An EMO Algorithm Using the Hypervolume Measure as Selection Criterion

Michael Emmerich*, Nicola Beume+, and Boris Naujoks+

+University of Dortmund, Chair of Systems Analysis,
44221 Dortmund, Germany
*University of Leiden, Leiden Institute for Advanced Computer Science,
2333 CA Leiden, NL
emmerich@liacs.nl, {nicola.beume, boris.naujoks}@cs.uni-dortmund.de
http://ls11-www.cs.uni-dortmund.de/
www.liacs.nl

Abstract. The hypervolume measure is one of the most frequently applied measures for comparing the results of evolutionary multiobjective optimization algorithms (EMOA). The idea to use this measure for selection is self-evident. A steady-state EMOA will be devised, that combines concepts of non-dominated sorting with a selection operator based on the hypervolume measure. The algorithm computes a well distributed set of solutions with bounded size thereby focussing on interesting regions of the Pareto front(s). By means of standard benchmark problems the algorithm will be compared to other well established EMOA. The results show that our new algorithm achieves good convergence to the Pareto front and outperforms standard methods in the hypervolume covered.
We also studied the applicability of the new approach in the important field of design optimization. In order to reduce the number of time consuming precise function evaluations, the algorithm will be supported by approximate function evaluations based on Kriging metamodels. First results on an airfoil redesign problem indicate a good performance of this approach, especially if the computation of a small, bounded number of well-distributed solutions is desired.

1 Introduction

Pareto optimization [1, 2] has become a well established technique for detecting interesting solution candidates for multiobjective optimization problems. It enables the decision maker to filter efficient solutions and to discover trade-offs between opposing objectives among these solutions. Provided a set of objective functions $f_{1,...,n} : \mathbb{S} \to \mathbb{R}$ defined on some search space $\mathbb{S}$ to be minimized, in Pareto optimization the aim is to detect the *Pareto-optimal set* $M = \{\mathbf{x} \in \mathbb{S} | \nexists \mathbf{x}' \in \mathbb{S} : \mathbf{x}' \prec \mathbf{x}\}$, or at least a good approximation to this set.

In practice, the decision maker wishes to evaluate only a limited number of Pareto-optimal solutions. This is due to the limited amount of time for examining the applicability of the solutions to be realized in practice. Typically these

solutions should include extremal solutions as well as solutions that are located in parts of the solution space, where balanced trade-offs can be found.

A measure for the quality of a non-dominated set is the hypervolume measure or $\mathcal{S}$ metric [3]. Until now, research mainly focussed on two approaches to utilize the $\mathcal{S}$ metric for multiobjective optimization: Fleischer [4] suggested to recast the multiobjective optimization problem to a single objective one by maximizing the $\mathcal{S}$ metric of a finite set of non-dominated points. Knowles et al. utilized the $\mathcal{S}$ metric within an archiving strategy for EMOA [5, 6].

Going one step further, our aim was to construct an algorithm in which the $\mathcal{S}$ metric governs the selection operator of an EMOA in order to find a set of solutions well distributed on the Pareto front. The basic idea of this EMOA is to integrate new points in the population, if replacing a member increases the hypervolume covered by the population. Moreover, we aimed at an algorithm that can easily be parallelized and is simple and transparent. It should be extendable by problem specific features, like approximate function evaluations. Thus, a steady-state $(\mu + 1)$-EMOA, the so-called $\mathcal{S}$ *metric selection EMOA (SMS-EMOA)*, is proposed.

Notice that in contrast to Knowles et al. [6], we do not evaluate an archiving operator solely, but the dynamics of a complete EMOA based on $\mathcal{S}$ metric selection. In our opinion, the design of an EA suitable for a given problem or a series of test problems is a multiobjective task again. This way we look at archiving strategies as only one component of the whole EMOA.

The article is structured as follows: The hypervolume or $\mathcal{S}$ metric that is used in the selection of our algorithm is discussed first (section 2). Afterwards, the integration in an EMOA as well as some features are described (section 3). Section 4 deals with the performance on several test problems whereas the results achieved on a real world design problem are the topic of section 5, including results with approximate function evaluations. In particular, the coupling of our method to a metamodel assisted fitness function approximation tool is presented here. We close with a summary and an outlook to implied future tasks (section 6).

2   The Hypervolume Measure

The hypervolume measure or S metric was originally proposed by Zitzler and Thiele [3], who called it the *size of the space covered* or *size of dominated space*. Coello Coello, Van Veldhuizen and Lamont [2] described it as the Lebesgue measure Λ of the union of hypercubes a_i defined by a non-dominated point m_i and a reference point x_{ref}:

$$\mathcal{S}(M) := \Lambda(\{\bigcup_i a_i | m_i \in M\}) = \Lambda(\bigcup_{m \in M} \{x | m \prec x \prec x_{ref}\}). \tag{1}$$

Zitzler and Thiele note that this measure prefers convex regions to non-convex ones [3]. A major drawback was the computational time for recursively calculating the values of $\mathcal{S}$. Knowles and Corne [5] estimated $O(k^{n+1})$ with k being the number of solutions in the Pareto set and n being the number

of objectives. Furthermore, an accurate calculation of the $\mathcal{S}$ metric requires a normalized and positive objective space and a careful choice of the reference point. In [5, 7] Knowles and Corne gave an example with two Pareto fronts, A and B, in the two dimensional case. They showed either $\mathcal{S}(A) < \mathcal{S}(B)$ or $\mathcal{S}(B) < \mathcal{S}(A)$ depending on the choice of the reference point.

Despite these disadvantages, the $\mathcal{S}$ metric is currently the only unary quality measure that is complete with respect to weak out-performance, while also indicating with certainty that one set is not worse than another [6]. It was used in several comparative studies of EMOA, e.g. [8, 9, 10]. Quite recently, Fleischer [4] proofed that the maximum of $\mathcal{S}$ is a necessary and sufficient condition for a finite true Pareto front ($|PF_{true}| < \infty$):

$$PF_{known} = PF_{true} \iff \mathcal{S}(PF_{known}) = max(\mathcal{S}(PF_{known})). \qquad (2)$$

Moreover, he developed a method for computing the $\mathcal{S}$ metric of a set in polynomial time: $O(k^3 n^2)$ [4]. This algorithm led to the efficient integration of the $\mathcal{S}$ metric in archiving strategies [6].

In addition, the $\mathcal{S}$ metric of a set of non-dominated solutions is suggested as a mapping to a scalar value. Fleischer proposed the use of metaheuristics to optimize this scalar. His idea was to try simulated annealing (SA) resulting in a provable global convergent algorithm towards the true Pareto front [4].

3   The Algorithm

Our aim was to design an EMOA that covers a maximal hypervolume with a limited number of points. Furthermore, we wanted to diminish the problem of choosing the right reference point. Our SMS-EMOA combines ideas borrowed from other EMOA, like the well established NSGA-II [11] and archiving strategies presented by Knowles, Corne, and Fleischer [5, 6]. It is a steady-state evolutionary algorithm with constant population size that firstly uses non-dominated sorting as a ranking criterion. Secondly the hypervolume is applied as selection criterion to discard that individual, which contributes least hypervolume to the worst-ranked Pareto-optimal front.

3.1   Details of the SMS-EMOA

A basic feature of the SMS-EMOA is that it updates a population of individuals within a steady-state approach, i. e. by generating only one new individual in each iteration. The basic algorithm is described in algorithm 1. Starting with an initial population of μ individuals, a new individual is generated by means of random variation operators[1]. The individual enters the population, if replacing a member

---

[1] We employed the variation operators used by Deb et al. for their ϵ-MOEA algorithm [10]. These are the SBX recombination 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/.

increases the hypervolume covered by the population. By this rule, individuals may always enter, if they replace dominated individuals and therefore contribute to a higher quality of the population. Apparently, the selection criterion assures that no non-dominated individual is replaced by a dominated one.

Before we will further explicate this selection strategy, we will spend a few more words on the steady-state approach. A steady-state scheme seems to be well suited for our approach, since it can be easily parallelized, enables the algorithm to keep a high diversity, and allows for an efficient implementation of the selection based on the hypervolume measure.

Algorithm 1 SMS-EMOA

1: $P_0 \leftarrow \text{init}()$ /* Initialize random start population of μ individuals */
2: $t \leftarrow 0$
3: **repeat**
4:    $q_{t+1} \leftarrow \text{generate}(P_t)$ /* Generate one offspring by variation operators */
5:    $P_{t+1} \leftarrow \text{Reduce}(P_t \cup \{q_{t+1}\})$ /* Select μ individuals for the new population */
6:    $t \leftarrow t + 1$
7: **until** stop criterium reached

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. A variable population size might lead to single individual populations in the worst case and therefore to a crucial loss of diversity for succeeding populations. If the population size is kept constant, the population may also have to include dominated individuals. In order to decide, which individuals are eliminated in the selection, also preferences among the dominated solutions have to be established.

Algorithm 2 Reduce(Q)

1: $\{\mathcal{R}_1, \ldots, \mathcal{R}_I\} \leftarrow$ fast-nondominated-sort(Q)
2:                                      /* all I non-dominated fronts of Q */
3: $r \leftarrow \text{argmin}_{s \in \mathcal{R}_I}[\Delta_\mathcal{S}(s, \mathcal{R}_I)]$ /* detect element of $\mathcal{R}_I$ with lowest $\Delta_\mathcal{S}(s, \mathcal{R}_I)$ */
4: $Q' \leftarrow Q \setminus \{r\}$ /* eliminate detected element */
5: **return** Q'

Algorithm 2 describes the replacement procedure `Reduce` employed. In order to decide, which individuals are kept in the population, the concept of Pareto front ranking from the well-known NSGA-II is be adopted. First, the Pareto fronts with respect to the non-domination level (or rank) are computed using the `fast-nondominated-sort`-algorithm [11]. Afterwards, one individual is discarded from the worst ranked front. If this front comprises $|\mathcal{R}_I| > 1$ individuals, the individual $s \in \mathcal{R}_I$ is eliminated that minimizes

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

For the case of two objective functions, we take the points of the worst-ranked non-dominated front and sort them ascending 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. Here for $\mathcal{R}_I = \{s_1, \ldots, s_{|\mathcal{R}_I|}\}$, $\Delta_{\mathcal{S}}$ is calculated as follows:

$$\Delta_{\mathcal{S}}(s_i, \mathcal{R}_I) = (f_1(s_{i+1}) - f_1(s_i)) \cdot (f_2(s_{i-1}) - f_2(s_i)). \tag{4}$$

3.2   Theoretical Aspects of the SMS-EMOA

The runtime complexity of the hypervolume procedure in the case of two objective functions is governed by the sorting algorithm. It is $O(\mu \cdot \log \mu)$, if all points lie on one non-dominated front. For the case of more than two objectives, we suggest to use the algorithm of Fleischer to calculate the contributing hypervolume $\Delta_{\mathcal{S}}$ of each point (compare [6]). Here, the runtime complexity of SMS-EMOA is governed by the calculation of the hypervolume and is $O(\mu^3 n^2)$.

The advantage of the steady-state approach is that only subsets of size $(|\mathcal{R}_I| - 1)$ have to be considered. By greedily discarding the individual that minimizes $\Delta_{\mathcal{S}}(s, \mathcal{R}_I)$, it is guaranteed that the subset which covers the maximal hypervolume compared to all $|\mathcal{R}_I|$ possible subsets remains in the population (for a proof we refer to Knowles and Corne [5]). With regard to the replacement operator this also implies that the covered hypervolume of a population cannot decrease by application of the `Reduce` operator, i. e. for algorithm 1 we can state the invariant:

$$\mathcal{S}(P_t) \leq \mathcal{S}(P_{t+1}). \tag{5}$$

Note, that the basic algorithm presented here nearly fits into the generic algorithm scheme AA_{reduce} presented by Knowles et al. [5] 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 and is the most similar one to our algorithm among those presented in [5].

Knowles et al. showed that the AA_S strategy converges to a subset of the true Pareto front and therefore to a local optimum of the $\mathcal{S}$ metric value achievable with a bounded set of points. A local optimum means that no replacement of an archive solution with a new one would increase the archive's $\mathcal{S}$ metric net value. 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 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 our algorithm as well. Knowles et al. analyzed the quality of local optima and remarked in [5] that the points of local optima of the $\mathcal{S}$ metric are "well distributed".

Often stated criticisms of the hypervolume measure regard the crucial choice of the reference point and the scaling of the search space. Our method of determining the solution contributing least to the hypervolume is actually independent from the choice of the reference point. The reference point is only needed to calculate the hypervolume of extremal points of a front and can alternatively be
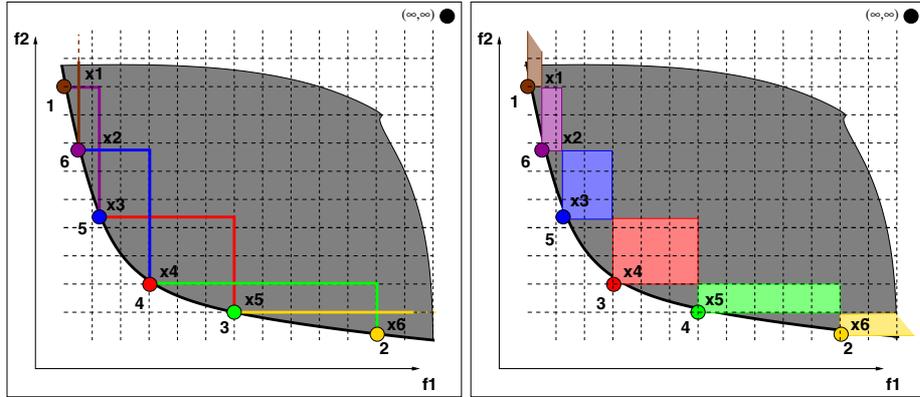
Fig. 1. Comparison of crowding distance sorting (left) and sorting by $\Delta_\mathcal{S}$ (right)

omitted, if extremal solution are to be kept anyway. Furthermore, our method is independent from the scaling of the objective space, in the sense that the order of solutions is not changed by multiplying the objective functions with a constant scalar vector.

3.3   Comparison of $\Delta_\mathcal{S}$ and the Crowding Distance

The similarity of the SMS-EMOA to the NSGA-II algorithm is noticable. The main differences between both procedures are the steady-state selection of the SMS-EMOA in contrast to the $(\mu + \mu)$ selection in NSGA-II and the different ranking of solutions located on the same Pareto front.

We would like to compare the crowding distance measure, that functions as ranking criterion for solutions of equal Pareto rank in NSGA-II, to the hypervolume based measure $\Delta_\mathcal{S}$. We recapitulate the definition of the crowding distance: It is defined as infinity for extremal solutions and as the sum of side lengths of the cuboid that touches neighboring solutions in case of non-extremal solutions on the Pareto front. It is meant to distribute solution points uniformly on the Pareto front. In contrast to this, the hypervolume measure is meant to distribute them in a way that maximizes the covered hypervolume.

In figure 1 a set R of non-dominated solutions is depicted in a two dimensional solution space. The left hand side figure shows the lines determining the ranking of solutions in the NSGA-II. The right hand side figure depicts the same solutions and their corresponding values of $\Delta_\mathcal{S}(s, R)$, which are given by the areas of the attached rectangles. Note that for the crowding distance, the value of a solution x_i depends on its neighbors and not directly on the position of the point itself, in contrast to $\Delta_\mathcal{S}(s, R)$. In both cases extremal solutions are ranked best, provided we choose a sufficiently large reference point for the hypervolume measure. Concerning the inner points of the front, x_5 (rank 3) outperforms x_4 (rank 4), if the crowding distance is used as a ranking criterion. On the other hand, x_4 (rank 3) outperforms x_5 (rank 4), if $\Delta_\mathcal{S}(s, R)$ is employed (right

figure). This indicates that good compromise solutions, which are located near knee-points of convex parts of the Pareto front are given better ranks in the SMS-EMOA than in the NSGA-II algorithm. Practically, solution x_5 is less interesting than solution x_4, since in the vicinity of x_5 little gains in objective f_2 can only be achieved at the price of large concession in objective f_1, which is not what is sought to be a well-balanced solution. Thus, the new method leads to more interesting solutions with fair trade-offs. It concentrates on knee-points without losing extremal points. This serves the practitioner who is mainly interested in a limited number of solutions on the Pareto front.

4   Test Problems

The SMS-EMOA from the last section was tested on several test problems from literature. We aimed at comparability to the papers of Deb and his coauthors presenting their ϵ-MOEA approach [9, 10]. That is why we also invoked the variation operators used for that approach. The test problems named ZDT1 to ZDT4 and ZDT6 from [10, 12] have been considered. For reasons of a clear overview, we copied the results for the hypervolume measure and the convergence achieved in [10] to table 1. This way, we compared our SMS-EMOA to NSGA-II, C-NSGA-II, SPEA2, and ϵ-MOEA.

4.1   Settings

We chose the parameters according to the ones given in [9, 10]. We set μ=100, calculated 20000 evaluations and used exactly the same variation operators as used for the ϵ-MOEA. The results of five runs are considered to create the values in table 1.

The hypervolume or $\mathcal{S}$ metric of the set of non-dominated points is calculated as described above, using the same reference point as in [9, 10]. The convergence measure is the average closest euclidean distance to a point of the true Pareto front as used in [10]. Note that the convergence measure is calculated concerning a set of 1000 equally distributed solution of the true Pareto front. Even an arbitrary point of the true Pareto front does not have a convergence value of 0, unless exactly equalling one of these 1000 points. Thus, the values are only comparable up to a certain degree of accuracy.

4.2   Results

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, two second 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. Building 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

Table 1. Results

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

$\mathcal{S}$ metric (with 6). So in conjunction, concerning this bundle of test problems, the SMS-EMOA can be regarded as the best one.

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 is a multi-modal function with multiple parallel Pareto fronts, whereas the best front is equivalent to that of ZDT1. On the basis of the given values from [9, 10], we assume that all algorithms achieved to jump above the second front with most solutions and aimed at the first front, like our SMS-EMOA. The worse values of the other algorithms seem to stem from disadvantageous distributions. ZDT2 has a smooth concave front and the SMS-EMOA covers most hypervolume, despite the criticism that the $\mathcal{S}$ metric favors convex regions. ZDT3 has a discontinuous Pareto 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 really better concerning the 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
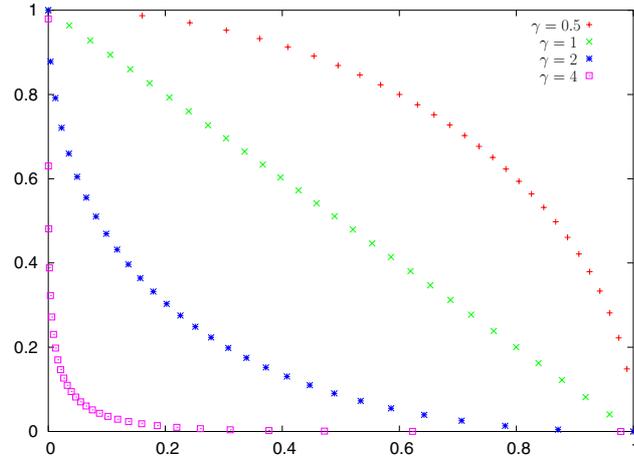
Fig. 2. This study visualizes results on the EBN problem family with Pareto fronts of different curvature computed by SMS-EMOA for a 20-dimensional search space

SMS-EMOA is ranked second on both measures, only outperformed by SPEA2, which shows apparently bad results on the other easier functions.

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. One should appreciate that our approach is a rather simple one with only one population and it is steady-state, resulting in a low selection pressure. Neither there are any special variation operators fitted to the selection strategy, nor it is tuned for performance in any way. All these facts would normally imply not that good results.

The good results in the convergence measure are maybe 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.

4.3   Distribution of Solutions

In order to get an impression of how the SMS-EMOA distributes solutions on Pareto fronts of different curvature, we conducted a study on simple but high dimensional test functions. The aim is to observe the algorithms behavior on convex, concave and linear Pareto fronts. For the study, we devised the following family of simple generic functions:

$$f_1(\mathbf{x}) = (\sum_{i=1}^{d} |x_i|)^{\gamma} d^{-\gamma}, \quad f_2(\mathbf{x}) := (\sum_{i=1}^{d} |x_i - 1|)^{\gamma} d^{-\gamma}, \quad \mathbf{x} \in [0,1]^d, \quad (6)$$

with d being the number of object variables. The ideal criterion vectors for these bicriterial problems (which we will abbreviate EBN) are given by $\mathbf{x}_1^* = (0, \ldots, 0)$,

$\mathbf{f}(\mathbf{x}_1^*) = (0,1)^T$ and $\mathbf{x}_2^* = (1,\ldots,1)$, $\mathbf{f}(\mathbf{x}_2^*) = (1,0)^T$. By the choice of the parameter γ the behavior of these functions can be adjusted. Parameter $\gamma = 1$ leads to a linear Pareto front, while $\gamma > 1$ yields convex fronts and $\gamma < 1$ concave ones.

Figure 2 shows that the solutions are not equally distributed on the Pareto front. The results demonstrate that the SMS-EMOA concentrates 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 always 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 emphasized. These are located near the angular point of the Pareto front. The results can be explained by the way the contributing hypervolume is defined and is discussed in the previous sections.

5   Design Optimization

A frequently addressed multiobjective design problem is the two-dimensional NACA redesign of an airfoil [13, 14]. Here, two target airfoils are given, each almost optimal for predefined flow conditions. A computational fluid dynamics (CFD) tool based on the solutions of Navier-Stokes equations calculates the properties, e.g. the pressure distribution of airfoils proposed by the coupled optimization technique. From these results, the differences in pressure distribution to the target airfoils are calculated and serve as the two objectives to minimize. The computation of objective function values based on CFD calculations are usually very time consuming with one evaluation typically taking several minutes, hence only a limited number of evaluations can be afforded. Here, we allow 1000 evaluations to stay comparable to previous studies on this test problem.

5.1   Integration of Fitness Function Approximations

We use Kriging metamodels [15] as fitness function approximation tools to accelerate the SMS-EMOA. The Kriging methods allows for a prediction of the objective function values for new design points $\mathbf{x}'$ from previously evaluated points stored in a database. Basically, Kriging is a distance based interpolation method. In addition to the predicted value, Kriging also provides a confidence value for each prediction. Based on the statistical assumption of Kriging, the predicted result $y(\mathbf{x}')$ and the confidence value $s(\mathbf{x}')$ can be interpreted as the mean value and standard deviation of a one-dimensional gaussian distribution describing the probability for the 'true' outcome of the evaluation. We refer to [15] for technical details of this procedure and the statistical assumptions about the continuous random process that – as it is assumed – generated the landscape $y(\mathbf{x})$.

As Kriging itself tends to be time consuming for a large number of training points, Kriging models are only build from the $2d$ nearest neighbors of each point, where d denotes the dimension of the search space.

Algorithm 3 Metamodel-assisted SMS-EMOA

1: $P_0 \leftarrow \text{init}()$ /* Initialize and evaluate start population of μ individuals */
2: $D \leftarrow P_0$ /* Initialize database of precisely evaluated solutions */
3: $t \leftarrow 0$
4: **repeat**
5:     Draw s_t randomly out of P_t
6:     $a_i \leftarrow \text{mutate}(s_t), i = 1, \ldots, \lambda$ /* Generate λ solutions via mutation */
7:     $\text{approximate}(D, a_1, \ldots, a_\lambda)$ /* Approximate results with local metamodels */
8:     $q_{t+1} \leftarrow \text{filter}(a_1, \ldots, a_\lambda)$ /* Detect 'best' approximate solution */
9:     evaluate q_{t+1} /* Evaluate selected solution precisely */
10:    $D \leftarrow D \cup \{q_{t+1}\}$
11:    $P_{t+1} \leftarrow \text{Reduce}(P_t \cup \{q_{t+1}\})$ /* Select new population of μ individuals */
12:    $t \leftarrow t + 1$
13: **until** stop criterion reached

The new method is depicted in algorithm 3. In order to make extensive use of approximate evaluations, it proved to be a good strategy, to produce a surplus of λ individuals by mutation of the same parent individual. For these new individuals an approximation is computed by means of the local metamodel. The filter procedure selects the most promising solution then. The chosen solution gets evaluated precisely and is considered for the `Reduce` method in the SMS-EMOA. This ensures that only precisely evaluated solutions enter the population P and that the amount of approximations employed can be scaled by the user. All precisely evaluated solutions enter a database, so they can subsequently be considered for the metamodeling procedure.

The basic idea of the filter algorithm is to devise a criterion based on the approximate evaluation of a search point. Criteria for the integration of approximations in EMOA have already been suggested in [14]. Here, confidence interval boxes in the solution space were calculated as $l_i = \hat{y}_i - \omega \hat{s}_i$ and $u_i = \hat{y}_i + \omega \hat{s}_i$, $i = 1, \ldots, n$, where n is the number of objectives and ω is a confidence factor that can be used to scale the confidence level. An illustrative example for approximations with Kriging and confidence interval boxes in a 2-D solution space is given in figure 3.

Among the criteria introduced in [14], two criteria seemed to be of special interest: First, the predicted result from the Kriging method, the mean value of the confidence interval box, is considered as a surrogate for the objective function value. This corresponds to the frequently employed approach to use merely the estimated function values as surrogates for the true objective functions and thus ignore the degree of uncertainty for these approximations. The second criterion goes one step further and upvalues those points with a high degree of uncertainty, by using the lower bound edge $\hat{\mathbf{y}} - \omega \hat{\mathbf{s}}$ of the interval boxes instead of its center $\hat{\mathbf{y}}$ for the prediction. This offers us a best case estimation for the solution.

Both surrogate points are employed to evaluate a criterion based on the $\mathcal{S}$ metric that is used for sorting the candidate solutions. For the mean value surrogate this is the most likely improvement (MLI) in hypervolume for population P when selecting $\mathbf{x}$:
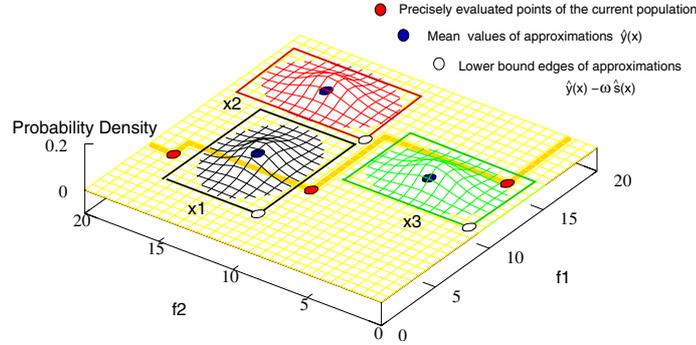
Fig. 3. Filtering of approximate solutions: Within the mean value criterion only x_3 is pre-selected while within the lower bound criterion the contributing hypervolume values of x_1 and x_3 are computed

$$\text{MLI}(\mathbf{x}) = \mathcal{S}(P_t \cup \{\hat{\mathbf{y}}(\mathbf{x})\}) - \mathcal{S}(P_t) \qquad (7)$$

and for the lower bound edge this is the potential improvement in hypervolume (LBI), that reads:

$$\text{LBI}(\mathbf{x}) = \mathcal{S}(P_t \cup \{\hat{\mathbf{y}} - \omega \hat{\mathbf{s}}(\mathbf{x})\}) - \mathcal{S}(P_t). \qquad (8)$$

It may occur that all values of the criterion are zero, if all surrogate points are dominated by the old population. In that case, the Pareto fronts of lower dominance level are considered for computing the values of MLI or LBI, respectively.

For the metamodel-assisted SMS-EMOA the user has to choose the parameters ω and λ. If the lower bound criterion is used, the choice of ω determines the degree of global search by the metamodel. For high values of ω the search focuses more on the unexplored regions of the search space.

5.2   Results

Like on the test problems, the SMS-EMOA provided very good and encouraging results on the design optimization problem. For this test series, we collected five runs for each setting again. We considered SMS-EMOA without fitness function approximation as well as the metamodel-assisted SMS-EMOA with mean value and lower bound criterion as described above.

For reasons of comparability, we utilized a method to average Pareto fronts from [13]. In short, parallel lines are drawn through the corresponding region of the search space. From the Pareto front of each run, the points with the shortest distance to these lines are considered for the calculation of the averaged front.

In the left hand part of figure 4 the different dotted sets describe three of the five Pareto fronts received from the different runs utilizing SMS-EMOA without Kriging. The line represents the received averaged Pareto front. This front is additionally copied to the right hand side figure for the reason of easier comparability. That figure compares the averaged fronts received using SMS-EMOA with and without fitness function approximations. In addition, a prior result, the best
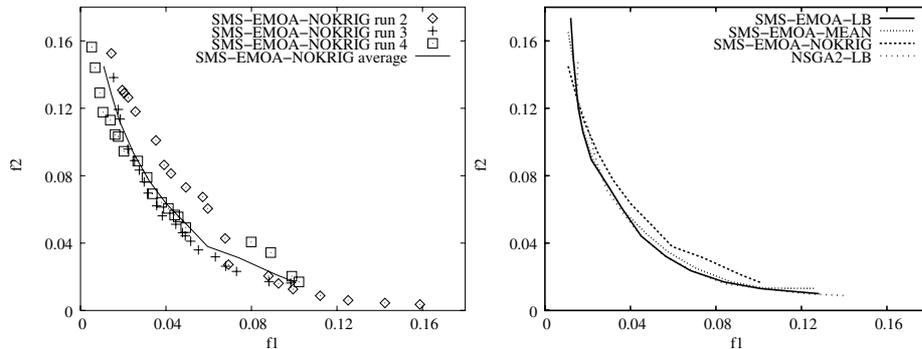
Fig. 4. The left hand side shows three of five runs used for averaging and the corresponding averaged front. The right hand side part compares SMS-EMOA without Kriging, using Kriging with lower bound (LB) and mean value (MEAN) criterion, next to NSGA-II using Kriging with lower bound criterion

one from the investigation presented in [14] coupling metamodeling techniques with multiobjective optimization is also included in the figure. This result stems from NSGA-II runs with Kriging and lower bound criterion within 1000 exact evaluations as well.

The points on the Pareto fronts achieved using fitness function approximations are much better distributed than the ones obtained without. In the left figure, each received Pareto front is biased towards one special region. In the runs utilizing fitness function approximations no focuses can be recognized. The solutions are more equally distributed all over the Pareto front, with the aspired higher density in regions with fair trade-offs as discussed above. The reason are the thousands of preevaluations that are used to find promising regions of the search space to place exact evaluations. Compared to the results with Kriging the runs without Kriging seem not to tap their full potential due to the too small amount of evaluations.

A clear superiority of the algorithms utilizing metamodels can be recognized. The averaged front without metamodel integration is the worst front all over the search space except for the upper left corner, the extreme f_2 flank of the front. In most other regions the SMS-EMOA with lower bound criterion seems to be better than the other algorithms shortly followed by the old results from NSGA-II with lower bound criterion. The SMS-EMOA with mean value criterion yielded the worst front with metamodel integration.

In the extreme f_2 flank of the front the results seem to be turned upside down. Here, the averaged front from runs without model integration achieved the best results. The left hand side of the figure, however suggests that this might be an effect of the averaging technique. It seems to be that one run achieves outstanding results here, which leads to an unbalanced average point that is better than the averaged points of the other algorithms. This extreme effect could be avoided by averaging over more than five runs which is a small and statistically not significant number of course.

Notice, that the lower bound approximation technique yielded better results than the mean value approximation again. This was also observed in [14] and seems to be a general achievement, where more attention should be drawn to.

6   Summary and Outlook

The SMS-EMOA has been devised in this work, which is a promising algorithm for Pareto optimization, especially if a small, limited number of solutions is desired and areas with balanced trade-offs shall be emphasized. The results on academic test problems show that the algorithm is rather competitive to established EMO algorithms like SPEA2 and NSGA-II regarding the convergence measure. It clearly outperforms these methods, if the $\mathcal{S}$ metric is considered as performance measure.

Compared to many other EMOA the new approach is simple and efficient for the two objective case. The selection and variation procedures do not interfere with an extra archive and the number of strategy parameters is very low (population size and reference point). Instead of specifying a reference point the SMS-EMOA can also work with an infinite reference point.

The focus of the performance assessment was on the two objective case. We demonstrated for this case that the approach is of special elegance, since its implementation is quite simple and the update of the population can be computed efficiently. Future research will have to deal with the performance assessment for three and more objectives and for constraint problems.

For a real world airfoil design problem Kriging metamodels have been employed to save time consuming precise function evaluations. The results indicate that these techniques can be used to further enhance the performance of the SMS-EMOA.

Acknowledgements

References

1. Deb, K.: Multi-Objective Optimization using Evolutionary Algorithms. John Wiley & Sons, Chichester, UK (2001)
2. Coello Coello, C.A., Van Veldhuizen, D.A., Lamont, G.B.: Evolutionary Algorithms for Solving Multi-Objective Problems. Kluwer Academic Publishers, New York (2002)
3. Zitzler, E., Thiele, L.: Multiobjective Optimization Using Evolutionary Algorithms—A Comparative Study. In Eiben, A.E., ed.: Parallel Problem Solving from Nature V, Amsterdam, Springer-Verlag (1998) 292–301

 4. Fleischer, M.: The Measure of Pareto Optima. Applications to Multi-objective Metaheuristics. In Fonseca, C.M., Fleming, P.J., Zitzler, E., Deb, K., Thiele, L., eds.: Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003, Faro, Portugal, Springer. Lecture Notes in Computer Science. Volume 2632 (2003) 519–533

 5. Knowles, J., Corne, D.: Properties of an Adaptive Archiving Algorithm for Storing Nondominated Vectors. IEEE Transactions on Evolutionary Computation **7** (2003) 100–116

 6. Knowles, J.D., Corne, D.W., Fleischer, M.: Bounded Archiving using the Lebesgue Measure. In: Proceedings of the 2003 Congress on Evolutionary Computation (CEC'2003). Volume 4., Canberra, Australia, IEEE Press (2003) 2490–2497

 7. Knowles, J., Corne, D.: On Metrics for Comparing Nondominated Sets. In: Congress on Evolutionary Computation (CEC'2002). Volume 1., Piscataway, New Jersey, IEEE Service Center (2002) 711–716

 8. Zitzler, E.: Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications. PhD thesis, Swiss Federal Institute of Technology (ETH), Zurich, Switzerland (1999)

 9. Deb, K., Mohan, M., Mishra, S.: Towards a Quick Computation of Well-Spread Pareto-Optimal Solutions. In Fonseca, C.M., Fleming, P.J., Zitzler, E., Deb, K., Thiele, L., eds.: Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003, Faro, Portugal, Springer. Lecture Notes in Computer Science. Volume 2632 (2003) 222–236

10. Deb, K., Mohan, M., Mishra, S.: A Fast Multi-objective Evolutionary Algorithm for Finding Well-Spread Pareto-Optimal Solutions. KanGAL report 2003002, Indian Institute of Technology, Kanpur, India (2003)

11. Deb, K., Pratap, A., Agarwal, S., Meyarivan, T.: A Fast and Elitist Multiobjective Genetic Algorithm: NSGA–II. IEEE Transactions on Evolutionary Computation **6** (2002) 182–197

12. Zitzler, E., Deb, K., Thiele, L.: Comparison of Multiobjective Evolutionary Algorithms: Empirical Results. Evolutionary Computation **8** (2000) 173–195

13. Naujoks, B., Willmes, L., T.Bäck, W.Haase: Evaluating multi-criteria evolutionary algorithms for airfoil optimisation. In Guervós, J.J.M., Adamidis, P., Beyer, H.G., Fernández-Villacañas, J.L., Schwefel, H.P., eds.: Parallel Problem Solving from Nature – PPSN VII, Proc. Seventh Int'l Conf., Granada, Berlin, Springer (2002) 841–850

14. Emmerich, M., Naujoks, B.: Metamodel assisted multiobjective optimisation strategies and their application in airfoil design. In Parmee, I.C., ed.: Adaptive Computing in Design and Manufacture VI, London, Springer (2004) 249–260

15. Sacks, J., Welch, W.J., Mitchell, W.J., Wynn, H.P.: Design and analysis of computer experiments. Statistical Science **4** (1989) 409–435