

Bringing Order to Special Cases of Klee’s Measure Problem

Karl Bringmann

Max Planck Institute for Informatics

Abstract. Klee’s Measure Problem (KMP) asks for the volume of the union of n axis-aligned boxes in \mathbb{R}^d . Omitting logarithmic factors, the best algorithm has runtime $\mathcal{O}^*(n^{d/2})$ [Overmars,Yap’91]. There are faster algorithms known for several special cases: CUBE-KMP (where all boxes are cubes), UNITCUBE-KMP (where all boxes are cubes of equal side length), HYPERVOLUME (where all boxes share a vertex), and k -GROUNDED (where the projection onto the first k dimensions is a HYPERVOLUME instance).

In this paper we bring some order to these special cases by providing reductions among them. In addition to the trivial inclusions, we establish HYPERVOLUME as the easiest of these special cases, and show that the runtimes of UNITCUBE-KMP and CUBE-KMP are polynomially related. More importantly, we show that any algorithm for one of the special cases with runtime $T(n, d)$ implies an algorithm for the general case with runtime $T(n, 2d)$, yielding the first non-trivial relation between KMP and its special cases. This allows to transfer W[1]-hardness of KMP to all special cases, proving that no $n^{o(d)}$ algorithm exists for any of the special cases assuming the Exponential Time Hypothesis. Furthermore, assuming that there is no *improved* algorithm for the general case of KMP (no algorithm with runtime $\mathcal{O}(n^{d/2-\varepsilon})$) this reduction shows that there is no algorithm with runtime $\mathcal{O}(n^{\lfloor d/2 \rfloor / 2 - \varepsilon})$ for any of the special cases. Under the same assumption we show a tight lower bound for a recent algorithm for 2-GROUNDED [Yıldız,Suri’12].

1 Introduction

Klee’s measure problem (KMP) asks for the volume of the union of n axis-aligned boxes in \mathbb{R}^d , where d is considered to be a constant. This is a classic problem with a long history [2, 9, 10, 14, 16, 17]. The fastest algorithm has runtime $\mathcal{O}(n^{d/2} \log n)$ for $d \geq 2$, given by Overmars and Yap [16], which was slightly improved to $n^{d/2} 2^{\mathcal{O}(\log^* n)}$ by Chan [9]. Thus, for over twenty years there has been no improvement over the runtime bound $n^{d/2}$. As already expressed in [9], one might conjecture that no *improved* algorithm for KMP exists, i.e., no algorithm with runtime $\mathcal{O}(n^{d/2-\varepsilon})$ for some $\varepsilon > 0$.

However, no matching lower bound is known, not even under reasonable complexity theoretic assumptions. The best unconditional lower bound is $\Omega(n \log n)$ for any dimension d [10]. Chan [9] proved that KMP is W[1]-hard by giving

a reduction to the k -Clique problem. Since his reduction has $k = d/2$, we can transfer runtime lower bounds from k -Clique to KMP, implying that there is no $n^{o(d)}$ algorithm for KMP assuming the Exponential Time Hypothesis (see [15]). However, this does not determine the correct constant in the exponent. Moreover, Chan argues that since no “purely combinatorial” algorithm with runtime $\mathcal{O}(n^{k-\varepsilon})$ is known for Clique, it might be that there is no such algorithm with runtime $\mathcal{O}(n^{d/2-\varepsilon})$ for KMP, but this does not rule out faster algorithms using, e.g., fast matrix multiplication techniques.

Since no further progress was made for KMP for a long time, research turned to the study of special cases. Over the years, the following special cases have been investigated. For each one we list the asymptotically fastest results.

- **Cube-KMP:** Here the given boxes are cubes, not necessarily all with the same side length. This case can be solved in time $\mathcal{O}(n^{(d+2)/3})$ for $d \geq 2$ [6]. In dimension $d = 3$ this has been improved to $\mathcal{O}(n \log^4 n)$ by Agarwal [1]. In dimensions $d \leq 2$ even the general case can be solved in time $\mathcal{O}(n \log n)$, the same bound clearly applies to this special case. As described in [6], there are simple reductions showing that the case of cubes is roughly the same as the case of “ α -fat boxes”, where all side lengths of a box differ by at most a constant factor α .
- **Unitcube-KMP:** Here the given boxes are cubes, all of the same side length. This is a specialization of CUBE-KMP, so all algorithms from above apply. The combinatorial complexity of a union of unit cubes is $\mathcal{O}(n^{\lfloor d/2 \rfloor})$ [5]. Using this, there are algorithms with runtime $\mathcal{O}(n^{\lfloor d/2 \rfloor} \text{polylog } n)$ [13] and $\mathcal{O}(n^{\lfloor d/2 \rfloor - 1 + \frac{1}{\lfloor d/2 \rfloor}} \text{polylog } n)$ [8]. Again, there is a generalization to “ α -fat boxes of roughly equal size”, and any algorithm for UNITCUBE-KMP can be adapted to an algorithm for this generalization [6].
- **Hypervolume:** Here all boxes have a common vertex. Without loss of generality, we can assume that they share the vertex $(0, \dots, 0) \in \mathbb{R}^d$ and lie in the positive orthant $\mathbb{R}_{\geq 0}^d$. This special case is of particular interest for practice, as it is used as an indicator of the quality of a set of points in the field of Evolutionary Multi-Objective Optimization [3, 12, 19, 20]. Improving upon the general case of KMP, there is an algorithm with runtime $\mathcal{O}(n \log n)$ for $d = 3$ [4]. The same paper also shows an unconditional lower bound of $\Omega(n \log n)$ for $d > 1$, while $\#P$ -hardness in the number of dimensions was shown in [7]. Recently, an algorithm with runtime $\mathcal{O}(n^{(d-1)/2} \log n)$ for $d \geq 3$ was presented in [18].
- **k -Grounded:** Here the projection of the input boxes to the first k dimensions is a HYPERVOLUME instance, where $0 \leq k \leq d$, the other coordinates are arbitrary. This rather novel special case appeared in [18], where an algorithm with runtime $\mathcal{O}(n^{(d-1)/2} \log^2 n)$ for $d \geq 3$ was given for 2-GROUNDED.

Note that for none of these special cases W[1]-hardness is known, so there is no larger lower bound than $\Omega(n \log n)$ (for constant or slowly growing d), not even under reasonable complexity theoretic assumptions. Also note that there are trivial inclusions of some of these special cases: Each special case can be seen

as a subset of all instances of the general case. As such subsets, the following inclusions hold.

- $\text{UNITCUBE-KMP} \subseteq \text{CUBE-KMP} \subseteq \text{KMP}$.
- $(k + 1)\text{-GROUNDED} \subseteq k\text{-GROUNDED}$ for all k .
- $d\text{-GROUNDED} = \text{HYPERVOLUME}$ and $0\text{-GROUNDED} = \text{KMP}$.

Clearly, for such an inclusion it holds that any algorithm for the superset is also an algorithm for the subset, with the same runtime. Moreover, any lower bound for the subset is also a lower bound for the superset. This allows to transfer some results listed above to other special cases, e.g., the CUBE-KMP algorithm with runtime $\mathcal{O}(n^{(d+2)/3})$ also applies to UNITCUBE-KMP .

1.1 Our results

We present several reductions among the above four special cases and the general case of KMP . They provide bounds on the runtimes needed for these variants and, thus, yield some order among the special cases.

Our first reduction relates HYPERVOLUME and UNITCUBE-KMP .

Theorem 1. *If there is an algorithm for UNITCUBE-KMP with runtime $T_{\text{UNITCUBE-KMP}}(n, d)$, then there is an algorithm for HYPERVOLUME with runtime*

$$T_{\text{HYPERVOLUME}}(n, d) \leq \mathcal{O}(T_{\text{UNITCUBE-KMP}}(n, d)).$$

Note that if HYPERVOLUME were a subset of UNITCUBE-KMP , then the same statement would hold, with the constant hidden by the \mathcal{O} -notation being 1. Hence, this reduction can nearly be seen as an inclusion. Also note that this reduction allows to transfer runtime bounds from UNITCUBE-KMP and CUBE-KMP to HYPERVOLUME , e.g., there is an algorithm for HYPERVOLUME with runtime $\mathcal{O}(n^{(d+2)/3})$ for $d \geq 2$. Moreover, together with the trivial inclusions this reduction establishes HYPERVOLUME as the easiest of all studied special cases. In particular, lower bounds for HYPERVOLUME also holds for all other special cases.

Corollary 1. *For all studied special cases, HYPERVOLUME , UNITCUBE-KMP , CUBE-KMP , and $k\text{-GROUNDED}$ (for any $0 \leq k \leq d$), we have the unconditional lower bound $\Omega(n \log n)$ for any $d > 1$.*

One can find contradicting statements regarding the feasibility of a reduction as in Theorem 1 in the literature. On the one hand, existence of such a reduction has been mentioned in [18]. On the other hand, a newer paper [11] contains this sentence: “Better bounds have been obtained for the KMP on unit cubes ..., but reducing the hypervolume indicator to such problems is not possible in general.” In any case, a proof of such a statement cannot be found anywhere in the literature.

Our second reduction substantiates the intuition that the special cases CUBE-KMP and UNITCUBE-KMP are very similar, by showing that their runtimes

differ by at most a factor of $\mathcal{O}(n)$. Recall that $\text{UNITCUBE-KMP} \subseteq \text{CUBE-KMP}$ was one of the trivial inclusions. Clearly, it implies that if there is an algorithm for CUBE-KMP with runtime $T(n, d)$, then there is an algorithm for UNITCUBE-KMP with the same runtime. We prove an inequality in the other direction in the following theorem.

Theorem 2. *If there is an algorithm for UNITCUBE-KMP with runtime $T_{\text{UNITCUBE-KMP}}(n, d)$, then there is an algorithm for CUBE-KMP with runtime*

$$T_{\text{CUBE-KMP}}(n, d) \leq \mathcal{O}(n \cdot T_{\text{UNITCUBE-KMP}}(n, d)).$$

Our third and last reduction finally allows to show lower bounds for all special cases. We show an inequality between the general case of KMP and $2k\text{-GROUNDED}$, in the opposite direction than the trivial inclusions. For this, we have to increase the dimension in which we consider $2k\text{-GROUNDED}$.

Theorem 3. *If there is an algorithm for $2k\text{-GROUNDED}$ in dimension $d + k$ with runtime $T_{2k\text{-GROUNDED}}(n, d + k)$, then there is an algorithm for KMP in dimension d with runtime*

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2k\text{-GROUNDED}}(n, d + k)).$$

Note that, if we set $k = d$, the special case $2k\text{-GROUNDED}$ in $d + k$ dimensions becomes HYPERVOLUME in $2d$ dimensions. Since we established HYPERVOLUME as the easiest variant, the above reduction allows to transfer $\text{W}[1]$ -hardness from the general case to all special cases. Since the dimension is increased only by a constant factor, even the tight lower bound on the runtime can be transferred to all special cases.

Corollary 2. *There is no $n^{o(d)}$ algorithm for any of the special cases HYPERVOLUME , UNITCUBE-KMP , CUBE-KMP , and $k\text{-GROUNDED}$, assuming the Exponential Time Hypothesis.*

We get more precise lower bounds if we assume that no improved algorithm exists for KMP (no algorithm with runtime $\mathcal{O}(n^{d/2-\varepsilon})$). Since any algorithm for HYPERVOLUME in dimension d yields an algorithm for KMP in dimension $\lfloor d/2 \rfloor$ by Theorem 3, and since HYPERVOLUME is the easiest special case, we get the following corollary.

Corollary 3. *If there is no improved algorithm for KMP , then there is no algorithm with runtime $\mathcal{O}(n^{\lfloor d/2 \rfloor / 2 - \varepsilon})$ for any of HYPERVOLUME , UNITCUBE-KMP , CUBE-KMP , and $k\text{-GROUNDED}$, for any $\varepsilon > 0$.*

This shows a lower bound for all studied special cases. Note that there is, however, a wide gap to the best known upper bound of $\mathcal{O}(n^{(d+2)/3})$ for HYPERVOLUME , UNITCUBE-KMP , and CUBE-KMP . Also note that, other than the unconditional lower bound of $\Omega(n \log n)$ for HYPERVOLUME , there were no lower bounds for any of these special cases before.

Furthermore, setting $k = 1$, Theorem 3 immediately implies that the recent algorithm for 2-GROUNDED with runtime $\mathcal{O}(n^{(d-1)/2} \log^2 n)$ [18] is optimal (apart from logarithmic factors and if there is no improved algorithm for KMP).

Corollary 4. *If there is no improved algorithm for KMP, then there is no algorithm for 2-GROUNDED with runtime $\mathcal{O}(n^{(d-1)/2-\varepsilon})$ for any $\varepsilon > 0$.*

An alternative way of stating this result would be that any algorithm faster than the one from [18] for 2-GROUNDED would yield an improved algorithm for KMP.

To simplify our runtime bounds, in some proofs we use the following technical lemma. Informally, it states that for any k -GROUNDED algorithm with runtime $T(n, d)$ we have $T(\mathcal{O}(n), d) \leq \mathcal{O}(T(n, d))$. Intuitively, this should hold, since any reasonable algorithm for this problem has runtime $T(n, d) = \mathcal{O}(n^d)$, i.e., bounded by a polynomial in n , and for polynomial T this simplification is valid. Note that in this paper we hide by the \mathcal{O} -notation any functions depending solely on d .

Lemma 1. *Fix $0 \leq k \leq d$ and $c > 1$. If there is an algorithm for k -GROUNDED with runtime $T_{k\text{-GROUNDED}}(n, d)$ then there is another algorithm for k -GROUNDED with runtime $T'_{k\text{-GROUNDED}}(n, d)$ satisfying*

$$T'_{k\text{-GROUNDED}}(cn, d) \leq \mathcal{O}(T_{k\text{-GROUNDED}}(n, d)).$$

Due to space constraints, the proofs of this and other statements can be found in the appendix, for many results we only provide proof sketches.

1.2 Notation and Organization

A *box* is a set of the form $B = [a_1, b_1] \times \dots \times [a_d, b_d] \subset \mathbb{R}^d$, $a_i, b_i \in \mathbb{R}$, $a_i \leq b_i$. A *cube* is a box with all side lengths equal, i.e., $|b_1 - a_1| = \dots = |b_d - a_d|$. Moreover, a KMP instance is simply a set M of n boxes. In CUBE-KMP all these boxes are cubes, and in UNITCUBE-KMP all these boxes are cubes of common side length. In HYPERVOLUME, all input boxes share the vertex $(0, \dots, 0) \in \mathbb{R}^d$, i.e., each input box is of the form $B = [0, b_1] \times \dots \times [0, b_d]$. In k -GROUNDED, the projection of each input box to the first k dimensions is a HYPERVOLUME box, meaning that each input box is of the form $B = [a_1, b_1] \times \dots \times [a_d, b_d]$ with $a_1 = \dots = a_k = 0$.

We write the usual Lebesgue measure of a set $A \subseteq \mathbb{R}^d$ as $\text{VOL}(A)$. For sets $R, A \subseteq \mathbb{R}^d$ we write $\text{VOL}_R(A) := \text{VOL}(R \cap A)$, the volume of A restricted to R . For a KMP instance M we let $\mathcal{U}(M) := \bigcup_{B \in M} B$. To shorten notation we write $\text{VOL}(M) := \text{VOL}(\mathcal{U}(M))$ and $\text{VOL}_R(M) := \text{VOL}(R \cap \mathcal{U}(M))$.

In the next section we present the proof of Theorem 1. In Section 3 we prove Theorem 2. The proof of Theorem 3 is split into Section 4 and Section 5: We first give the reduction for 2-GROUNDED (again split into the case $d = 1$ and a generalization to larger dimensions) and then generalize this result to $2k$ -GROUNDED, $k > 1$. We close with an extensive list of open problems.

2 Hypervolume \leq Unitcube-KMP

In this section we provide a short sketch of a proof of Theorem 1 by giving a reduction from HYPERVOLUME to UNITCUBE-KMP.

Given an instance of HYPERVOLUME, let Δ be the largest coordinate of any box. We extend all boxes to cubes of side length Δ , yielding a UNITCUBE-KMP instance. In this process, we make sure that the new parts of each box will not lie in the positive orthant $\mathbb{R}_{\geq 0}^d$, but in the other orthants, as depicted in Figure 1. This means that the volume of the newly constructed cubes - restricted to $\mathbb{R}_{\geq 0}^d$ - is the same as the volume of the input boxes. To compute this restricted volume, we compute the volume of the constructed UNITCUBE-KMP instance once with and once without an additional cube $C = [0, \Delta]^d$. From this we can infer the volume of the input HYPERVOLUME instance.




Fig. 1: Construction in the proof of Theorem 1.

3 Unitcube-KMP \geq Cube-KMP

In this section we provide a short sketch of the proof of Theorem 2 by giving a reduction from CUBE-KMP to UNITCUBE-KMP.

Given a CUBE-KMP instance, let C be the cube with smallest side length. We will compute the *contribution* v of C , i.e., the volume of space that is contained in C but no other cube. Having this, we can delete C and recurse on the remaining boxes. Adding up yields the total volume of the input instance.

To compute v , we modify each cube such that it becomes a cube of C 's sidelength and its restriction to C stays the same, as depicted in Figure 2. Applying this construction to all input boxes, we get a UNITCUBE-KMP instance that, inside C , looks the same as the input CUBE-KMP instance. Computing the volume of this new instance once with and once without C allows to infer v .




Fig. 2: Construction in the proof of Theorem 2.

4 2-Grounded \geq KMP

We first show the reduction of Theorem 3 for 2-GROUNDED, i.e., we show $T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2\text{-GROUNDED}}(n, d + 1))$ by giving a reduction from KMP to 2-GROUNDED. This already implies Corollary 4 and lays the foundations for the complete reduction given in the next section.

We begin by showing the reduction for $d = 1$. As a second step we show how to generalize this to larger dimensions.

4.1 Dimension $d = 1$

We want to give a reduction from KMP in 1 dimension to 2-GROUNDED in 2 dimensions. Note that the latter is the same as HYPERVOLUME in 2 dimensions. Let M be an instance of KMP in 1 dimension, i.e., a set of n intervals in \mathbb{R} . We will reduce the computation of $\text{VOL}(M)$ to two instances of 2-GROUNDED.

Denote by $x_1 < \dots < x_m$ the endpoints of all intervals in M (if all endpoints are distinct then $m = 2n$). We can assume that $x_1 = 0$ after translation. Consider the boxes

$$A_i := [m - i - 1, m - i] \times [x_i, x_{i+1}]$$

in \mathbb{R}^2 for $1 \leq i \leq m - 1$, as depicted in Figure 3. Denote the union of these boxes by A . Note that the volume of box A_i is the same as the length of the interval $[x_i, x_{i+1}]$. This means that we took the chain of intervals $\{[x_i, x_{i+1}]\}$ and made it into a staircase of boxes $\{A_i\}$, where each box has the same volume as the corresponding interval.




Fig. 3: The left hand side depicts all endpoints $0 = x_1 \leq \dots \leq x_6$ of a 1-dimensional KMP instance. An input interval I is indicated. The right hand side shows the result of our transformation. Each interval $[x_i, x_{i+1}]$ to the left corresponds to a box A_i to the right. The interval I gets mapped to the box C_I . The shaded regions depict the set A (■, the union of all A_i) and the set T_0 (◻).

Now consider an interval $I = [x_j, x_k] \in M$. We construct the box

$$C_I := [0, m - j] \times [0, x_k],$$

also shown in Figure 3. Then C_I contains the boxes A_i with $j \leq i < k$ and (its interior) has no common intersection with any other box A_i . This is easily seen as $A_i \subseteq C_I$ iff $m - i \leq m - j$ and $x_{i+1} \leq x_k$. Hence, for any interval $I \in M$ we constructed a box C_I that contains exactly those boxes A_i whose corresponding interval $[x_i, x_{i+1}]$ is contained in I , or in other words

$$\begin{aligned} [x_i, x_{i+1}] \subseteq I &\Leftrightarrow A_i \subseteq C_I, \\ \text{VOL}([x_i, x_{i+1}] \cap I) &= \text{VOL}(A_i \cap C_I). \end{aligned}$$

From these properties it follows that the volume of C_I restricted to A is the same as the length of I , i.e.,

$$\text{VOL}_A(C_I) = \text{VOL}(I).$$

Furthermore, considering the whole set M of intervals, the interval $[x_i, x_{i+1}]$ is contained in some interval in M iff the box A_i is contained in some box in $C_M := \{C_I \mid I \in M\}$. This yields

$$\text{VOL}(M) = \text{VOL}_A(C_M).$$

It remains to reduce the computation of $\text{VOL}_A(C_M)$ to two 2-GROUNDED instances. For this we consider

$$T_0 := \bigcup_{1 \leq i \leq m} C_{[x_i, x_i]}.$$

Informally speaking, T_0 consists of all points “below” A , as depicted in Figure 3. Note that no set A_j is contained in T_0 . Moreover, we consider the set $T_1 := T_0 \cup A$. Observe that we can write

$$T_1 = \bigcup_{1 \leq i \leq m-1} C_{[x_i, x_{i+1}]},$$

since $A_i \subseteq C_{[x_i, x_{i+1}]}$. Note that both sets T_0 and T_1 are unions of $\mathcal{O}(n)$ 2-GROUNDED boxes. Informally, T_0 is the maximum 2-GROUNDED instance that has $\text{VOL}_A(T_0) = 0$, and T_1 is the minimum 2-GROUNDED instance with $\text{VOL}_A(T_1) = \text{VOL}(A)$. Now, we can compute $\text{VOL}_A(C_M)$ as follows.

Lemma 2. *In the above situation we have*

$$\text{VOL}_A(C_M) = \text{VOL}(A) + \text{VOL}(T_0 \cup \mathcal{U}(C_M)) - \text{VOL}(T_1 \cup \mathcal{U}(C_M)).$$

Proof. Set $U := \mathcal{U}(C_M)$. Using $T_0 \subseteq T_1$ and $A = T_1 \setminus T_0$ in a sequence of simple transformations, we get

$$\begin{aligned} \text{VOL}(T_1 \cup U) - \text{VOL}(T_0 \cup U) &= \text{VOL}((T_1 \cup U) \setminus (T_0 \cup U)) \\ &= \text{VOL}((T_1 \setminus T_0) \setminus U) \\ &= \text{VOL}(A \setminus U) \\ &= \text{VOL}(A) - \text{VOL}(A \cap U) \\ &= \text{VOL}(A) - \text{VOL}_A(C_M), \end{aligned}$$

which proves the claim. □

Note that $\text{VOL}(A) = \sum_i \text{VOL}(A_i) = \sum_i |x_{i+1} - x_i| = |x_m - x_1|$ is trivial. Also note that both sets T_0 and T_1 are the union of $\mathcal{O}(n)$ 2-GROUNDED boxes, so that $\text{VOL}(T_b \cup \mathcal{U}(C_M))$ can be seen as a 2-GROUNDED instance of size $\mathcal{O}(n)$, for both $b \in \{0, 1\}$. Hence, we reduced the computation of the input instance's volume $\text{VOL}(M)$ to $\text{VOL}_A(C_M)$ and further to the 2-GROUNDED instances $\text{VOL}(T_0 \cup \mathcal{U}(C_M))$ and $\text{VOL}(T_1 \cup \mathcal{U}(C_M))$.

As we have to sort the given intervals first, we get

$$T_{\text{KMP}}(n, 1) \leq \mathcal{O}(T_{2\text{-GROUNDED}}(\mathcal{O}(n), 2) + n \log n).$$

Note that this inequality alone gives no new information, as already Klee [14] showed that $T_{\text{KMP}}(n, 1) \leq \mathcal{O}(n \log n)$. However, we get interesting results when we generalize this reduction to higher dimensions in the next section.

4.2 Larger Dimensions

In this section we sketch that the reduction from the last section easily carries over to larger dimensions, yielding a reduction from KMP in d dimensions to 2-GROUNDED in $d+1$ dimensions. This implies $T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2\text{-GROUNDED}}(n, d+1))$.

Assume we are given a KMP instance M in dimension d . The idea is that we use the dimension doubling reduction from the last section on the first dimension and leave all other dimensions untouched. More precisely, for a box $B \in M$ let $\pi_1(B)$ be its projection onto the first dimension and let $\pi_*(B)$ be its projection onto the last $d-1$ dimensions, so that $B = \pi_1(B) \times \pi_*(B)$. Now follow the reduction from the last section on the instance $M' := \{\pi_1(B) \mid B \in M\}$. This yields sets A, T_0, T_1 , and a box C_I for each $I \in M'$.

We set $C_B := C_{\pi_1(B)} \times \pi_*(B)$ and $C_M = \{C_B \mid B \in M\}$. A possible way of generalizing A would be to set $A'' := A \times \mathbb{R}^{d-1}$. Then we would be interested in $\text{VOL}_{A''}(C_M)$, which can be seen to be exactly $\text{VOL}(M)$. This definition of A'' is, however, not simple enough, as it is not a difference of 2-GROUNDED instances (unlike $A = T_1 \setminus T_0$). To give a different definition, assume (after translation) that all coordinates of the input instance are nonnegative and let Δ be the maximal coordinate in any dimension. We set $A' := A \times [0, \Delta]^{d-1}$ and still get the same volume $\text{VOL}_{A'}(C_M) = \text{VOL}(M)$. This allows to generalize T_0 and T_1 to $T'_0 := T_0 \times [0, \Delta]^{d-1}$ and $T'_1 := T_1 \times [0, \Delta]^{d-1}$, while still having

$$\text{VOL}_{A'}(C_M) = \text{VOL}(A') + \text{VOL}(T'_0 \cup \mathcal{U}(C_M)) - \text{VOL}(T'_1 \cup \mathcal{U}(C_M)).$$

Note that T'_0 and T'_1 are also a union of $\mathcal{O}(n)$ 2-GROUNDED boxes, so a volume such as $\text{VOL}(T'_0 \cup \mathcal{U}(C_M))$ can be seen as a 2-GROUNDED instance. This completes the reduction and yields the time bound

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2\text{-GROUNDED}}(\mathcal{O}(n), d+1) + n \log n).$$

Using the lower bound $\Omega(n \log n)$ of Corollary 1 we can hide the additional $n \log n$ in the first summand. Moreover, first using the technical Lemma 1 we can finally simplify this to the statement of Corollary 4,

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2\text{-GROUNDED}}(n, d+1)).$$

5 2k-Grounded \geq KMP

It is left to show the full version of Theorem 3, i.e., to give a reduction from KMP in dimension d to $2k$ -GROUNDED in dimension $d+k$. A full proof of this can be found in the full version, here we only present an outline.

The first steps of generalizing the reduction from the last section to the case $k > 1$ are straight forward. We want to use the dimension doubling reduction from Section 4.1 on each one of the first k dimensions. For any box $B \in \mathbb{R}^d$ denote its projection onto the i -th dimension by $\pi_i(B)$, $1 \leq i \leq k$, and its projection onto dimensions $k+1, \dots, d$ by $\pi_*(B)$. We use the reduction from Section 4.1 on each dimension $1 \leq i \leq k$, i.e., on each instance $M^{(i)} := \{\pi_i(B) \mid B \in M\}$, yielding sets $A^{(i)}, T_0^{(i)}, T_1^{(i)}$, and a box $C_I^{(i)}$ for each $I \in M^{(i)}$. For a box $B \in M$ we now define $C_B := C_{\pi_1(B)}^{(1)} \times \dots \times C_{\pi_k(B)}^{(k)} \times \pi_*(B)$.

This is a box in \mathbb{R}^{d+k} , it is even a $2k$ -GROUNDED box, as its projection onto the first $2k$ coordinates has the vertex $(0, \dots, 0)$. Let $C_M := \{C_B \mid B \in M\}$. Setting $A := A^{(1)} \times \dots \times A^{(k)} \times [0, \Delta]^{d-k}$, we now can show the following.

Lemma 3. *We have $\text{VOL}(M) = \text{VOL}_A(C_M)$.*

The hard part of the reduction that remains to show is that the right hand side of this can indeed be computed using $2k$ -GROUNDED calls, although A is a non-trivial region.

For $1 \leq i \leq k$ and $b \in \{0, 1\}$ we set $\tilde{T}_b^{(i)} := \Omega^{2(i-1)} \times T_b^{(i)} \times \Omega^{d+k-2i}$. This set in Ω^{d+k} consists of all points x whose projection to dimensions $2i-1$ and $2i$ is contained in $T_b^{(i)}$. Note that each set $\tilde{T}_b^{(i)}$ can be written as the union of $\mathcal{O}(n)$ $2k$ -GROUNDED boxes, since $T_b^{(i)}$ is the union of $\ell = \mathcal{O}(n)$ 2-GROUNDED boxes in \mathbb{R}^2 , i.e., $T_b^{(i)} = \bigcup_{j=1}^{\ell} C_j$, so that we may write $\tilde{T}_b^{(i)} = \bigcup_{j=1}^{\ell} \Omega^{2(i-1)} \times C_j \times \Omega^{d-k-2i}$. Thus, we can use an algorithm for $2k$ -GROUNDED to compute any volume of the form $\text{VOL}(\tilde{T}_b^{(i)} \cup V)$, where V is a union of $\mathcal{O}(n)$ $2k$ -GROUNDED boxes.

Furthermore, define for $S \subseteq [k]$

$$D_S := \left(\bigcup_{i \in S} \tilde{T}_1^{(i)} \right) \cup \bigcup_{i \in [k] \setminus S} \tilde{T}_0^{(i)}.$$

Note that $D_S \subseteq D_{S'}$ holds for $S \subseteq S'$. We can express A using the sets D_S as shown by the following lemma.

Lemma 4. *We have $A = \bigcap_{1 \leq i \leq k} D_{\{i\}} \setminus D_{\emptyset}$.*

Moreover, each D_S can be written as the union of $\mathcal{O}(n)$ $2k$ -GROUNDED instances, since the same was true for the sets $\tilde{T}_b^{(i)}$. Hence, we can use an algorithm for $2k$ -GROUNDED to compute the volume

$$H_S := \text{VOL}(D_S \cup \mathcal{U}(C_M)).$$

Finally, we show that we can compute $\text{VOL}_A(C_M)$ from the H_S by an interesting usage of the inclusion-exclusion principle, finishing the reduction.

Lemma 5. *We have $\text{VOL}_A(C_M) = \text{VOL}(A) + \sum_{S \subseteq [k]} (-1)^{|S|} H_S$.*

6 Conclusion

We presented reductions between the special cases CUBE-KMP, UNITCUBE-KMP, HYPERVOLUME, and k -GROUNDED of Klee’s measure problem. These reductions imply statements about the runtime needed for these problem variants. We established HYPERVOLUME as the easiest among all studied special cases, and showed that the variants CUBE-KMP and UNITCUBE-KMP have polynomially related runtimes. Moreover, we presented a reduction from the general case of KMP to $2k$ -GROUNDED. This allows to transfer $W[1]$ -hardness from KMP to all special cases, proving that no $n^{o(d)}$ algorithm exists for any of the special cases assuming the Exponential Time Hypothesis. Moreover, assuming that no improved algorithm exists for KMP, we get a tight lower bound for a recent algorithm for 2-GROUNDED, and a lower bound of roughly $n^{(d-1)/4}$ for all other special cases. Thus, we established some order among the special cases of Klee’s measure problem.

Our results lead to a number of open problems, both asking for new upper and lower bounds:

- Is there a polynomial relation between HYPERVOLUME and UNITCUBE-KMP, similar to CUBE-KMP and UNITCUBE-KMP, or do both problems have significantly different runtimes?
- Show that no improved algorithm exists for KMP, e.g., assuming the Strong Exponential Time Hypothesis, as has been done for the Dominating Set problem, see [15]. Or give an improved algorithm.
- Assuming that no improved algorithm for KMP exists, we know that the optimal runtimes of HYPERVOLUME and CUBE-KMP/UNITCUBE-KMP are of the form $n^{c_d \cdot d \pm \mathcal{O}(1)}$, with $c_d \in [1/4, 1/3]$. Determine the correct value of c_d .
- Generalize the $\mathcal{O}(n^{(d-1)/2} \log^2 n)$ algorithm for 2-GROUNDED [18] to an $\mathcal{O}(n^{(d-k)/2 + o(1)})$ algorithm for $2k$ -GROUNDED. This would again be optimal by Theorem 3.
- We showed the relation $T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2k\text{-GROUNDED}}(n, d + k))$. Show an inequality in the opposite direction, i.e., a statement of the form $T_{k\text{-GROUNDED}}(n, d) \leq \mathcal{O}(T_{\text{KMP}}(n, d'))$ with $d' < d$.

Bibliography

- [1] P. K. Agarwal. An improved algorithm for computing the volume of the union of cubes. In *Proc. 2010 Annual Symposium on Computational Geometry (SoCG'10)*, pp. 230–239, New York, NY, USA, 2010. ACM.
- [2] J. L. Bentley. Algorithms for Klee’s rectangle problems, 1977. Department of Computer Science, Carnegie Mellon University, Unpublished notes.
- [3] N. Beume, B. Naujoks, and M. T. M. Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181:1653–1669, 2007.

- [4] N. Beume, C. M. Fonseca, M. López-Ibáñez, L. Paquete, and J. Vahrenhold. On the complexity of computing the hypervolume indicator. *IEEE Trans. Evolutionary Computation*, 13:1075–1082, 2009.
- [5] J.-D. Boissonnat, M. Sharir, B. Tagansky, and M. Yvinec. Voronoi diagrams in higher dimensions under certain polyhedral distance functions. *Discrete & Computational Geometry*, 19:485–519, 1998.
- [6] K. Bringmann. An improved algorithm for Klee’s measure problem on fat boxes. *Computational Geometry: Theory and Applications*, 45:225–233, 2012.
- [7] K. Bringmann and T. Friedrich. Approximating the least hypervolume contributor: NP-hard in general, but fast in practice. *Theoretical Computer Science*, 425:104–116, 2012.
- [8] T. M. Chan. Semi-online maintenance of geometric optima and measures. *SIAM J. Comput.*, 32:700–716, 2003.
- [9] T. M. Chan. A (slightly) faster algorithm for Klee’s measure problem. *Computational Geometry: Theory and Applications*, 43:243–250, 2010.
- [10] M. L. Fredman and B. W. Weide. On the complexity of computing the measure of $\bigcup[a_i, b_i]$. *Commun. ACM*, 21:540–544, 1978.
- [11] A. P. Guerreiro, C. M. Fonseca, and M. T. M. Emmerich. A fast dimension-sweep algorithm for the hypervolume indicator in four dimensions. In *Proc. 24th Canadian Conference on Computational Geometry (CCCG’12)*, pp. 77–82, 2012.
- [12] C. Igel, N. Hansen, and S. Roth. Covariance matrix adaptation for multi-objective optimization. *Evolutionary Computation*, 15:1–28, 2007.
- [13] H. Kaplan, N. Rubinfeld, M. Sharir, and E. Verbin. Counting colors in boxes. In *Proc. 18th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA’07)*, pp. 785–794, 2007.
- [14] V. Klee. Can the measure of $\bigcup[a_i, b_i]$ be computed in less than $O(n \log n)$ steps? *American Mathematical Monthly*, 84:284–285, 1977.
- [15] D. Lokshtanov, D. Marx, and S. Saurabh. Lower bounds based on the exponential time hypothesis, The Complexity Column by V. Arvind. *Bulletin of the EATCS*, 105:41–72, 2011.
- [16] M. H. Overmars and C.-K. Yap. New upper bounds in Klee’s measure problem. *SIAM J. Comput.*, 20:1034–1045, 1991.
- [17] J. van Leeuwen and D. Wood. The measure problem for rectangular ranges in d -space. *J. Algorithms*, 2:282–300, 1981.
- [18] H. Yıldız and S. Suri. On Klee’s measure problem for grounded boxes. In *Proc. 2012 Annual Symposium on Computational Geometry (SoCG’12)*, pp. 111–120, New York, NY, USA, 2012. ACM.
- [19] E. Zitzler and S. Künzli. Indicator-based selection in multiobjective search. In *Proc. 8th International Conference on Parallel Problem Solving from Nature (PPSN VIII)*, Vol. 3242 of LNCS, pp. 832–842. Springer, 2004.
- [20] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Trans. Evolutionary Computation*, 3:257–271, 1999.

A 2k-Grounded \geq KMP

While Section 5 contains an outline of the proof of Theorem 3, in this section we present the details of this proof. Recall that we want to give a reduction from KMP in dimension d to $2k$ -GROUNDED in dimension $d + k$.

Let M be a KMP instance. After translation we can assume that in every dimension the minimal coordinate among all boxes in M is 0. Denoting the largest coordinate of any box by Δ we thus have $B \subseteq [0, \Delta]^d = \Omega^d$ for all $B \in M$, where $\Omega := [0, \Delta]$.

The first steps of generalizing the reduction from the last section to the case $k > 1$ are straight forward. We want to use the dimension doubling reduction from Section 4.1 on each one of the first k dimensions. For any box $B \in \mathbb{R}^d$ denote its projection onto the i -th dimension by $\pi_i(B)$, $1 \leq i \leq k$, and its projection onto dimensions $k + 1, \dots, d$ by $\pi_*(B)$. We use the reduction from Section 4.1 on each dimension $1 \leq i \leq k$, i.e., on each instance $M^{(i)} := \{\pi_i(B) \mid B \in M\}$, yielding sets $A^{(i)}, T_0^{(i)}, T_1^{(i)}$, and a box $C_I^{(i)}$ for each $I \in M^{(i)}$. We assume that all these sets are contained in $[0, \Delta]^2 = \Omega^2$, meaning that all coordinates are upper bounded by Δ (this holds after possibly increasing the Δ we chose before).

For a box $B \in M$ we now define

$$C_B := C_{\pi_1(B)}^{(1)} \times \dots \times C_{\pi_k(B)}^{(k)} \times \pi_*(B).$$

This is a box in \mathbb{R}^{d+k} , it is even a $2k$ -GROUNDED box, as its projection onto the first $2k$ coordinates has the vertex $(0, \dots, 0)$. Set $C_M := \{C_B \mid B \in M\}$. As shown by the following lemma, we want to determine the volume $\text{VOL}_A(C_M)$ in

$$A := A^{(1)} \times \dots \times A^{(k)} \times [0, \Delta]^{d-k}.$$

Lemma 3. *In the above situation we have*

$$\text{VOL}(M) = \text{VOL}_A(C_M).$$

Proof. Denote by $x_1^{(i)} < \dots < x_{m_i}^{(i)}$ the coordinates of all boxes in M in the i -th dimension. We can express $\text{VOL}(M)$ in terms of the boxes

$$E_{j_1, \dots, j_d} := [x_{j_1}^{(1)}, x_{j_1+1}^{(1)}] \times \dots \times [x_{j_d}^{(d)}, x_{j_d+1}^{(d)}],$$

for $1 \leq j_i < m_i$. Since each such box is either completely included in some box in M or does not contribute to $\text{VOL}(M)$, we have

$$\text{VOL}(M) = \sum_{j_1, \dots, j_d} [E_{j_1, \dots, j_d} \subseteq \mathcal{U}(M)] \cdot \text{VOL}(E_{j_1, \dots, j_d}),$$

where $[X]$ is 1 if X is true, and 0 otherwise.

Recall from the reduction in Section 4.1 that $A^{(i)} = A_1^{(i)} \cup \dots \cup A_{m_i-1}^{(i)}$, and there is a one-to-one correspondence between intervals $[x_j^{(i)}, x_{j+1}^{(i)}]$ and 2-dimensional boxes $A_j^{(i)}$, in particular both have the same volume. This carries

over to a one-to-one correspondence between d -dimensional boxes E_{j_1, \dots, j_d} and $(d+k)$ -dimensional boxes

$$E'_{j_1, \dots, j_d} := A_{j_1}^{(1)} \times \dots \times A_{j_k}^{(k)} \times [x_{j_{k+1}}^{(k+1)}, x_{j_{k+1}+1}^{(k+1)}] \times \dots \times [x_{j_d}^{(d)}, x_{j_d+1}^{(d)}],$$

in particular both have the same volume.

Additionally, recall that an interval $I \in M^{(i)}$ includes $[x_j^{(i)}, x_{j+1}^{(i)}]$ if and only if the 2-dimensional box $C_I^{(i)}$ contains $A_j^{(i)}$. If I does not include $[x_j^{(i)}, x_{j+1}^{(i)}]$, then $\text{VOL}(I \cap [x_j^{(i)}, x_{j+1}^{(i)}]) = 0$, and we also have $\text{VOL}(C_I^{(i)} \cap A_j^{(i)}) = 0$. Hence, we have for any $B \in M$ that $E_{j_1, \dots, j_d} \subseteq B$ if and only if $E'_{j_1, \dots, j_d} \subseteq C_B$, which implies that $E_{j_1, \dots, j_d} \subseteq \mathcal{U}(M)$ if and only if $E'_{j_1, \dots, j_d} \subseteq \mathcal{U}(C_M)$. Furthermore, E'_{j_1, \dots, j_d} is either included in $\mathcal{U}(C_M)$ or does not contribute to $\text{VOL}(C_M)$. In total, we get

$$\begin{aligned} \text{VOL}(M) &= \sum_{j_1, \dots, j_d} [E_{j_1, \dots, j_d} \subseteq \mathcal{U}(M)] \cdot \text{VOL}(E_{j_1, \dots, j_d}) \\ &= \sum_{j_1, \dots, j_d} [E'_{j_1, \dots, j_d} \subseteq \mathcal{U}(C_M)] \cdot \text{VOL}(E'_{j_1, \dots, j_d}) = \text{VOL}_A(C_M), \end{aligned}$$

which completes the proof. \square

Unfortunately, the set A is not simply a difference of two $2k$ -GROUNDED instances. Thus, the hard part is to reduce the computation of $\text{VOL}_A(C_M)$ to $2k$ -GROUNDED instances, which we will do in the remainder of this section. For $1 \leq i \leq k$ and $b \in \{0, 1\}$ we set

$$\tilde{T}_b^{(i)} := \Omega^{2(i-1)} \times T_b^{(i)} \times \Omega^{d+k-2i}.$$

This set in Ω^{d+k} consists of all points x whose projection to dimensions $2i-1$ and $2i$ is contained in $T_b^{(i)}$. Note that each set $\tilde{T}_b^{(i)}$ can be written as the union of $\mathcal{O}(n)$ $2k$ -GROUNDED boxes, since $T_b^{(i)}$ is the union of $\ell = \mathcal{O}(n)$ 2-GROUNDED boxes in \mathbb{R}^2 , i.e., $T_b^{(i)} = \bigcup_{j=1}^{\ell} C_j$, so that we may write $\tilde{T}_b^{(i)} = \bigcup_{j=1}^{\ell} \Omega^{2(i-1)} \times C_j \times \Omega^{d-k-2i}$. Thus, we can use an algorithm for $2k$ -GROUNDED to compute any volume of the form $\text{VOL}(\tilde{T}_b^{(i)} \cup V)$, where V is a union of $\mathcal{O}(n)$ $2k$ -GROUNDED boxes.

Furthermore, define for $S \subseteq [k]$

$$D_S := \left(\bigcup_{i \in S} \tilde{T}_1^{(i)} \right) \cup \bigcup_{i \in [k] \setminus S} \tilde{T}_0^{(i)}.$$

Note that $D_S \subseteq D_{S'}$ holds for $S \subseteq S'$. We can express A using the sets D_S as shown by the following lemma.

Lemma 4. *In the above situation we have*

$$A = \bigcap_{1 \leq i \leq k} D_{\{i\}} \setminus D_{\emptyset}.$$

Proof. We have $D_\emptyset = \bigcup_{i \in [k]} \tilde{T}_0^{(i)}$ and $\tilde{T}_0^{(i)} \subseteq \tilde{T}_1^{(i)}$ for all i , implying

$$D_{\{i\}} = \tilde{T}_1^{(i)} \cup D_\emptyset.$$

This yields

$$\bigcap_{1 \leq i \leq k} D_{\{i\}} \setminus D_\emptyset = \left(\bigcap_{1 \leq i \leq k} \tilde{T}_1^{(i)} \right) \setminus \bigcup_{i \in [k]} \tilde{T}_0^{(i)}.$$

A point $x = (x_1, \dots, x_{d+k}) \in \mathbb{R}^{d+k}$ is in the set on the right hand side if and only if it has the following three properties:

- $x_i \in [0, \Delta]$ for all $1 \leq i \leq d+k$,
- (x_{2i-1}, x_{2i}) lies in $T_1^{(i)}$ for all $1 \leq i \leq k$,
- (x_{2i-1}, x_{2i}) does not lie in $T_0^{(i)}$ for all $1 \leq i \leq k$.

Since $A^{(i)} = T_1^{(i)} \setminus T_0^{(i)}$, this description captures exactly $A^{(1)} \times \dots \times A^{(k)} \times [0, \Delta]^{d-k} = A$, finishing the proof. \square

Moreover, each D_S can be written as the union of $\mathcal{O}(n)$ $2k$ -GROUNDED instances, since the same was true for the sets $\tilde{T}_b^{(i)}$. Hence, we can use an algorithm for $2k$ -GROUNDED to compute the volume

$$H_S := \text{VOL}(D_S \cup \mathcal{U}(C_M)).$$

Next we show that we can compute $\text{VOL}_A(C_M)$ from the H_S by an interesting usage of the inclusion-exclusion principle.

Lemma 5. *In the above situation we have*

$$\text{VOL}_A(C_M) = \text{VOL}(A) + \sum_{S \subseteq [k]} (-1)^{|S|} H_S.$$

Proof. In this proof we write for short $U := \mathcal{U}(C_M)$. We clearly have

$$\text{VOL}(A) - \text{VOL}_A(U) = \text{VOL}(A \setminus U). \quad (1)$$

We first show

$$\text{VOL}(A \setminus U) = \sum_{\emptyset \neq S \subseteq [k]} (-1)^{|S|+1} (H_S - H_\emptyset), \quad (2)$$

and simplify the right hand side later. Using $A = \bigcap_{1 \leq i \leq k} D_{\{i\}} \setminus D_\emptyset$ (Lemma 4) and the inclusion-exclusion principle we arrive at

$$\begin{aligned} \text{VOL}(A \setminus U) &= \text{VOL} \left(\bigcap_{1 \leq i \leq k} D_{\{i\}} \setminus (D_\emptyset \cup U) \right) \\ &= \sum_{\emptyset \neq S \subseteq [k]} (-1)^{|S|+1} \text{VOL} \left(\bigcup_{i \in S} D_{\{i\}} \setminus (D_\emptyset \cup U) \right). \end{aligned}$$

Note that $D_S = \bigcup_{i \in S} D_{\{i\}}$, so that the above equation simplifies to

$$\text{VOL}(A \setminus U) = \sum_{\emptyset \neq S \subseteq [k]} (-1)^{|S|+1} \text{VOL}\left(D_S \setminus (D_\emptyset \cup U)\right) \quad (3)$$

Using the definition of H_S and $D_\emptyset \subseteq D_S$ for any $S \subseteq [k]$, we get

$$\begin{aligned} H_S - H_\emptyset &= \text{VOL}(D_S \cup U) - \text{VOL}(D_\emptyset \cup U) \\ &= \text{VOL}((D_S \cup U) \setminus (D_\emptyset \cup U)) \\ &= \text{VOL}(D_S \setminus (D_\emptyset \cup U)). \end{aligned}$$

Plugging this into (3) yields (2).

Observe that

$$\sum_{\emptyset \neq S \subseteq [k]} (-1)^{|S|+1} (-H_\emptyset) = -H_\emptyset.$$

This allows to further simplify (2) to

$$\text{VOL}(A \setminus U) = \sum_{S \subseteq [k]} (-1)^{|S|+1} H_S.$$

Plugging this into (1) yields the desired equation. \square

As $\text{VOL}(A) = \Delta^{d-k} \cdot \prod_{1 \leq i \leq k} \text{VOL}(A^{(i)})$ is trivial, we have reduced the computation of $\text{VOL}_A(C_M)$ to $2^k = \mathcal{O}(1)$ instances of $2k$ -GROUNDED, each consisting of $\mathcal{O}(n)$ boxes. During the construction of these instances we need to sort the coordinates, so that we need additional time $\mathcal{O}(n \log n)$. This yields

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2k\text{-GROUNDED}}(\mathcal{O}(n), d+k) + n \log n).$$

Because of the lower bound from Corollary 1, we have $T_{2k\text{-GROUNDED}}(\mathcal{O}(n), d+k) = \Omega(n \log n)$, so we can hide the second summand in the first,

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T_{2k\text{-GROUNDED}}(\mathcal{O}(n), d+k)).$$

We may use the technical Lemma 1 to get rid of the inner \mathcal{O} : This lemma guarantees an algorithm with runtime $T'_{2k\text{-GROUNDED}}(n, d)$ such that we get

$$T_{\text{KMP}}(n, d) \leq \mathcal{O}(T'_{2k\text{-GROUNDED}}(\mathcal{O}(n), d+k)) \leq \mathcal{O}(T_{2k\text{-GROUNDED}}(n, d+k)).$$

This finishes the proof.

B Proof of Lemma 1

In this section we prove the technical Lemma 1.

Lemma 1. Fix $0 \leq k \leq d$ and $c > 1$. If there is an algorithm for k -GROUNDED with runtime $T_{k\text{-GROUNDED}}(n, d)$ then there is another algorithm for k -GROUNDED with runtime $T'_{k\text{-GROUNDED}}(n, d)$ satisfying

$$T'_{k\text{-GROUNDED}}(cn, d) \leq \mathcal{O}(T_{k\text{-GROUNDED}}(n, d)).$$

Proof. Let M be an instance of k -GROUNDED of size $|M| = n$. Denote by $z_1 \leq \dots \leq z_{2n}$ the coordinates in the first dimension of all boxes in M . Let $a_1 := z_{(1-\alpha)n}$ and $b_1 := z_{(1+\alpha)n}$, where $\alpha := 1/(3d)$. Denote by $\{x_1 \leq a_1\}$ the set of all points $x = (x_1, \dots, x_d) \in \mathbb{R}^d$ with $x_1 \leq a_1$, similarly for $\{a_1 \leq x_1 \leq b_1\}$ and $\{x_1 \geq b_1\}$. We consider the three KMP instances

$$\begin{aligned} M_a &:= \{B \cap \{x_1 \leq a_1\} \mid B \in M\}, \\ M_b &:= \{B \cap \{x_1 \geq b_1\} \mid B \in M\}, \\ M' &:= \{B \cap \{a_1 \leq x_1 \leq b_1\} \mid B \in M\}. \end{aligned}$$

Note that all three instances can be seen as k -GROUNDED instances: Projected onto the first k dimensions, all boxes in M_a share the vertex $(0, \dots, 0)$, all boxes in M_b share the vertex $(b_1, 0, \dots, 0)$, and all boxes in M' share the vertex $(a_1, 0, \dots, 0)$, so after translation they all share the vertex $(0, \dots, 0)$. Moreover, $\{x_1 < a_1\}$ and $\{x_1 > b_1\}$ contain only $(1-\alpha)n$ coordinates of boxes in M in the first dimension. Hence, there are at most $(1-\alpha)n$ boxes intersecting $\{x_1 < a_1\}$, so after deleting boxes with volume 0 we get $|M_a| \leq (1-\alpha)n$, similarly for M_b . This reasoning does not work for M' , it might even be that all n boxes are present in M' : If a box has left coordinate smaller than a_1 and right coordinate larger than b_1 , then none of its coordinates is seen in $\{a_1 \leq x_1 \leq b_1\}$, although it has non-empty intersection with $\{a_1 \leq x_1 \leq b_1\}$. However, such a box in M' is trivial in the first dimension: its coordinates in the first dimension are simply $[a_1, b_1]$. If all boxes in M' were trivial in the first dimension, then M' would clearly be simpler than the input instance. Although this is not the case, we can bound the number of boxes in M' that are non-trivial in the first dimension: Since there are at most $2\alpha n$ coordinates z_i in $[a_1, b_1]$, all but at most $2\alpha n$ boxes in M' are trivial in the first dimension. Thus, also M' is easier than the input instance M , in a certain sense.

Note that $\text{VOL}(M) = \text{VOL}(M_a) + \text{VOL}(M_b) + \text{VOL}(M')$ and M_a, M_b are strictly easier than M , as they contain at most $(1-\alpha)n$ boxes. We have to simplify M' further. For this, we use the same construction as above (on M and dimension 1) on M' and dimension 2, i.e., we split by coordinates in dimension 2 at a_2 and b_2 . This yields three k -GROUNDED instances. Two of them contain at most $(1-\alpha)n$ boxes. The third one, M'' , may contain up to n boxes. However, all but at most $4\alpha n$ of these boxes are trivial in the first and second dimension, meaning that their projection onto the first 2 dimensions is $[a_1, b_1] \times [a_2, b_2]$.

Iterating this reduction d times yields $2d$ instances of k -GROUNDED containing at most $(1-\alpha)n$ points and one instance M^* that may contain up to n boxes. However, all but at most $2d\alpha n$ of these boxes are trivial in all d dimensions, meaning that they are equal to $[a_1, b_1] \times \dots \times [a_d, b_d]$. Since

all boxes in M^* are contained in $[a_1, b_1] \times \dots \times [a_d, b_d]$, if any such trivial box exists, the volume of M^* is trivial. Otherwise M^* only contains at most $2d\alpha n = \frac{2}{3}n \leq (1 - 1/(3d))n = (1 - \alpha)n$ boxes. Thus, we have reduced the computation of $\text{VOL}(M)$ to $2d + 1$ instances of k -GROUNDED with at most $(1 - \alpha)n$ boxes each. The reduction itself can be made to run in $\mathcal{O}(n)$ time. Hence, if we solve the reduced problems by an algorithm with runtime $T_{k\text{-GROUNDED}}(n, d)$, then we get an algorithm with runtime $T'_{k\text{-GROUNDED}}(n, d)$ satisfying

$$T'_{k\text{-GROUNDED}}(n, d) \leq (2d + 1)T_{k\text{-GROUNDED}}((1 - \alpha)n, d) + \mathcal{O}(n).$$

As every algorithm for k -GROUNDED at least has to read its whole input, we can hide the $\mathcal{O}(n)$ by the first term,

$$T'_{k\text{-GROUNDED}}(n, d) \leq \mathcal{O}(T_{k\text{-GROUNDED}}((1 - \alpha)n, d)),$$

or,

$$T'_{k\text{-GROUNDED}}(n/(1 - \alpha), d) \leq \mathcal{O}(T_{k\text{-GROUNDED}}(n, d)).$$

Repeating this construction an appropriate number of times we can increase the constant $1/(1 - \alpha)$ to any constant $c > 1$, while the factor on the right hand side is still bounded by a constant. This finally yields an algorithm with runtime satisfying

$$T''_{k\text{-GROUNDED}}(cn, d) \leq \mathcal{O}(T_{k\text{-GROUNDED}}(n, d)).$$

□