The Hypervolume Indicator Revisited: On the Design of Pareto-compliant Indicators Via Weighted Integration

Eckart Zitzler, Dimo Brockhoff, and Lothar Thiele

Computer Engineering (TIK), ETH Zurich
{zitzler,brockho,thiele}@tik.ee.ethz.ch
http://www.tik.ee.ethz.ch/sop/

Abstract. The design of quality measures for approximations of the Pareto-optimal set is of high importance not only for the performance assessment, but also for the construction of multiobjective optimizers. Various measures have been proposed in the literature with the intention to capture different preferences of the decision maker. A quality measure that possesses a highly desirable feature is the hypervolume measure: whenever one approximation completely dominates another approximation, the hypervolume of the former will be greater than the hypervolume of the latter. Unfortunately, this measure—as any measure inducing a total order on the search space—is biased, in particular towards convex, inner portions of the objective space. Thus, an open question in this context is whether it can be modified such that other preferences such as a bias towards extreme solutions can be obtained. This paper proposes a methodology for quality measure design based on the hypervolume measure and demonstrates its usefulness for three types of preferences.

1 Motivation

Using the hypervolume of the dominated portion of the objective space as a measure for the quality of Pareto set approximations has received more and more attention in recent years. The reason is that this measure has two important advantages over other set measures:

1. It is sensitive to any type of improvements, i.e., whenever an approximation set A dominates another approximation set B, then the measure yields a strictly better quality value for the former than for the latter set [23].
2. As a result from the first property, the hypervolume measure guarantees that any approximation set A that achieves the maximally possible quality value for a particular problem contains all Pareto-optimal objective vectors [5].

So far, this is the only measure known in the literature on evolutionary multi-criterion optimization that possesses these properties.

The hypervolume measure—or *hypervolume indicator* [23]—was first proposed and employed in [21,22] where it was denoted as 'size of the space covered'; later,

S. Obayashi et al. (Eds.): EMO 2007, LNCS 4403, pp. 862–876, 2007.

also other terms such as 'hyperarea metric' [14], 'S-metric' [18], and 'Lebesgue measure' [11,5] were used. On the one hand, the hypervolume indicator is meanwhile among the most popular measures for the performance assessment of multiobjective optimizers and in this context it has been subject to several theoretical investigations [8,5,23,15]. On the other hand, there are some studies that discuss the usage of this measure for multiobjective search [10,20,4] and in particular the issue of fast hypervolume calculation has been a focus of research recently [16,17,6,1].

Despite the aforementioned advantages of the hypervolume indicator, it inevitably has its biases. There is some freedom with respect to the choice of the reference point, but nevertheless it represents only one particular class of preference information that may not be appropriate in specific situations. This discussion directly leads to the question of whether it is possible to design quality measures that (i) share the two above properties of the hypervolume indicator, while (ii) standing for a different type of preferences of the decision maker. The fact that besides the hypervolume no other measures with these properties are known indicates that the formalization of *arbitrary* preferences in terms of a quality measure may be difficult. However, not being aware of such measures does not imply that such indicators do not exist.

This paper presents a first step to tackle this issue: it demonstrates that novel quality measures with the aforementioned properties can be designed and proposes a general design methodology on the basis of the hypervolume indicator. In detail, the key contributions are:

- A generalized definition of the hypervolume indicator using attainment functions [2] that can be used for any type of dominance relation;
- A weighted-integration approach to directly manipulate and control the influence of certain regions in the objective space for the hypervolume indicator;
- Three new example set measures for biobjective problems that provide the same sensitivity as the hypervolume indicator, but represent different types of preference information: (i) the preference of extreme solutions, (ii) the preference of predefined reference points, and (iii) bias towards one of the objectives.

The usefulness of the methodology and the three proposed measures is demonstrated on selected test problems.

2 Mathematical Framework

2.1 Preliminaries

Without loss of generalization, we consider a maximization problem with n objective functions $f_i : X \to (0,1)^n$, $1 \leq i \leq n$. Requiring the objective values to lay between 0 and 1 instead of using $\mathbb{R}^n$ as objective space simplifies the following discussions, but does not represent a serious limitation as there exists a bijective mapping from $\mathbb{R}$ into the open interval $(0,1) \subset \mathbb{R}$. The objective functions map a solution $\boldsymbol{x} \in X$ in the *decision space* to an objective vector $f(\boldsymbol{x}) = (f_1(\boldsymbol{x}), \ldots, f_n(\boldsymbol{x})) \in (0,1)^n$ in the *objective space* $Z = (0,1)^n$.

In the following, the (weak) Pareto-dominance relation $\succeq$ is used as a preference relation on the search space X indicating that a solution $\boldsymbol{x}$ is at least as good as a solution $\boldsymbol{y}$ ($\boldsymbol{x} \succeq \boldsymbol{y}$) if and only if $\forall 1 \leq i \leq n : f_i(\boldsymbol{x}) \geq f_i(\boldsymbol{y})$.[1] This relation can be canonically extended to sets of solutions where a set $A \subseteq X$ weakly dominates a set $B \subseteq X$ ($A \succeq B$) iff $\forall \boldsymbol{y} \in B\ \exists \boldsymbol{x} \in A : \boldsymbol{x} \succeq \boldsymbol{y}$ [23]. Note that any other type of dominance relation, e.g., based on arbitrary convex cones [13], could be used as well, and the considerations in this paper apply to any dominance relation.

Given the preference relation $\succeq$, we consider the optimization goal to identify a set of solutions that approximates the set of Pareto-optimal solutions and ideally this set is not strictly dominated by any other approximation set. For reasons of simplicity though, we assume that the outcome of a multiobjective optimizer is a set of objective vectors, also denoted as *approximation set*, and the set of all possible objective vector sets is denoted as $\Omega := 2^Z$. Therefore, we will also use the symbol $\succeq$ for objective vectors and objective vector sets, although it is originally defined on X. In practice, one always obtains a set of decision vectors instead of objective vectors, but most often only the objective vectors are considered to evaluate the quality of a solution set.

Since the generalized weak Pareto dominance relation $\succeq$ defines only a partial order on Ω, there may be incomparable sets in Ω which may cause difficulties with respect to search and performance assessment. These difficulties become more serious as the number of objectives in the problem formulation increases, see [3] for details. One way to circumvent this problem is to define a total order on Ω which guarantees that any two objective vector sets are mutually comparable. To this end, quality indicators have been introduced that assign, in the simplest case, each approximation set a real number, i.e., a (unary) indicator I is a function $I : \Omega \to \mathbb{R}$, cf. [23]. One important feature an indicator should have is *Pareto compliance* [9], i.e., it must not contradict the order induced by the Pareto dominance relation. In detail, this means that whenever $A \succeq B \wedge B \not\succeq A$, then the indicator value of A must not be worse than the indicator value of B. A stricter version of compliance would be to require that $A \succeq B\ \wedge\ B \not\succeq A$ implies that the indicator value of A is strictly better than the indicator value of B (if better means a higher indicator value):

$$A \succeq B\ \wedge\ B \not\succeq A \Rightarrow I(A) > I(B)$$

So far, the hypervolume indicator has been the only known indicator with this property.

[1] If $\boldsymbol{x} \succeq \boldsymbol{y}$, we say $\boldsymbol{x}$ *weakly dominates* $\boldsymbol{y}$. Two solutions $\boldsymbol{x}$ and $\boldsymbol{y}$ are called *incomparable* if neither weakly dominates the other one. If for two solutions $\boldsymbol{x}$ and $\boldsymbol{y}$ both $\boldsymbol{x} \succeq \boldsymbol{y}$ and $\boldsymbol{y} \not\succeq \boldsymbol{x}$ holds, we say that $\boldsymbol{x}$ is strictly better than $\boldsymbol{y}$ or $\boldsymbol{x}$ *strictly dominates* $\boldsymbol{y}$ ($\boldsymbol{x} \succ \boldsymbol{y}$). A solution $\boldsymbol{x}^* \in X$ is called *Pareto optimal* if there is no other $\boldsymbol{x} \in X$ that strictly dominates $\boldsymbol{x}^*$.

2.2 The Hypervolume Indicator

The classical definitions of the hypervolume indicator are based on volumes of polytopes [22] or hypercubes [5] and assume that Pareto dominance is the underlying preference relation. Here, we give a generalized definition based on attainment functions that allows to consider arbitrary dominance relations.

The attainment function [2] gives, roughly speaking, for each objective vector in Z the probability that it is weakly dominated by the outcome of a particular multiobjective optimizer. As only single sets are considered here, we can take a slightly simplified definition of the attainment function:

Definition 1 (Attainment function for an objective vector set). *Given a set $A \in \Omega$, the attainment function $\alpha_A : [0,1]^n \rightarrow \{0,1\}$ for A is defined as*

$$\alpha_A(\boldsymbol{z}) := \begin{cases} 1 & \text{if } A \succeq \{\boldsymbol{z}\} \\ 0 & \text{else} \end{cases}$$

for $\boldsymbol{z} \in Z$.

This definition is illustrated for weak Pareto dominance in Fig. 1 and applies to any type of dominance relation.

The concept of attainment functions can now be used to give a formal definition of the well known hypervolume indicator. It is simply defined as the volume of the objective space enclosed by the attainment function and the axes.

Definition 2 (Hypervolume indicator). *The hypervolume indicator I_H^* with reference point $(0,\ldots,0)$ can be formulated via the attainment function as*

$$I_H^*(A) := \int_{(0,\ldots,0)}^{(1,\ldots,1)} \alpha_A(\boldsymbol{z}) d\boldsymbol{z}$$

where A is any objective vector set in Ω.

In the following section, we will give a rough overview about the new concepts that are introduced in the paper and illustrate how $I_H^*(A)$ can be modified to incorporate preference information without violating compliance to Pareto dominance.

3 Introductory Example and Outline of the Proposed Approach

The attainment function, the integration over which gives the hypervolume for a given set A, is a binary function: all weakly dominated objective vectors are assigned 1, while the remaining objective vectors are assigned 0. That means all weakly dominated objective vectors have the same weight and contribute equally to the overall indicator value.

The main idea behind the approach proposed in this paper is to give different weights to different regions in the objective space. This can be achieved by

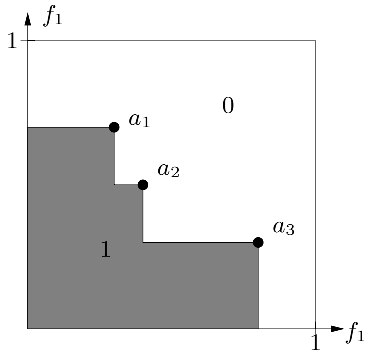

Fig. 1. Illustration of the attainment function α_A for $A = \{\boldsymbol{a}_1, \boldsymbol{a}_2, \boldsymbol{a}_3\}$ in the two-dimensional case

defining a weight distribution over the objective space such that the value that a particular weakly dominated objective vector contributes to the overall indicator value can be any real value greater than 0—provided the integral over the resulting function still exists. To this end, we introduce a weight distribution function $w : Z \rightarrow \mathbb{R}^+$, and the hypervolume is calculated as the integral over the product of the weight distribution function and the attainment function:

$$I_H^w(A) := \int_{(0,\dots,0)}^{(1,\dots,1)} w(\boldsymbol{z}) \cdot \alpha_A(\boldsymbol{z}) d\boldsymbol{z}$$

As will be shown later, thereby the basic hypervolume indicator can be modified such that (a) the compliance to Pareto dominance is preserved and (b) preference information can be flexibly introduced.

To see what the effect of different weight distribution functions is on the behavior of the corresponding modified hypervolume indicator I_H^w, it is helpful to consider equi-indicator surfaces. An equi-indicator surface $S(I, K)$ for a given indicator function I and an indicator value K is defined as the set of points $\boldsymbol{z} \in Z$ that each has an indicator value K, i.e.:

$$S(I, K) = \{\boldsymbol{z} \in Z \,:\, I(\{z\}) = K\}$$

In other words, the equi-indicator surface represents the indicator field for approximation sets with a single element.

If we consider a uniform weight distribution function with $w(\boldsymbol{z}) = 1$ for $\boldsymbol{z} \in Z$, we obtain the standard hypervolume indicator I_H^*. In this case, the equi-indicator surfaces looks for $n = 2$ as depicted in Fig. 2a; from the representation of these curves, one can conclude that the standard hypervolume indicator has convex equi-indicator surfaces and therefore implicitly introduces a preference towards solutions close the the diagonal. For instance, consider the point $(0.5, 0.5)$ located on the diagonal. To obtain the same indicator value for a point not on the diagonal, the degradation in one objective needs to be compensated by a larger

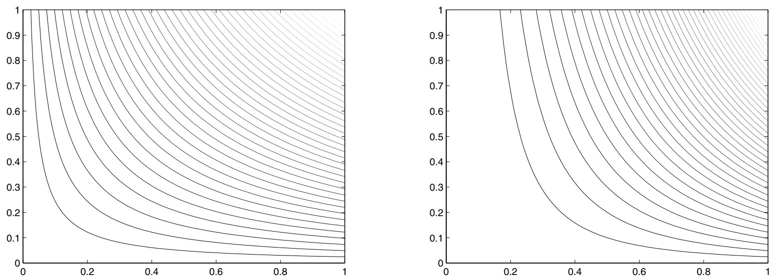

(a) The hypervolume indicator (b) A biased indicator

Fig. 2. Equi-indicator surfaces for simple indicators in the biobjective case. The abscissa in these two-dimensional examples denotes f_1 and the ordinate f_2. Figure (a) shows (a sample of) surfaces for the hypervolume indicator I_H^* (weight distribution function $w((z_1, z_2)) = 1$), Figure (b) illustrates a biased, modified indicator with weight distribution function $w((z_1, z_2)) = z_1$. Points on one equipotential curve share the same indicator value.

improvement in the other objective, e.g., $(0.25, 1)$ instead of $(0.25, 0.75)$ where degradation and improvement would be both 0.25.

If we now change the weight distribution function to $w(\boldsymbol{z}) = z_1$ with $\boldsymbol{z} = (z_1, z_2, \ldots, z_n)$, then in the biobjective case the equi-indicator surfaces shown in Fig. 2b are obtained. Obviously, solutions with objective vectors that have large components in the direction of z_1 are preferred.

Another possibility is to impose special emphasis on the border of the objective space, see Fig. 3a. The objective vectors in the 'center' of the objective space have weight 1, while the objective vectors on the axes are assigned a substantially larger weight.[2] The corresponding equi-indicator surfaces are shown in Fig. 3b. Here, the bias of the original hypervolume indicator for a single solution towards the diagonal is removed by putting more emphasis on the areas close to the coordinate axes.

The above two examples illustrate how weight distribution functions on the objective space can be used to change the bias of the hypervolume indicator. Based on these informal observations, we will describe the underlying methodology next.

4 Methodology: The Weighted-Integration Approach

The main concept of the approach proposed in this paper is to extend the basic hypervolume indicator by a weight distribution function $w : [0, 1]^n \rightarrow \mathbb{R}^+$ which serves to emphasize certain regions of the objective space:

[2] Since the borders have zero width, they will actually not influence the integral; therefore, dirac-type functions need to be used to make the border weights effective.

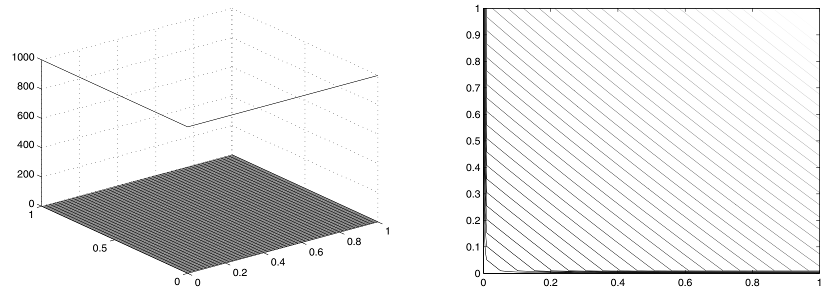

(a) Weight distribution function with emphasis on the coordinate axes

(b) (Sample of the) equi-indicator surfaces for the corresponding indicator

Fig. 3. Weight distribution function (left) and corresponding indicator (right) when stressing on coordinate axes

Definition 3 (Generalized Hypervolume Indicator). *The generalized hypervolume indicator I_H^w with weight distribution function $w : [0,1]^n \to \mathbb{R}^+$ is defined as the weighted integral*

$$I_H^w(A) := \int_{(0,\ldots,0)}^{(1,\ldots,1)} w(\boldsymbol{z}) \cdot \alpha_A(\boldsymbol{z}) d\boldsymbol{z}$$

where A is an approximation set in Ω.

If using this indicator as the basis for optimization algorithms or performance assessment tools, it would be important to know whether it is compliant with the concept of Pareto-dominance. This property will be shown next.

Theorem 1. *Let w be a weight distribution function $w : [0,1]^n \to \mathbb{R}^+$ such that the corresponding generalized hypervolume indicator I_H^w is well-defined for all $A \in \Omega$. Then for any two arbitrary approximation sets $A \in \Omega$ and $B \in \Omega$, it holds*

$$A \succeq B \ \wedge \ B \not\succeq A \Rightarrow I_H^w(A) > I_H^w(B).$$

Proof. If we have $A \succeq B \ \wedge \ B \not\succeq A$, then the following two conditions hold: $\forall \boldsymbol{y} \in B \ \exists \boldsymbol{x} \in A : \boldsymbol{x} \succeq \boldsymbol{y}$ and $\exists \boldsymbol{x} \in A \ \ \nexists \boldsymbol{y} \in B : \boldsymbol{y} \succeq \boldsymbol{x}$. Now we can easily see that the attainment functions of A and B satisfy $(\alpha_A(\boldsymbol{z}) = 1) \Rightarrow (\alpha_B(\boldsymbol{z}) = 1)$ as $A \succeq B$. Every point in the objective space that is weakly dominated by some element in B is also weakly dominated by some element in A. In addition, as $B \not\succeq A$ there are some points in the objective space that are weakly dominated by points in A but not weakly dominated by points in B. Therefore, there exists a region $R \subset Z$ with $(\alpha_A(\boldsymbol{z}) = 1) \wedge (\alpha_B(\boldsymbol{z}) = 0)$ for $\boldsymbol{z} \in R$; in particular:

$$\int_{(0,\ldots,0)}^{(1,\ldots,1)} (\alpha_A(\boldsymbol{z}) - \alpha_B(\boldsymbol{z})) d\boldsymbol{z} > 0$$

Using the definition of the generalized hypervolume indicator and noting that $w(\boldsymbol{z}) > 0$, we find $I_H^w(A) > I_H^w(B)$. □

In order to simplify the definition of weight distribution functions and to avoid the use of dirac-type functions, we use a slightly different representation of the generalized hypervolume indicator where line segments can be used to establish emphasis on zero-width regions such as axes. Every line segment l_i is specified by a start point $\boldsymbol{s}_i \in Z$, an end point $\boldsymbol{e}_i \in Z$, and a corresponding weight distribution function $\overline{w}_i : [0,1] \to \mathbb{R}_0^+$. Using these notation, we can rewrite the generalized hypervolume indicator according to Def. 3 as follows

$$I_H^{w,\overline{w}_1,\overline{w}_2,\ldots,\overline{w}_L}(A) := \int_{(0,\ldots,0)}^{(1,\ldots,1)} w(\boldsymbol{z}) \cdot \alpha_A(\boldsymbol{z}) \cdot d\boldsymbol{z} +$$

$$\sum_{i \in \{1,2,\ldots,L\}} \int_0^1 \overline{w}_i(\boldsymbol{z}) \cdot \alpha_A(\boldsymbol{s}_i + t \cdot (\boldsymbol{e}_i - \boldsymbol{s}_i)) \cdot dt$$

Assuming that the weight distribution functions are chosen such that all integrals are well-defined, it is easy to see that the property proven in Theorem 1 is preserved.

In the following, we will discuss three examples of useful weight distribution functions that will also be used for experimental results.

1. The first weight distribution function is the sum of two exponential functions in direction of the axes:

$$w^{ext}(\boldsymbol{z}) = (e^{20 \cdot z_1} + e^{20 \cdot z_2})/(2 \cdot e^{20})$$

   with $L = 0$. The effect is an indicator with preference of extremal solutions. Because of the weight distribution function's steep slope near the two axes, a Pareto front approximation with solutions crowded near the axes yield a larger indicator value than a population with solutions in the interior region of the objective space where the weight distribution function contribute less to the indicator value.
2. The second weight distribution function focuses on the second objective by using an exponential function in f_2-direction:

$$w^{asym}(\boldsymbol{z}) = e^{20 \cdot z_2}/e^{20}$$

   In addition, the following line segment with a constant weight distribution function on the f_1-axis is used:

$$\overline{w}_1^{asym}(\boldsymbol{z}) = 400, \quad \boldsymbol{s}_1 = (0,0), \quad \boldsymbol{e}_1 = (1,0)$$

   where $L = 1$. This combination results in an indicator preferring solutions with extreme f_2 values and an additional solution near the f_1 axis. The additional line segment along the f_1 axis used here instead of an additional exponential function in f_1 direction yields only a single additional solution lying near the f_1 axis instead of many solutions with large f_1 value as with the weight distribution function defined above.

3. Often, a decision maker has some idea which points in the search space are the most important ones. With the third weight distribution function, we can integrate such information into a Pareto-compliant indicator. A point (a, b) of interest, also called reference point, can be chosen in advance. The weight distribution function defined below will direct the search of indicator based algorithms towards this point. Multiple reference points can be considered simultaneously by adding up the corresponding distinct weight distribution functions.

   The following weight distribution function is based on a ridge-like function through the reference point (a, b), parallel to the diagonal:

$$w^{ref}(\boldsymbol{z}) = \begin{cases} c + \frac{(2-((2(x-a))^2+(2(y-b))^2))}{(0.01+(2(x-a)-2(y-b))^2)} & \text{if } |z_1 - a| < 0.5 \wedge |z_2 - b| < 0.5 \\ c & \text{else} \end{cases}$$

   with $L = 0$. The constant c should be chosen small in comparison to the values of the ridge; here, we use $c = 10^{-5}$.

The computation of the generalized hypervolume indicator is based on the representation described above. It first partitions the whole unit hypercube $[0, 1]^n$ into smaller hyperrectangles based on the objective vectors contained in the set A, and then the weighted volumes of these hyperrectangles are added. To this end, the above weight distribution functions have been symbolically integrated using a commercial symbolic mathematics tool.

5 Proof-of-Principle Results

5.1 Simple Indicator-Based Optimization Algorithm

For the experimental validation of the weighted-integration approach, a simple indicator-based evolutionary algorithm (SIBEA) is considered that uses similar concepts as proposed in [10,20,4,7]. As the purpose of this section is to show the influence of preference information which has been incorporated into the generalized hypervolume indicator and not to compare different optimization algorithms, methods to improve the convergence rate such as fitness-based mating selection are not taken into account.

SIBEA (Simple Indicator-Based Evolutionary Algorithm)

Input: population size μ; number of generations N; indicator function I;

Output: approximation of Pareto-optimal set A;

Step 1 (Initialization): Generate an initial set of decision vectors P of size μ; set the generation counter $m := 0$.

Step 2 (Environmental Selection): Iterate the following three steps until the size of the population does no longer exceed μ:

1. Rank the population using Pareto dominance and determine the set of individuals $P' \subseteq P$ with the worst rank.

2. For each solution $\boldsymbol{x} \in P'$ determine the loss $d(\boldsymbol{x})$ w.r.t. the indicator I if it is removed from P', i.e., $d(\boldsymbol{x}) := I(P') - I(P' \setminus \{x\})$.
3. Remove the solution with the smallest loss $d(\boldsymbol{x})$ from the population P (ties are broken randomly).

Step 3 (Termination): If $m \geq N$ then set $A := P$ and stop; otherwise set $m := m + 1$.

Step 4 (Mating): Randomly select elements from P to form a temporary mating pool Q of size μ. Apply variation operators such as recombination and mutation to the mating pool Q which yields Q'. Set $P := P + Q'$ (multi-set union) and continue with Step 2.

As to the environmental selection step, an issue are dominated individuals in the population: they never lead to a change in the indicator value which is entirely determined by the nondominated front of the population. Therefore, the population is first partitioned into fronts (Step 2.1) using the dominance rank (number of dominating individuals)[3], and only individuals located in the worst front are investigated for deletion.

Furthermore, we consider two scaling variants to obtain the maximum effect of the weighted integral: online and offline scaling. In the online variant, the objective function values are scaled to the interval $[0, 1]$ within each generation; to guarantee that boundary solutions contribute a non-zero hypervolume to the overall indicator value, for each axis a line segment with a constant weight distribution function is added. The offline variant does not scale the objective function values but the weighting distribution function. In detail, the weighted integral is only computed over and scaled to the region of the Pareto front, which needs to be known in advance. Since any approximation set outside this region would yield an indicator value of zero, the standard hypervolume indicator value, down-scaled such that is does not interfere with the weighted integral, is added.

5.2 Experiments

We now show how the three weight distribution functions defined above influence the search process of SIBEA for three biobjective test problems. For each weight distribution function, we derive two indicators, one for the online scaling method and one for offline scaling, resulting in six different indicators overall. We name the corresponding indicators I_H^{ext}, I_H^{asym}, and I_H^{ref} respectively, and distinguish between the online and the offline version. The same holds for the usual hypervolume indicator I_H, where we also distinguish between the two scaling methods.

The test functions ZDT1, ZDT3, and ZDT6, cf. [19], are optimized by a SIBEA run with population size 20 for 1000 generations.[4] Note, that the ZDT

[3] A nondominating sorting could be used as well.

[4] The individuals are coded as real vectors with 30 (ZDT1 and ZDT3) and 10 (ZDT6 decision variables, where the SBX-20 operator is used for recombination and a polynomial distribution for mutation. The recombination and mutation probabilities were set to 1.0, according to [3].

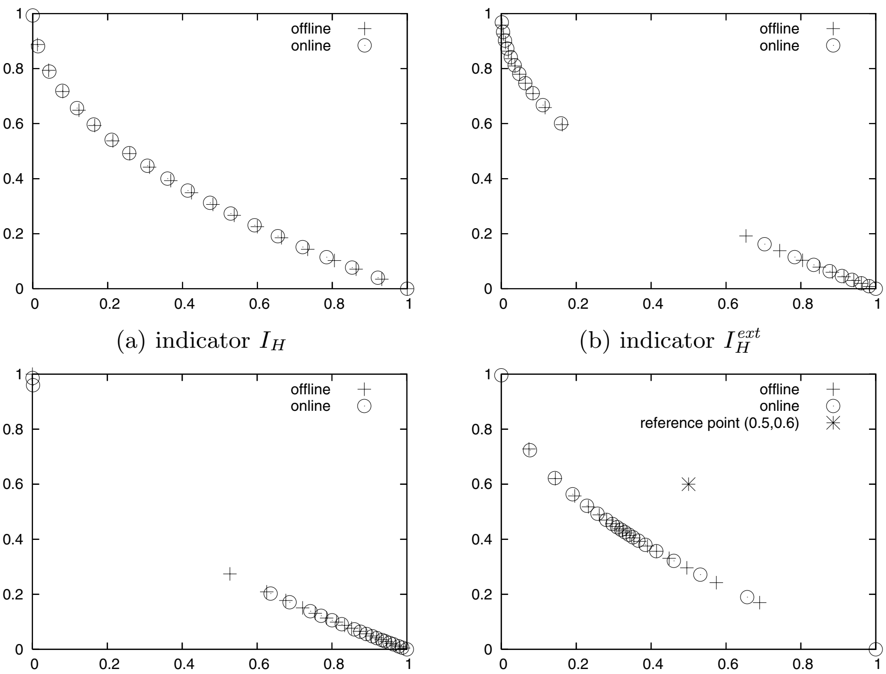

(a) indicator I_H (b) indicator I_H^{ext}

(c) indicator I_H^{asym} (d) indicator I_H^{ref}

Fig. 4. Pareto front approximations for the three different indicators based on weight distribution functions on the function ZDT1. For reference, the generated Pareto front approximation using the usual hypervolume indicator I_H is given in (a). The two scaling methods are plotted for comparison.

functions are to be minimized. Thus, an internal transformation is performed, independent whether the online or offline scaling is enabled.

The figures Fig 4, Fig. 5, and Fig. 6 show the computed Pareto front approximations after 1000 generations for the three ZDT functions and the three indicators I_H^{ext}, I_H^{asym}, and I_H^{ref} with both scaling methods. The reference point for I_H^{ref} is chosen as $(0.5, 0.6)$ for ZDT1 and ZDT6 and as $(0.5, 1.2)$ for ZDT3[5]. The approximation derived with the established hypervolume indicator I_H is also shown as golden reference.

The experiments show two main aspects. Firstly, the behavior of the evolutionary algorithm is similar for all three problems if always the same indicator is used—independent of the front shape and the scaling method used. With the indicator I_H^{ext} the solutions accumulate near the extremal points. When using the indicator I_H^{asym}, mainly the f_2 values are minimized. Due to the additional weight on the line segment, at least one solution with large f_1 value is also kept in the population if I_H^{asym} is used. With the indicator I_H^{ref}, the population moves towards

[5] The reference point is changed for ZDT3 due to the larger Pareto-optimal front of the ZDT3 problem.

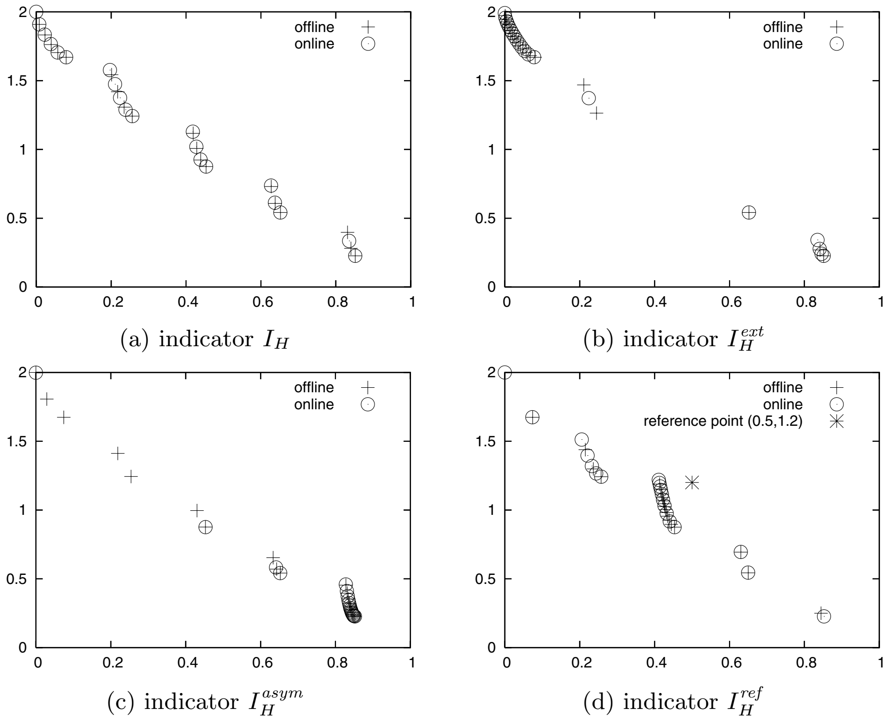

(a) indicator I_H (b) indicator I_H^{ext}

(c) indicator I_H^{asym} (d) indicator I_H^{ref}

Fig. 5. Pareto front approximations for the three different indicators and the two scaling methods on the function ZDT3. For reference, the generated Pareto front approximation using the usual hypervolume indicator I_H is given in (a). Due to the larger Pareto-optimal front, the reference point for I_H^{ref} is chosen as $(0.5, 1.2)$.

the predefined reference point $(0.5, 0.6)$, and $(0.5, 1.2)$ respectively. Secondly, the weighted-integration approach seems to be feasible for designing new Pareto-compliant indicators according to specific preferences. The simple indicator-based algorithm was indeed attracted to those regions in the objective space that were particularly emphasized by means of large weight values.

When comparing the two scaling variants, online and offline, only slight differences can be observed with the test cases studied in this paper. Online scaling has the advantage that the preferences are always adapted according to the current shape of the Pareto front approximation. However, thereby the actual global indicator changes during the run and potentially cycles can occur during the optimization process—a phenomenon, cf. [12], that emerges with most algorithms. Cycling is not necessarily a problem in the biobjective case, but as the number of objectives increases, it is likely that this behavior causes difficulties. The alternative is offline scaling. Here, the indicator remains fixed and can be used for comparing the outcomes of different methods. The drawback of this approach is the requirement that domain knowledge is available: either about the location of the Pareto front or about regions of interest. This problem holds basically for all types of indicators.

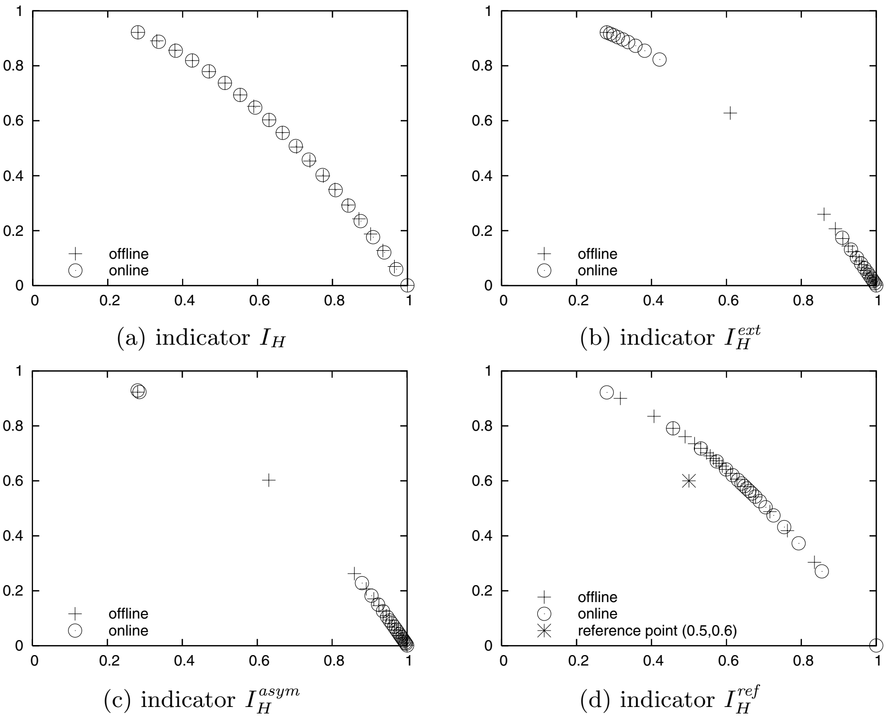

Fig. 6. Pareto front approximations for the three different indicators based on weight distribution functions on the function ZDT6. Plot (a) shows the generated Pareto front approximation using the usual hypervolume indicator I_H for comparison. The reference point for I_H^{ref} is chosen as $(0.5, 0.6)$.

6 Discussion

This paper has introduced a novel methodology to design Pareto-compliant indicators on the basis of the hypervolume indicator. Different preferences can be integrated, while an important property of the hypervolume indicator, sensitivity to dominance, is preserved. This is insofar an important result as up to now the pure hypervolume indicator has been the only one with this property. The possibility to design dominance-sensitive and Pareto-compliant indicators that can guide the search towards extreme solutions or reference points offers therefore more flexibility to tune the search with respect to the decision maker's preferences. We have demonstrated how this approach works and can be used for three example indicators in biobjective scenarios. As expected, the outcomes reflected the encoded preferences.

The presented methodology offers new ways for multiobjective optimization and performance assessment. However, this paper is just a first step in this direction, and the capabilities as well as the limitations of the weighted-integration approach need to be explored and require more research. In particular, the following considerations may point to interesting future research topics:

– The presented new indicators are designed for biobjective problems, but clearly one is interested in general indicators for n objectives; the first two indicators can be easily extended to higher dimensions, but for the ridge-based indicators this extension is not straight forward. The definition of general indicator classes for an arbitrary number of objectives will be one of the next steps to take.
– The efficient computation of the generalized hypervolume indicator based on weight distribution functions is especially an issue, if it is hard to obtain a function for the integral in closed form; here, numerical approximation might be a solution, although it is unclear how such an approach could work in practice.
– Whether novel indicators require new algorithms is an open issue; this holds in particular when other dominance relations based an arbitrary convex cones are used.

Acknowledgements

The authors would like to thank Stefan Bleuler, Tim Hohm, and Simon Künzli for valuable discussions. Dimo Brockhoff has been supported by the Swiss National Science Foundation (SNF) under grant 112079.

References

1. N. Beume and G. Rudolph. Faster S-Metric Calculation by Considering Dominated Hypervolume as Klee's Measure Problem. Technical Report CI-216/06, Sonderforschungsbereich 531 Computational Intelligence, Universität Dortmund, July 2006. shorter version published at IASTED International Conference on Computational Intelligence (CI 2006).
2. V. G. da Fonseca, C. M. Fonseca, and A. O. Hall. Inferential performance assessment of stochastic optimisers and the attainment function. In E. Zitzler, K. Deb, L. Thiele, C. A. C. Coello, and D. Corne, editors, *Evolutionary Multi-Criterion Optimization, First International Conference, EMO 2001*, pages 213–225. Springer-Verlag, Berlin, 2001.
3. K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable test problems for evolutionary multi-objective optimization. In A. Abraham, R. Jain, and R. Goldberg, editors, *Evolutionary Multiobjective Optimization: Theoretical Advances and Applications*, chapter 6, pages 105–145. Springer, 2005.
4. M. Emmerich, N. Beume, and B. Naujoks. An EMO Algorithm Using the Hypervolume Measure as Selection Criterion. In *Evolutionary Multi-Criterion Optimization (EMO). Third International Conference*, pages 62–76. Springer-Verlag Berlin, 2005.
5. M. Fleischer. The measure of Pareto optima. Applications to multi-objective metaheuristics. In C. M. Fonseca, P. J. Fleming, E. Zitzler, K. Deb, and L. Thiele, editors, *Evolutionary Multi-Criterion Optimization. Second International Conference, EMO 2003*, pages 519–533, Faro, Portugal, April 2003. Springer. Lecture Notes in Computer Science. Volume 2632.
6. C. M. Fonseca, L. Paquete, and M. López-Ibáñez. An Improved Dimension-Sweep Algorithm for the Hypervolume Indicator. In *IEEE Congress on Evolutionary Computation (CEC 2006)*, pages 1157–1163, Sheraton Vancouver Wall Centre Hotel, Vancouver, BC Canada, July 2006. IEEE Press.

7. C. Igel, N. Hansen, and S. Roth. The multi-objective variable metric evolution strategy, part i. Technical report, Institut für Neuroinformatik, Ruhr-Universität Bochum, Germany, 2005. IRINI-2005-04.
8. J. Knowles and D. Corne. On metrics for comparing non-dominated sets. In *Congress on Evolutionary Computation (CEC 2002)*, pages 711–716, Piscataway, NJ, 2002. IEEE Press.
9. J. Knowles, L. Thiele, and E. Zitzler. A Tutorial on the Performance Assessment of Stochastic Multiobjective Optimizers. TIK Report 214, Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Feb. 2006.
10. J. D. Knowles, D. W. Corne, and M. Fleischer. Bounded Archiving using the Lebesgue Measure. In *Proceedings of the 2003 Congress on Evolutionary Computation (CEC'2003*, pages 2490–2497, Canberra, Australia, Dec. 2006. IEEE Press.
11. M. Laumanns, G. Rudolph, and H.-P. Schwefel. Approximating the Pareto Set: Concepts, Diversity Issues, and Performance Assessment. Technical Report CI-72/99, 1999.
12. M. Laumanns, L. Thiele, K. Deb, and E. Zitzler. Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary Computation*, 10(3):263–282, 2002.
13. K. Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer, Boston, 1999.
14. D. A. V. Veldhuizen. *Multiobjective Evolutionary Algorithms: Classifications, Analyses, and New Innovations*. PhD thesis, Graduate School of Engineering, Air Force Institute of Technology, Air University, June 1999.
15. L. While. A New Analysis of the LebMeasure Algorithm for Calculating Hypervolume. In *Evolutionary Multi-Criterion Optimization. Third International Conference, EMO 2005*, pages 326–340, Guanajuato, México, Mar. 2005. Springer.
16. L. While, L. Bradstreet, L. Barone, and P. Hingston. Heuristics for Optimising the Calculation of Hypervolume for Multi-objective Optimisation Problems. In *IEEE Congress on Evolutionary Computation (CEC'2005)*, pages 2225–2232, IEEE Service Center, Edinburgh, Scotland, Sept. 2005. IEEE Press.
17. L. While, P. Hingston, L. Barone, and S. Huband. A Faster Algorithm for Calculating Hypervolume. *IEEE Transactions on Evolutionary Computation*, 10(1):29–38, Feb. 2006.
18. E. Zitzler. *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. PhD thesis, Swiss Federal Institute of Technology (ETH) Zürich, Switzerland, 1999.
19. E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation*, 8(2):173–195, 2000.
20. E. Zitzler and S. Künzli. Indicator-based selection in multiobjective search. In *Proc. 8th International Conference on Parallel Problem Solving from Nature (PPSN VIII)*, Birmingham, UK, September 2004.
21. E. Zitzler and L. Thiele. Multiobjective Optimization Using Evolutionary Algorithms - A Comparative case study. In *PPSN-V*, pages 292–301, Amsterdam, Sept. 1998.
22. E. Zitzler and L. Thiele. Multiobjective Evolutionary Algorithms: A Comparative Case Study and the Strength Pareto Approach. *IEEE Transactions on Evolutionary Computation*, 3(4):257–271, 1999.
23. E. Zitzler, L. Thiele, M. Laumanns, C. M. Foneseca, and V. Grunert da Fonseca. Performance assessment of multiobjective optimizers: An analysis and review. *IEEE Transactions on Evolutionary Computation*, 7(2):117–132, 2003.