# Approximating the Least Hypervolume Contributor: NP-Hard in General, But Fast in Practice 

Karl Bringmann ${ }^{1}$ and Tobias Friedrich ${ }^{2}$<br>${ }^{1}$ Universität des Saarlandes, Saarbrücken, Germany<br>${ }^{2}$ International Computer Science Institute, Berkeley, CA, USA


#### Abstract

The hypervolume indicator is an increasingly popular set measure to compare the quality of two Pareto sets. The basic ingredient of most hypervolume indicator based optimization algorithms is the calculation of the hypervolume contribution of single solutions regarding a Pareto set. We show that exact calculation of the hypervolume contribution is \#P-hard while its approximation is NP-hard. The same holds for the calculation of the minimal contribution. We also prove that it is NP-hard to decide whether a solution has the least hypervolume contribution. Even deciding whether the contribution of a solution is at most $(1+\varepsilon)$ times the minimal contribution is NP-hard. This implies that it is neither possible to efficiently find the least contributing solution (unless $\mathbf{P}=\mathbf{N P}$ ) nor to approximate it (unless $\mathbf{N P}=\mathbf{B P P}$ ).

Nevertheless, in the second part of the paper we present a very fast approximation algorithm for this problem. We prove that for arbitrarily given $\varepsilon, \delta>0$ it calculates a solution with contribution at most $(1+\varepsilon)$ times the minimal contribution with probability at least $(1-\delta)$. Though it cannot run in polynomial time for all instances, it performs extremely fast on various benchmark datasets. The algorithm solves very large problem instances which are intractable for exact algorithms (e.g., 10000 solutions in 100 dimensions) within a few seconds.


## 1 Introduction

Multi-objective optimization deals with the task of optimizing several objective functions at the same time. As these functions are often conflicting, we cannot aim for a single optimal solution but for a set of Pareto optimal solutions. Unfortunately, the Pareto set frequently grows exponentially in the problem size. In this case, it is not possible to compute the whole front efficiently and the goal is to compute a good approximation of the Pareto front.

There are many indicators to measure the quality of a Pareto set, but there is only one widely used that is strictly Pareto compliant 22], namely the hypervolume indicator. Strictly Pareto compliant means that given two Pareto sets $A$ and $B$ the indicator values $A$ higher than $B$ if the Pareto set $A$ dominates the Pareto set $B$. The hypervolume (HYP) measures the volume of the dominated portion of the objective space. It was first proposed and employed for multi-objective optimization by Zitzler and Thiele 20].

It has become very popular recently and several algorithms have been developed to calculate it. The first one was the Hypervolume by Slicing Objectives (HSO) algorithm which was suggested independently by Zitzler [19] and Knowles [9]. To improve its runtime on practical instances, various speed up heuristics of HSO have been suggested [16, 18]. The currently best asymptotic runtime of $O\left(n \log n+n^{d / 2}\right)$ is obtained by Beume and Rudolph 2] by an adaption of Overmars and Yap's algorithm [11] for Klee's Measure Problem [8].

From a geometric perspective, the hypervolume indicator is just measuring the volume of the union of a certain kind of boxes in $\mathbb{R}_{\geq 0}^{d}$, namely of boxes which share the reference point $\sqrt[1]{1}$ as a common point. We will use the terms point and box interchangeably for solutions as the dominated volume of a point defines a box and vice versa. Given a set $M$ of $n$ points in $\mathbb{R}^{d}$, we define the hypervolume of $M$ to be

$$
\operatorname{HYP}(M):=\operatorname{vOL}\left(\bigcup_{\left(x_{1}, \ldots, x_{d}\right) \in M}\left[0, x_{1}\right] \times \ldots \times\left[0, x_{d}\right]\right)
$$

In [4] the authors have proven that it is $\# \mathbf{P}$-hard ${ }^{2}$ in the number of dimension to calculate HYP precisely. Therefore, all hypervolume algorithms must have an exponential runtime in the number of objectives (unless $\mathbf{P}=\mathbf{N P}$ ). Without the widely accepted assumption $\mathbf{P} \neq \mathbf{N P}$, the only known lower bound for any $d$ is $\Omega(n \log n)$ [3]. Note that the worst-case combinatorial complexity (i.e., the number of faces of all dimensions on the boundary of the union) of $\Theta\left(n^{d}\right)$ does not imply any bounds on the computational complexity.

Though the \#P-hardness of HYP dashes the hope for an exact subexponential algorithm, there are a few estimation algorithms [1, 4] for approximating the hypervolume based on Monte Carlo sampling. However, the only approximation algorithm with proven bounds is presented in [4]. There, the authors describe an FPRAS for HYP which gives an $\varepsilon$-approximation of the hypervolume with probability $(1-\delta)$ in time $\mathcal{O}\left(\log (1 / \delta) n d / \varepsilon^{2}\right)$.

## New Complexity Results

We will now describe a few problems related to the calculation of the hypervolume indicator and state our results. For this, observe that calculating the hypervolume itself is actually not necessary in a hypervolume-based evolutionary multi-objective optimizer as the algorithm actually only has to find a box with the minimal contribution to the hypervolume.

The contribution of a box $x \in M$ to the hypervolume of a set $M$ of boxes is the volume dominated by $x$ and no other element of $M$. We define the contribution $\operatorname{CON}(M, x)$ of $x$ to be

$$
\operatorname{CON}(M, x):=\operatorname{HYP}(M)-\operatorname{HYP}(M \backslash x)
$$

[^0]In Section 2 we show that this problem is \#P-hard to solve exactly. Furthermore, approximating CON by a factor of $2^{d^{1-\varepsilon}}$ is NP-hard for any $\varepsilon>0$. Hence, CON is not approximable. Note that this is no contradiction to the above-mentioned FPRAS for HYPas an approximation of HYP yields no approximation of CON.

As a hypervolume-based optimizer is only interested in the box with the minimal contribution, we also consider the following problem. Given a set $M$ of $n$ boxes in $\mathbb{R}^{d}$, find the least contribution of any box in $M$, that is,

$$
\operatorname{MINCON}(M):=\min _{x \in M} \operatorname{CON}(M, x)
$$

The reduction in Section 2 shows that MINCON is \# $\mathbf{P}$-hard and not approximable, even if we know the box which is the least contributor.

Both mentioned problems can be used to find the box contributing the least hypervolume, but their hardness does not imply hardness of the problem itself, which we are trying to solve, namely calculating which box has the least contribution. Therefore we also examine the following problem. Given a set $M$ of $n$ boxes in $\mathbb{R}^{d}$, we want to find a box with the least contribution in $M$, that is,

$$
\mathrm{LC}(M):=\underset{x \in M}{\operatorname{argmin}} \operatorname{CON}(M, x) .
$$

If there are multiple boxes with the same (minimal) contribution, we are, of course, satisfied with any of them. In Section 2 we prove that this problem is NP-hard to decide.

However, for practical purposes it most often suffices to solve a relaxed version of the above problem. That is, we just need to find a box which contributes not much more than the minimal contribution, meaning that is is only a $(1+\varepsilon)$ factor away. If we then throw out such a box, we have an error of at most $\varepsilon$. We will call this $\varepsilon-\mathrm{LC}(M)$ as it is an "approximation" of the problem LC. Given a set $M$ of $n$ boxes in $\mathbb{R}^{d}$ and $\varepsilon>0$, we want to find a box with contribution at most $(1+\varepsilon)$ times the minimal contribution of any box in $M$, that is,

$$
\operatorname{CON}(M, \varepsilon-\mathrm{LC}(M)) \leq(1+\varepsilon) \operatorname{MINCON}(M)
$$

The final result of Section 2 is the NP-hardness of $\varepsilon$-LC. This shows, that there is no way of computing the least contributor efficiently, and even no way to approximate it.

## New Approximation Algorithm

In Section 3 we will give a "practical" algorithm for determining a small contributor. Technically speaking, it solves the following problem we call $\varepsilon-\delta-\mathrm{LC}(M)$ : Given a set $M$ of $n$ boxes in $\mathbb{R}^{d}, \varepsilon>0$ and $\delta>0$, with probability at least $1-\delta$ find a box with contribution at most $(1+\varepsilon) \operatorname{MINCON}(M)$.

$$
\operatorname{Pr}[\mathrm{CON}(M, \varepsilon-\delta-\mathrm{LC}(M)) \leq(1+\varepsilon) \operatorname{MINCON}(M)] \geq 1-\delta
$$

As we will be able to choose $\delta$ arbitrarily, solving this problem is of high practical interest. By the NP-hardness of $\varepsilon$-LC there is no way of solving $\varepsilon-\delta$-LC
efficiently, unless $\mathbf{N P}=\mathbf{B P P}$. This means, our algorithm cannot run in polynomial time for all instances. Its runtime depends on some hardness measure H (cf. Section 3.2), which is an intrinsic property of the given input, but generally unbounded, i.e., not bounded by some function in $n$ and $d$.

However, in Section 4 we show that our algorithm is practically very fast on various benchmark datasets, even for dimensions completely intractable for exact algorithms like $d=100$ for which we can solve instances with $n=10000$ points within seconds. This implies a huge shift in the practical usability of the hypervolume indicator.

## 2 Hardness of Approximation

In this section we first show hardness of approximating MINCON, which we will use afterwards to show hardness of LC and $\varepsilon-\mathrm{LC}$. We will reduce \#MON-CNF to MINCON, which is the problem of counting the number of satisfying assignments of a Boolean formula in conjunctive normal form in which all variables are unnegated. While the problem of deciding satisfiability of such formula is trivial, counting the number of satisfying assignments is \#P-hard and even approximating it by a factor of $2^{d^{1-\varepsilon}}$ for any $\varepsilon>0$ is NP-hard, where $d$ is the number of variables (see Roth [14] for a proof).

Theorem 1. MINCON is \# $\mathbf{P}$-hard and approximating it by a factor of $2^{d^{1-\varepsilon}}$ is NP-hard for any $\varepsilon>0$.

Proof. To show the theorem, we reduce \#MON-CNF to MINCON. Let $\square\left(a_{1}, \ldots, a_{d}\right)$ denote a box $\left[0, a_{1}\right] \times \ldots \times\left[0, a_{d}\right]$. Let $f=\bigwedge_{k=1}^{n} \bigvee_{i \in C_{k}} x_{i}$ be a monotone Boolean formula given in CNF with $C_{k} \subseteq[d]:=\{1, \ldots, d\}$, for $k \in[n], d$ the number of variables, $n$ the number of clauses. First, we construct a box $A_{k}=\square\left(a_{1}^{k}, \ldots, a_{d}^{k}, 2^{d}+2\right) \subseteq \mathbb{R}^{d+1}$ for each clause $C_{k}$ with one vertex at the origin and the opposite vertex at $\left(a_{1}^{k}, \ldots, a_{d}^{k}, 2^{d}+2\right)$, where we set

$$
a_{i}^{k}=\left\{\begin{array}{ll}
1, & \text { if } i \in C_{k} \\
2, & \text { otherwise }
\end{array}, \quad i \in[d] .\right.
$$

Additionally, we need a box $B=\square(2, \ldots, 2,1) \subseteq \mathbb{R}^{d+1}$ and set $M=$ $\left\{A_{1}, \ldots, A_{n}, B\right\}$. Since we can assume without loss of generality that no clause is dominated by another, meaning $C_{i} \nsubseteq C_{j}$ for every $i \neq j$, every box $A_{k}$ uniquely overlaps a region $\left[x_{1}, x_{1}+1\right] \times \ldots \times\left[x_{d}, x_{d}+1\right] \times\left[1,2^{d}+2\right]$ with $x_{i} \in\{0,1\}, i \in[d]$, so that the contribution of every box $A_{k}$ is greater than $2^{d}$ and the contribution of $B$ is at most $2^{d}$, so that $B$ is indeed the least contributor.

Observe that the contribution of $B$ to $\operatorname{HYP}(M)$ can be written as a union of boxes of the form $B_{x_{1}, \ldots, x_{d}}=\left[x_{1}, x_{1}+1\right] \times \cdots \times\left[x_{d}, x_{d}+1\right] \times[0,1]$ with $x_{i} \in\{0,1\}, i \in[d]$. Moreover, $B_{x_{1}, \ldots, x_{d}}$ is not a subset of the contribution of $B$ to $\operatorname{HYP}(M)$ iff it is a subset of $\bigcup_{k=1}^{n} A_{k}$ iff it is a subset of some $A_{k}$ iff we have $a_{i}^{k} \geq x_{i}+1$ for $i \in[d]$ iff $a_{i}^{k}=2$ for all $i$ with $x_{i}=1$ iff $i \notin C_{k}$ for all $i$ with $x_{i}=1$ iff $\left(x_{1}, \ldots, x_{d}\right)$ satisfies $\bigwedge_{i \in C_{k}} \neg x_{i}$ for some $k$ iff $\left(x_{1}, \ldots, x_{d}\right)$
satisfies the negated formula $\bar{f}=\bigvee_{k=1}^{n} \bigwedge_{i \in C_{k}} \neg x_{i}$. This implies that $B_{x_{1}, \ldots, x_{d}}$ is a subset of the contribution of $B$ iff $\left(x_{1}, \ldots, x_{d}\right)$ satisfies $f$. Hence, since $\operatorname{voL}\left(B_{x_{1}, \ldots, x_{d}}\right)=1$, we have $\operatorname{MINCON}(M)=\operatorname{CON}(M, B)=\mid\left\{\left(x_{1}, \ldots, x_{d}\right) \in\right.$ $\{0,1\}^{d} \mid\left(x_{1}, \ldots, x_{d}\right)$ satisfies $\left.f\right\} \mid$. Thus a polynomial time algorithm solving MINCON $(M)$ would result in a polynomial time algorithm for \#MON-CNF, which proves the claim.

Note that the reduction from above implies that MINCON is \#P-hard and NP-hard to approximate even if the least contributor is known. Moreover, since we constructed boxes with integer coordinates in $\left[0,2^{d}+2\right]$ a number of $b=$ $\mathcal{O}\left(d^{2} n\right)$ bits suffices to represent all $d+1$ coordinates of the $n+1$ constructed points. Hence, MINCON is hard even if all coordinates are integral. We define as input size $b+n+d$, where $b$ is the number of bits in the input. We will use this result in the next proof. Also note that the same hardness for CON follows immediately, as it is hard to compute $\operatorname{CON}(M, B)$ as constructed above.

By reducing MINCON to LC, one can now show NP-hardness of LC. We skip this proof and directly prove NP-hardness of $\varepsilon$-LC by using the hardness of approximating MINCON in the following theorem.

Theorem 2. $\varepsilon$-LC is NP-hard for any constant $\varepsilon$. More precisely, it is NP-hard for $(1+\varepsilon)$ bounded from above by $2^{d^{1-c}-1}$ for some $c>0$.
Proof. We reduce MINCON to $\varepsilon$-LC. Let $M$ be a set of $n$ boxes in $\mathbb{R}^{d}$, i.e., a problem instance of MINCON represented by a number of $b$ bits, so that the input size is $b+n+d$.

As discussed above, we can assume that the coordinates are integral. We can further assume that $d \geq 2$ as MINCON is trivial for $d=1$. The minimal contribution of $M$ might be 0 , but this occurs if and only if one box in $M$ dominates another. As the latter can be checked in polynomial time, we can without loss of generality also assume that $\operatorname{MiNCON}(M)>0$.

Now, let $V$ be the volume of the bounding box of all the boxes in $M$, i.e., the product of all maximal coordinates in the $d$ dimensions. We know that $V$ is an integer with $1 \leq V \leq 2^{b}$, as there are only $b$ bits in the input.

We now define a slightly modified set of boxes:

$$
\begin{aligned}
A & =\left\{\square\left(a_{1}+2 V, a_{2}, \ldots, a_{d}\right) \mid \square\left(a_{1}, \ldots, a_{d}\right) \in M\right\}, \\
B & =\square(2 V, \ldots, 2 V), \\
C_{\lambda} & =\square(1, \ldots, 1,2 V+\lambda), \\
M_{\lambda} & =A \cup\{B\} \cup\left\{C_{\lambda}\right\} .
\end{aligned}
$$

The boxes in $A$ are the boxes of $M$, but shifted along the $x_{1}$-axis. By definition, $a_{i} \leq V, i \in[d]$ for all $\square\left(a_{1}, \ldots, a_{d}\right) \in M$. The contribution to $\operatorname{HYP}\left(M_{\lambda}\right)$ of a box in $A$ is the same as the contribution to $\operatorname{HYP}(M)$ of the corresponding box in $M$ as the additional part is overlapped by the "blocking" box $B$. Also note that the contribution of a box in $A$ is less or equal than $V$.

The box $B$ uniquely overlaps at least the space $[V, 2 V] \times \ldots \times[V, 2 V]$ (as every coordinate of a point in $M$ is less than equal to $V$ ) which has volume at
least $V$. Hence, $B$ is never the least contributor of $M_{\lambda}$. The box $C_{\lambda}$ then has a contribution of $\operatorname{VOL}([0,1] \times \ldots \times[0,1] \times[2 V, 2 V+\lambda])=\lambda$, so that $C_{\lambda}$ is a least contributor iff $\lambda$ is less than or equal to the minimal contribution of any box in $A$ to $\operatorname{HYP}\left(M_{\lambda}\right)$ which holds iff we have $\lambda \leq \operatorname{MINCON}(M)$.

Since we can decide, whether $C_{\lambda}$ is the least contributor, by one call to $\mathrm{LC}\left(M_{\lambda}\right)$, we can do kind of a binary search on $\lambda$. As we are interested in a multiplicative approximation, we search for $\kappa:=\log _{2}(\lambda)$ to be the largest value less than equal to $\log _{2}(\operatorname{MINCON}(M))$, where $\kappa$ now is an integer in the range $[0, b]$. As we can only answer $\varepsilon$-LC-queries we cannot do exact binary search. But we can still follow its lines, recurring on the left half of the current interval, if for the median value $\kappa_{\mathrm{m}}$ we get $\varepsilon-\mathrm{LC}\left(M_{\lambda_{\mathrm{m}}}\right)=C_{\lambda_{\mathrm{m}}}$, where $\lambda_{\mathrm{m}}=2^{\kappa_{\mathrm{m}}}$, and on the right half, if we get any other result.

The incorrectness of $\varepsilon$-LC may misguide our search, but since we have $\operatorname{CON}(M, \varepsilon-\mathrm{LC}(M)) \leq(1+\varepsilon) \operatorname{MINCON}(M)$ it can give a wrong answer (i.e., not the least contributor) only if we have $(1+\varepsilon)^{-1} \operatorname{MINCON}(M) \leq 2^{\kappa} \leq$ $(1+\varepsilon) \operatorname{MINCON}(M)$. Outside of this interval our search goes perfectly well. Thus, after the binary search, i.e, after at most $\left\lceil\log _{2}(b)\right\rceil$ many calls to $\varepsilon$-LC, we end up at a value $\kappa$ which is either inside the above interval (in which case we are satisfied) or the largest integer smaller than $\log _{2}\left((1+\varepsilon)^{-1} \operatorname{MINCON}(M)\right)$ or the smallest integer greater than $\log _{2}((1+\varepsilon) \operatorname{MINCON}(M))$. Hence, we have $\kappa \leq \log _{2}((1+\varepsilon) \operatorname{MINCON}(M))+1$ implying $\lambda=2^{\kappa} \leq 2(1+\varepsilon) \operatorname{MINCON}(M)$. Analogously, we get $\lambda=2^{\kappa} \geq \operatorname{MINCON}(M) /(2(1+\varepsilon))$. Therefore after $\mathcal{O}(\log (b))$ many calls to $\varepsilon-\mathrm{LC}$ we get a $2(1+\varepsilon)$ approximation of $\operatorname{MINCON}(M)$. Since this is NP-hard for $2(1+\varepsilon)$ bounded from above by $2^{d^{1-c}}$ for some $c>0$, we showed NP-hardness of $\varepsilon$-LC in this case. Note that this includes any constant $\varepsilon$.

The NP-hardness of $\varepsilon$-LC not only implies NP-hardness of LC, but also the non-existence of an efficient algorithm for $\varepsilon-\delta$-LC unless $\mathbf{N P}=\mathbf{B P P}$. The above proof also gives a very good intuition about the problem $\varepsilon$-LC: As we can approximate the minimal contribution by a small number of calls to $\varepsilon$-LC, there cannot be a much faster way to solve $\varepsilon$-LC but to approximate the contributions - approximating at least the least contribution can be only a factor of $\mathcal{O}(\log (b))$ slower than solving $\varepsilon$-LC. This motivates the algorithm we present in the next section, which tries to approximate the contributions of the various boxes.

## 3 Practical Approximation Algorithm

The last section ruled out the possibility of a worst case efficient algorithm for computing or approximating the least contributor. Nevertheless, we are now presenting an algorithm $\mathcal{A}$ that is "safe" and has a good practical runtime, but no polynomial worst case runtime (as this is not possible). By "safe" we mean that it provably solves $\varepsilon-\delta$-LC, i.e., it holds that

$$
\operatorname{Pr}[\operatorname{CON}(M, \mathcal{A}(M, \varepsilon, \delta)) \leq(1+\varepsilon) \operatorname{MINCON}(M)] \geq 1-\delta
$$

As the algorithm is going to approximate the contributions, we cannot avoid $\varepsilon$ and solve LC directly, as with no $(1+\Delta)$-approximation, for any $\Delta>0$ we can decide whether two contributions are equal or just nearly equal (and in the latter case which one is greater). We consider an $\varepsilon$ around $10^{-2}$ or $10^{-3}$ as sufficient for typical instances. This implies for most instances that we return the correct result as there are no two small contributions which are only a $(1+\varepsilon)$-factor apart. For the remaining cases we return at least a box which has contribution at most $(1+\varepsilon)$ times the minimal contribution, which means we make an "error" of $\varepsilon$.

Additionally, the algorithm is going to be a randomized Monte Carlo algorithm, which is why we need the $\delta$ and do not always return the correct result. However, we will be able to set $\delta=10^{-6}$ or even $\delta=10^{-12}$ without increasing the runtime overly. In the following we will describe algorithm $\mathcal{A}$, prove its correctness and describe its runtime.

### 3.1 The Algorithm $\mathcal{A}$

Our algorithm works as follows. For each box $A$ it determines the minimal bounding box of the space that is uniquely overlapped by the box. To do so we start with the box $A$ itself. Then we iterate over all other boxes $B$. If $B$ dominates $A$ in all but one dimension, then we can cut the bounding box in the non-dominated dimension. This can be realized in $\mathcal{O}\left(d n^{2}\right)$.

Having the bounding box $\mathrm{BB}_{A}$ of the contribution of $A$ we start to sample randomly in it. For each random point we determine if it is uniquely dominated by $A$. If we checked $\operatorname{noSamples}(A)$ random points and $\operatorname{noSuccSamples}(A)$ of them were uniquely dominated by $A$, then the contribution of $A$ is about $\widetilde{V}_{A}:=\frac{\operatorname{NoSuccSamples}(A)}{\operatorname{NoSamples}(A)} \operatorname{vol}\left(\mathrm{BB}_{A}\right)$, where $\operatorname{vol}\left(\mathrm{BB}_{A}\right)$ denotes the volume of the bounding box of the contribution of $A$. Additionally, we can give an estimate of the deviation of $\widetilde{V}_{A}$ from $V_{A}$, the correct contribution of $A$ (i.e., $\left.V_{A}=\operatorname{CON}(M, A)\right)$ : Using Chernoff's inequality we get that for

$$
\begin{equation*}
\Delta(A):=\sqrt{\frac{\log (2 n / \delta)}{2 \operatorname{NOSAMPLES}(A)}} \operatorname{vOL}\left(\mathrm{BB}_{A}\right) \tag{1}
\end{equation*}
$$

the probability that $V_{A}$ deviates from $\widetilde{V}_{A}$ by more than $\Delta(A)$ is small enough.
We would like to sample in the bounding boxes in parallel such that every $\widetilde{V}_{A}$ deviates about the same $\Delta$. We do this by initializing $\Delta$ arbitrarily (e.g., $\Delta=1$ ) and then in every iteration decrease $\Delta$ by some factor (e.g., $\frac{1}{2}$ ) and sample in each bounding box until we have $\Delta(A) \leq \Delta$. If we then have at any point two boxes $A$ and $B$ with

$$
\begin{equation*}
\widetilde{V}_{A}-\Delta(A)>\widetilde{V}_{B}+\Delta(B) \tag{2}
\end{equation*}
$$

we can with good probability assume that $A$ is not a least contributor as we would need to have $\widetilde{V}_{A}-V_{A}>\Delta(A)$ or $V_{B}-\widetilde{V}_{B}>\Delta(B)$ for $A$ having a less contribution than $B$ (which is necessary for $A$ being the least contributor). Hence, whenever such a situation occurs we can delete $A$ from our race, meaning
that we do not have to sample in its bounding box anymore. Note that we never have to compare two arbitrary boxes, but only a box $A$ to the currently smallest box $\widetilde{L C}$, i.e., the box with $\widetilde{V_{\overparen{L C}}}$ minimal.

We can run this race, deleting boxes if their contribution is clearly too much by the above selection equation until either there is just one box left, in which case we have found the least contributor, or until we have reached a point where we have approximated all contributions well enough. Given an abortion criterion $\varepsilon$ we can just return $\widetilde{L C}$ (the box with currently smallest approximated contribution) when we have

$$
0<\frac{\widetilde{V}_{\widetilde{L C}}+\Delta(\widetilde{L C})}{\widetilde{V}_{A}-\Delta(A)} \leq 1+\varepsilon
$$

for any box $A \neq \widetilde{L C}$ still in the race. If this equation holds, then we can be quite sure that any box has contribution at least $\frac{1}{1+\varepsilon} V_{\widetilde{L C}}$, and, similarly, all other boxes that are still in the race, too. So, after all, we have solved $\varepsilon-\delta-\mathrm{LC}$.

Due to space limitations, the proof that the described algorithm is indeed correct has been removed. It can be found in Section 3.3 of [5].

### 3.2 Runtime

As discussed above, our algorithm needs a runtime of at least $\Omega\left(d n^{2}\right)$. This seems to be the true runtime on many practical instances (cf. Section 4). However, by Theorem 2 we cannot hope for a matching upper bound. In this section we present an upper bound on the runtime depending on some characteristics of the input.

For an upper bound, observe that we have to approximate each box $A$ up to $\Delta=\left(V_{A}-\operatorname{MINCON}(M)\right) / 4$ to be able to delete it with high probability: At this point, $\widetilde{V}_{A} \geq V_{A}-\Delta$ and $\widetilde{V}_{B} \leq V_{B}+\Delta$, for $B$ a least contributor, so that $\widetilde{V}_{A}-\widetilde{V}_{B} \geq 2 \Delta$ with probability at least $1-\delta / n$. Similarly, we can show that the expected value of $\Delta$ where we delete box $A$ is $\Omega\left(V_{A}-\operatorname{MINCON}(M)\right)$. By equation (1) we observe that we need a number of

$$
\frac{\log (2 n / \delta) \operatorname{vOL}\left(\mathrm{BB}_{A}\right)^{2}}{2 \Omega\left(V_{A}-\operatorname{MINCON}(M)\right)^{2}}=\mathcal{O}\left(\frac{\log (n / \delta) \operatorname{voL}\left(\mathrm{BB}_{A}\right)^{2}}{\left(V_{A}-\operatorname{MINCON}(M)\right)^{2}}\right)
$$

samples to delete box $A$ on average. For the least contributor $L C$, we need $\mathcal{O}\left(\frac{\log (n / \delta) \operatorname{voL}\left(\mathrm{BB}_{L C}\right)^{2}}{(\sec -\min (V)-\operatorname{MINCON}(M))^{2}}\right)$ many samples until we have finally deleted all other boxes, where $\sec -\min (V)$ denotes the second smallest contribution of any box in $M$. Since each sample takes runtime $\mathcal{O}(d n)$ and everything besides the sampling takes much less runtime, we get an overall runtime of $\mathcal{O}(d n(n+\log (n / \delta) \mathrm{H}))$, where

$$
\mathrm{H}=\frac{\operatorname{voL}\left(\mathrm{BB}_{L C}\right)^{2}}{(\sec -\min (V)-\operatorname{MINCON}(M))^{2}}+\sum_{L C \neq A \in S} \frac{\operatorname{voL}\left(\mathrm{BB}_{A}\right)^{2}}{\left(V_{A}-\operatorname{MINCON}(M)\right)^{2}}
$$

```
\(\overline{\text { Algorithm 1. } \mathcal{A}(M, \varepsilon, \delta) \text { solves } \varepsilon-\delta-\mathrm{LC}(M) \text { for a set } M \text { of } n \text { boxes in } \mathbb{R}^{d}}\)
and \(\varepsilon, \delta>0\), i.e., it determines a box \(x \in M\) s.t. \(\operatorname{Pr}[\operatorname{CON}(M, x) \leq(1+\)
ع) \(\operatorname{MINCON}(M)] \geq 1-\delta\).
    determine the bounding boxes \(\mathrm{BB}_{A}\) for all \(A \in M\)
    initialize \(\operatorname{NoSamples}(A)=\operatorname{NoSuccSamples}(A)=0\) for all \(A \in M\)
    initialize \(\Delta\)
    set \(S:=M\)
    repeat
        set \(\Delta:=\Delta / 2\)
        for all \(A \in S\) do
            repeat
            sample a random point in \(\mathrm{BB}_{A}\)
            increase \(\operatorname{noSamples}(A)\) and possibly noSuccSamples \((A)\)
            update \(\widetilde{V}_{A}\) and \(\Delta(A)\)
            until \(\Delta(A) \leq \Delta\)
        od
        set \(\widetilde{L C}:=\operatorname{argmin}\left\{\widetilde{V}_{A} \mid A \in S\right\}\)
        for all \(A \in S\) do
            if \(\widetilde{V}_{A}-\Delta(A)>\widetilde{V}_{\widetilde{L C}}+\Delta(\widetilde{L C})\) then
            \(S:=S \backslash\{A\}\)
            od
        od
    until \(|S|=1\) or \(0<\frac{\tilde{V}_{\widetilde{L C}}+\Delta(\widetilde{L C})}{\widetilde{V_{A}}-\Delta(A)} \leq 1+\varepsilon\) for any \(\widetilde{L C} \neq A \in S\)
    return \(\widetilde{L C}\)
```

is a certain measure of hardness of the input. This value is unbounded and can even be undefined if there is no unique least contributor. In this case our abortion criterion comes into play: With probability $(1-\delta)$ after approximating every contribution up to $\Delta=\frac{\varepsilon}{4+2 \varepsilon} \operatorname{MINCON}(M)$ we have $\widetilde{V}_{L C} \leq V_{L C}+\Delta$, thus $\widetilde{V}_{\widetilde{L C}} \leq V_{L C}+\Delta$, and $\widetilde{V}_{A} \geq V_{L C}-\Delta$ for every other box $A$ still in the race. Then we conclude

$$
\frac{\widetilde{V}_{\widetilde{L C}}+\Delta(\widetilde{L C})}{\widetilde{V}_{A}-\Delta(A)} \leq \frac{V_{L C}+2 \Delta}{V_{L C}-2 \Delta}=\frac{1+2 \frac{\varepsilon}{4+2 \varepsilon}}{1-2 \frac{\varepsilon}{4+2 \varepsilon}}=1+\varepsilon
$$

for every box $\widetilde{L C} \neq A \in S$. Hence, the above defined value for $\Delta$ suffices to enforce abortion. Since we get this $\Delta$ after $\operatorname{NoSAMPLES}(A)=\frac{\log (2 n / \delta) \operatorname{voL}\left(\mathrm{BB}_{A}\right)^{2}}{2\left(\frac{\varepsilon}{4+2 \varepsilon} \operatorname{MINCON}(M)\right)^{2}}$ samples, this yields another upper bound for the overall number of samples, a still unbounded but always finite value:

$$
\mathcal{O}\left(\frac{\log (n / \delta)}{\varepsilon^{2} \operatorname{MINCON}(M)^{2}} \sum_{A \in M} \operatorname{voL}\left(\mathrm{BB}_{A}\right)^{2}\right)
$$

However, for the random testcases that we consider in Section 4 the above defined hardness H is a more realistic measure of runtime as there are never two
identical contributions and not too many equally small contributions. There one observes values for H that roughly lie in the interval $[n, 10 n$ ].

## 4 Experimental Analysis

To demonstrate the performance of the described approximation algorithm for the hypervolume contribution, we have implemented it and measured its performance on different datasets. To yield a practically relevant algorithm, we have implemented several heuristical improvements which are described in detail in Section 3.4 of [5]. The most important for the correct interpretation of the experiments is that we use a classical exact algorithm for small $n$ and $d$. We now first describe the used benchmark datasets and then our results.

### 4.1 Datasets

We used five different fronts similar to the DTLZ test suite 7]. As we do not want to compare the hypervolume algorithms for point distributions specific to different optimizers like NSGA-II [6] or SPEA2 [21], we have sampled the points from different surfaces randomly. This allows full scalability of the datasets in the number of points and the number of dimensions.

To define the datasets, we use random variables with two different distributions. Simple uniformly distributed random variables are provided by the buildin random number generator rand() of $\mathrm{C}++$. To get random variables with a Gaussian distribution, we used the polar form of the Box-Muller transformation as described in 13].

Linear Dataset: The first dataset consists of points $\left(x_{1}, x_{2}, \ldots, x_{d}\right) \in[0,1]^{d}$ with $\sum_{i=1}^{d} x_{i}=1$. They are obtained by generating $d$ Gaussian random variables $y_{1}, y_{2}, \ldots, y_{d}$ and then using the normalized points

$$
\left(x_{1}, x_{2}, \ldots, x_{n}\right):=\frac{\left(\left|y_{1}\right|,\left|y_{2}\right|, \ldots,\left|y_{n}\right|\right)}{\left|y_{1}\right|+\left|y_{2}\right|+\ldots+\left|y_{d}\right|}
$$

Spherical Dataset: To obtain uniformly distributed points $\left(x_{1}, x_{2}, \ldots, x_{d}\right) \in$ $[0,1]^{d}$ with $\sum_{i=1}^{d} x_{i}^{2}=1$ we follow the method of Muller [10]. That is, we generate $d$ Gaussian random variables $y_{1}, y_{2}, \ldots, y_{d}$ and take the points

$$
\left(x_{1}, x_{2}, \ldots, x_{n}\right):=\frac{\left(\left|y_{1}\right|,\left|y_{2}\right|, \ldots,\left|y_{n}\right|\right)}{\sqrt{y_{1}^{2}+y_{2}^{2}+\ldots+y_{d}^{2}}}
$$

Concave Dataset: Analogously to the spherical dataset we choose points $\left(x_{1}, x_{2}, \ldots, x_{d}\right) \in[0,1]^{d}$ with $\sum_{i=1}^{d} \sqrt{x_{i}}=1$. For this, we generate again $d$ Gaussian random variables $y_{1}, y_{2}, \ldots, y_{d}$ and use the points

$$
\left(x_{1}, x_{2}, \ldots, x_{n}\right):=\frac{\left(\left|y_{1}\right|,\left|y_{2}\right|, \ldots,\left|y_{n}\right|\right)}{\left(\sqrt{\left|y_{1}\right|}+\sqrt{\left|y_{2}\right|}+\ldots+\sqrt{\left|y_{d}\right|}\right)^{2}} .
$$



Fig. 1. Visualization of the first three datasets


Fig. 2. Experimental results for $d=3$


Fig. 3. Experimental results for $d=10$


Fig. 4. Experimental results for $d=100$


Fig. 5. Experimental results for random datasets with $d=5$


Fig. 6. Experimental results for random datasets with $d=100$

For $d=3$, the surface of the dataset is shown in Figure 1. Additionally to random points lying on a lower-dimensional surface, we have also examined the following two datasets with points sampled from the actual space similar to the random dataset examined by While et al. 17].

Random Dataset 1: We first draw $n$ uniformly distributed points from $[0,1]^{d}$ and then replace all dominated points by new random points until we have a set of $n$ nondominated points.

Random Dataset 2: Very similar to the previous dataset, we choose random points until there are no dominated points. The only difference is that this time the points are not drawn uniformly, but Gaussian distributed in $\mathbb{R}^{d}$ with mean 1.

Note that the last two datasets are far from being uniformly distributed. The points of the first set all have at least one coordinate very close to 1 while the points of the second set all have at least one coordinate which is significantly above the mean value. This makes their computation for many points (e.g., $n \geq 100$ ) in small dimensions (e.g., $d \leq 5$ ) computationally very expensive as it becomes more and more unlikely to sample a nondominated point.

### 4.2 Comparison

We have implemented our algorithm in $\mathrm{C}++$ and compared it with the available implementations of HSO by Eckart Zitzler [19] and BR by Nicola Beume [2]. We did not add any further heuristics to both exact algorithms as all published heuristics do not improve the asymptotic runtime and even a speedup of a few magnitudes does not change the picture significantly.

It would be better to compare our approximation algorithm with other approximation algorithms instead of exact algorithms. However, the only other published approximation algorithm seems to be 1], which is not publicly available yet. Another reason is that all available optimization algorithms based on the hypervolume indicator use exact calculations and hence our speedup is carried over to them.

All experiments were run on a cluster of 100 machines with two 2.4 GHz AMD Opteron processors, operating in 32-bit mode, running Linux. For our approximation algorithm we used the parameters $\delta=10^{-6}$ and $\varepsilon=10^{-2}$. The code used is available upon request and will be distributed from the homepage of the second author.

Figure $2 \sqrt{6}$ show double-logarithmic plots of the runtime for different datasets and number of dimensions. The shown values are the median of 100 runs each. To illustrate the occurring deviations below and above the median, we also plotted all measured runtimes as ligther single points in the back. As both axes are scaled logarithmically, also the examined problem sizes are distributed logarithmically. That is, we only calculated Pareto sets of size $n$ if $n \in\{\lfloor\exp (k / 100)\rfloor \mid k \in \mathbb{N}\}$. We examined dimensions $d=3,10,100$ for the first three datasets and $d=5,100$ for the last two datasets.

Independent of the number of solutions and dimension, we always observed that, unless $n \leq 10$, our algorithm outperformed HSO and BR substantially. On the used machines this means that only if the calculation time was insignificant (say, below $10^{-4}$ seconds), the exact algorithm could compete. On the other hand, the much lower median of our algorithm also comes with a much higher empirical standard deviation and interquartile range. In fact, we observed that the upper quartile can be up to five times slower than the median (for the especially degenerated random dataset 1 ). The highest ratio observed between the maximum runtime and the average runtime is 66 (again for the random dataset 1). This behavior is represented in the plots by the spread of lighter datapoints in the back of the median. However, there are not too many outliers and even their runtime outperforms HSO and BR. The non-monotonicity of our algorithm around $n=10$ for $d=10$ is caused by the approximate for the runtimes of the exact algorithms.

For larger dimensions the advantage of our approximation algorithm becomes tremendous. For $d=100$ we observed that within 100 seconds our algorithm could solve all problems with less than 6000 solutions while HSO an BR could not solve any problem for a population of 6 solutions in the same time. For example for 7 solutions on the 100-dimensional linear front, HSO needed 13 minutes, BR 7 hours while our algorithm terminated within 0.5 milliseconds.

## 5 Conclusions

We have proven that most natural questions about the hypervolume contribution which are relevant for evolutionary multi-objective optimizers are not only computationally hard to decide, but also hard to approximate. On the other hand, we have presented a new approximation algorithm which works extremely fast for all tested practical instances. It can solve efficiently large high-dimensional instances ( $d \geq 10, n \geq 100$ ) which are intractable for all previous exact algorithms and heuristics.

It would be very interesting to compare the algorithms on further datasets. We believe that only when two solutions have contributions of very close value, our algorithm slows down. For practical instances this should not matter as it simply occurs too rarely - but this conjecture should be substantiated by some broader experimental study in the future.

## References

[1] Bader, J., Zitzler, E.: HypE: An Algorithm for Fast Hypervolume-Based ManyObjective Optimization. TIK Report 286, Institut für Technische Informatik und Kommunikationsnetze, ETH Zürich (2008)
[2] Beume, N., Rudolph, G.: Faster S-metric calculation by considering dominated hypervolume as Klee's measure problem. In: Proc. Second International Conference on Computational Intelligence (IASTED 2006), pp. 233-238 (2006)
[3] Beume, N., Fonseca, C., López-Ibáñez, M., Paquete, L., Vahrenhold, J.: On the complexity of computing the hypervolume indicator. Technical report CI-235/07, Technical University of Dortmund (2007)
[4] Bringmann, K., Friedrich, T.: Approximating the volume of unions and intersections of high-dimensional geometric objects. In: Hong, S.-H., Nagamochi, H., Fukunaga, T. (eds.) ISAAC 2008. LNCS, vol. 5369, pp. 436-447. Springer, Heidelberg (2008a), http://arxiv.org/abs/0809.0835
[5] Bringmann, K., Friedrich, T.: Approximating the least hypervolume contributor: NP-hard in general, but fast in practice (2008a), http://arxiv.org/abs/0812.2636
[6] Deb, K., Agrawal, S., Pratap, A., Meyarivan, T.: A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: NSGA-II. In: Deb, K., Rudolph, G., Lutton, E., Merelo, J.J., Schoenauer, M., Schwefel, H.-P., Yao, X. (eds.) PPSN 2000. LNCS, vol. 1917, pp. 849-858. Springer, Heidelberg (2000)
[7] Deb, K., Thiele, L., Laumanns, M., Zitzler, E.: Scalable multi-objective optimization test problems. In: Proc. IEEE Congress on Evolutionary Computation (CEC 2002), pp. 825-830 (2002)
[8] Klee, V.: Can the measure of $\bigcup\left[a_{i}, b_{i}\right]$ be computed in less than $O(n \log n)$ steps? American Mathematical Monthly 84, 284-285 (1977)
[9] Knowles, J.D.: Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization. PhD thesis, Department of Computer Science, University of Reading, UK (2002)
[10] Muller, M.E.: A note on a method for generating points uniformly on $n$ dimensional spheres. Commun. ACM 2, 19-20 (1959)
[11] Overmars, M.H., Yap, C.-K.: New upper bounds in Klee's measure problem. SIAM J. Comput. 20, 1034-1045 (1991); announced at 29th Annual Symposium on Foundations of Computer Science (FOCS 1988)
[12] Papadimitriou, C.M.: Computational Complexity. Addison-Wesley Publishing Company, Reading (1994)
[13] Press, W.H., Teukolsky, S.A., Vetterling, W.T., Flannery, B.P.: Numerical Recipes: The Art of Scientific Computing, 2nd edn. Cambridge University Press, Cambridge (1992)
[14] Roth, D.: On the hardness of approximate reasoning. Artif. Intell. 82, 273-302 (1996)
[15] Valiant, L.G.: The complexity of computing the permanent. Theor. Comput. Sci. 8, 189-201 (1979)
[16] While, R.L., Bradstreet, L., Barone, L., Hingston, P.: Heuristics for optimizing the calculation of hypervolume for multi-objective optimization problems. In: Proc. IEEE Congress on Evolutionary Computation (CEC 2005), pp. 2225-2232 (2005)
[17] While, R.L., Hingston, P., Barone, L., Huband, S.: A faster algorithm for calculating hypervolume. IEEE Trans. Evolutionary Computation 10, 29-38 (2006)
[18] Zhou, X., Mao, N., Li, W., Sun, C.: A fast algorithm for computing the contribution of a point to the hypervolume. In: Proc. Third International Conference on Natural Computation (ICNC 2007), vol. 4, pp. 415-420 (2007)
[19] Zitzler, E.: Hypervolume metric calculation, 2001. Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Switzerland, ftp://ftp.tik.ee.ethz.ch/pub/people/zitzler/hypervol.c
[20] Zitzler, E., Thiele, L.: Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. IEEE Trans. Evolutionary Computation 3, 257-271 (1999); announced at 5th International Conference Parallel Problem Solving from Nature (PPSN V)
[21] Zitzler, E., Laumanns, M., Thiele, L.: SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In: Giannakoglou, K., et al. (eds.) Proc. Evolutionary Methods for Design, Optimisation and Control with Application to Industrial Problems (EUROGEN 2001), pp. 95-100. International Center for Numerical Methods in Engineering (CIMNE) (2002)
[22] Zitzler, E., Thiele, L., Laumanns, M., Fonseca, C.M., da Fonseca, V.G.: Performance assessment of multiobjective optimizers: an analysis and review. IEEE Trans. Evolutionary Computation 7, 117-132 (2003)


[^0]:    ${ }^{1}$ Without loss of generality we assume the reference point to be $0^{d}$.
    ${ }^{2} \# \mathbf{P}$ is the analog of $\mathbf{N P}$ for counting problems. For details see either the original paper by Valiant 15] or the standard textbook on computational complexity by Papadimitriou 12].

