# Don't be Greedy when Calculating Hypervolume Contributions 

Karl Bringmann<br>Universität des Saarlandes<br>Postfach 151150<br>66041 Saarbrücken, Germany

Tobias Friedrich<br>International Computer Science Institute<br>1947 Center St., Suite 600<br>94704 Berkeley, CA, USA


#### Abstract

Most hypervolume indicator based optimization algorithms like SIBEA [Zitzler et al. 2007], SMS-EMOA [Beume et al. 2007], or MO-CMA-ES [Igel et al. 2007] remove the solution with the smallest loss with respect to the dominated hypervolume from the population. This is usually iterated $\lambda$ times until the size of the population no longer exceeds a fixed size $\mu$. We show that there are populations such that the contributing hypervolume of the $\lambda$ solutions chosen by this greedy selection scheme can be much smaller than the contributing hypervolume of an optimal set of $\lambda$ solutions. Selecting the optimal $\lambda$-set implies calculating $\binom{\mu+\lambda}{\mu}$ conventional hypervolume contributions, which is considered computationally too expensive. We present the first hypervolume algorithm which calculates directly the contribution of every set of $\lambda$ solutions. This gives an additive term of $\binom{\mu+\lambda}{\mu}$ in the runtime of the calculation instead of a multiplicative factor of $\binom{\mu+\lambda}{\mu}$. Given a population of size $n=\mu+\lambda$, our algorithm can calculate a set of $\lambda \geqslant 1$ solutions with minimal $d$-dimensional hypervolume contribution in time $\mathcal{O}\left(n^{d / 2} \log n+n^{\lambda}\right)$ for $d>2$. This improves all previously published algorithms by a factor of order $n^{\min \{\lambda, d / 2\}}$ for $d>3$. Therefore even if we remove the solutions one by one greedily as usual, we gain a speed up factor of $n$ for all $d>3$.


Categories and Subject Descriptors: F. 2 [Theory of Computation]: Analysis of Algorithms and Problem Complexity

General Terms: Theory, Algorithms, Performance

## 1. INTRODUCTION

How to compare Pareto sets lies at the heart of research in multi-objective optimization. One measure that has been the subject of much recent study in evolutionary multi-objective optimization is the "hypervolume indicator"

[^0](HYP). It measures the volume of the dominated portion of the objective space. The hypervolume metric is of exceptional interest as it possesses the highly desirable feature of strict Pareto compliance [24]. That is, considering two Pareto sets $A$ and $B$, the hypervolume indicator values $A$ higher than $B$ if the Pareto set $A$ dominates the Pareto set $B$. This property makes it well suited to many-objective problems.

The hypervolume was first introduced for performance assessment in multi-objective optimization by Zitzler and Thiele [23]. Later on it was used to guide the search in various hypervolume-based evolutionary optimizers $[4,11,12$, $13,21,22]$. Since then, several algorithms for calculating the hypervolume have been developed. The first one was the Hypervolume by Slicing Objectives (HSO) algorithm which was suggested independently by Zitzler [20] and Knowles [14]. To improve its runtime on practical instances, various speed up heuristics of HSO have been suggested [17, 19]. The currently best asymptotic runtime of $\mathcal{O}\left(n \log n+n^{d / 2}\right)$ by Beume and Rudolph [5] is based on Overmars and Yap's algorithm for Klee's Measure Problem [16]. There are also various algorithms for small dimensions [11, 15].

So far, 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 the boxes) of $\Theta\left(n^{d}\right)$ does not imply any lower bounds on the computational complexity. It has recently been shown by the authors [7] that the calculation of HYP is \#P-hard, which implies that all hypervolume algorithms must have a superpolynomial runtime in the number of objectives or boxes unless $\mathbf{P}=\mathbf{N P}$. [7] also presents an FPRAS (fully polynomial-time randomized approximation scheme) which gives an $\varepsilon$-approximation of the hypervolume indicator with probability $1-\delta$ in time $\mathcal{O}\left(\log (1 / \delta) n d / \varepsilon^{2}\right)$. Though this algorithm gives a very fast approximation in time (linear in $n$ and $d$ ) for the hypervolume, it is important to note that this is not an approximation of the contributing hypervolume. Even the approximation of the latter is NP-hard as has recently been shown in [8].

Let us now look at a hypervolume-based algorithm that maintains a population (set of solutions) $M \subseteq \mathbb{R}_{\geqslant 0}^{d}$ of size $\mu$. Then the hypervolume ${ }^{1}$ is defined as

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

[^1]with $\operatorname{VOL}(\cdot)$ being the usual Lebesgue measure. To avoid an unbounded population, the number of solutions in the population is usually fixed to a certain threshold and with every new Pareto optimal solution another one has to be removed. Most hypervolume based algorithms remove the solution $x \in M$ with the smallest contribution ${ }^{2}$ [4, 9, 21]
$$
\operatorname{CON}_{M}(x):=\operatorname{HYP}(M)-\operatorname{HYP}(M \backslash\{x\}) .
$$

The contribution can be seen as the measure of the space that is dominated by $x$, but no other point in $M$. A point $\left(x_{1}, \ldots, x_{d}\right) \in \mathbb{R}^{d}$ dominates a point $\left(y_{1}, \ldots, y_{d}\right) \in \mathbb{R}^{d}$, iff we have $x_{i} \leq y_{i}$ for $i=1, \ldots, d$. The solution with minimal contribution can be calculated easily by $\mu+1$ hypervolume calculations.

The problem is that often $\lambda$ solutions should be removed at once. In this case one aims for a set $S \subseteq M$ of size $\lambda$, i.e., $S$ is a $\lambda$-set or $\lambda$-subset of $M$, such that

$$
\operatorname{CON}_{M}(S):=\operatorname{HYP}(M)-\operatorname{HYP}(M \backslash S)
$$

is minimized. To calculate the optimal $\lambda$-set $S_{\text {opt }}^{\lambda}$ which has the smallest joint contribution requires $\binom{\mu+\lambda}{\mu}$ calculations of $\operatorname{HYP}(\cdot)$ which is considered computationally too expensive $[1,4]$. This is why all current hypervolume based optimization algorithms just calculate a greedy $\lambda$-set $S_{\text {greedy }}^{\lambda}$ by starting with $S_{\text {greedy }}^{0}:=\emptyset$ and then iteratively setting

$$
S_{\text {greedy }}^{\lambda+1}:=S_{\text {greedy }}^{\lambda} \cup\left\{\underset{x \in M \backslash S_{\text {greedy }}^{\lambda}}{\operatorname{argmin}} \operatorname{CON}_{M \backslash S_{\text {greedy }}^{\lambda}}(x)\right\} .
$$

In Section 2 we will show that $\operatorname{CON}\left(S_{\text {greedy }}^{\lambda}\right) / \operatorname{CON}\left(S_{\text {opt }}^{\lambda}\right)$ can be arbitrarily large. That is, the error introduced by choosing the $\lambda$-set greedily instead of optimally can misguide the search process. We also report on Pareto fronts with significant differences between both $\lambda$-sets in the DTLZ library [10].

This observation motivates the algorithm introduced in Section 3. It is an adaption of Overmars and Yap's algorithm which allows the direct computation of the contribution of every set of $\lambda$ solutions. This avoids calculating $\binom{\mu+\lambda}{\mu}$ conventional hypervolume calculations by maintaining the volume of the contribution of every set of $\lambda$ solutions during the calculation. This yields an additive term of $\binom{\mu+\lambda}{\mu}$ in the runtime of the calculation instead of a multiplicative factor of $\binom{\mu+\lambda}{\mu}$. For a population of size $n=\mu+\lambda$ and $d>2$, the algorithm calculates $S_{\text {opt }}^{\lambda}$ in time $\mathcal{O}\left(n^{d / 2} \log n+n^{\lambda}\right)$, which improves all previously published algorithms by a factor of order $n^{\min \{\lambda, d / 2\}}$.

This has two consequences. First, even if we remove only $\lambda=1$ solutions in every step, this is a speed up by a factor of order $n$. This means we can determine the box with least contribution, i.e., the greedy solution, in time of order $n^{d / 2}$ instead of the usual $\mathcal{O}\left(n^{d / 2+1}\right)$. Second, for $\lambda \leqslant d / 2$ the asymptotic runtime is independent of $\lambda$. Therefore, using $\lambda=d / 2$ instead of the commonly used $\lambda=1$ (greedy) gives the same asymptotic runtime and yields the same or potentially even smaller contributions of the calculated $\lambda$-sets.

Note that we will assume throughout the paper, that $d$ and $\lambda$ are constant, i.e., $n$ is the only growing parameter for

[^2]the asymptotic analysis. This is no real restriction as both of them appear in the exponent of the resulting runtime and hence superconstant values for them would immediately make the algorithm inefficient.

## 2. GREEDY IS BAD

For $\lambda=1$, the greedy $\lambda$-set is optimal, i.e., $S_{\text {greedy }}^{\lambda}=S_{\text {opt }}^{\lambda}$. However, it is known that for $\lambda \geqslant 2$ the greedy algorithm is not able to construct the optimal solution set in general. For example Bradstreet et al. [6] present a three-dimensional example of $n=6$ points where for $\lambda=2$ the contribution of the greedy $\lambda$-set $S_{\text {greedy }}^{\lambda}$ is $12.5 \%$ larger than the optimal $\lambda$-set $S_{\text {opt }}^{\lambda}$. In this section we show that the $\lambda$-set $S_{\text {greedy }}^{\lambda}$ found by the greedy algorithm can have an arbitrarily larger contribution than the optimal $\lambda$-set $S_{\text {opt }}^{\lambda}$ for all $\lambda \geqslant 2$. Let $\kappa$ denote the ratio between the contribution of $S_{\text {greedy }}^{\lambda}$ and $S_{\mathrm{opt}}^{\lambda}$. For many sets $M, \kappa$ is either one or very close to one. Next, we prove that for given $\kappa$, dimension $d \geqslant 3$ and number of boxes $n>d$ there is a set of solutions $M$ of size $n$ such that the ratio between $\operatorname{CON}_{M}\left(S_{\text {greedy }}^{\lambda}\right)$ and $\mathrm{CON}_{M}\left(S_{\text {opt }}^{\lambda}\right)$ is larger than $\kappa$ for all $2 \leq \lambda \leq d$. Additionally, we show that $\kappa>1$ also holds for some fronts from the DTLZ library [10].

Lemma 1. For all $\kappa \geqslant 1, d \geqslant 3$ and $n>d$ there is a set $M \subseteq \mathbb{R}_{\geqslant 0}^{d}$ with $|M|=n$ such that $\operatorname{CON}_{M}\left(S_{\text {greedy }}^{\lambda}\right) / \operatorname{CON}_{M}\left(S_{\text {opt }}^{\lambda}\right)>\kappa$ for all $2 \leqslant \lambda \leqslant d$.

Proof. We first assume $n=d+1$. Let $\varepsilon:=1 /\left(2 \kappa d^{2}\right)$ and $\sigma:=d^{2} \varepsilon^{2}$. We choose $M=\left\{q, p_{1}, p_{2}, \ldots, p_{d}\right\}$ with

$$
\begin{aligned}
q & =(1+\varepsilon, 1+\varepsilon, \ldots, 1+\varepsilon), \\
p_{1} & =(1+\varepsilon+\sigma, 1,1, \ldots, 1), \\
p_{2} & =(1,1+\varepsilon+\sigma, 1, \ldots, 1), \\
& \ldots \\
p_{d} & =(1,1, \ldots, 1,1+\varepsilon+\sigma) .
\end{aligned}
$$

A three-dimensional example is shown in Figure 1. As the space dominated by $p_{i}$ but no other point in $M$ is exactly

$$
[0,1]^{i-1} \times[1+\varepsilon, 1+\varepsilon+\sigma] \times[0,1]^{d-i}
$$

we have $\operatorname{CON}_{M}\left(p_{i}\right)=\sigma$. The point $q$ dominates $[0,1+\varepsilon]^{d}$, but $[0,1]^{d}$ and the $d$ rectangular regions of the form $[0,1] \times$ $\cdots \times[0,1] \times[1,1+\varepsilon] \times[0,1] \times \cdots \times[0,1]$ are dominated by the other points in $M$, too, so that we have

$$
\begin{aligned}
\operatorname{CON}_{M}(q) & =(1+\varepsilon)^{d}-1^{d}-d \cdot \varepsilon \\
& =\sum_{i=2}^{d}\binom{d}{i} \varepsilon^{i} .
\end{aligned}
$$

Similarly, we get

$$
\begin{gathered}
\operatorname{CON}_{M}(\{\underbrace{p_{i_{1}}, p_{i_{2}}, \ldots, p_{i_{\lambda}}}_{\lambda \text { different } p_{i} \text { 's }}\})=\lambda \sigma, \\
\operatorname{CON}_{M}(\{q, \underbrace{p_{i_{1}}, \ldots, p_{i_{(\lambda-1)}}}_{\lambda-1 \text { different } p_{i} \text { 's }}\})=\operatorname{CON}_{M}(q)+(\lambda-1)(\varepsilon+\sigma),
\end{gathered}
$$

| Test case | $\boldsymbol{d}$ | $\boldsymbol{n}$ | $\boldsymbol{\lambda}$ | $\frac{\operatorname{CON}\left(\boldsymbol{S}_{\text {greedy }}^{\boldsymbol{\lambda}}\right)-\mathbf{C O N}\left(\boldsymbol{S}_{\mathrm{opt}}^{\boldsymbol{\lambda}}\right)}{\mathbf{C O N}\left(\boldsymbol{S}_{\mathrm{opt}}^{\boldsymbol{\lambda}}\right)}$ |
| :--- | :---: | :---: | :---: | :---: |
| DTLZLinearShape.3d.front.50pts | 3 | 50 | 5 | $4.23 \%$ |
| DTLZLinearShape.3d.front.10pts | 3 | 10 | 9 | $5.76 \%$ |
| DTLZSphereShape.3d.front.50pts | 3 | 50 | 6 | $4.19 \%$ |
| DTLZSphereShape.3d.front.50pts | 3 | 50 | 7 | $8.57 \%$ |
| DTLZDiscontinuousShape.5d.front.20pts | 5 | 20 | 8 | $2.60 \%$ |
| DTLZDegenerateShape.8d.front.10pts | 8 | 10 | 3 | $11.31 \%$ |
| DTLZDegenerateShape.6d.front.10pts | 6 | 10 | 3 | $32.64 \%$ |

Table 1: Some examples of populations on fronts from the DTLZ library [10] where the contributions of the greedy $\lambda$-set $S_{\text {greedy }}^{\lambda}$ and the optimal $\lambda$-set $S_{\mathrm{opt}}^{\lambda}$ deviate significantly.
where we used that after picking $q$ every $p_{i}$ dominates a portion of space with volume $\varepsilon+\sigma$. Since

$$
\begin{aligned}
\operatorname{CON}_{M}(q) & <\sum_{i=0}^{\infty}\binom{d}{i+2} \varepsilon^{i+2}<\sum_{i=0}^{\infty} \frac{d^{i+2} \varepsilon^{i+2}}{(i+2)!} \\
& =\varepsilon^{2} d^{2} \sum_{i=0}^{\infty} \frac{d^{i} \varepsilon^{i}}{(i+2)!}<\varepsilon^{2} d^{2} \sum_{i=0}^{\infty} \frac{1}{(i+2)!} \\
& =\varepsilon^{2} d^{2}(e-2)<\varepsilon^{2} d^{2}=\operatorname{CON}_{M}\left(p_{i}\right),
\end{aligned}
$$

the greedy algorithm chooses the $\lambda$-set

$$
S_{\text {greedy }}^{\lambda}=\{q, \underbrace{p_{i_{1}}, p_{i_{2}}, \ldots, p_{i_{(\lambda-1)}}}_{(\lambda-1) \text { different } p_{i}{ }^{\prime} \text { 's }}\}
$$

though the optimal $\lambda$-set is

$$
S_{\mathrm{opt}}^{\lambda}=\{\underbrace{p_{i_{1}}, p_{i_{2}}, \ldots, p_{i_{\lambda}}}_{\lambda \text { different } p_{i}{ }^{\prime} \mathrm{s}}\} .
$$

Therefore for all $\lambda \leqslant d$,

$$
\begin{aligned}
\frac{\operatorname{CON}_{M}\left(S_{\text {greedy }}^{\lambda}\right)}{\operatorname{CON}_{M}\left(S_{\text {opt }}^{\lambda}\right)} & =\frac{\operatorname{CON}_{M}(q)+(\lambda-1)(\varepsilon+\sigma)}{\lambda \sigma} \\
& >\frac{(\lambda-1) \varepsilon}{\lambda \sigma}>\frac{1}{2 d^{2} \varepsilon}=\kappa
\end{aligned}
$$

This shows the claim for $n=d+1$. To prove the remaining case $n>d+1$, we take the set $M$ from above, shift all points in $M$ by 1 along the first dimension and add some extra boxes, each one contributing too much to be chosen by the greedy or the optimal algorithm. For this we define

$$
\begin{aligned}
M^{\prime} & =\left\{\left(x_{1}+1, x_{2}, \ldots, x_{d}\right) \mid\left(x_{1}, \ldots, x_{d}\right) \in M\right\} \\
B & =\bigcup_{1 \leq i<n-d}\{(1, C \cdot i, C \cdot(n-d-i), C, \ldots, C)\}, \\
N & =M^{\prime} \cup B
\end{aligned}
$$

where $C>2(1+\varepsilon+\sigma)^{2}$ is a sufficiently large number. As $|M|=d+1, N$ contains exactly $n$ boxes. Furthermore, each point in $B$ uniquely dominates the rectangular region
$[0,1] \times[C(i-1), C i] \times[C(n-d-i-1), C(n-d-i)] \times[0, C]^{d-3}$,
when considering only points from $B$. Hence, we have $\operatorname{CON}_{B}(x) \geq C^{d-1}$ for all $x \in B$. The contribution of $x$ in $N$ can be smaller than the contribution of $x$ in $B$. However, the contribution cannot decrease by more than the overall hypervolume of $M^{\prime}$ as every point lying in the one contributing space but not in the other one has to be dominated by a point in $M^{\prime}$. Therefore $\operatorname{CON}_{N}(x) \geq \operatorname{CON}_{B}(x)-\operatorname{HYP}\left(M^{\prime}\right)$. We further know that the hypervolume of $M^{\prime}$ is bounded
from above by $2(1+\varepsilon+\sigma)^{d}$ as each $i$-th coordinate of a point in $M^{\prime}$ is less than $1+\varepsilon+\sigma$ for $2 \leq i \leq d$ and less than $2+\varepsilon+\sigma<2+2 \varepsilon+2 \sigma$ for $i=1$. Hence,

$$
\begin{aligned}
\operatorname{CON}_{N}(x) & \geq C^{d-1}-2(1+\varepsilon+\sigma)^{d} \\
& >2^{d-1}(1+\varepsilon+\sigma)^{2 d-2}-2(1+\varepsilon+\sigma)^{d} \\
& \geq 4(1+\varepsilon+\sigma)^{d}-2(1+\varepsilon+\sigma)^{d} \\
& \geq 2(1+\varepsilon+\sigma)^{d} \\
& \geq \operatorname{HYP}\left(M^{\prime}\right) .
\end{aligned}
$$

This implies that it is better to remove from $N$ all elements in $M^{\prime}$ than to remove one element in $B$. Therefore none of $S_{\mathrm{opt}}^{\lambda}$ or $S_{\text {greedy }}^{\lambda}$ can contain a point in $B$ for $2 \leq \lambda \leq d$. Moreover, the contribution of an element $x \in M$ to $\bar{M}$ is the same as the contribution of the corresponding element $x^{\prime} \in M^{\prime}$ to $N$, as the boxes in $B$ cut away all additional dominated space (every box in $B$ dominates $[0,1] \times[0,1+$ $\left.\varepsilon+\sigma]^{d-1}\right)$. Hence

$$
\begin{aligned}
\operatorname{CON}_{N}\left(S_{\text {opt }}^{\lambda}(N)\right) & =\operatorname{CON}_{M}\left(S_{\text {opt }}^{\lambda}(M)\right) \\
\operatorname{CON}_{N}\left(S_{\text {greedy }}^{\lambda}(N)\right) & =\operatorname{CON}_{M}\left(S_{\text {greedy }}^{\lambda}(M)\right)
\end{aligned}
$$

for $2 \leq \lambda \leq d$, which implies that their ratio is at least $\kappa$ as shown for the case $n=d+1$.

To validate that significant differences indeed occur in real data sets, we have calculated the greedy and the optimal $\lambda$-set contribution for some populations on fronts from the DTLZ library [10]. To allow an easy verification of our results we used the populations generated by [18] available from http://www.wfg.csse.uwa.edu.au/hypervolume/. For different $\lambda$ 's we calculated for all populations the greedy $\lambda$-set $S_{\text {greedy }}^{\lambda}$ and the optimal $\lambda$-set $S_{\text {opt }}^{\lambda}$. We observed relative differences of up to one third between calculating the contribution greedily or optimally. Some representative numbers of larger deviations between both contributions are shown in Table 1. Note that even a small difference between the contribution of the greedy and the optimal $\lambda$-set can misguide the search process.

## 3. ALGORITHM

Consider a set $S$ of boxes in $\mathbb{R}^{d}, n:=|S|$. Throughout this chapter, we use the term $\lambda$-set for a subset $T \subseteq S$ with $|T|=\lambda$. We also say that $T$ is a $\lambda$-subset of $S$. Similarly, we use $\lambda^{\leq}$-set for denoting any set $T \subseteq S$ with $|T| \leq \lambda$, or say that $T$ is a $\lambda^{\leq}$-subset of $S$.

The optimal $\lambda$-set $T^{*}$ of $S$ is a $\lambda$-set with $\operatorname{CON}_{S}\left(T^{*}\right)$ minimal among all $\lambda$-sets $T$. This set $T^{*}$ is the set we would like
to discard from our solution set $S$. The task of finding $T^{*}$ can easily be accomplished by computing $\operatorname{HYP}(S \backslash T)$ for all $\lambda$-sets $T$ (and $T=\emptyset$ ), as $\operatorname{CON}_{S}(T)=\operatorname{HYP}(S)-\operatorname{HYP}(S \backslash T)$. The main idea of our algorithm is that for doing this we do not have to compute these hypervolume measures independently, but can "parallelize" the execution of the algorithm the currently fastest hypervolume algorithm is based on: the algorithm of Overmars and Yap [16]. Therefore, we present their ideas in short and give a sketch of our changes afterwards.

The general framework of the algorithm is the same as the one of Bentley [2]: do a space sweep along dimension $d$ stopping at each endpoint of a box in decreasing order and, inserting the new box, solve the dynamic ( $d-1$ )-dimensional measure problem. The latter is the same as the problem of computing the hypervolume, but the boxes are given one by one and we have to output the current hypervolume after each box. Bentley's original approach to this dynamic problem took $\mathcal{O}\left(n^{d-2} \log n\right)$. Those $(d-1)$-dimensional measures then have to be multiplied by the length of the interval in dimension $d$ we overleaped and summed up to get the overall hypervolume of $S$. Confer [2] for details and correctness of this formula. Note that Bentley solved a more general problem than computing the hypervolume. In our context we, other than Bentley, never have to delete boxes, as all boxes have the same lower $d$-th coordinate 0 (as they all share the origin as a joint corner).

For the tree approach by Overmars and Yap [16] to the dynamic problem we need some more terminology. For a point $x \in \mathbb{R}^{d}$ we denote its $i$-th coordinate by $x_{i}$. In general, we will denote an $d$-dimensional object as $d$-object if we want to emphasize that it is an object lying in $\mathbb{R}^{d}$. We consider a $d$-box $B$ to be a set $\left[0, b_{1}\right] \times \ldots \times\left[0, b_{d}\right]$, so that we can think of $B$ also as the point $\left(b_{1}, \ldots, b_{d}\right) \in \mathbb{R}^{d}$. We will use this dualism often, speaking of boxes and points being the same. Moreover, we consider a $(d-1)$-region $R$, or just region for short, to be a set $\left[a_{1}, b_{1}\right] \times \ldots \times\left[a_{d-1}, b_{d-1}\right]$, i.e., a rectangular region in $\mathbb{R}^{d-1}$. For such a region we define $R^{*}$ to be the $d$-region $R \times[0, U]$, where $U:=\max \left\{x_{d} \mid\right.$ $x \in S\}$ is a fixed upper bound for the $d$-th coordinates of the points in $S$. Furthermore, we speak of the projected box $B_{\pi}$ by dropping the $d$-th coordinate of the box $B$, i.e., $B_{\pi}=\left[0, b_{1}\right] \times \ldots \times\left[0, b_{d-1}\right]$. Also, for a set $S$ of points (or boxes) in $\mathbb{R}^{d}$ we denote by $S_{\pi}$ the set of all projected boxes $\left\{\left(x_{1}, \ldots, x_{d-1}\right) \mid\left(x_{1}, \ldots, x_{d}\right) \in S\right\}$.

Definition 1. A $d$-box $B$ is said to partially cover a (d-1)-region $R$ if the boundary of $B_{\pi}$ intersects the interior of $R . B$ is said to (fully) cover $R$ if $R \subseteq B_{\pi}$.

We speak of the two $i$-boundaries of a box $B$ being the two sets $\left\{x \in B \mid x_{i}=0\right\}$ and $\left\{x \in B \mid x_{i}=b_{i}\right\}$. Additionally, we let the $i$-interval of a box or region to be its projection on the $x_{i}$-axis. Then we speak of a box $B$ being an $i$-pile of the region $R$, if for each $1 \leq j \leq d-1, j \neq i$ the $j$-interval of $R$ is fully contained in the $j$-interval of $B$. Less formally that means that within a region $R$ all $j$-intervals of $i$-piles of $R$ only differ for $j=i$. We consider a set $S$ of $d$-boxes to form a trellis in the region $R$, if each box in $S$ is an $i$-pile for some $1 \leq i \leq d-1$ in $R$. See Figure 2 for an illustration of a general trellis as defined by Overmars and Yap. At last, we will need a restricted hypervolume: For any region $R$ and finite point set $T \subset \mathbb{R}^{d}$ we define
$\operatorname{HYP}_{R^{*}}(T):=\operatorname{voL}\left(R^{*} \cap \bigcup_{x \in T}\left[0, x_{1}\right] \times \ldots \times\left[0, x_{d}\right]\right)$, which is the hypervolume dominated by $T$ restricted to $R^{*}$.

To calculate the volume efficiently, Overmars and Yap [16] cleverly use a partitioning of the $(d-1)$-dimensional space by an orthogonal partition tree. There each node $u$ is associated to a (leaf-)region $R_{u}$ The root is associated to a region $R_{\text {root }}$ with $R_{\text {root }}^{*}$ being a bounding box for all the boxes in $S$. We will in this paper always assume $R_{\text {root }}^{*}$ to be the box $\mathbf{B B}=\left[0, \mathbf{B B}_{1}\right] \times \ldots \times\left[0, \mathbf{B B}_{d}\right]$ with $\mathbf{B B}_{i}=\max \left\{x_{i} \mid x \in S\right\}$. Additionally, the associated region to each node splits up to the two children covering and intersection free, i.e., if $\ell(u)$ and $r(u)$ are the two children of $u$, then $R_{\ell(u)} \cup R_{r(u)}=R_{u}$ and $R_{\ell(u)} \cap R_{r(u)}$ has zero Lebesgue measure. In every leaf $\ell$ they require that any box in $S$ (the problem instance) that partially covers $R_{\ell}$ is a pile for $R_{\ell}$, so that the boxes in any leaf form a trellis.

Overmars and Yap [16] show how to build such a tree with several nice properties. Among others, they prove (in their Lemma 4.2) the following three properties which will be useful in the remainder.

Lemma 2. The partition tree built by the algorithm of Overmars and Yap has the following properties:

- The depth of the tree is $\mathcal{O}(\log n)$.
- Each projection of a box in $S$ partially covers $\mathcal{O}\left(n^{(d-2) / 2}\right)$ leaf-regions.
- Each projection of $a$ box in $S$ partially covers $\mathcal{O}\left(n^{(d-2) / 2} \log n\right)$ regions of inner nodes.

Note, we build this tree only for the first $d-1$ dimensions, while Overmars and Yap solve the problem in $d$ dimensions, which explains the difference in the statements.

For inserting or deleting a box in this tree, one only has to update the measure in each of the influenced regions. This can be done in constant time for internal nodes and in logarithmic time for the leafs, as we will see in the next but one section. Hence, this tree helps to efficiently determine the dynamic measure.

## Streaming variant

Overmars and Yap [16] also present a streaming variant of their algorithm. It uses less space but elsewise performs the same operations as the tree variant, just in a different ordering. The tree variant can be seen as a sweep in "time" (being the $d$-th coordinate), where we insert a box into the tree when it is reached in time. We can rearrange this the following way: We traverse the tree, and for each leaf we sweep the "time" inserting all boxes that influence the current leaf. In other words, we do not perform every insertion one by one on the whole tree structure, but look at the leafregions and perform all the insertions that will influence the region at hand. This rearrangement is possible as we know all insertion times beforehand. The benefit of the latter variant is that we do not have to explicitly store the whole tree structure: As Overmars and Yap managed to simulate the splitting of an inner node just by looking at the boxes that influence the associated region, we just need the tree structure implicitly, reducing the amount of storage to $\mathcal{O}(n)$. This variant fits better our purpose of a practical algorithm.

## Trellises

What is left is how to deal with the leaf-regions of the tree. Overmars and Yap [16] saw that maintaining the


Figure 1: A three-dimensional example of a set $M \subseteq \mathbb{R}_{\geqslant_{0}}^{3}$ such that the greedy $\lambda$-set $S_{\text {greedy }}^{\lambda}(M)$ gives a much higher contribution than the optimal $\lambda$-set $S_{\mathrm{opt}}^{\lambda}(M)$.
measure of a projected trellis dynamically can be done in $\mathcal{O}(\log n)$ : Consider a region $R$, i.e., a rectangle with side lengths $L_{1}, \ldots, L_{d-1}$. Furthermore, consider the $i$-piles of this region: by projecting them onto dimension $i$, we can determine their measure by solving a 1 -dimensional measure problem with overlapping intervals, which can be maintained in $\mathcal{O}(\log n)$ per update by an interval tree. This way, we get $M_{i}$, the measure of the 1-dimensional problem of the $i$-piles. Then we can compute the measure of the projected trellis easily as

$$
\prod_{i=1}^{d-1} L_{i}-\prod_{i=1}^{d-1}\left(L_{i}-M_{i}\right)
$$

as explained in [16].
Beume and Rudolph [5] noticed that in the case of the hypervolume indicator the measures $M_{i}$ can be maintained even in constant time, since we do not delete boxes and the interval overlapped by the $i$-piles is always of the form $[0, r]$ for some $r \in \mathbb{R}$, so that we just have to save the largest right end of such an interval, which can be updated in $\mathcal{O}(1)$.

## Sketch of our algorithm

Roughly speaking, the algorithm of Overmars and Yap can be summarized as follows:

- By building the partition tree, compute for all leafregions $R$ the hypervolume $\operatorname{HYP}_{R^{*}}(S)$ of the space dominated by $S$ restricted to $R^{*}$.
- Sum up these volumes.

The crucial observation is that the space dominated by $S$ restricted to $R^{*}$ as considered in the first step forms a trellis, whose hypervolume can be determined efficiently.

As sketched at the beginning of Section 3, we want to compute all the hypervolumes $\operatorname{HYP}(S \backslash T)$ for $\lambda$-sets $T$ in parallel. Observing that we can use the same partitioning tree for $S \backslash T$ as for $S$, we can come up with a simple adaption of the above method:


Figure 2: A 2-dimensional trellis for arbitrary boxes as in Overmars and Yap [16]. There, a trellis consists of long vertical (pairwise disjoint) rectangles superposed on long horizontal (pairwise disjoint) rectangles. The dotted rectangle around the trellis shows the corresponding region $R$.

- By building the partition tree, compute for all leafregions $R$ the hypervolume $\operatorname{HYP}_{R^{*}}(S \backslash T)$ for all $\lambda$-subsets $T$ of $S$.
- Sum up these volumes independently to get $\operatorname{HYP}(S \backslash T)$ for each $T$.

One way to carry out the first step is described in the following. For this, we compute for every leaf-region $R$ and all $\lambda^{\leq}$-sets $U$ the following volumes $M_{U}^{R}$.

Definition 2. The volume $M_{U}^{R}$ denotes the volume of the space in $R^{*}$ that is not dominated by $S \backslash U$, but is dominated by every $S \backslash W$ for $W \subsetneq U$.

This way every point in $R^{*}$ not dominated by $S \backslash T$ is counted in exactly one of the measures $M_{U}^{R}$ with $U \subseteq T$. We get

$$
\begin{equation*}
\operatorname{HYP}_{R^{*}}(S \backslash T)=\operatorname{voL}\left(R^{*}\right)-\sum_{U \subseteq T} M_{U}^{R} \tag{1}
\end{equation*}
$$

Using this and postponing the summation for each $T$ we can restate the method now as follows:

- Compute the measures $M_{U}^{R}$ for all leaf-regions $R$ and $\lambda \leq$-sets $U$
- Sum up these measures to get $M_{U}:=\sum_{R} M_{U}^{R}$ (the sum goes over all leaf-regions $R$ )
- For each $\lambda$-set $T$ compute $\operatorname{HYP}(S \backslash T)=\operatorname{vol}\left(R_{\text {root }}^{*}\right)-$ $\sum_{U \subseteq T} M_{U}$

The subtle point making this method superior to the naïve approach is that most of the $M_{U}^{R}$, s are actually zero. For seeing this, let us take a closer look on how one would compute the values $M_{U}^{R}$ : Inside $R^{*}$ we can do a space sweep along dimension $d$, just as in the algorithm of Bentley [2], stopping at each of the $d$-th coordinates of the points in $S$ in decreasing order. At the stop for $x \in S$, we have to insert the box $x$ and compute a ( $d-1$ )-dimensional measure, namely the volume $M_{U}^{R, x}$ :

Definition 3. The volume $M_{U}^{R, x}$ denotes the volume of the $(d-1)$-dimensional space in $R$ that is not dominated by $S_{\pi}^{x} \backslash U_{\pi}$, but is dominated by every $S_{\pi}^{x} \backslash W_{\pi}$ for $W \subsetneq U$, where $S^{x}=\left\{y \in S \mid y_{d} \geq x_{d}\right\}$ denotes the set of already inserted boxes.

We multiply this measure by the covered distance $x_{d}-x_{d}^{N}$ in the $d$-th dimension, where $x_{d}^{N}$ is the $d$-th coordinate of the next stop. This is summed up over all stops to get the measure $M_{U}^{R}$.

This way, we reduced the computation of $M_{U}^{R}$ to $M_{U}^{R, x}$, which is a measure inside the first $d-1$ dimensions of a trellis, which have a fairly simple geometric structure, as depicted in Figure 3. In the picture we can see for each part, to which measure $M_{U}^{R, x}$ it corresponds. The set $\{A, D\}$ for example marks the space corresponding to $M_{\{A, D\}}^{R, x}$. There we can also verify equation (1) (reduced to $d-1$ dimensions): The hypervolume of the space dominated by all but the points $A$ and $D$ equals the total volume of $R$ minus the volumes of the parts marked with $\{A, D\},\{A\},\{D\}$ and $\emptyset$, i.e., the subsets of $\{A, D\}$.

Moreover, observe that $M_{U}^{R, x}$ is non-zero only if $U$ contains the largest $\ell_{i} i$-piles in $R$ for some $\ell_{i}$, but no other $i$ pile, $i=1, \ldots, d-1$. In the example we may choose $\emptyset,\{A\}$, $\{A, B\}$ or $\{A, B, C\}$ being the 1-piles contained in $U$ to get a non-zero $M_{U}^{R, x}$, but not, e.g., only $\{B\}$. As we need to compute $M_{U}$ only for $\lambda^{\leq}$-sets $U$, we additionally condition on $|U|=\sum_{i=1}^{d-1} \ell_{i} \leq \lambda$. But then there are at most as many nonzero $M_{U}^{R, x}$,s, as there are $(d-1)$-tuples $\left(\ell_{1}, \ldots, \ell_{d-1}\right) \in \mathbb{N}_{0}^{d-1}$ with $\sum_{i=1}^{d-1} \leq \lambda$, which is a constant number for $d$ and $\lambda$ being constant. This is why there is only a constant number of non-zero $M_{U}^{R, x}$, s for $R$ and $x$ fixed, which implies that there is only a small number of non-zero $M_{U}^{R}$ 's for $R$ fixed. As we will see in the next section in detail, we can even determine those non-zero values quickly, even in the same asymptotic runtime as we need for the standard algorithm of Overmars and Yap. Computing the hypervolumes $\operatorname{HYP}(S \backslash T)$ for all $\lambda$-sets $T$ can then be accomplished by summing up all $M_{U}$ with $U \subseteq T$, as pointed out above, from which we can compute $\operatorname{CON}_{S}(T)=\operatorname{HYP}(S)-\operatorname{HYP}(S \backslash T)$, as we get $\operatorname{HYP}(S)=\operatorname{vOL}\left(R_{\text {root }}^{*}\right)-M_{\emptyset}$ for free, and thus quickly determine the optimal $\lambda$-set.

## Details of the algorithm

## Computemeasures

One obvious arising problem concerns boxes that fully cover a region of some inner node of the tree. In such a case Overmars and Yap collapse the interval in dimension $d$, where the region is fully covered, into a single moment, memorizing the deleted volume, and recur. We may not do this, as the fully covering box may be in the set $T$ we disregard, so that in $\operatorname{HYP}(S \backslash T)$ the region is not fully covered. This is why we do not collapse any intervals, but have to pass the fully covering boxes to the recursive calls, so that we can deal with them in the leaf-nodes. Note that the runtime analysis of Overmars and Yap's algorithm does not rely at any point on collapsing intervals, which is why we get the same asymptotic runtime. It does, however, rely on the fact, that inside a leaf-node we spend time $\mathcal{O}\left(\left|S^{\prime}\right| \log n\right)$ and in an inner node $\mathcal{O}\left(\left|S^{\prime}\right|\right)$, where $S^{\prime}$ is the set of boxes in $S$ that partially cover the region at hand. Hence, we may not pass all fully covering boxes to the recursive calls to be inside this time bound. Luckily, any measure $M_{U}^{R}$ is zero, if the fully covering boxes of $R$ contained in $U$ are not the $\ell$ largest ones for some $\ell$, i.e., have largest $d$-th coordinate among all fully covering boxes. This stems from the fact that for fully covering boxes $x \in U, y \notin U$ with $d$-th coordinates $x_{d} \leq y_{d}$


Figure 3: The first two dimensions of a 3-dimensional trellis consisting of the hypervolume-boxes defined by $S=$ $\{A, B, C, D, E, F\}$. Here, a trellis consists of boxes that cover the region completely in each of the $(d-1)$ dimensions except one. The dotted rectangle indicates the corresponding region $R$.
we have $\operatorname{HYP}_{R^{*}}(S \backslash U)=\operatorname{HYP}_{R^{*}}(S \backslash(U \backslash\{x\}))$, so that $M_{U}^{R}=0$. This is why we need to pass only the up to $\lambda+1$ largest covering boxes to the recursive calls. Since this is a constant number, we do not increase the runtime in an inner node or leaf-node asymptotically.

The streaming variant of [16] is essentially the same as the algorithm ComputeMeasures (cf. Algorithm 1), only that we added the set Cov containing the up to $\lambda+1$ largest covering boxes, i.e., out of the set $C$ of boxes that fully cover the region $R$ at hand (including covering boxes of any parent region) we save the $\min \{\lambda+1,|C|\}$ many ones with greatest $d$-th coordinate. This set is updated determining the set $U \subseteq S$ of boxes fully covering the region $R$, where $S$ is the current set of boxes. Everything after determining the set Cov' is copied from Overmars and Yap: We proceed by computing the measures in a trellis, if the remaining boxes $S^{\prime}$ form one, and by splitting the region $R$ into two regions $R_{1}, R_{2}$ and recursing, otherwise. For splitting we need the sets $S_{1}^{\prime}$ and $S_{2}^{\prime}$, where $S_{1}^{\prime}$ is the set of all boxes in $S^{\prime}$ that have a 1 - or 2 - or $\ldots$ or $(i-1)$-boundary in $R$ and $S_{2}^{\prime}$ is the set of boxes in $S^{\prime}$ that do not have such a boundary in $R$. For details of this splitting method confer [16]. Note that we never split a region along dimension $d$ as is implied by the use of Overmars and Yaps splitting method and our definition of trellis (which considers only the first $d-1$ dimensions).

The procedure ComputeMeasuresTrellis will need the boxes $S^{\prime}$ to be sorted by $d$-th coordinate. This can be achieved easily by sorting the boxes before the first call of ComputeMeasures and maintaining this ordering during all steps of ComputeMeasures, without increasing the overall asymptotic runtime. Hence, we may assume in the following, that $S^{\prime}$ in the input of ComputeMeasuresTrellis is sorted.

Computing the set $\mathrm{Cov}^{\prime}$ can be done in $\mathcal{O}(|S|)$, assuming $\lambda$ to be constant. Hence, as long as we provide a Compute-MeasuresTrellis-function which runs in $\mathcal{O}\left(\left|S^{\prime}\right| \log n\right)$ for the set of boxes $S^{\prime}$, which is exactly the same runtime as Overmars and Yap's, we do not increase the overall runtime of $\mathcal{O}\left(n^{d / 2} \log n\right)$ of their algorithm.

For determining the needed storage consider the following trick of Overmars and Yap [16]: If we save the boxes in $U$ we can reconstruct the old $S$ by joining $U$ and $S^{\prime}$. Hence, we can send $S^{\prime}$ down the recursion, not copying it; we just have to reconstruct it at the end of the recursion call. This way,

```
Algorithm 1 ComputeMeasures ( \(R, S, i, \mathrm{Cov}, \lambda\) ) com-
putes the measures \(M_{U}\) of the \(\lambda^{\leq}\)-subsets \(U\) of the set of
boxes \(S\) in the region \(R^{*} \subseteq \mathbb{R}^{d}\), where \(i\) is the current split-
ting dimension and Cov is a set containing the up to \(\lambda+1\)
largest covering boxes.
    discard boxes in \(S\) not influencing \(R\)
    determine the set \(U \subseteq S\) of boxes fully covering \(R\)
    \(S^{\prime}:=S \backslash U\)
    determine the new set \(\operatorname{Cov}^{\prime} \subseteq \operatorname{Cov} \cup U\)
    if the boxes in \(S^{\prime}\) form a trellis in \(R\) then
        ComputeMeasuresTrellis \(\left(R, S^{\prime}, \operatorname{Cov}^{\prime}, \lambda\right)\)
    else
        determine the sets \(S_{1}^{\prime}\) and \(S_{2}^{\prime}\) (as defined on page )
        if \(S_{1}^{\prime} \neq \emptyset\) then
            split \(R\) into \(R_{1}, R_{2}\) along the median \(i\)-boundary in
            \(S_{1}^{\prime}\)
            ComputeMeasures \(\left(R_{1}, S^{\prime}, i, \operatorname{Cov}^{\prime}, \lambda\right)\)
            ComputeMeasures \(\left(R_{2}, S^{\prime}, i, \operatorname{Cov}^{\prime}, \lambda\right)\)
    else if \(S_{2}^{\prime}\) contains more than \(\sqrt{n} i\)-boundaries then
            split \(R\) into \(R_{1}, R_{2}\) along the median \(i\)-boundary in
            \(S_{2}^{\prime}\)
            ComputeMeasures \(\left(R_{1}, S^{\prime}, i, \operatorname{Cov}^{\prime}, \lambda\right)\)
            ComputeMeasures \(\left(R_{2}, S^{\prime}, i, \operatorname{Cov}^{\prime}, \lambda\right)\)
    else
        \(\operatorname{ComputeMeasures}\left(R, S^{\prime}, i+1, \operatorname{Cov}^{\prime}, \lambda\right)\)
    od
od
```

no box is saved at two places at any time, so that the overall space for the sets $S$ and $U$ is just $\mathcal{O}(n)$. Since the size of Cov is $\lambda$ and thus constant, we can save it normally, getting an additional space needed of $\mathcal{O}(\lambda \log n)$, as the recursion depth equals the depth of the partition tree which is $\mathcal{O}(\log n)$ by Lemma 2, so that we overall need a storage of $\mathcal{O}(n)$. Note that if we would not follow this trick we had a storage of $\mathcal{O}(n \log n)$, i.e., $\mathcal{O}(n)$ in each of the $\mathcal{O}(\log n)$ levels of the recursion.

## ComputeMeasuresTrellis

To complete the description of ComputeMeasures we have to provide the procedure ComputeMeasuresTrellis, which will report the measures $M_{U}^{R}$. These measure will then be summed up to get $M_{U}$ from which we can directly compute the hypervolume dominated by any $S \backslash T$, for $T$ a $\lambda$-subset of $S$, as sketched in the preceeding section. One way to compute these measures is given in Algorithm 2.

There, at first, we remove all boxes from the set of boxes $S$ that do not influence the current region $R$ at all. The variable $x_{d}^{L}$ is going to be the $d$-th coordinate of the last box inserted and initialized to $\mathbf{B B}_{d}$. We will maintain an ordered list of the up to $\lambda+1$ largest $i$-piles of the $(d-1)$-region $R$ for each $i$, i.e., the $i$-piles with greatest $i$-th coordinate. Those will be the boxes $A_{j}^{i}, 1 \leq i<d, 1 \leq j \leq \lambda+1$, which are undefined initially and get updated every time we insert a box, so that $A_{1}^{i}$ is the greatest $i$-pile, $A_{2}^{i}$ the second greatest and so on. We use $(A)_{i}$ to denote the maximal $i$-th coordinate of a point in a box $A$, i.e., $A_{i}$ (viewed as a point). For simplicity, we define $(A)_{i}$ to be 0 if $A$ is undefined and $A_{0}^{i}=R$ to be the region itself for each $i$.

Going further, we add Cov to $S$, which we need, since each fully covering box is a pile of $R$, and those boxes are

```
Algorithm 2 ComputeMeasuresTrellis \((R, S, \operatorname{Cov}, \lambda)\)
computes the measures \(M_{U}^{R}\) for each \(\lambda^{\leq}\)-subset \(U\) of \(S\),
where the boxes in \(S\) restricted to \(R^{*} \subseteq \mathbb{R}^{d}\) form a trellis
and Cov is a set containing the up to \(\lambda+1\) largest covering
boxes.
    discard boxes in \(S\) not influencing \(R\)
    set \(A_{j}^{i}:=\operatorname{undef}(1 \leq i \leq d-1,1 \leq j \leq \lambda+1)\)
    set \(A_{0}^{i}:=R(1 \leq i \leq d-1)\)
    \(S:=S \cup \operatorname{Cov} \cup\{(0, \ldots, 0)\}\)
    \(x_{d}^{L}:=\mathbf{B B}_{d}\)
    initialize \(M_{U}^{R}\), s to 0
    for all \(x \in S\) ordered by decreasing \(x_{d}\) do
        for all \(\left(k_{1}, \ldots, k_{d-1}\right) \in \mathbb{N}_{0}^{d-1}\) with \(\sum_{i=1}^{d-1} k_{i} \leq \lambda\) and
        \(A_{k_{i}}^{i}\) defined for all \(i\) do
            \(U:=\left\{A_{j}^{i} \mid 1 \leq i<d\right.\) and \(\left.1 \leq j \leq k_{i}\right\}\)
            \(M_{U}^{R, x}:=\prod_{i=1}^{d-1}\left(\left(A_{k_{i}}^{i}\right)_{i}-\left(A_{k_{i}+1}^{i}\right)_{i}\right)\)
            \(M_{U}^{R}:=M_{U}^{R}+\left(x_{d}^{L}-x_{d}\right) \cdot M_{U}^{R, x}\)
        od
        if \(x\) is a \(k\)-pile: update \(A_{j}^{k}(1 \leq j \leq \lambda+1)\)
        \(x_{d}^{L}:=x_{d}\)
    od
    for all \(U\) with non-zero \(M_{U}^{R}\) do
        \(M_{U}:=M_{U}+M_{U}^{R}\)
    od
```

not already in $S$. Since we need the set $S$ sorted according to $d$-th coordinate and $S$ was sorted in the beginning, we have to insert the points in Cov into $S$ properly, which can be done in $\mathcal{O}(|S|)$ as $|\operatorname{Cov}|=\mathcal{O}(1)$. Additionally, we need to add the dummy point $(0, \ldots, 0)$ to $S$, as we want to sweep along the entire $d$-interval of $R^{*}$, i.e., we want to end at 0 .

Now, we go through all the boxes in $S$ ordered by $d$-th coordinate in decreasing order. For each point $x \in S$, we go through all the tuples $\left(k_{1}, \ldots, k_{d-1}\right) \in \mathbb{N}_{0}^{d-1}$ with $\sum_{i=1}^{d-1} k_{i} \leq \lambda$, but only those, for which $A_{k_{i}}^{i}$ is not undefined. Each such tuple corresponds to a set $U=\left\{A_{j}^{i} \mid 1 \leq\right.$ $i<d$ and $\left.1 \leq j \leq k_{i}\right\}$, where all occurring $A_{j}^{i}$ are defined for the condition mentioned before. We then compute the ( $d-1$ )-dimensional measure not covered by $S^{x} \backslash U$, but by $S^{x} \backslash W$ for any $W \subsetneq U$, where $S^{x}=\left\{y \in S \mid y_{d}>x_{d}\right\}$ denotes the set of the already inserted boxes. This measure is $M_{U}^{R, x}=\prod_{i=1}^{d-1}\left(\left(A_{k_{i}}^{i}\right)_{i}-\left(A_{k_{i}+1}^{i}\right)_{i}\right)$. It has to be multiplied by the length of the interval of the $d$-th coordinate we are currently regarding, which is $x_{d}^{L}-x_{d}$ (as $x_{d}^{L}$ was the last $d$-th coordinate of an insertion). The resulting product has to be added to $M_{U}^{R}$. We implicitly initialize the measures $M_{U}^{R}$ to 0 . Also, we do not want to explicitly save each value $M_{U}^{R}$, as most of them are zero, but save only the non-zero ones. Both points can be achieved by using a dynamic hash-table, that contains $U$ and $M_{U}^{R}$ iff $M_{U}^{R}$ is non-zero.

Afterwards, we determine the number $k, 1 \leq k<d$, for which $x$ is a $k$-pile in $R$. If this number is not unique, which can only happen if the box fully covers $R$, assign an arbitrary $1 \leq k<d$. Then we update the largest $k$-piles $A_{j}^{k}, 1 \leq j \leq$ $\lambda+1$, i.e., we insert $x$ at the correct position, shifting all smaller ones by one position.

In the end we report the computed measures $M_{U}^{R}$, i.e., we add them to $M_{U}$. Here, again, we will implicitly initializing each $M_{U}$ with 0 (before the start of ComputeMeasures) and save $U$ and $M_{U}$ in a dynamic hash-table for every nonzero $M_{U}$.

Considering the runtime we see that everything inside the main loop can be done in constant time: Since $d$ and $\lambda$ are considered to be constant, we have to update a constant number of boxes $A_{j}^{i}$. Furthermore, there are at most $(\lambda+1)^{d-1}$ many tuples $\left(k_{1}, \ldots, k_{d-1}\right)$, since every entry lies between 0 and $\lambda$. Since we can view this as assigning at most $\lambda$ many ones to $d-1$ many buckets, the number of tuples is also bounded from above by $\sum_{i=0}^{\lambda}(d-1)^{i}<(d-1)^{\lambda+1}$. All we do with such a tuple can be done in constant time for the same reason, which establishes, that ComputeMeasuresTrellis runs in $\mathcal{O}(|S|)$. Observe that this is even better than Overmars and Yaps runtime of $\mathcal{O}(|S| \log n)$, so that we definitely get their overall asymptotic runtime of $\mathcal{O}\left(n^{d / 2} \log n\right)$.

## Correctness

We now show that the above methods are indeed correct.

Lemma 3. The measures $M_{U}$ computed by ComputeMeasures satisfy the following equation for any $\lambda^{\leq}$-set $T$ of $S$ :

$$
\operatorname{HYP}(S \backslash T)=\operatorname{voL}(\boldsymbol{B B})-\sum_{U \subseteq T} M_{U}
$$

Proof. The described algorithm partitions the bounding box BB into a number of leaf-regions that contain trellises. Since we sum up over all of those regions, all we have to show is that for each region $R$ for which we call ComputeMeasuresTrellis it holds that

$$
\operatorname{HYP}_{R^{*}}(S \backslash T)=\operatorname{voL}\left(R^{*}\right)-\sum_{U \subseteq T} M_{U}^{R}
$$

Now, inside $R^{*}$ we sweep along the $d$-th dimension considering intervals $\left[x_{d}, x_{d}^{L}\right.$ ], where the boxes influencing the ( $d-1$ )-dimensional measures stay the same, and sum up weighted by $x_{d}^{L}-x_{d}$. Since we start with $x_{d}^{L}=\mathbf{B B}_{d}$ and end at the dummy point $(0, \ldots, 0)$ with $x_{d}=0$ our sweep indeed covers the interval $\left[0, \mathbf{B B}_{d}\right]$. Hence, the summation along dimension $d$ is correct as long as it holds that at the stop for $x \in S$, with $S^{x}=\{y \in S$ | $y_{d}>x_{d}$ and $y$ partially covers $\left.R\right\}$ the set of already inserted boxes:

$$
\begin{equation*}
\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}\right)=\operatorname{voL}(R)-\sum_{U \subseteq T} M_{U}^{R, x} \tag{2}
\end{equation*}
$$

Note that we are here dealing with $(d-1)$-dimensional volumes, which is why we used the projected set of boxes $S_{\pi}^{x}$ and $T_{\pi}$.

Let $T$ be of the form as in the pseudo code, i.e., $T=\left\{A_{j}^{i} \mid\right.$ $1 \leq i<d$ and $\left.1 \leq j \leq k_{i}\right\}$ for some $\left(k_{1}, \ldots, k_{d-1}\right) \in \mathbb{N}_{0}^{d-1}$ with $\sum_{i=1}^{d-1} k_{i} \leq \lambda$ and all $A_{k_{i}}^{i}$ defined. We compute (nonzero) measures $M_{U}^{R, x}$ only for subsets $U \subseteq T$ of the form $\left\{A_{j}^{i} \mid 1 \leq i<d\right.$ and $\left.1 \leq j \leq \ell_{i}\right\}$ for some $\left(\ell_{1}, \ldots, \ell_{d-1}\right) \in$
$\mathbb{N}_{0}^{d-1}$ with $\ell_{i} \leq k_{i}$ for all $i$. Thus, we have:

$$
\begin{aligned}
& \sum_{U \subseteq T} M_{U}^{R, x} \\
& =\sum_{\substack{\left(\ell_{1}, \ldots, \ell_{d-1}\right) \in \mathbb{N}^{d-1} \\
\ell_{i} \leq k_{i} \text { for all }{ }_{i}}} M_{\left\{A_{j}^{i} \mid 1 \leq i<d \text { and } 1 \leq j \leq \ell_{i}\right\}}^{R, x} \\
& =\sum_{\substack{\left(\ell_{1}, \ldots, \ell_{d-1}\right) \in \mathbb{N}_{0}^{d-1} \\
\ell_{i} \leq k_{i} \text { for all } i}} \prod_{i=1}^{d-1}\left(\left(A_{\ell_{i}}^{i}\right)_{i}-\left(A_{\ell_{i}+1}^{i}\right)_{i}\right) \\
& =\prod_{i=1}^{d-1} \sum_{\ell_{i}=0}^{k_{i}}\left(\left(A_{\ell_{i}}^{i}\right)_{i}-\left(A_{\ell_{i}+1}^{i}\right)_{i}\right) \\
& =\prod_{i=1}^{d-1}\left(\left(A_{0}^{i}\right)_{i}-\left(A_{k_{i}+1}^{i}\right)_{i}\right) \\
& =\prod_{i=1}^{d-1}\left(R_{i}-\left(A_{k_{i}+1}^{i}\right)_{i}\right) .
\end{aligned}
$$

There, we denote by $R_{i}$ the maximal $i$-th coordinate of a point in $R$. Observe that $S_{\pi}^{x} \backslash T_{\pi}$ is of a very simple form (restricted to $R$ ): It forms a trellis where the maximal $i$-pile is $A_{k_{i}+1}^{i}$. This means that the space inside $R$ not overlapped by this trellis is a rectangle with side lengths $R_{i}-\left(A_{k_{i}+1}^{i}\right)_{i}$ for $i=1, \ldots, d-1$, so that we established the equality

$$
\sum_{U \subseteq T} M_{U}^{R, x}=\operatorname{vOL}(R)-\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}\right)
$$

Note that the above argument also makes sense if $A_{k_{i}+1}^{i}$ is not defined, since then we have $\left(A_{k_{i}+1}^{i}\right)_{i}=0$. This gives us correctness for sets $T$ of the aforementioned form.

If, on the other hand, $T$ is not of the indicated form, then it has some maximal subset $T^{\prime} \subseteq T$, which is of this form, i.e., $T^{\prime}=\left\{A_{j}^{i} \mid 1 \leq i<d\right.$ and $\left.1 \leq j \leq k_{i}\right\}$ for some $\left(k_{1}, \ldots, k_{d-1}\right) \in \mathbb{N}_{0}^{d-1}$. Since $T^{\prime}$ is maximal, either $A_{k_{i}+1}^{i}$ is not contained in $T$ or it is not defined. In both cases every box in $T_{\pi} \backslash T_{\pi}^{\prime}$ is included in some box in $S_{\pi}^{x} \backslash T_{\pi}^{\prime}$, so that those boxes do not influence the measure in $R$, i.e., we have $\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}\right)=\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}^{\prime}\right)$. Also, we will report a measure for a set $U \subseteq T$ only if $U \subseteq T^{\prime}$ by construction, so that we have, using the former case:

$$
\begin{aligned}
\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}\right) & =\operatorname{HYP}_{R}\left(S_{\pi}^{x} \backslash T_{\pi}^{\prime}\right) \\
& =\operatorname{voL}(R)-\sum_{U \subseteq T^{\prime}} M_{U}^{R, x} \\
& =\operatorname{voL}(R)-\sum_{U \subseteq T} M_{U}^{R, x}
\end{aligned}
$$

This shows the desired equality. It also implies that any box which is not among the largest $\lambda$ in one dimension does not contribute to our measures at all, which also makes clear why we only need the $\lambda+1$ largest covering boxes in Cov.

## Putting everything together

After we computed the measures $M_{U}$, we can compute the actual contribution of a $\lambda$-set $T$ easily, using Lemma 3 . We
have:

$$
\begin{aligned}
\operatorname{CON}_{S}(T) & =\operatorname{HYP}(S)-\operatorname{HYP}(S \backslash T) \\
& =\left(\operatorname{voL}(\mathbf{B B})-M_{\emptyset}\right)-\left(\operatorname{voL}(\mathbf{B B})-\sum_{U \subseteq T} M_{U}\right) \\
& =\sum_{\emptyset \neq U \subseteq T} M_{U}
\end{aligned}
$$

This confirms the following combined procedure:

```
Algorithm 3 ComputeOptimalSubset \((S, \lambda)\) computes
the optimal \(\lambda\)-subset \(T\) of the set of boxes \(S\) in \(\mathbb{R}^{d}\).
    initialize \(M_{U}\) 's to 0
    ComputeMeasures(BB, \(S, 1, \emptyset, \lambda\) )
    return \(\operatorname{argmin}\left\{\sum_{\emptyset \neq U \subseteq T} M_{U}|T \subseteq S,|T|=\lambda\}\right.\)
```

Since $T$ has size at most $\lambda$, it has at most $2^{\lambda}$ subsets, which is a constant. Hence, given the measures $M_{U}$ we can compute the contribution of all $\lambda$-sets in $\mathcal{O}\left(n^{\lambda}\right)$ (as their number is bounded by this), so that we get an overall runtime of $\mathcal{O}\left(n^{d / 2} \log n+n^{\lambda}\right)$ for ComputeOptimalSubset. Its correctness follows directly from Lemma 3.

As mentioned in the preceding section, we need a big hashtable to store the non-zero measures $M_{U}$. Since there are $\mathcal{O}\left(n^{\lambda}\right) \lambda^{\leq}$-subsets of a size- $n$ set, there are $\mathcal{O}\left(n^{\lambda}\right)$ entries in the hash. On the other hand, by Lemma 2 each $d$-box partially covers $\mathcal{O}\left(n^{(d-2) / 2}\right)$ many leaf-regions. As ComputeMeasuresTrellis runs in $\mathcal{O}\left(\left|S^{\prime}\right|\right)$, where $S^{\prime}$ is the set of boxes in $S$ that partially cover the region $R$ at hand, we report $\mathcal{O}\left(\left|S^{\prime}\right|\right)$ non-zero measures $M_{U}^{R}$ in each region. This way, we get an upper bound of $\mathcal{O}\left(n^{d / 2}\right)$ reported non-zero measures. Since there are at most this many entries in the hash-table, ComputeOptimalSubset needs a space of $\mathcal{O}\left(\min \left(n^{d / 2}, n^{\lambda}\right)\right)$, using a dynamically growing hash-table.

## 4. DISCUSSION

We have presented an algorithm which calculates the optimal $\lambda$-set $S_{\text {opt }}^{\lambda}(M)$ of a population of size $n=|M|$ in time $\mathcal{O}\left(n^{d / 2} \log n+n^{\lambda}\right)$ for $d>2$. For $d>3$ this improves all previously published algorithms by a factor of order $n^{\min \{\lambda, d / 2\}}$.

For small $\lambda(\lambda \leqslant d / 2)$ the algorithm gives an improvement in the runtime of the calculation of the hypervolume by a factor of order $n^{\lambda}$. Hence even for the greedy calculation of $S_{\text {opt }}^{1}=S_{\text {greedy }}^{1}$, we have a speed up by a factor of $n$.

For very large $\lambda$ the algorithm might still be intractable. It is open whether this can be avoided. Our algorithm allows the calculation of $S_{\mathrm{opt}}^{\lambda}$ in the same time as $S_{\mathrm{opt}}^{1}$ if $\lambda \leqslant d / 2$. We therefore suggest the following compromise between $S_{\text {opt }}^{\lambda}$ and $S_{\text {greedy }}^{\lambda}$ for large $\lambda$ :

$$
\begin{array}{rr}
S_{\text {comp }}^{\lambda}(M):=S_{\text {opt }}^{\lambda}(M) & \text { for all } \lambda \leqslant d / 2, \\
S_{\text {comp }}^{\lambda+d / 2}(M):=S_{\text {comp }}^{\lambda}(M) \cup S_{\text {opt }}^{d / 2}\left(M \backslash S_{\text {comp }}^{\lambda}(M)\right) \\
& \text { for all } \lambda>d / 2 .
\end{array}
$$

As $\operatorname{CON}\left(S_{\text {greedy }}^{\lambda}\right) \geqslant \operatorname{CON}\left(S_{\text {comp }}^{\lambda}\right) \geqslant \operatorname{CON}\left(S_{\text {opt }}^{\lambda}\right)$ for all $\lambda$, above improved greedy algorithm returns $\lambda$-sets with the same or maybe smaller contributions than the classical greedy algorithm within the same asymptotic runtime.

## Acknowledgements

The authors would like to thank the anonymous reviewers for several useful suggestions regarding the presentation of this paper. This work was partially supported by a postdoctoral fellowship from the German Academic Exchange Service (DAAD).

## References

[1] J. Bader and E. Zitzler. HypE: An Algorithm for Fast Hypervolume-Based Many-Objective Optimization. TIK Report 286, Institut für Technische Informatik und Kommunikationsnetze, ETH Zürich, Nov. 2008.
[2] J. L. Bentley. Algorithms for Klee's rectangle problems, 1977. Department of Computer Science, Carnegie Mellon University, Unpublished notes.
[3] N. Beume, C. Fonseca, M. López-Ibáñez, L. Paquete, and J. Vahrenhold. On the complexity of computing the hypervolume indicator. Technical report CI-235/07, Technical University of Dortmund, December 2007.
[4] N. Beume, B. Naujoks, and M. Emmerich. SMSEMOA: Multiobjective selection based on dominated hypervolume. European Journal of Operational Research, 181(3):1653-1669, 2007.
[5] N. Beume and G. Rudolph. Faster S-metric calculation by considering dominated hypervolume as Klee's measure problem. In Proc. Second International Conference on Computational Intelligence (IASTED '06), pages 233-238, 2006.
[6] L. Bradstreet, L. Barone, and L. While. Maximising hypervolume for selection in multi-objective evolutionary algorithms. In Proc. IEEE Congress on Evolutionary Computation (CEC '06), pages 6208-6215, 2006.
[7] K. Bringmann and T. Friedrich. Approximating the volume of unions and intersections of high-dimensional geometric objects. In Proc. 19th International Symposium on Algorithms and Computation (ISAAC '08), volume 5369 of Lecture Notes in Computer Science, pages 436447. Springer, 2008.
[8] K. Bringmann and T. Friedrich. Approximating the least hypervolume contributor: NP-hard in general, but fast in practice. In Proc. 5th International Conference on Evolutionary Multi-Criterion Optimization (EMO '09), 2009. To appear, available from http://arxiv.org/abs/0812.2636.
[9] D. Brockhoff and E. Zitzler. Improving hypervolumebased multiobjective evolutionary algorithms by using objective reduction methods. In Proc. IEEE Congress on Evolutionary Computation (CEC '07), pages 20862093, 2007.
[10] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In Proc. IEEE Congress on Evolutionary Computation (CEC '02), pages 825-830, 2002.
[11] M. Emmerich, N. Beume, and B. Naujoks. An EMO algorithm using the hypervolume measure as selection criterion. In Proc. Third International Conference on Evolutionary Multi-Criterion Optimization (EMO '05), pages 62-76, 2005.
[12] C. Igel, N. Hansen, and S. Roth. Covariance matrix adaptation for multi-objective optimization. Evol. Comput., 15(1):1-28, 2007.
[13] J. Knowles, D. Corne, and M. Fleischer. Bounded archiving using the Lebesgue measure. In Proc. IEEE Congress on Evolutionary Computation (CEC'03), volume 4, pages 2490-2497, Dec 2003.
[14] J. D. Knowles. Local-Search and Hybrid Evolutionary Algorithms for Pareto Optimization. PhD thesis, Department of Computer Science, University of Reading, UK, 2002.
[15] B. Naujoks, N. Beume, and M. T. M. Emmerich. Multi-objective optimisation using S-metric selection: application to three-dimensional solution spaces. In Proc. IEEE Congress on Evolutionary Computation (CEC '05), pages 1282-1289, 2005.
[16] M. H. Overmars and C.-K. Yap. New upper bounds in Klee's measure problem. SIAM J. Comput., 20(6):10341045, 1991. Announced at 29th Annual Symposium on Foundations of Computer Science (FOCS '88).
[17] R. L. While, L. Bradstreet, L. Barone, and P. Hingston. Heuristics for optimizing the calculation of hypervolume for multi-objective optimization problems. In Proc. IEEE Congress on Evolutionary Computation (CEC '05), pages 2225-2232, 2005.
[18] R. L. While, P. Hingston, L. Barone, and S. Huband. A faster algorithm for calculating hypervolume. IEEE Trans. Evolutionary Computation, 10(1):29-38, 2006.
[19] X. Zhou, N. Mao, W. Li, and C. Sun. A fast algorithm for computing the contribution of a point to the hypervolume. In Proc. Third International Conference on Natural Computation (ICNC '07), volume 4, pages 415420, 2007.
[20] E. Zitzler. Hypervolume metric calculation, 2001. Computer Engineering and Networks Laboratory (TIK), ETH Zurich, Switzerland, ftp://ftp.tik.ee.ethz.ch/pub/people/zitzler/hypervol.c.
[21] E. Zitzler, D. Brockhoff, and L. Thiele. The hypervolume indicator revisited: On the design of Pareto-compliant indicators via weighted integration. In Proc. Fourth International Conference on Evolutionary Multi-Criterion Optimization (EMO '07), volume 4403 of Lecture Notes in Computer Science, pages 862876. Springer, 2007.
[22] E. Zitzler and S. Künzli. Indicator-based selection in multiobjective search. In Proc. 8th International Conference Parallel Problem Solving from Nature (PPSN VIII), volume 3242 of Lecture Notes in Computer Science, pages 832-842. Springer, 2004.
[23] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength Pareto approach. IEEE Trans. Evolutionary Computation, 3(4):257-271, 1999. Announced at 5th International Conference Parallel Problem Solving from Nature (PPSN V).
[24] E. Zitzler, L. Thiele, M. Laumanns, C. M. Fonseca, and V. G. da Fonseca. Performance assessment of multiobjective optimizers: an analysis and review. IEEE Trans. Evolutionary Computation, 7(2):117-132, 2003.


[^0]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
    FOGA'09, January 9-11, 2009, Orlando, Florida, USA.
    Copyright 2009 ACM 978-1-60558-414-0/09/01 ...\$5.00.

[^1]:    ${ }^{1}$ Without loss of generality we assume the reference point to be $0^{d}$. Also, we are dealing with maximization problems throughout the paper.

[^2]:    ${ }^{2}$ Throughout the paper, we will use $\operatorname{CON}(x):=\operatorname{CON}_{M}(x)$, $S_{\mathrm{opt}}^{\lambda}:=S_{\mathrm{opt}}^{\lambda}(M)$ and $S_{\text {greedy }}^{\lambda}:=S_{\text {greedy }}^{\lambda}(M)$ if there is no ambiguity in the choice of $M$.

