\Delta_{\mathcal{S}}$ is calculated as follows ($i = 2, \dots, |\mathcal{R}_v| - 1$):

$$\Delta_{\mathcal{S}}(s_i, \mathcal{R}_v) = (f_1(s_{i+1}) - f_1(s_i)) \cdot (f_2(s_{i-1}) - f_2(s_i)). \quad (6)$$

A fast algorithm for computing all contributions in a three objective space is presented in [16]. It has a runtime of $O(\mu^3)$ and is based on a two-dimensional projection of the set of points. The f_1 – f_2 plane is divided into a grid according to the coordinates of all $\mu + 1$ solution vectors and the reference point. For each grid cell c_{ij} ($i = 1, \dots, \mu + 1, j = 1, \dots, \mu + 1$), the $\mu + 1$ values of f_3 are computed at the corner point of the cell, which dominates it in the f_1 – f_2 plane. Among these values, the best value $h_1(c_{ij})$ and second best value $h_2(c_{ij})$ are determined. Given this information, which can be stored in two matrices of dimension $(\mu + 1) \times (\mu + 1)$, the exclusive contribution of the i th solution vector ($i = 1, \dots, \mu + 1$) can be calculated as follows. For any grid cell c_{ij} with $h_1(c_{ij})$ value determined by the i th solution, we determine the volume v_{ij} as the product of the surface area of c_{ij} and the difference $|h_1(c_{ij}) - h_2(c_{ij})|$. Accumulating all these values of v_{ij} yields to the exclusive contribution of the i th solution to the dominated hypervolume. The runtime is $O(\mu^3)$ because the number of cells is quadratic and the calculations for each cell value are done in linear time. It is likely possible to further tune this algorithm, i.e. by using incremental updates. However, for the computations performed in this paper the algorithm performed reasonably fast. Another potential advantage is that the algorithm is non-recursive and uses simple data structures. Thus it is relatively easy to implement/debug.

An algorithm that computes the whole hypervolume of a set of points is the LebMeasure algorithms by Fleischer [5]. While [17] recently showed, that the runtime complexity is exponential in the number of objectives in the worst case, contrary to the runtime complexity mistakenly stated in [5].

The HSO algorithm (Hypervolume by Slicing Objectives) described by Knowles [18] and Zitzler [19] is exponential in the number of objectives as well. It is easier to implement, uses simpler data structures, and needs less space. Recently While et al. [20] presented heuristics for accelerating HSO. For a set of l points, HSO processes objectives one after the other and thereby divides the problem into l sets with a dimension reduced by one in each iteration. The runtime differs with the order in which the objectives are processed. The heuristics are applied as a pre-processing of HSO and calculate a promising order of objectives which reliably avoids the worst possible case (but not reliably the worst case complexity) of HSO. While et al. [20] demonstrated for some DTLZ functions the significant runtime reduction of 25–98% compared to the basic HSO.

2.4. Comparison to other EMOA

2.4.1. ESP

Huband et al. [12] incorporated an approximation of the hypervolume contribution into a selection operator resembling SPEA2 [21], replacing the common nearest neighbour criterion. They involved the “product of the one-dimensional lengths to the next worse objective function value in the front for each objective” [12] as selection criterion in case of equal fitness values. In the two-objective case, this exactly corresponds to the \mathcal{S} metric contribution (cf. Eq. (6)). In case of more than two objectives, this procedure provides only a lower bound of the hypervolume value. The great advantage is the reduced runtime of $O(k)$ with k being the number of solutions on the front. The selection is performed using a $(\mu + \lambda)$ -ES selection scheme, setting $\mu = \lambda = 100$. Selection with equal numbers of parents and offsprings is very common in evolutionary multiobjective optimisation; though evidence for the advantage of this setting is missing up to now. As variation operator,

the ESP algorithm employs usual operators from evolution strategies coupled with a mutation probability p . Huband et al. choose $p = 1/n$ as the probability to apply mutation for each decision variable.

ESP has been compared to several well-known EMOA like NSGA-II and SPEA2 on a set of test functions (ZDT1 to ZDT4 and ZDT6 [22]). In [12], the authors account the probabilistic mutation operator to be responsible for the obtained quality of ESP.

2.4.2. IBEA

The indicator-based EA (IBEA), proposed by Zitzler and Künzli [13], is a general framework to incorporate quality indicators in EMOA. Next to a hypervolume based indicator, they also incorporated a binary additive ϵ -indicator. The fitness of a solution is based on the sum of the indicator values resulting from pairwise comparisons to all other solutions.

They utilised ordinary EMOA variation operators for different test functions, i.e. SBX recombination and polynomial mutation on the real valued test functions under investigation (ZDT6, KUR, DTLZ2 and DTLZ6). Again, a $(\mu + \mu)$ selection scheme and a specially scaled value of the objective and indicator values were incorporated in the improved adaptive IBEA. Both proposed variants yielded results “significantly better than SPEA2 and NSGA-II with respect to the performance indicators” [13].

2.4.3. The AA_{reduce} scheme

The SMS-EMOA presented here nearly fits into the generic algorithm scheme AA_{reduce} presented by Knowles and Corne [8] within the context of archiving strategies. The archiving strategy called AA_S uses the contributing hypervolume of the non-dominated points to determine the worst solution. Knowles and Corne showed that the AA_S strategy converges to a subset of the true Pareto front. Provided that the population size in the SMS-EMOA equals the archive size in AA_S and only non-dominated solutions are concerned, the AA_S strategy is equivalent to our selection method. If dominated solutions are considered as well, the SMS-EMOA population contains even more solutions than the AA_S archive. Thus, the proof of convergence holds for SMS-EMOA as well.

2.4.4. NSGA-II

The NSGA-II by Deb et al. [14] is an improved version of the preceding NSGA. It features non-dominated-sorting and a $(\mu + \mu)$ selection scheme. The secondary ranking criterion for solutions on the same front is called crowding distance. Its task corresponds to $\Delta_{\mathcal{F}}$ in the SMS-EMOA. For a solution, the crowding distance is defined as the sum of the side lengths of the cuboid through the neighbouring solutions on the front. For extremal solutions it is defined as infinity. The crowding distance value of a solution depends on its neighbours and not directly on the position of the point itself, in contrast to the $\Delta_{\mathcal{F}}$ measure.

The crowding distance is meant to distribute solutions uniformly on the fronts. In contrast to this, the hypervolume measure is meant to distribute them in a way that approximately maximises the covered hypervolume. This indicates that good compromise solutions, which are located in parts of the front with high curvature, are valued higher in SMS-EMOA than in NSGA-II.

3. Studies on artificial testfunctions

3.1. Distribution of solutions

A study was performed to get an impression of how the SMS-EMOA distributes solutions on Pareto fronts of different curvature, e.g. convex, concave and linear Pareto fronts. We conducted a study on simple two-objective test functions with a high dimensional decision space. The following family of generic functions, named EBN, is devised:

$$f_1(\mathbf{x}) := \left(\sum_{i=1}^n |x_i| \right)^{\gamma} n^{-\gamma}, \quad f_2(\mathbf{x}) := \left(\sum_{i=1}^n |x_i - 1| \right)^{\gamma} n^{-\gamma}, \quad \mathbf{x} \in [0, 1]^n, \quad (7)$$

with n being the number of decision variables. The curvature is adjusted by the parameter γ : $\gamma = 1$ leads to a linear Pareto front, while $\gamma > 1$ yields convex fronts and $\gamma < 1$ concave ones. With $y_1 := f_1(\mathbf{x})$ and $y_2 := f_2(\mathbf{x})$ the precise Pareto fronts are analytically given by

$$y_2 = (1 - y_1^{(1/\gamma)})^\gamma. \quad (8)$$

All points in $[0, 1]^n$ are Pareto-optimal solutions [23].

Fig. 2 shows that the solutions are not uniformly distributed on the Pareto front, but distributed in a way that maximises the covered hypervolume. The results demonstrate that the SMS-EMOA favours solutions in regions where the Pareto front has knee-points and captures the regions with fair trade-offs between different objectives. The regions with unbalanced trade-offs, located on the flanks of the Pareto front, are covered with less density, although extremal solutions are maintained. On the linear Pareto front, the points get uniformly distributed. In case of a concave Pareto front the regions with fair trade-offs are emphasised. These are located near the angular point of the Pareto front. This serves the practitioner who is mainly interested in a limited number of solutions on the Pareto front to get an overview of possible compromise solutions.

3.2. Settings for benchmarks

The SMS-EMOA was tested on several test problems from literature. We aimed at comparability to frequently cited papers of Deb et al. presenting the ϵ -MOEA approach [9,10]. Therefore, we also invoked the same variation operators, namely SBX recombination and polynomial mutation with the parametrisation given in [9,10]. The test problems named ZDT1 to ZDT4 and ZDT6 from [10,22] have been considered and are presented in Section 3.3, as well as the functions DTLZ1 to DTLZ4 [24] (see Section 3.4). This way, we compared our SMS-EMOA to NSGA-II, C-NSGA-II, SPEA2, and ϵ -MOEA.

The same parameter settings as in the benchmark from Deb et al. are invoked [9,10]. 20,000 function evaluations are calculated for the ZDT functions and 30,000 on all DTLZ functions except for DTLZ3, where 100,000 function evaluations are used. The population size is set to 100 for all functions, unlike in [10] where a size of 200 individuals was applied for DTLZ3. For DTLZ4, ten runs are executed and five for all other DTLZ and ZDT functions. Like in [9], only ‘successful’ runs are used to calculate the metric values. A run is termed successful, if the solution set is distributed over the whole true Pareto front. The quality of the algorithms is compared considering the \mathcal{S} metric and another performance measure, the *convergence metric* like in [9,10]. For every point in the solution set, the euclidean distance to the nearest point on the Pareto front is measured in the solution space. The convergence metric is the arithmetic mean of these distances. It is only applicable if the Pareto front is known a priori, as it is the case for many benchmark problems. The con-




Fig. 2. This study visualises results on the EBN problem family with Pareto fronts of different curvature computed by SMS-EMOA for a 20-dimensional decision space.

Table 1
Results of several EMOA on functions ZDT1 to ZDT4 and ZDT6

	Algorithm	Convergence measure			\mathcal{S} metric		
		Average	Standard deviation	r	Average	Standard deviation	r
ZDT1	NSGA-II	0.00054898	6.62e−05	4	0.8701	3.85e−04	6
	C-NSGA-II	0.00061173	7.86e−05	5	0.8713	2.25e−04	3
	SPEA2	0.00100589	12.06e−05	6	0.8708	1.86e−04	4
	ϵ -MOEA	0.00039545	1.22e−05	1	0.8702	8.25e−05	5
	SMS-EMOA	0.00044394	2.88e−05	2	0.8721	2.26e−05	1
	SMS-EMOA dp	0.00044872	4.35e−05	3	0.8721	1.59e−05	1
ZDT2	NSGA-II	0.00037851	1.88e−05	1	0.5372	3.01e−04	6
	C-NSGA-II	0.00040011	1.91e−05	2	0.5374	4.42e−04	4
	SPEA2	0.00082852	11.38e−05	6	0.5374	2.61e−04	4
	ϵ -MOEA	0.00046448	2.47e−05	5	0.5383	6.39e−05	3
	SMS-EMOA	0.00041004	2.34e−05	3	0.5388	3.60e−05	1
	SMS-EMOA dp	0.00041923	2.94e−05	4	0.5388	1.77e−05	1
ZDT3	NSGA-II	0.00232321	13.95e−05	4	1.3285	1.72e−04	4
	C-NSGA-II	0.00239445	12.30e−05	5	1.3277	9.82e−04	6
	SPEA2	0.00260542	15.46e−05	6	1.3276	2.54e−04	5
	ϵ -MOEA	0.00175135	7.45e−05	3	1.3287	1.31e−04	3
	SMS-EMOA	0.00057233	5.81e−05	2	1.3295	2.11e−05	1
	SMS-EMOA dp	0.00053864	3.27e−05	1	1.3295	2.33e−05	1
ZDT4	NSGA-II	0.00639002	0.0043	5	0.8613	0.00640	3
	C-NSGA-II	0.00618386	0.0744	4	0.8558	0.00301	5
	SPEA2	0.00769278	0.0043	6	0.8609	0.00536	4
	ϵ -MOEA	0.00259063	0.0006	2	0.8509	0.01537	6
	SMS-EMOA	0.00251878	0.0014	1	0.8677	0.00258	1
	SMS-EMOA dp	0.00289158	0.0019	3	0.8660	0.00324	2
ZDT6	NSGA-II	0.07896111	0.0067	5	0.3959	0.00894	6
	C-NSGA-II	0.07940667	0.0110	6	0.3990	0.01154	5
	SPEA2	0.00573584	0.0009	1	0.4968	0.00117	1
	ϵ -MOEA	0.06792800	0.0118	4	0.4112	0.01573	4
	SMS-EMOA	0.05043192	0.0217	2	0.4354	0.02957	2
	SMS-EMOA dp	0.06569538	0.0106	3	0.4148	0.01414	3

vergence metric does only account for the precision of the approximation to the Pareto front but not for the distribution of points. Thus, it is recommended to use it in combination with other measures [1].

For a clear overview, the best results for the \mathcal{S} metric and the convergence achieved in [10,9] have been copied to Tables 1 and 2.

3.3. Two-dimensional functions (ZDT1 to ZDT4 and ZDT6)

The following results are depicted in Table 1. The results of the SMS-EMOA with the modified selection operator using the number of dominating points (cf. Section 2.2) are given in the lines named ‘SMS-EMOA dp’. The results of the two SMS-EMOA variants are very similar and they often share the same rank. This is a very positive result since the quality of the SMS-EMOA is maintained though algorithmic complexity and runtime have been reduced. In the following, we start addressing the achievements for each test functions separately, but do not further distinguish the results for the SMS-EMOA variants.

ZDT1 has a smooth convex Pareto front where the SMS-EMOA is ranked best on the \mathcal{S} metric and near to the best concerning the convergence measure. ZDT4 provides multiple parallel local fronts and the Pareto front is equivalent to that of ZDT1. On the basis of the given values from [9,10], we assume that all algorithms jumped across the second best front with most solutions and aimed at the best front, like SMS-EMOA. The worse ranked values of the other algorithms seem to stem from disadvantageous distributions. ZDT2 has a smooth concave front and the SMS-EMOA covers most hypervolume. ZDT3 has a discontinuous Pareto

Table 2
SMS-EMOA and best results from [9] of other EMOA

Algorithm		Convergence measure		\mathcal{S} metric	
		Average	Standard deviation	Average	Standard deviation
DTLZ1	SPEA2	0.0033377	3.54e−02	0.315981	6.98e−04
	ϵ -MOEA	0.00245	9.52e−05	0.298487	NC
	SMS-EMOA	0.0029175	1.72e−04	0.316930	5.30e−05
	SMS-EMOA dp	0.0028909	1.06e−04	0.316936	8.38e−05
DTLZ2	C-NSGA-II	0.00986	8.8e−04	NC	NC
	SMS-EMOA	0.0063652	3.20e−04	0.757911	4.49e−05
	SMS-EMOA dp	0.0065383	5.12e−04	0.757994	4.74e−05
DTLZ3	ϵ -MOEA	0.0122290	2.23e−03	NC	NC
	SMS-EMOA	0.0071626	5.94e−04	0.755294	2.22e−03
	SMS-EMOA dp	0.0069858	4.28e−04	0.755443	7.9e−04
DTLZ4	ϵ -MOEA (6)	0.0097755	2.0e−04	NC	NC
	SMS-EMOA (4)	0.0065006	3.39e−04	0.757949	8.65e−05
	SMS-EMOA dp (5)	0.0065193	4.42e−04	0.757967	3.87e−05

front that consists of five slightly convex parts. Here, the SMS-EMOA is a little bit better concerning the \mathcal{S} metric than the second ranked ϵ -MOEA and significantly better with respect to convergence. ZDT6 has a concave Pareto front that is equivalent to that of ZDT2, except for the differences that the front is truncated to a smaller range and that points are non-uniformly spaced. Here, the SMS-EMOA is ranked second on both measures, outperformed by SPEA2, which in contrast shows apparently bad results on the other more simple problems.

The SMS-EMOA is ranked best concerning the \mathcal{S} metric in all functions except for ZDT6. Concerning the convergence measure, it has two first ranks, two second ranks and one third rank. According to the sum of ranks of the two measures on each function, one can state that the SMS-EMOA provides best results on all considered functions, except for ZDT6, where it is outperformed by SPEA2. Comparing the sum of the achieved ranks of each measure shows that our algorithm obtains best results concerning both the convergence measure (with 9) and the \mathcal{S} metric (with 6). In order to test the significance of the results a Student t -test has been applied. The results of the SMS-EMOA were tested for significance to all strategies that performed worse. The results with respect to the \mathcal{S} metric are significant with a significance level of more than 90%. Except for the ZDT6 test problem, even a higher significance level of 95% holds. Also for the convergence measure, the SMS-EMOA is better than the next best strategy with a significance level of more than 90%, except for ZDT4, where the comparison to CNSGA-2 and ϵ -EMOA does not lead to a sufficient significance level.

3.4. Three-dimensional functions (DTLZ1 to DTLZ4)

The results for the DTLZ functions are presented in Table 2. For easier comparison, the best values of those published in [9] are copied into the table. The abbreviation ‘NC’ stands for ‘not computed’, indicating missing values in the literature [9].

The results indicate that both SMS-EMOA variants clearly outperform the aforementioned EMOA concerning the \mathcal{S} metric value. Nearly the same holds for the convergence, where only the ϵ -MOEA applied to DTLZ1 is better. Just like on the ZDT functions, the comparison of the SMS-EMOA variants shows no significant differences. This confirms that the hypervolume-based selection is successful within different frameworks.

The SMS-EMOA solutions for functions DTLZ1 and DTLZ2 are visualised in Fig. 3. The median run concerning the \mathcal{S} metric values is shown, respectively. As there are no observably differences between the selection variants, only results received with the basic SMS-EMOA are displayed.

The Pareto front of DTLZ1 is a beveled, triangular plane in the first quadrant. The solutions of the SMS-EMOA variants lay exactly on this plane. They are distributed uniformly and extremal solutions at the edges




Fig. 3. Solutions of SMS-EMOA for DTLZ1 (top) and DTLZ2 (bottom). The lattice adumbrates the Pareto front, respectively.

of the plane are contained (cf. Fig. 3). The functions DTLZ2 to DTLZ4 share the same Pareto front, which is the eighth of the unit sphere, bounded to the first quadrant. There are no major differences between the results of the SMS-EMOA variants, again. Fig. 3 shows that the solutions for DTLZ2 are not equally distributed over the whole Pareto front. An adequate amount of boundary solutions are positioned along the eighth sphere. The remaining solutions are uniformly distributed over the inner part of the sphere surface, leaving

an obvious gap to the boundary solutions. The function DTLZ3 has several local fronts parallel to the global one. All runs of the SMS-EMOA variants reached the global Pareto front and their convergence values are substantially better than the best EMOA studied in previous work (cf. Table 2). The distribution of solutions corresponds to the one on DTLZ2. The difficulty for function DTLZ4 is that the first two object variables are raised to the power of 100. Hence, there is a bias towards the special solutions resulting in sets located solely in the f_1 – f_2 -plane or the f_1 – f_3 -plane. Just like in [10], only the successful runs are considered for the metric values and their number out of 10 performed runs is given in brackets. The results of the successful runs are similar to those of the solutions of DTLZ2 and DTLZ3.

On almost all DTLZ functions the obtained results for the convergence metric and \mathcal{S} metric are significantly better for the SMS-EMOA than for the other approaches (with a significance level of 99% using a Student t -test). Only on DTLZ1, the significance of the better \mathcal{S} metric values is not given and the ϵ -MOEA achieves better convergence values than the SMS-EMOA.

In conjunction, concerning this bundle of test problems, the SMS-EMOA can be regarded as the best evolutionary multiobjective algorithm. The outstanding performance concerning the \mathcal{S} metric is a very encouraging result even though good results seem to be natural because of the use of the \mathcal{S} metric as selection criterion and quality measure. One should admit that our approach is a rather simple one with only one population and only a few parameters. It is steady-state, resulting in a low selection pressure per generation. Neither there are any special variation operators fitted to the selection strategy, nor it is tuned for performance in any way.

The results in the convergence measure are even more surprising. Especially on the function that are supposed to be more difficult, the SMS-EMOA achieves very good results. A possible explanation might be that a population of well-distributed points is able to sample individuals with larger improvement. Further investigations are required to clarify this topic.

4. Application to airfoil design optimisation

Airfoil design applications are frequently addressed multiobjective real-world problems. The work on these applications is pushed by aeronautic industry because new solutions promise huge potential savings as well as better and more secure flight characteristics. Moreover, research in aeronautics receives governmental support in order to push up-to-date research in the related countries.

The multiobjective point of view is introduced by contradicting demands made to airfoils. On the one hand, it should provide a maximum lift to ensure save and stable flight qualities and a low energy consumption during the starting phase. On the other hand, a minimum drag is wanted during the cruising flight allowing for minimal energy consumption. Many more qualities influence the complex interplay of parameters to ensure best flight characteristics.

In this paper, we exemplarily deal with two airfoil design test cases. A computational fluid dynamics (CFD) tool based on the solutions of Navier–Stokes equations computes the properties of the airfoils proposed by the optimisation technique. The parameterisation of the airfoils is done using Bezier-splines. One computation of an objective function value by time consuming CFD tools typically takes several minutes or more. Hence, only a limited number of evaluations can be afforded. We allow 1000 evaluations to stay comparable to previous studies on these test problems.

An airfoil redesign test case is the application investigated firstly. The two objectives defined are the differences in pressure distribution between target airfoils that are nearly optimal for given flow conditions and the

Table 3
Design conditions for the NACA redesign test case and the RAE drag minimisation test case

	NACA		RAE		
	Take-off	Cruise	Cruise	Off-design 1	Off-design 2
Mach no.	0.20	0.77	0.734	0.754	0.680
Reynolds no.	5×10^6	10^7	6.5×10^6	6.2×10^6	5.7×10^6
Angle of attack	10.8	1.0	2.8	2.8	1.8
Transition	3%	3%	3%	3%	11%

airfoil proposed by the optimisation technique. The target airfoils are NACA 0012 and NACA 4412, why we call this one the ‘NACA test case’. These flow conditions correspond to the take-off phase (high lift required) and the cruising phase (low drag desired) during an aircraft flight and can be taken from the left hand part of Table 3. For both flow conditions, the pressure distribution around the airfoil is calculated to derive the differences to the two target airfoils. The integrated differences serve as fitness function values for the test case. Due to the parameterisation in use, this test case features 18 degrees of freedom. A more detailed description can be found in [7].

Within the second test case, three drag coefficients are to be minimised. The considered airfoil type is the RAE 2822 airfoil, hence we will term the corresponding test case ‘RAE test case’. The different drag coefficients stem from three different flow conditions, each yielding different values for drag, lift and pitching moment. The conditions can be taken from the right hand part of Table 3. The task is to minimise the drag coefficients while not losing lift and keeping the pitching moment within a 2% range. Furthermore, geometrical constraints have to be met for the thickness of the airfoil measured at two particular points: the maximum thickness, its leading edge radius, and its trailing edge angle.

The geometrical information about a proposed airfoil can be received from the evaluation software just after the airfoil shape is generated. The whole time-consuming procedure of grid generation, solving the flow describing equations, and all post-processing tasks are not required to receive this information. Therefore, the geometrical constraints are treated differently from the aerodynamic ones, which require the costly evaluation. Variation operators are applied up to 1000 times unless a feasible solution subject to the geometrical constraints is obtained. This was done in order to achieve a higher ratio of feasible individuals. If no feasible solution is found after 1000 samples, the geometrical constraints will be treated like all other constraints. Due to the parameterisation in use here, the RAE test case features six degrees of freedom.

4.1. Fitness function approximations

We use Kriging metamodels as described by Sacks et al. [25] as fitness function approximation tools (FFA) to accelerate the SMS-EMOA in the presence of time consuming objective function evaluations. For a new design point $\mathbf{x} \in \mathbb{S}$, the Kriging metamodel allows for a prediction $\hat{y}_i(\mathbf{x}) (i \in \{1, \dots, m\})$ of the objective function value $f_i(\mathbf{x})$ from the set of previously evaluated (neighbouring) points stored in a database. Moreover, it estimates the uncertainty of the prediction by means of a standard deviation $\hat{s}_i(\mathbf{x})$.

In each generation, the *metamodel-assisted SMS-EMOA* [7] generates a surplus of individuals by mutations of one chosen individual. From these individuals, we retrieve the most promising solution by means of the metamodel. The chosen individual is then evaluated exactly using the precise evaluation function and considered in the $(\mu + 1)$ -selection. An optimistic guess of the objective function values, $\hat{y}_i(\mathbf{x}) - 2\hat{s}_i(\mathbf{x})$, $(i = 1, \dots, m)$ serves as a ranking criterion in the selection. The eligible individual is the one with the potentially highest value of $\Delta_{\mathcal{J}}$ [7] based on this optimistic guess. This selection criterion will be termed lower bound criterion in the following. Another variant is to apply simply the predicted function value for estimating the contribution in hypervolume, thereby neglecting its associated uncertainty measure. The latter strategy will be referred to as mean value criterion.²

4.2. Results

The SMS-EMOA provided very good and encouraging results on the theoretical test problems as well as on the design optimisation tasks. For the current study on airfoils, we collected five runs for each setting.

4.2.1. NACA results

Here, we considered SMS-EMOA without fitness function approximation (FFA) as well as the metamodel-assisted SMS-EMOA with mean value and lower bound criterion as described above. The points of the sets achieved using fitness function approximations are much better distributed than the ones obtained without.

² Whenever the potential increase of the hypervolume measure is zero for all solutions, the algorithm performs a random choice.

While runs without FFA are always biased towards one special region in the solution space, runs utilising FFA do not show any special crowding. The solutions are more equally distributed all over the Pareto front, with a desired higher density in regions with fair trade-offs as discussed above. The reason is the huge amount of pre-evaluations that are used to find promising regions of the search space to locate individuals that are to be evaluated exactly. Compared to the results with FFA, the runs without FFA do not seem to tap into their full potential probably due to the fact of too few evaluations.

The resulting sets without FFA are the worst fronts of all presented ones except for one special region. In the other regions, the SMS-EMOA with lower bound criterion seems to be better than the other algorithms shortly followed by old results from NSGA-II with lower bound criterion. This is the best result from [26] where different EMOA with and without FFA were studied. We therefore conclude, that the SMS-EMOA with lower bound criterion yields better results than all EMOA tested in [26].

The SMS-EMOA with mean value criterion provided the worst results utilising FFA within this study. The observation that the lower bound approximation technique offers better results than the mean value approximation was made in [26] as well. This seems to be a general achievement that deserves further attention.

4.2.2. RAE results

The results presented here have been obtained by the basic SMS-EMOA and are compared with results from an earlier publication, where a standard NSGA-II implementation was considered [27]. NSGA-II (with SBX, polynomial mutation and a population size of 20 individuals) was the best algorithm in that publication. We therefore conclude, that algorithms outperforming the NSGA-II results, provide better results than all other algorithmic variants studied in [27].

The NSGA-II results (termed NSGA-II-SBP-20 in Fig. 4) are compared to results from SMS-EMOA featuring the same variation operators and the same population size (SMS-SBP-20). Furthermore, SMS-EMOA was tested with standard variation operators from evolution strategies.³ This setting was tested with a population size of 10 and 20 individuals (SMS-ES-10, SMS-ES-20), respectively.

Fig. 4 presents the non-dominated solutions from the union of all five received Pareto front approximations per setting in the left hand part. The right hand part gives a median attainment surface consisting of all points from a Pareto front approximation that are weakly dominated by at least two of the other Pareto front approximations obtained with the same strategy.

In the top row of Fig. 4, the projections of the three-dimensional solutions into the f_1 – f_2 -plane are shown. The results demonstrate, that the fitness function components f_1 and f_2 are highly correlated to each other. Nearly all solutions appear to improve the baseline design (marked with a circle in the upper right corner). Furthermore, solutions received from one strategy seem to dominate each other. This is an artifact caused by the projection of the three-dimensional points. The high correlation between the fitness function components f_1 and f_2 causes the projections in planes f_1 – f_3 (middle row) and f_2 – f_3 (bottom row) to look similar. In the f_1 – f_3 -plane (middle row in Fig. 4), the variance in the results received from NSGA-II is lower than the ones from SMS-EMOA. The results of SMS-EMOA show a wider distribution, which is strictly in accordance with the intentions followed within the development of the algorithm.

From the right hand part of Fig. 4 (median attainment surface) it is hardly possible to distinguish the different results. The left hand part (the best obtained solutions) allow for a much better distinction. Here, a clear best setting can be recognised, which is SMS-EMOA, featuring ES variation operators and a population size of 10 individuals. The second best setting is the aforementioned one with a population size of 20 individuals. Both other strategies can be ranked equally with an advantage for the SMS-EMOA variant in distribution of solutions and quality in the third fitness function component (f_3). The NSGA-II behaves a bit better with respect to the other fitness function components (f_1 and f_2). Note, that SMS-ES-10 is significantly better than all other strategies with respect to all characteristics mentioned above.

³ We invoked one step size for each search direction, two randomly chosen parents to generate one offspring, intermediate recombination on the step sizes, discrete recombination on the object parameters and a simple Gaussian mutation operator for step sizes and object parameters according to Beyer and Schwefel [28].




Fig. 4. Scatter plots from the RAE 2822 drag minimisation test case. The top row shows the projection of the three-dimensional Pareto front approximations into the f_1 – f_2 -plane. The middle row shows the projection in the f_1 – f_3 -plane and the bottom row the projection in the f_2 – f_3 -plane. Four different settings are depicted: NSGA-II with SBX crossover, polynomial mutation and a population size of 20 individuals next to SMS-EMOA with the same parameterisation and the same variation operators as well as SMS-EMOA with standard ES variation operators with a population size of 10 and 20 individuals. The left column provides the non-dominated solutions of the union of solutions obtained in five runs per setting, while the right column presents the median attainment surfaces obtained.

The superiority of these results may be explained by the use of the ES variation operators in contrast to the SBX crossover and the polynomial mutation as well as the smaller population size. In concordance with our experience the SBX and polynomial mutation perform well on the theoretical test cases under investigation. Real-world problems, however, seems to be better addressed with the ES variation operators. The benefit of the smaller population size is detectable in all plots. This is probably due to the higher reproduction probability of each individual in smaller populations and the resulting higher selection pressure.

From an aerodynamic point of view, it is important to improve the baseline design with respect to all given flight conditions; this means finding solutions dominating the baseline design. Only one such result has been calculated until now, using the metamodel assisted NSGA-II as discussed in [27]. The ability to produce such good results is due to the fact that a large number of pre-evaluations is comprised. Our results show that

SMS-EMOA is always able to simultaneously improve all objectives of the baseline design. This can only be explained by the favourable hypervolume selection incorporated in SMS-EMOA.

5. Summary and prospects

Multiple results on standard benchmark problems indicate that the SMS-EMOA is well-suited for Pareto optimisation with two and three objectives. It clearly outperforms established techniques such as SPEA2, ϵ -MOEA, and NSGA-II, both in convergence and in \mathcal{S} metric values received. The examples reveal that results are well-distributed on the Pareto front. Boundaries of the Pareto front and regions around knee points are favoured. A distinguishing feature of the SMS-EMOA is that it is well-suited for approximating Pareto sets by a small number of individuals. This is often desired in practice.

A new selection criterion has been developed, which replaces the hypervolume-based selection in the presence of dominated solutions in the population. Using this new concept, the rating of a solution depends on the number of solutions dominating it. Hence, it focuses the evolutionary search towards less explored regions near the Pareto front. Both variants of the SMS-EMOA (the basic algorithm and the coupling with the number of dominated points criterion) led to better results than obtained by conventional strategies. This indicates that the hypervolume selection mechanism, performs well, regardless of the particular details of the selection scheme. This indicates a kind of robustness with regard to the special selection method chosen.

Besides the comparison on state-of-the-art benchmarks, the SMS-EMOA has been tested on challenging real-world applications, namely the optimisation of airfoils within two or three different flight conditions. Not only the distributions of solutions could be improved, but also solutions that clearly dominate the baseline design have been found. Additionally, SMS-EMOA was coupled with a metamodeling tool. This tool provides fast fitness function approximations to save costly exact evaluation in unfavourable regions. The SMS-EMOA was the first algorithm which always produced solutions dominating the baseline one. Moreover this algorithm produces the solutions robustly.

Future research should envisage the study of further algorithmic variants and a deepened analysis of strategy parameters. By now, the computational effort to compute the \mathcal{S} metric values is very high for more than three objectives. Therefore, the SMS-EMOA is hardly applicable to higher dimensional problems that normally desire a large number of function evaluations. Nevertheless, many real-world applications allow only a very limited number of evaluations due to costly simulations that govern the runtime of the optimisation process. In these cases, the runtime of the operations performed within the EMOA can almost be neglected, and the SMS-EMOA is certainly a well-suited optimiser.

Acknowledgements

This work was supported by the Deutsche Forschungsgemeinschaft (DFG) as part of the *Collaborative Research Center* ‘Computational Intelligence’ (SFB 531) as well as the research project SCHW 361/15-1 ‘Ein Verfahren zur Optimierung von aus Mehrkomponenten bestehenden Schiffsantrieben’. M. Emmerich acknowledges financial support of NWO within the HEFBOOM Project. Furthermore, the authors thank André Deutz for valuable comments and corrections.

References

- [1] K. Deb, Multi-Objective Optimization using Evolutionary Algorithms, John Wiley & Sons, Chichester, UK, 2001.
- [2] C.A. Coello Coello, D.A. Van Veldhuizen, G.B. Lamont, Evolutionary Algorithms for Solving Multi-Objective Problems, Kluwer Academic Publishers, New York, 2002.
- [3] E. Zitzler, L. Thiele, Multiobjective Optimization Using Evolutionary Algorithms – A Comparative Study, in: Parallel Problem Solving from Nature V, Springer, Amsterdam, 1998, pp. 292–301.
- [4] E. Zitzler, Evolutionary algorithms for multiobjective optimization: methods and applications, Ph.D. Thesis, Swiss Federal Institute of Technology (ETH), Zürich, Switzerland, 1999.
- [5] M. Fleischer, The Measure of Pareto Optima. Applications to Multi-objective Metaheuristics, Evolutionary Multi-Criterion Optimization (EMO 2003), LNCS, vol. 2632, Springer, Berlin, 2003, pp. 519–533.

- [6] J.D. Knowles, D.W. Corne, M. Fleischer, Bounded Archiving using the Lebesgue Measure, Congress on Evolutionary Computation (CEC'2003), vol. 4, IEEE Press, Piscataway, NJ, 2003, pp. 2490–2497.
- [7] M. Emmerich, N. Beume, B. Naujoks, An EMO Algorithm Using the Hypervolume Measure as Selection Criterion, Evolutionary Multi-Criterion Optimization (EMO 2005), LNCS, vol. 3410, Springer, Berlin, 2005, pp. 62–76.
- [8] J. Knowles, D. Corne, Properties of an adaptive archiving algorithm for storing nondominated vectors, *IEEE Transactions on Evolutionary Computation* 7 (2) (2003) 100–116.
- [9] K. Deb, M. Mohan, S. Mishra, Towards a Quick Computation of Well-Spread Pareto-Optimal Solutions, Evolutionary Multi-Criterion Optimization (EMO 2003), LNCS, vol. 2632, Springer, Berlin, 2003, pp. 222–236.
- [10] K. Deb, M. Mohan, S. Mishra, A fast multi-objective evolutionary algorithm for finding well-spread Pareto-optimal solutions, KanGAL report 2003002, Indian Institute of Technology, Kanpur, India, 2003.
- [11] E. Zitzler, L. Thiele, M. Laumanns, C.M. Fonseca, V. G. da Fonseca, Performance assessment of multiobjective optimizers: An analysis and review, TIK-Report 139, Swiss Federal Institute of Technology (ETH), Zürich, 2002.
- [12] S. Huband, P. Hingston, L. While, L. Barone, An evolution strategy with probabilistic mutation for multi-objective optimisation Congress on Evolutionary Computation (CEC 2003), vol. 4, IEEE Press, Piscataway, NJ, 2003, pp. 2284–2291.
- [13] E. Zitzler, S. Künzli, Indicator-based selection in multiobjective search, in: *Parallel Problem Solving from Nature (PPSN 2004)*, Springer, Berlin, 2004, pp. 832–842.
- [14] K. Deb, A. Pratap, S. Agarwal, T. Meyarivan, A fast and elitist multiobjective genetic algorithm: NSGA-II, *IEEE Transactions on Evolutionary Computation* 6 (2) (2002) 182–197.
- [15] D. Goldberg, *Genetic Algorithms in Search, Optimization and Machine Learning*, Addison-Wesley, Reading, MA, 1989.
- [16] B. Naujoks, N. Beume, M. Emmerich, Multi-objective optimisation using *S*-metric selection: Application to three-dimensional solution spaces, Congress on Evolutionary Computation (CEC 2005), vol. 2, IEEE Press, Piscataway, NJ, 2005, pp. 1282–1289.
- [17] L. While, A New Analysis of the LebMeasure Algorithm for Calculating the Hypervolume, Evolutionary Multi-Criterion Optimization (EMO 2005), LNCS, vol. 3410, Springer, Berlin, 2005, pp. 326–340.
- [18] J. Knowles, Local search and hybrid evolutionary algorithms for Pareto optimization, Ph.D. Thesis, University of Reading, Reading, UK, 2002.
- [19] E. Zitzler, Hypervolume metric calculation, Computer Engineering and Networks Laboratory (TIK), Zürich, 2001. Available from: <ftp://ftp.tik.ee.ethz.ch/pub/people/zitzler/hypervol.c>.
- [20] L. While, L. Bradstreet, L. Barone, P. Hingston, Heuristics for Optimising the Calculation of Hypervolume for Multi-objective Optimisation Problems, Congress on Evolutionary Computation (CEC 2005), vol. 3, IEEE Press, Piscataway, NJ, 2005, pp. 2225–2232.
- [21] E. Zitzler, M. Laumanns, L. Thiele, SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization, in: *Evolutionary Methods for Design, Optimisation, and Control, CIMNE*, Barcelona, Spain, 2002, pp. 95–100.
- [22] E. Zitzler, K. Deb, L. Thiele, Comparison of multiobjective evolutionary algorithms: Empirical results, *Evolutionary Computation* 8 (2) (2000) 173–195.
- [23] M. Emmerich, A rigorous analysis of two bi-criteria problem families with scalable curvature of the Pareto fronts, Technical Report LIACS TR 2005-05, Leiden Institute on Advanced Computer Science, Leiden, NL, 2005.
- [24] K. Deb, L. Thiele, M. Laumanns, E. Zitzler, Scalable Multi-objective Optimization Test Problems, Congress on Evolutionary Computation (CEC 2002), vol. 1, IEEE Press, Piscataway, NJ, 2002, pp. 825–830.
- [25] J. Sacks, W.J. Welch, W.J. Mitchell, H.P. Wynn, Design and analysis of computer experiments, *Statistical Science* 4 (1989) 409–435.
- [26] M. Emmerich, B. Naujoks, Metamodel assisted multiobjective optimisation strategies and their application in airfoil design, in: I.C. Parmee (Ed.), *Adaptive Computing in Design and Manufacture VI (ACDM 2004)*, Springer, London, 2004, pp. 249–260.
- [27] M. Emmerich, B. Naujoks, Metamodel-assisted multi-objective optimisation with implicit constraints and their application in airfoil design, in: *Proc. Int. Conf. ERCOFTAC'04, Athens (CD-ROM)*, CIMNE, Barcelona, Spain, 2004.
- [28] H.-G. Beyer, H.-P. Schwefel, Evolution strategies – A comprehensive introduction, *Natural Computing* 1 (1) (2002) 3–52.