

Generic Postprocessing via Subset Selection for Hypervolume and Epsilon-Indicator^{*}

Karl Bringmann¹, Tobias Friedrich², and Patrick Klitzke³

¹ Max Planck Institute for Informatics, Saarbrücken, Germany

² Friedrich-Schiller-Universität Jena, Jena, Germany

³ Universität des Saarlandes, Saarbrücken, Germany

Abstract. Most biobjective evolutionary algorithms maintain a population of fixed size μ and return the final population at termination. During the optimization process many solutions are considered, but most are discarded. We present two generic postprocessing algorithms which utilize the archive of all non-dominated solutions evaluated during the search. We choose the best μ solutions from the archive such that the hypervolume or ε -indicator is maximized. This postprocessing costs no additional fitness function evaluations and has negligible runtime compared to most EMOAs.

We experimentally examine our postprocessing for four standard algorithms (NSGA-II, SPEA2, SMS-EMOA, IBEA) on ten standard test functions (DTLZ 1–2,7, ZDT 1–3, WFG 3–6) and measure the average quality improvement. The median decrease of the distance to the optimal ε -indicator is 95%, the median decrease of the distance to the optimal hypervolume value is 86%. We observe similar performance on a real-world problem (wind turbine placement).

1 Introduction

Biobjective optimization aims at minimizing (or maximizing) a two-dimensional fitness function $f: \mathcal{X} \rightarrow \mathbb{R}^2$. As the two objectives f_1 and f_2 are typically contradicting, the outcome of the optimization is a set of incomparable solutions describing a Pareto front. Multiobjective evolutionary algorithms (MOEA) typically maintain a set of solutions called population during the optimization. The simplest MOEAs (like SEMO [13, 15, 16]) keep all non-dominated solutions in the population. As the Pareto front of a biobjective fitness function can be exponential in the input size [11], this results in exponential runtimes of SEMO for such fitness functions [7, 14]. More advanced MOEAs therefore avoid keeping all non-dominated solutions in the population and assume some upper limit on the size of the population.

^{*} The research leading to these results has received funding from the Australian Research Council (ARC) under grant agreement DP140103400 and from the European Union Seventh Framework Programme (FP7/2007–2013) under grant agreement no 618091 (SAGE). K.B. is a recipient of the *Google Europe Fellowship in Randomized Algorithms*, and this research is supported in part by this Google Fellowship.

With a population of fixed maximum size, MOEAs have to decide in each step which solutions to keep in the population and which solutions to remove. This can be done based on various measures like the hypervolume contribution [24], ε -approximation, crowding distance, or many others. Independent of the specific measure, the archiving algorithm of a MOEA has to solve an online problem: It has to decide which solutions to remove without knowing what new solutions will be generated in the future. While the MOEA has to decide ‘online’ which solutions to keep, an ‘offline’ algorithm would have access to all points generated during the optimization process and would just choose the best set from this archive. It is known that a MOEA which in each iteration keeps these μ solutions in the population that maximize the hypervolume, can still only reach a final hypervolume which is a factor μ smaller than achieved by the optimal choice of μ solutions from the whole archive (in the worst case) [3].

This shows that MOEAs potentially lose a lot of information by dropping solutions during the optimization. We suggest to make best use of the accumulated information by collecting all search points seen during the optimization in an archive. After the MOEA has stopped, we suggest to do a postprocessing that selects the best subset (of size μ) from the whole archive. This costs no additional fitness evaluations and should not reduce the quality of the reported final population. Depending on the ultimate aim of the optimization, the problem solved by this postprocessing is known as the *Hypervolume Subset Selection Problem* (HYPSSP) [1, 6, 17] or *ε -Indicator Subset Selection Problem* (EPS SSP) [6, 19, 21]. Until recently the best known runtimes for these subset selection problems were quadratic in the archive size n . As a single run of a MOEA can easily produce $n = 10^5$ or more non-dominated search points, a runtime of $\mathcal{O}(n^2)$ is prohibitively large. The postprocessing therefore only becomes tractable due to two recently published quasi-linear algorithms by the authors [6], specifically, an algorithm for HYPSSP with runtime $\mathcal{O}(n(\mu + \log n))$ and one for EPS SSP with runtime $\mathcal{O}(n \log n)$.

In this paper, we are going to investigate the effect of these two postprocessing algorithms for four standard algorithms on 10 common multi-objective optimization test functions and one real-world problem of optimizing the placement of wind turbines [20, 22].

Quality Measures. There are several ways to measure the quality of solution sets. We focus on two measures. Our first metric is the hypervolume indicator. It measures the volume of the objective space dominated by the set of solutions relative to a reference point. Its main disadvantage is its high computational complexity in higher dimensions [2, 4]. Our second metric is the ε -indicator. For measuring how well a set P approximates another set R , the (additive) ε -indicator returns the minimal ε by which we have to increase all points in P in all coordinates so that every point in R is dominated by some point in P . The disadvantage of this notion is that its computation requires knowing the Pareto front, as we would like to plug in the Pareto front for R . In contrast to the hypervolume indicator, it can therefore not directly be used to guide the search [5].

2 Preliminaries

We consider biobjective minimization problems, where a vector-valued function $f = (f_1, f_2): \mathcal{X} \rightarrow \mathbb{R}^2$ is minimized with respect to the weak Pareto dominance relation $\preceq$. We will mainly work in the *objective space* $f(\mathcal{X})$ and say that a point $p = (p_1, p_2) \in \mathbb{R}^2$ *weakly dominates* another point $q = (q_1, q_2) \in \mathbb{R}^2$ (denoted as $p \preceq q$) iff $p_1 \leq q_1$ and $p_2 \leq q_2$. The aim of most MOEAs is to approximate the *Pareto front* $\mathcal{F} = \{f(x) \mid x \in \mathcal{X}: \nexists y \in \mathcal{X}: f(y) \preceq f(x) \wedge f(x) \not\preceq f(y)\}$.

Hypervolume Indicator. For a set of points $P \subset \mathbb{R}^2$, the *hypervolume indicator* is defined as the volume of the set of points that are weakly dominated by solutions in P and at the same time weakly dominate a given reference point $r \in \mathbb{R}^2$, that is, $\mathcal{I}_{\text{hyp}}(P) := \mathcal{I}_{\text{hyp}}(P, r) := \lambda(\{z \in \mathbb{R}^2 \mid \exists a \in P: a \preceq z \preceq r\})$, where λ is the Lebesgue measure.

The *Hypervolume Subset Selection Problem* (HYPSSP) is then defined as follows: Given a set $P \subset \mathbb{R}^2$ of size n , $r \in \mathbb{R}^2$, and $\mu \in \mathbb{N}$, compute a subset $P^* \subseteq P$ of size at most μ that maximizes $\mathcal{I}_{\text{hyp}}(P^*, r)$. We write the result of this problem as $\text{HYPSSP}(P, r, \mu)$.

ε -Indicator. How well a point $p = (p_1, p_2) \in \mathbb{R}^2$ approximates another point $r = (r_1, r_2) \in \mathbb{R}^2$ in the objective space can be measured by the minimal number ε by which we have to decrease p in both coordinates so that it dominates r . More formally, we set $\mathcal{I}_{\text{eps}}(p, r) := \max\{p_1 - r_1, p_2 - r_2\}$. This can be used to define how well a set $P \subset \mathbb{R}^2$ approximates a set $R \subset \mathbb{R}^2$: The ε -indicator is defined as $\mathcal{I}_{\text{eps}}(P, R) := \max_{r \in R} \min_{p \in P} \mathcal{I}_{\text{eps}}(p, r)$. This denotes the minimal number ε by which we have to decrease all points in P in both coordinates so that every point in R is dominated by some point in P .

The ε -Indicator Subset Selection Problem (EPS SSP) is defined as follows: Given a set $P \subset \mathbb{R}^d$ of size n , $R \subset \mathbb{R}^d$ of size m , and $\mu \in \mathbb{N}$, compute a subset $P^* \subseteq P$ of size at most μ that minimizes $\mathcal{I}_{\text{eps}}(P^*, R)$. We write the result of this problem as $\text{EPS SSP}(P, R, \mu)$.

3 Postprocessing

Consider any EMOA with population size μ running until it performed n fitness evaluations. Let P be the final population after n fitness evaluations and A be the archive of all n solutions that were evaluated during the run. We describe our postprocessing in the objective space, i.e., we let $P, A \subset \mathbb{R}^2$.

Hypervolume. For $\mathcal{I}_{\text{hyp}}$ we may pick any reference point $r \in \mathbb{R}^2$. The general optimization goal is then to find a population P^* of μ Pareto optimal points that maximize the hypervolume, i.e., $P^* = \text{HYPSSP}(\mathcal{F}, r, \mu)$. Unfortunately, the Pareto front is unknown. To overcome this problem, we introduce the assumption that the archive converges to the Pareto front (as has been done in AGE [5] and is implicit in the design of most EMOAs). Thus, our $\mathcal{I}_{\text{hyp}}$ -postprocessing computes the set of μ points maximizing the hypervolume among all points in the archive,

$$\text{PP}_{\text{hyp}}(A, \mu) := \text{HYPSSP}(A, r, \mu).$$

Typically the hypervolume of $\text{PP}_{\text{hyp}}(A, \mu)$ should be larger than the hypervolume of P , so that our postprocessing improves the quality. In Section 5 we will see an experimental evaluation of this claim. In any case, since $P \subseteq A$ we have $\mathcal{I}_{\text{hyp}}(\text{PP}_{\text{hyp}}(A, \mu)) \geq \mathcal{I}_{\text{hyp}}(P)$, so our postprocessing does not decrease the quality of the result.

ε -Indicator. In case of $\mathcal{I}_{\text{eps}}$ the general optimization goal is to find a population P^* of μ Pareto optimal points that optimally approximate the Pareto front, i.e., $P^* = \text{EPS SSP}(\mathcal{F}, \mathcal{F}, \mu)$. Again, $\mathcal{F}$ is unknown and we assume that the archive converges to the Pareto front. Thus, our $\mathcal{I}_{\text{eps}}$ -postprocessing computes the set of μ points among all points in the archive that best approximate the archive,

$$\text{PP}_{\text{eps}}(A, \mu) := \text{EPS SSP}'(A, A, \mu).$$

Here, the prime in $\text{EPS SSP}'$ hides a minor modification that improves the experimental results, namely that we choose a population P among all μ -subsets of A that include the leftmost point of A and the bottommost point of A , i.e., the single-objective optima. Intuitively, this is necessary since these extrema are needed to cover the boundaries of the Pareto front. We remark that one can compute $\text{EPS SSP}'(A, A, \mu)$ with a minor modification of [6].

Again, typically $\mathcal{I}_{\text{eps}}(\text{PP}_{\text{eps}}(A, \mu), \mathcal{F})$ should be smaller than $\mathcal{I}_{\text{eps}}(P, \mathcal{F})$, so that our postprocessing improves the quality, and we will examine this claim experimentally. A noteworthy difference to the hypervolume case is that our postprocessing for $\mathcal{I}_{\text{eps}}$ does not come with the guarantee that quality cannot deteriorate, in fact, we will see in Section 5 that worsenings can happen but are rare.

Complexity. For a population size of μ and n fitness evaluations, NSGA-II, SPEA2, and IBEA have a running time of $\mathcal{O}(n\mu \log \mu)$, while SMS-EMOA has a running time of $\mathcal{O}(n\mu^2)$. Our postprocessing takes time $\mathcal{O}(n(\mu + \log n))$ for $\mathcal{I}_{\text{hyp}}$ and $\mathcal{O}(n \log n)$ for $\mathcal{I}_{\text{eps}}$ [6], which is comparable to the runtime of typical EMOAs. In our experiments, the postprocessing tends to be even faster, since we only have to store the non-dominated points of the archive, which are much less than n .

4 Experimental Setup

Implementation and Hardware. All presented algorithms have been implemented in Java using the jMetal framework [10] and run on a compute cluster with 128 nodes, each having two Intel Xeon E5620 @ 2.40GHz. The code is available at <http://docs.theinf.uni-jena.de/code/ssp.zip>.

Benchmark Problems and EMOAs. We compared the improvement gained by the postprocessing for the well established EMOAs NSGA-II [8], SPEA2 [26], SMS-EMOA [12], and IBEA [23]. As test functions we used DTLZ1, DTLZ2, DTLZ7 from DTLZ [9] with 7 variables, ZDT1–3 from ZDT [25] with 30 variables, and WFG3–6 from the [18] with 4 variables. We chose these benchmarks, because explicit expressions for their Pareto fronts are readily available. For measuring the hypervolume, we choose the reference point $r = (11, 11)$.

Additionally to the standard test functions, we used a simulation of a wind turbine placement function [20], which optimizes for the maximum power and the minimum perimeter of the convex hull formed by the turbine positions with 30 turbines on a discrete area of size 3000×3000 .

All experimental results (medians, quartiles, ...) that we will report in the next section are based on 700 independent runs for the benchmark problems and 100 independent runs for the turbine. As population size we used 100 for the benchmark problems and 10 for the turbine.

Quality Evaluation. We want to compare the quality $\mathcal{I}_{\text{hyp}}(P, r)$ of a population P with the optimal hypervolume $\text{OPT}_{\text{hyp}} := \text{HYPSSP}(\mathcal{F}, \mu)$ of any μ points on the Pareto front. For measuring the proximity to the front, we measure $\mathcal{I}_{\text{hyp}}^{\Delta}(P, r) := \text{OPT}_{\text{hyp}} - \mathcal{I}_{\text{hyp}}(P, r)$ in our experiments. However, as $\mathcal{F}$ is infinite it seems impossible to compute OPT_{hyp} . Instead of $\mathcal{F}$ we therefore consider a set $\mathcal{F}' \subseteq \mathcal{F}$ of m points placed equidistantly along the front, which we can compute because we know an explicit expression for $\mathcal{F}$ for the chosen benchmark problems. Now we simply replace $\mathcal{F}$ by $\mathcal{F}'$ in the definition of OPT_{hyp} to obtain an approximation. In case of $\mathcal{I}_{\text{eps}}$, for the same reasons we cannot directly compute $\text{OPT}_{\text{eps}} := \text{EPS SSP}(\mathcal{F}, \mathcal{F}, \mu)$. Moreover, we have the additional difficulty that we cannot evaluate the quality $\mathcal{I}_{\text{eps}}(P, \mathcal{F})$ of a population P . Again, we replace $\mathcal{F}$ by its finite approximation $\mathcal{F}'$ and obtain approximations for the optimum and the quality of a population. We use $\mathcal{I}_{\text{eps}}^{\Delta}(P, \mathcal{F}) := \text{OPT}_{\text{eps}} - \mathcal{I}_{\text{eps}}(P, \mathcal{F})$ as quality measure.

In our experiments we choose $m = 10^6$, which makes the error smaller than any of our reported values, and we ignore this error from now on. Note that it is infeasible to make m much larger, since the runtime for computing the approximation of, e.g., OPT_{hyp} is $\mathcal{O}(m(\mu + \log m))$.

Statistics. Additionally to calculating the (median) quality with and without postprocessing, we also perform a non-parametric test on the significance of the observed behavior. For this, we use the Wilcoxon-Mann-Whitney two-sample rank-sum test at the 95% confidence level.

5 Experimental Results

Test functions. The results of our experimental study on standard test functions are presented in Figure 1. The tables show the median of the indicators $\mathcal{I}_{\text{hyp}}^{\Delta}$ and $\mathcal{I}_{\text{eps}}^{\Delta}$ after 100 000 fitness evaluations. Our postprocessing improves (or does not worsen) the hypervolume and ε -indicator for all functions and algorithms. In all but six (out of 80 combinations) the improvement is statistically significant at the 95% confidence level. In fact, the median hypervolume is always increased and the distance to the optimal hypervolume $\mathcal{I}_{\text{hyp}}^{\Delta}$ therefore decreased. In 37 out of 40 cases the median $\mathcal{I}_{\text{hyp}}^{\Delta}$ decreases by more than 0.01%. The median reduction of $\mathcal{I}_{\text{hyp}}^{\Delta}$ for all 40 combinations of algorithms and functions is -85.6% (mean -60.5%). On the other hand, the median $\mathcal{I}_{\text{eps}}^{\Delta}$ could be decreased by more than 0.01% in 36 out of 40 cases. The median reduction of $\mathcal{I}_{\text{eps}}^{\Delta}$ is -95.2% (mean -73.4%).

Function	NSGA-II	IBEA	SPEA2	SMS-EMOA
DTLZ1	$4.7 \cdot 10^{-4} \rightarrow 7.1 \cdot 10^{-5}$ (-85.0%)	$8.1 \cdot 10^{-2} \rightarrow 3.1 \cdot 10^{-4}$ (-99.6%)	$2.0 \cdot 10^{-1} \rightarrow 5.4 \cdot 10^{-5}$ (-99.9%)	$3.1 \cdot 10^{-5} \rightarrow 2.8 \cdot 10^{-5}$ (-10.1%)
DTLZ2	$1.9 \cdot 10^{-3} \rightarrow 1.3 \cdot 10^{-5}$ (-99.3%)	$1.3 \cdot 10^{-2} \rightarrow 1.1 \cdot 10^{-5}$ (-99.9%)	$1.7 \cdot 10^{-1} \rightarrow 1.2 \cdot 10^{-5}$ (-99.9%)	$1.9 \cdot 10^{-5} \rightarrow 1.1 \cdot 10^{-5}$ (-41.3%)
DTLZ7	$2.1 \cdot 10 \rightarrow 2.1 \cdot 10$ ($\pm 0\%$)	$3.4 \cdot 10 \rightarrow 3.4 \cdot 10$ (-0.04%)	$2.2 \cdot 10 \rightarrow 2.1 \cdot 10$ (-0.6%)	$2.1 \cdot 10 \rightarrow 2.1 \cdot 10$ ($\pm 0\%$)
ZDT1	$1.9 \cdot 10^{-3} \rightarrow 1.8 \cdot 10^{-5}$ (-99.1%)	$5.4 \cdot 10^{-2} \rightarrow 8.2 \cdot 10^{-5}$ (-99.8%)	$6.0 \cdot 10^{-2} \rightarrow 2.3 \cdot 10^{-5}$ (-99.9%)	$2.7 \cdot 10^{-5} \rightarrow 1.4 \cdot 10^{-5}$ (-46.9%)
ZDT2	$1.7 \cdot 10^{-3} \rightarrow 1.4 \cdot 10^{-5}$ (-99.2%)	$1.6 \cdot 10^{-2} \rightarrow 1.1 \cdot 10^{-5}$ (-99.9%)	$1.8 \cdot 10^{-1} \rightarrow 2.9 \cdot 10^{-5}$ (-99.9%)	$3.7 \cdot 10^{-5} \rightarrow 1.9 \cdot 10^{-5}$ (-47.3%)
ZDT3	$1.1 \cdot 10^{-3} \rightarrow 6.1 \cdot 10^{-6}$ (-99.4%)	$3.5 \cdot 10^{-2} \rightarrow 3.3 \cdot 10^{-5}$ (-99.9%)	$5.7 \cdot 10^{-2} \rightarrow 9.2 \cdot 10^{-6}$ (-99.9%)	$1.3 \cdot 10^{-5} \rightarrow 5.1 \cdot 10^{-6}$ (-59.3%)
WFG3	$2.9 \cdot 10^{-2} \rightarrow 1.2 \cdot 10^{-2}$ (-60.1%)	$2.9 \cdot 10^{-1} \rightarrow 7.7 \cdot 10^{-3}$ (-97.4%)	$8.1 \cdot 10^{-1} \rightarrow 1.2 \cdot 10^{-2}$ (-98.5%)	$7.3 \cdot 10^{-3} \rightarrow 7.2 \cdot 10^{-3}$ (-2.1%)
WFG4	$1.8 \cdot 10^{-2} \rightarrow 1.7 \cdot 10^{-3}$ (-90.5%)	$9.5 \cdot 10^{-2} \rightarrow 3.8 \cdot 10^{-5}$ (-99.9%)	1.2 $\rightarrow 1.2 \cdot 10^{-2}$ (-99.0%)	$1.2 \cdot 10^{-2} \rightarrow 1.2 \cdot 10^{-2}$ (-0.4%)
WFG5	1.7 $\rightarrow 1.6$ (-0.8%)	1.7 $\rightarrow 1.6$ (-2.7%)	2.8 $\rightarrow 1.7$ (-40.3%)	1.8 $\rightarrow 1.8$ ($\pm 0\%$)
WFG6	$1.9 \cdot 10^{-1} \rightarrow 1.8 \cdot 10^{-1}$ (-7.2%)	$4.0 \cdot 10^{-1} \rightarrow 2.0 \cdot 10^{-1}$ (-49.0%)	1.3 $\rightarrow 1.9 \cdot 10^{-1}$ (-86.2%)	$2.3 \cdot 10^{-1} \rightarrow 2.3 \cdot 10^{-1}$ (-0.03%)

(a) Distance to optimal hypervolume: $\mathcal{I}_{\text{hyp}}^{\Delta}(P) = \text{OPT}_{\text{hyp}} - \mathcal{I}_{\text{hyp}}(P)$.

Function	NSGA-II	IBEA	SPEA2	SMS-EMOA
DTLZ1	$4.1 \cdot 10^{-3} \rightarrow 2.3 \cdot 10^{-4}$ (-94.3%)	$2.0 \cdot 10^{-1} \rightarrow 7.5 \cdot 10^{-3}$ (-96.2%)	$2.5 \cdot 10^{-2} \rightarrow 2.2 \cdot 10^{-4}$ (-99.1%)	$3.5 \cdot 10^{-4} \rightarrow 1.7 \cdot 10^{-4}$ (-50.4%)
DTLZ2	$8.2 \cdot 10^{-3} \rightarrow 1.7 \cdot 10^{-4}$ (-97.9%)	$4.0 \cdot 10^{-2} \rightarrow 1.2 \cdot 10^{-4}$ (-99.7%)	$2.4 \cdot 10^{-2} \rightarrow 1.4 \cdot 10^{-4}$ (-99.4%)	$6.7 \cdot 10^{-4} \rightarrow 1.2 \cdot 10^{-4}$ (-82.3%)
DTLZ7	2.3 $\rightarrow 2.3$ ($\pm 0\%$)	3.6 $\rightarrow 3.6$ (-0.03%)	2.3 $\rightarrow 2.3$ (-0.2%)	2.3 $\rightarrow 2.3$ ($\pm 0\%$)
ZDT1	$8.3 \cdot 10^{-3} \rightarrow 1.7 \cdot 10^{-4}$ (-98.0%)	$4.0 \cdot 10^{-2} \rightarrow 1.0 \cdot 10^{-4}$ (-99.7%)	$2.2 \cdot 10^{-2} \rightarrow 1.5 \cdot 10^{-4}$ (-99.3%)	$5.7 \cdot 10^{-4} \rightarrow 1.2 \cdot 10^{-4}$ (-79.7%)
ZDT2	$8.1 \cdot 10^{-3} \rightarrow 1.5 \cdot 10^{-4}$ (-98.1%)	$4.1 \cdot 10^{-2} \rightarrow 1.8 \cdot 10^{-4}$ (-99.6%)	$2.5 \cdot 10^{-2} \rightarrow 1.5 \cdot 10^{-4}$ (-99.4%)	$5.5 \cdot 10^{-4} \rightarrow 1.3 \cdot 10^{-4}$ (-76.1%)
ZDT3	$6.0 \cdot 10^{-3} \rightarrow 1.0 \cdot 10^{-4}$ (-98.3%)	$2.4 \cdot 10^{-2} \rightarrow 3.8 \cdot 10^{-4}$ (-98.4%)	$2.5 \cdot 10^{-2} \rightarrow 1.2 \cdot 10^{-4}$ (-99.5%)	$1.3 \cdot 10^{-3} \rightarrow 6.6 \cdot 10^{-5}$ (-94.9%)
WFG3	$2.5 \cdot 10^{-2} \rightarrow 1.1 \cdot 10^{-3}$ (-95.5%)	$1.1 \cdot 10^{-1} \rightarrow 6.5 \cdot 10^{-4}$ (-99.4%)	$8.3 \cdot 10^{-2} \rightarrow 1.1 \cdot 10^{-3}$ (-98.7%)	$1.9 \cdot 10^{-3} \rightarrow 6.2 \cdot 10^{-4}$ (-66.7%)
WFG4	$2.4 \cdot 10^{-2} \rightarrow 5.2 \cdot 10^{-4}$ (-97.8%)	$8.7 \cdot 10^{-2} \rightarrow 2.3 \cdot 10^{-4}$ (-99.7%)	$1.0 \cdot 10^{-1} \rightarrow 5.0 \cdot 10^{-4}$ (-99.5%)	$2.2 \cdot 10^{-3} \rightarrow 1.7 \cdot 10^{-4}$ (-92.4%)
WFG5	$1.2 \cdot 10^{-1} \rightarrow 1.2 \cdot 10^{-1}$ ($\pm 0\%$)	$1.3 \cdot 10^{-1} \rightarrow 1.2 \cdot 10^{-1}$ (-10.3%)	$2.1 \cdot 10^{-1} \rightarrow 1.2 \cdot 10^{-1}$ (-41.0%)	$1.4 \cdot 10^{-1} \rightarrow 1.4 \cdot 10^{-1}$ ($\pm 0\%$)
WFG6	$3.3 \cdot 10^{-2} \rightarrow 8.7 \cdot 10^{-3}$ (-73.8%)	$1.2 \cdot 10^{-1} \rightarrow 9.8 \cdot 10^{-3}$ (-91.9%)	$1.1 \cdot 10^{-1} \rightarrow 8.7 \cdot 10^{-3}$ (-92.3%)	$1.2 \cdot 10^{-2} \rightarrow 1.0 \cdot 10^{-2}$ (-16.2%)

(b) Distance to optimal ε -Indicator: $\mathcal{I}_{\text{eps}}^{\Delta}(P) = \text{OPT}_{\text{eps}} - \mathcal{I}_{\text{eps}}(P, \mathcal{F})$.

Fig. 1. Medians after 100 000 evaluations. We show the value of the respective indicator before and after the postprocessing. We also give the factor by which the indicator got smaller (=better). Improvements of less than 0.01% are treated as equal and marked $\pm 0\%$. For each test function we marked the best indicator values before and after postprocessing in **blue**.

Figure 2 gives a more detailed view on how the indicators $\mathcal{I}_{\text{hyp}}^{\Delta}$ and $\mathcal{I}_{\text{eps}}$ decrease over time for some exemplary combinations of MOEAs and test functions. Shown are box-and-whisker plots that specify the median, quartiles, and whiskers from minimum to maximum. Figures 2a and 2b show the typical behaviour of our postprocessing for NSGA-II. It is interesting to compare Figure 2b with Figure 3a, which shows the size of the population and of the archive over time. This shows that our postprocessing starts to “kick in” at the time we are starting to throw away points, where the archive size diverges from the population size.

Figure 2c shows the visually largest improvement that our $\mathcal{I}_{\text{hyp}}$ -postprocessing was able to achieve for any benchmark problem and MOEA. Note that SPEA2 without postprocessing converges to a suboptimal hypervolume, but our $\mathcal{I}_{\text{hyp}}$ -postprocessing is able to correct this. Figure 2d shows a similar profitable situation for $\mathcal{I}_{\text{eps}}$ -postprocessing.

Since SMS-EMOA is hypervolume driven, one expects that $\mathcal{I}_{\text{hyp}}$ -postprocessing is less effective for this algorithm compared to other EMOAs. Indeed, Figure 1 confirms this expectation, specifically, the median improvement of $\mathcal{I}_{\text{hyp}}^{\Delta}$ for SMS-EMOA is only -6.1% (mean -20.8%). However, Figure 2e shows that even the hypervolume driven SMS-EMOA can benefit a lot from our $\mathcal{I}_{\text{hyp}}$ -postprocessing in some situations.




Fig. 2. Quality measures as a function of time (evaluations) before (in red) and after (in green) postprocessing for some exemplary combinations of MOEAs and test functions.




Fig. 3. Population size and runtime of NSGA-II on WFG3

In contrast to the $\mathcal{I}_{\text{hyp}}$ -postprocessing, the $\mathcal{I}_{\text{eps}}$ -postprocessing can (at least theoretically) make $\mathcal{I}_{\text{eps}}$ worse. However, Figure 2f shows the only case where we observed that the $\mathcal{I}_{\text{eps}}$ -postprocessing visually worsens the algorithm at some time, namely in a thin region from 2500 to 5000 fitness evaluations. After that point, the $\mathcal{I}_{\text{eps}}$ -postprocessing gives again a huge improvement.

Figure 3b shows the total runtime up to n fitness evaluations of NSGA-II, the runtime of the other algorithms is similar (with SMS-EMOA being somewhat slower). Moreover, the runtime of the postprocessing when started after n fitness evaluations is plotted. This shows that the runtime of both postprocessing algorithms is negligible compared to the runtime of the MOEA.

Additionally, we examined whether a hypervolume-based algorithm (like SMS-EMOA) achieves more hypervolume than a non-hypervolume-based algorithms (like NSGA-II, SPEA2) with $\mathcal{I}_{\text{hyp}}$ -postprocessing. To this end, we compared SMS-EMOA without postprocessing to the other three algorithms with postprocessing. The Wilcoxon-Mann-Whitney U-test at 95% confidence level showed that NSGA-II with $\mathcal{I}_{\text{hyp}}$ -postprocessing outperforms SMS-EMOA without postprocessing on 7 out of 10 test functions. The same holds for SPEA2, but not for IBEA. This shows that our $\mathcal{I}_{\text{hyp}}$ -postprocessing makes algorithms that do not aim at maximizing $\mathcal{I}_{\text{hyp}}$ very competitive, even compared to algorithms that directly optimize $\mathcal{I}_{\text{hyp}}$. As one might expect, we no longer observe this behavior if SMS-EMOA is also allowed postprocessing: SMS-EMOA with $\mathcal{I}_{\text{hyp}}$ -postprocessing achieves higher hypervolumes than all other algorithms with $\mathcal{I}_{\text{hyp}}$ -postprocessing (on more than half of the test functions).

Figure 4 shows an exemplary result for the turbine problem for SMS-EMOA. Note that we cannot plot $\mathcal{I}_{\text{hyp}}^{\Delta}$ as OPT_{hyp} is not known. We observe a significant improvement also for this real-world problem. Specifically, after 300 000 evaluations our $\mathcal{I}_{\text{hyp}}$ -postprocessing increased $\mathcal{I}_{\text{hyp}}$ by 14%.




Fig. 4. Experimental results for turbine: $\mathcal{I}_{\text{hyp}}$ (not $\mathcal{I}_{\text{hyp}}^{\Delta}$!) over time (evaluations) for SMS-EMOA before (in red) and after (in green) postprocessing

6 Conclusion

We studied two generic postprocessing methods which have the potential to improve the final output of any EMOA on any biobjective optimization problem. These methods choose the optimal subset of μ solutions from the archive of all solutions seen during the run of an EMOA such that the hypervolume or ε -indicator is optimized. This requires no additional fitness evaluations and therefore zero additional ‘optimization time’. Moreover, the computation time of our postprocessing methods is negligible compared to the computation time of typical EMOAs. We experimentally evaluated the quality of our postprocessing

on four standard EMOAs and ten standard test functions and one real-world problem. This showed that our postprocessing typically returns a set of solutions which is about 90% closer to the optimum than the regular outcome of the EMOAs.

References

- [1] Auger, A., Bader, J., Brockhoff, D., Zitzler, E.: Investigating and exploiting the bias of the weighted hypervolume to articulate user preferences. In: Genetic and Evolutionary Computation Conference, GECCO 2009, pp. 563–570 (2009)
- [2] Bringmann, K., Friedrich, T.: Approximating the volume of unions and intersections of high-dimensional geometric objects. *Computational Geometry: Theory and Applications* 43, 601–610 (2010)
- [3] Bringmann, K., Friedrich, T.: Convergence of hypervolume-based archiving algorithms II: Competitiveness. In: Genetic and Evolutionary Computation Conference, GECCO 2012, pp. 457–464 (2012)
- [4] Bringmann, K., Friedrich, T.: Parameterized average-case complexity of the hypervolume indicator. In: Genetic and Evolutionary Computation Conference, GECCO 2013, pp. 575–582 (2013)
- [5] Bringmann, K., Friedrich, T., Neumann, F., Wagner, M.: Approximation-guided evolutionary multi-objective optimization. In: 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011, pp. 1198–1203. IJCAI/AAAI (2011)
- [6] Bringmann, K., Friedrich, T., Klitzke, P.: Two-dimensional subset selection for hypervolume and epsilon-indicator. In: Genetic and Evolutionary Computation Conference, GECCO (2014)
- [7] Brockhoff, D., Friedrich, T., Hebbinghaus, N., Klein, C., Neumann, F., Zitzler, E.: On the effects of adding objectives to plateau functions. *IEEE Trans. Evolutionary Computation* 13(3), 591–603 (2009)
- [8] Deb, K., Pratap, A., Agrawal, S., Meyarivan, T.: A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evolutionary Computation* 6(2), 182–197 (2002)
- [9] Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable test problems for evolutionary multiobjective optimization. In: *Evolutionary Multiobjective Optimization, Advanced Information and Knowledge Processing*, pp. 105–145 (2005)
- [10] Durillo, J.J., Nebro, A.J., Alba, E.: The jMetal framework for multi-objective optimization: Design and architecture. In: *IEEE Congress on Evolutionary Computation, CEC 2010*, pp. 4138–4325 (2010)
- [11] Ehrgott, M.: *Multicriteria Optimization*, 2nd edn. Springer (2005)
- [12] Emmerich, M.T.M., Beume, N., Naujoks, B.: An EMO algorithm using the hypervolume measure as selection criterion. In: *3rd International Conference on Evolutionary Multi-Criterion Optimization, EMO 2005*, pp. 62–76 (2005)
- [13] Friedrich, T., He, J., Hebbinghaus, N., Neumann, F., Witt, C.: Approximating covering problems by randomized search heuristics using multi-objective models. *Evolutionary Computation* 18(4), 617–633 (2010)
- [14] Friedrich, T., Hebbinghaus, N., Neumann, F.: Plateaus can be harder in multi-objective optimization. *Theoretical Computer Science* 411(6), 854–864 (2010)
- [15] Giel, O.: Expected runtimes of a simple multi-objective evolutionary algorithm. In: *IEEE Congress on Evolutionary Computation, CEC 2003*, pp. 1918–1925 (2003)

- [16] Giel, O., Lehre, P.K.: On the effect of populations in evolutionary multi-objective optimisation. *Evolutionary Computation* 18(3), 335–356 (2010)
- [17] Glasmachers, T.: Optimized approximation sets of low-dimensional benchmark pareto fronts. In: 13th International Conference on Parallel Problem Solving from Nature, PPSN (2014)
- [18] Huband, S., Barone, L., While, R.L., Hingston, P.: A scalable multi-objective test problem toolkit. In: 3rd International Conference on Evolutionary Multi-Criterion Optimization, EMO 2005, pp. 280–295 (2005)
- [19] Ponte, A., Paquete, L., Figueira, J.R.: On beam search for multicriteria combinatorial optimization problems. In: 9th International Conference in Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, CPAIOR 2012, pp. 307–321 (2012)
- [20] Tran, R., Wu, J., Denison, C., Ackling, T., Wagner, M., Neumann, F.: Fast and effective multi-objective optimisation of wind turbine placement. In: Genetic and Evolutionary Computation Conference, GECCO 2013, pp. 1381–1388 (2013)
- [21] Vaz, D., Paquete, L., Ponte, A.: A note on the ε -indicator subset selection. *Theoretical Computer Science* 499, 113–116 (2013)
- [22] Wagner, M., Day, J., Neumann, F.: A fast and effective local search algorithm for optimizing the placement of wind turbines. *Renewable Energy* 51, 64–70 (2013)
- [23] Zitzler, E., Künzli, S.: Indicator-based selection in multiobjective search. In: Yao, X., et al. (eds.) PPSN VIII. LNCS, vol. 3242, pp. 832–842. Springer, Heidelberg (2004)
- [24] Zitzler, E., Thiele, L.: Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Trans. Evolutionary Computation* 3, 257–271 (1999)
- [25] Zitzler, E., Deb, K., Thiele, L.: Comparison of multiobjective evolutionary algorithms: Empirical results. *Evolutionary Computation* 8(2), 173–195 (2000)
- [26] Zitzler, E., Laumanns, M., Thiele, L.: SPEA2: Improving the strength Pareto evolutionary algorithm for multiobjective optimization. In: *Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems*, EUROGEN 2001, pp. 95–100 (2002)