

Solving k -SUM Using Few Linear Queries

Jean Cardinal^{*1}, John Iacono^{†2}, and Aurélien Ooms^{‡3}

- 1 Université libre de Bruxelles (ULB), Brussels, Belgium
jcardin@ulb.ac.be
- 2 New York University, New York, USA
esa2016submission@johniacono.com
- 3 Université libre de Bruxelles (ULB), Brussels, Belgium
aureooms@ulb.ac.be

Abstract

The k -SUM problem is given n input real numbers to determine whether any k of them sum to zero. The problem is of tremendous importance in the emerging field of complexity theory within P , and it is in particular open whether it admits an algorithm of complexity $O(n^c)$ with $c < \lceil \frac{k}{2} \rceil$. Inspired by an algorithm due to Meiser (1993), we show that there exist linear decision trees and algebraic computation trees of depth $O(n^3 \log^2 n)$ solving k -SUM. Furthermore, we show that there exists a randomized algorithm that runs in $\tilde{O}(n^{\lceil \frac{k}{2} \rceil + 8})$ time, and performs $O(n^3 \log^2 n)$ linear queries on the input. Thus, we show that it is possible to have an algorithm with a runtime almost identical (up to the +8) to the best known algorithm but for the first time also with the number of queries on the input a polynomial that is independent of k . The $O(n^3 \log^2 n)$ bound on the number of linear queries is also a tighter bound than any known algorithm solving k -SUM, even allowing unlimited total time outside of the queries. By simultaneously achieving few queries to the input without significantly sacrificing runtime vis-à-vis known algorithms, we deepen the understanding of this canonical problem which is a cornerstone of complexity-within- P .

We also consider a range of tradeoffs between the number of terms involved in the queries and the depth of the decision tree. In particular, we prove that there exist $o(n)$ -linear decision trees of depth $\tilde{O}(n^3)$ for the k -SUM problem.

1998 ACM Subject Classification F.2.2 Nonnumerical Algorithms and Problems

Keywords and phrases k -SUM problem, linear decision trees, point location, ε -nets

Digital Object Identifier 10.4230/LIPIcs.ESA.2016.25

1 Introduction

The k -SUM problem is defined as follows: given a collection of n real numbers decide whether any k of them sum to zero, where k is a constant. It is a fixed-parameter version of the subset-sum problem, a standard NP -complete problem. The k -SUM problem, and in particular the special case of 3SUM, has proved to be a cornerstone of the fine-grained complexity program aiming at the construction of a complexity theory for problems in P . In particular, there are deep connections between the complexity of k -SUM, the Strong

* Supported by the “Action de Recherche Concertée” (ARC) COPHYMA, convention number 4.110.H.000023.

† Research partially completed while on sabbatical at the Algorithms Research Group of the Département d’Informatique at the Université Libre de Bruxelles with support from a Fulbright Research Fellowship, the Fonds de la Recherche Scientifique – FNRS, and NSF grants CNS-1229185, CCF-1319648, and CCF-1533564.

‡ Supported by the Fund for Research Training in Industry and Agriculture (FRIA).



Exponential Time Hypothesis [30, 12], and the complexity of many other major problems in P [20, 7, 28, 29, 5, 2, 23, 25, 1, 3, 13].

It has been long known that the k -SUM problem can be solved in time $O(n^{\frac{k}{2}} \log n)$ for even k , and $O(n^{\frac{k+1}{2}})$ for odd k . Erickson [17] proved a near-matching lower bound in the k -linear decision tree model. In this model, the complexity is measured by the depth of a decision tree, every node of which corresponds to a query of the form $q_{i_1} + q_{i_2} + \dots + q_{i_k} \leq? 0$, where q_1, q_2, \dots, q_n are the input numbers. In a recent breakthrough paper, Grønlund and Pettie [23] showed that in the $(2k - 2)$ -linear decision tree model, where queries test the sign of weighted sums of up to $2k - 2$ input numbers, only $O(n^{\frac{k}{2}} \sqrt{\log n})$ queries are required for odd values of k . In particular, there exists a 4-linear decision tree for 3SUM of depth $\tilde{O}(n^{\frac{3}{2}})$ (here the notation \tilde{O} ignores polylogarithmic factors), while every 3-linear decision tree has depth $\Omega(n^2)$ [17]. This indicates that increasing the size of the queries, defined as the maximum number of input numbers involved in a query, can yield significant improvements on the depth of the minimal-height decision tree. Ailon and Chazelle [4] slightly extended the range of query sizes for which a nontrivial lower bound could be established, elaborating on Erickson’s technique.

It has been well established that there exist nonuniform polynomial-time algorithms for the subset-sum problem. One of them was described by Meiser [26], and is derived from a data structure for point location in arrangements of hyperplanes using the bottom vertex decomposition. This algorithm can be cast as the construction of a linear decision tree in which the queries have non-constant size.

1.1 Our results

In Section 3, we show the existence of an n -linear decision tree of depth $\tilde{O}(n^3)$ for k -SUM using a careful implementation of Meiser’s algorithm [26]. Although the high-level algorithm itself is not new, we refine the implementation and analysis for the k -SUM problem.¹ Meiser presented his algorithm as a general method of point location in m given n -dimensional hyperplanes that yielded a $\tilde{O}(n^4 \log m)$ -depth algebraic computation tree; when viewing the k -SUM problem as a point location problem, m is $O(n^k)$ and thus Meiser’s algorithm can be viewed as giving a $\tilde{O}(n^4)$ -depth algebraic computation tree. We show that while the original algorithm was cast as a nonuniform polynomial-time algorithm, it can be implemented in the linear decision tree model with an $\tilde{O}(n^3)$ upper bound. Moreover, this result implies the same improved upper bound on the depth of algebraic computation trees for the k -SUM problem, as shown in Appendix B.

There are two subtleties to this result. The first is inherent to the chosen complexity model: even if the number of queries to the input is small (in particular, the degree of the polynomial complexity is invariant on k), the time required to *determine which queries should be performed* may be arbitrary. In a naïve analysis, we show it can be trivially bounded by $\tilde{O}(n^{k+2})$. In Section 4 we present an algorithm to choose which decisions to perform whereby the running time can be reduced to $\tilde{O}(n^{\frac{k}{2}+8})$. Hence, we obtain an $\tilde{O}(n^{\frac{k}{2}+8})$ time randomized algorithm in the RAM model expected to perform $\tilde{O}(n^3)$ linear queries on the input².

¹ After submitting this manuscript, we learned from a personal communication with Hervé Fournier that a similar analysis appears in his PhD thesis [18] (in French).

² Grønlund and Pettie [23] mention the algorithms of Meyer auf der Heyde [27] and Meiser [26], and state “(…) it was known that all k -LDT problems can be solved by n -linear decision trees with depth $O(n^5 \log n)$ [26], or with depth $O(n^4 \log(nK))$ if the coefficients of the linear function are integers with

■ **Table 1** Complexities of our new algorithms for the k -SUM problem. The query size is the maximum number of elements of the input that can be involved in a single linear query. The number of blocks is a parameter that allows us to change the query size (see Section 5). The origin of the constant in the exponent of the time complexity is due to Lemma 9. We conjecture it can be reduced, though substantial changes in the analysis will likely be needed to do so.

	# blocks	query size	# queries	time
Theorem 3	1	n	$\tilde{O}(n^3)$	$\tilde{O}(n^{\lceil \frac{k}{2} \rceil + 8})$
Theorem 11	b	$k \lceil \frac{n}{b} \rceil$	$\tilde{O}(b^{k-4} n^3)$	$\tilde{O}(b^{\lfloor \frac{k}{2} \rfloor - 9} n^{\lceil \frac{k}{2} \rceil + 8})$
Corollary 12	$b = \Theta(\text{polylog}(n))$	$o(n)$	$\tilde{O}(n^3)$	$\tilde{O}(n^{\lceil \frac{k}{2} \rceil + 8})$
Corollary 13	$b = \Theta(n^\alpha)$	$O(n^{1-\alpha})$	$\tilde{O}(n^{3+(k-4)\alpha})$	$\tilde{O}(n^{(1+\alpha)\frac{k}{2} + 8.5})$

The second issue we address is that the linear queries in the above algorithm may have size n , that is, they may use all the components of the input. The lower bound of Erickson shows that if the queries are of minimal size, the number of queries cannot be a polynomial independent of k such as what we obtain, so non-minimal query size is clearly essential to a drastic reduction in the number of queries needed. This gives rise to the natural question as to what is the relation between query size and number of queries. In particular, one natural question is whether queries of size less than n would still allow the problem to be solved using a number of queries that is a polynomial independent of k . We show that this is possible; in Section 5, we introduce a range of algorithms exhibiting an explicit tradeoff between the number of queries and their size. Using a blocking scheme, we show that we can restrict to $o(n)$ -linear decision trees. We also give a range of tradeoffs for $O(n^{1-\alpha})$ -linear decision trees. Although the proposed algorithms still involve nonconstant-size queries, this is the first time such tradeoffs are explicitly tackled. Table 1 summarizes our results.

2 Definitions and previous work

2.1 Definitions

We consider the k -SUM problem for $k = O(1)$. In what follows, we use the notation $[n] = \{1, 2, \dots, n\}$.

► **Problem (k -SUM).** *Given an input vector $q \in \mathbb{R}^n$, decide whether there exists a k -tuple $(i_1, i_2, \dots, i_k) \in [n]^k$ such that $\sum_{j=1}^k q_{i_j} = 0$.*

The problem amounts to deciding in n -dimensional space, for each hyperplane H of equation $x_{i_1} + x_{i_2} + \dots + x_{i_k} = 0$, whether q lies on, above, or below H . Hence this indeed amounts to locating the point q in the arrangement formed by those hyperplanes. We emphasize that the set of hyperplanes depends only on k and n and not on the actual input vector q .

Linear degeneracy testing (k -LDT) is a generalization of k -SUM where we have arbitrary rational coefficients³ and an independent term in the equations of the hyperplanes.

► **Problem (k -LDT).** *Given an input vectors $q \in \mathbb{R}^n$ and $\alpha \in \mathbb{Q}^n$ and constant $c \in \mathbb{Q}$ decide whether there exists a k -tuple $(i_1, i_2, \dots, i_k) \in [n]^k$ such that $c + \sum_{j=1}^k \alpha_j q_{i_j} = 0$.*

absolute value at most K [27]. Unfortunately these decision trees are not efficiently constructible. The time required to determine *which* comparisons to make is exponential." We prove that the trees can have depth $\tilde{O}(n^3)$ and that the whole algorithm can run in randomized polynomial-time.

³ The usual definition of k -LDT allows arbitrary *real* coefficients. However, the algorithm we provide for Lemma 8 needs the vertices of the arrangement of hyperplanes to have rational coordinates.

Our algorithms apply to this more general problem with only minor changes.

The s -linear decision tree model is a standard model of computation in which several lower bounds for k -SUM have been proven. In the decision tree model, one may ask well-defined questions to an oracle that are answered “yes” or “no.” For s -linear decision trees, a well-defined question consists of testing the sign of a linear function on at most s numbers q_{i_1}, \dots, q_{i_s} of the input q_1, \dots, q_n and can be written as

$$c + \alpha_1 q_{i_1} + \dots + \alpha_s q_{i_s} \stackrel{?}{\leq} 0.$$

Each question is defined to cost a single unit. All other operations can be carried out for free but may not examine the input vector q . We refer to n -linear decision trees simply as linear decision trees.

In this paper, we consider algorithms in the standard integer RAM model with $\Theta(\log n)$ -size words, but in which the input $q \in \mathbb{R}^n$ is accessible *only* via a linear query oracle. Hence we are not allowed to manipulate the input numbers directly. The complexity is measured in two ways: by counting the total number of queries, just as in the linear decision tree model, and by measuring the overall running time, taking into account the time required to determine the sequence of linear queries. This two-track computation model, in which the running time is distinguished from the query complexity, is commonly used in results on comparison-based sorting problems where analyses of both runtime and comparisons are of interest (see for instance [31, 10, 11]).

2.2 Previous Results

The seminal paper by Gajentaan and Overmars [20] showed the crucial role of 3SUM in understanding the complexity of several problems in computational geometry. Since then, there has been an enormous amount of work focusing on the complexity of 3SUM and this problem is now considered a key tool of complexity-within-P [20, 7, 28, 6, 29, 5, 2, 23, 25, 1, 3, 13]. The current conjecture is that no $O(n^{2-\delta})$ -time algorithm exists for 3SUM. It has been known for long that k -SUM is $W[1]$ -hard. Recently, it was shown to be $W[1]$ -complete by Abboud et al. [1].

In Erickson [17], it is shown that we cannot solve 3SUM in subquadratic time in the 3-linear decision tree model:

► **Theorem 1** (Erickson [17]). *The optimal depth of a k -linear decision tree that solves the k -LDT problem is $\Theta(n^{\lceil \frac{k}{2} \rceil})$.*

The proof uses an adversary argument which can be explained geometrically. As we already observed, we can solve k -LDT problems by modeling them as point location problems in an arrangement of hyperplanes. Solving one such problem amounts to determining which cell of the arrangement contains the input point. The adversary argument of Erickson [17] is that there exists a cell having $\Omega(n^{\lceil \frac{k}{2} \rceil})$ boundary facets and in this model point location in such a cell requires testing each facet.

Ailon and Chazelle [4] study s -linear decision trees to solve the k -SUM problem when $s > k$. In particular, they give an additional proof for the $\Omega(n^{\lceil \frac{k}{2} \rceil})$ lower bound of Erickson [17] and generalize the lower bound for the s -linear decision tree model when $s > k$. Note that the exact lower bound given by Erickson [17] for $s = k$ is $\Omega((nk^{-k})^{\lceil \frac{k}{2} \rceil})$ while the one given by Ailon and Chazelle [4] is $\Omega((nk^{-3})^{\lceil \frac{k}{2} \rceil})$. Their result improves therefore the lower bound for $s = k$ when k is large. The lower bound they prove for $s > k$ is the following:

► **Theorem 2** (Ailon and Chazelle [4]). *The depth of an s -linear decision tree solving the k -LDT problem is*

$$\Omega\left((nk^{-3})^{\frac{2k-s}{2\lceil\frac{s-k+1}{2}\rceil}(1-\epsilon_k)}\right),$$

where $\epsilon_k > 0$ tends to 0 as $k \rightarrow \infty$.

This lower bound breaks down when $k = \Omega(n^{1/3})$ or $s \geq 2k$ and the cases where $k < 6$ give trivial lower bounds. For example, in the case of 3SUM with $s = 4$ we only get an $\Omega(n)$ lower bound.

As for upper bounds, Baran et al. [6] gave subquadratic Las Vegas algorithms for 3SUM on integer and rational numbers in the circuit RAM, word RAM, external memory, and cache-oblivious models of computation. The idea of their approach is to exploit the parallelism of the models, using linear and universal hashing.

More recently, Grønlund and Pettie [23] proved the existence of a linear decision tree solving the 3SUM problem using a strongly subquadratic number of linear queries. The classical quadratic algorithm for 3SUM uses 3-linear queries while the decision tree of Grønlund and Pettie uses 4-linear queries and requires $O(n^{3/2}\sqrt{\log n})$ of them. Moreover, they show that their decision tree can be used to get better upper bounds for k -SUM when k is odd.

They also provide two subquadratic 3SUM algorithms. A deterministic one running in $O(n^2/(\log n/\log \log n)^{2/3})$ time and a randomized one running in $O(n^2(\log \log n)^2/\log n)$ time with high probability. These results refuted the long-lived conjecture that 3SUM cannot be solved in subquadratic time in the RAM model.

Freund [19] and Gold and Sharir [21] later gave improvements on the results of Grønlund and Pettie [23]. Freund [19] gave a deterministic algorithm for 3SUM running in $O(n^2 \log \log n / \log n)$ time. Gold and Sharir [21] gave another deterministic algorithm for 3SUM with the same running time and shaved off the $\sqrt{\log n}$ factor in the decision tree complexities of 3SUM and k -SUM given by Grønlund and Pettie.

Meyer auf der Heide [27] gave the first point location algorithm to solve the knapsack problem in the linear decision tree model in polynomial time. He thereby answers a question raised by Dobkin and Lipton [15, 16], Yao [32] and others. However, if one uses this algorithm to locate a point in an arbitrary arrangement of hyperplanes the running time is increased by a factor linear in the greatest coefficient in the equations of all hyperplanes. On the other hand, the complexity of Meiser's point location algorithm is polynomial in the dimension, logarithmic in the number of hyperplanes and does not depend on the value of the coefficients in the equations of the hyperplanes. A useful complete description of the latter is also given by Bürgisser et al. [9] (Section 3.4).

3 Query complexity

In this section and the next, we prove the following first result.

► **Theorem 3.** *There exist linear decision trees of depth at most $O(n^3 \log^2 n)$ solving the k -SUM and the k -LDT problems. Furthermore, for the two problems there exists an $\tilde{O}(n^{\lceil\frac{k}{2}\rceil+8})$ Las Vegas algorithm in the RAM model expected to perform $O(n^3 \log^2 n)$ linear queries on the input.*

3.1 Algorithm outline

For a fixed set of hyperplanes \mathcal{H} and given input vertex q in \mathbb{R}^n , Meiser's algorithm allows us to determine the cell of the arrangement $\mathcal{A}(\mathcal{H})$ that contains q in its interior (or that *is* q if q is a 0-cell of $\mathcal{A}(\mathcal{H})$), that is, the positions $\sigma(H, q) \in \{-, 0, +\}$ of q with respect to all hyperplanes $H \in \mathcal{H}$. In the k -SUM problem, the set \mathcal{H} is the set of $\Theta(n^k)$ hyperplanes with equations of the form $x_{i_1} + x_{i_2} + \dots + x_{i_k} = 0$. These equations can be modified accordingly for k -LDT.

We use standard results on ε -nets. By combining a theorem due to Blumer et al. [8] with the results of Meiser [26]⁴, it is possible to construct an ε -net \mathcal{N} for the range space defined by hyperplanes and simplices using a random uniform sampling on \mathcal{H} .

► **Theorem 4.** *For all real numbers $\varepsilon > 0, c \geq 1$, if we choose at least $cn^2 \log n \varepsilon^{-1} \log \varepsilon^{-1}$ hyperplanes of \mathcal{H} uniformly at random and denote this selection \mathcal{N} then for any simplex intersected by more than $\varepsilon|\mathcal{H}|$ hyperplanes of \mathcal{H} , with probability $1 - 2^{-\Omega(c)}$, at least one of the intersecting hyperplanes is contained in \mathcal{N} .*

The contrapositive states that if no hyperplane in \mathcal{N} intersects a given simplex, then with high probability the number of hyperplanes of \mathcal{H} intersecting the simplex is at most $\varepsilon|\mathcal{H}|$.

We can use this to design a prune and search algorithm as follows:

- (A) construct an ε -net \mathcal{N} ,
- (B) compute the cell C of $\mathcal{A}(\mathcal{N})$ containing the input point q in its interior,
- (C) construct a simplex S inscribed in C and containing q in its interior,
- (D) recurse on the hyperplanes of \mathcal{H} intersecting the interior of S .

Proceeding this way with a constant ε guarantees that at most a constant fraction ε of the hyperplanes remains after the pruning step, and thus the cumulative number of queries made to determine the enclosing cell at each step is $O(n^2 \log n \log |\mathcal{H}|)$ when done in a brute-force way. However, we still need to explain how to find a simplex S inscribed in C and containing q in its interior. This procedure corresponds to the well-known *bottom vertex decomposition* (or *triangulation*) of a hyperplane arrangement [22, 14].

3.2 Finding a simplex

In order to simplify the exposition of the algorithm, we assume, without loss of generality, that the input numbers q_i all lie in the interval $[-1, 1]$. This assumption is justified by observing that we can normalize all the input numbers by the largest absolute value of a component of q . One can then see that every linear query on the normalized input can be implemented as a linear query on the original input. A similar transformation can be carried out for the k -LDT problem. This allows us to use bounding hyperplanes of equations $x_i = \pm 1, i \in [n]$. We denote by \mathcal{B} this set of hyperplanes. Hence, if we choose a subset \mathcal{N} of the hyperplanes, the input point is located in a bounded cell of the arrangement $\mathcal{A}(\mathcal{N} \cup \mathcal{B})$. Note that $|\mathcal{N} \cup \mathcal{B}| = O(|\mathcal{N}|)$ for all interesting values of ε .

We now explain how to construct S under this assumption. The algorithm can be sketched as follows. (Recall that $\sigma(H, p)$ denotes the relative position of p with respect to the hyperplane H .)

⁴ Note that Meiser used an older result due to Haussler and Welzl [24] and got an extra $\log n$ factor in the size of the ε -net.

Algorithm 1 (Constructing S).

Input: A point q in $[-1, 1]^n$, a set \mathcal{I} of hyperplanes not containing q , and a set \mathcal{E} of hyperplanes in general position containing q , such that the cell

$$C = \{ p: \sigma(H, p) = \sigma(H, q) \text{ or } \sigma(H, p) = 0 \text{ for all } H \in (\mathcal{I} \cup \mathcal{E}) \}$$

is a bounded polytope. The value $\sigma(H, q)$ is known for all $H \in (\mathcal{I} \cup \mathcal{E})$.

Output: A simplex $S \in C$ that contains q in its interior (if it is not a point), and all vertices of which are vertices of C .

0. If $|\mathcal{E}| = n$, return q .
1. Determine a vertex ν of C .
2. Let q' be the projection of q along $\nu\vec{q}$ on the boundary of C . Compute $\mathcal{I}_\theta \subseteq \mathcal{I}$, the subset of hyperplanes in \mathcal{I} containing q' . Compute $\mathcal{I}_\tau \subseteq \mathcal{I}_\theta$, a maximal subset of those hyperplanes such that $\mathcal{E}' = \mathcal{E} \cup \mathcal{I}_\tau$ is a set of hyperplanes in general position.
3. Recurse on q' , $\mathcal{I}' = \mathcal{I} \setminus \mathcal{I}_\theta$, and \mathcal{E}' , and store the result in S' .
4. Return S , the convex hull of $S' \cup \{\nu\}$.

Step 0 is the base case of the recursion: when there is only one point left, just return that point. This step uses no query.

We can solve step 1 by using linear programming with the known values of $\sigma(H, q)$ as linear constraints. We arbitrarily choose an objective function with a gradient non-orthogonal to all hyperplanes in \mathcal{I} and look for the optimal solution. The optimal solution being a vertex of the arrangement, its coordinates are independent of q , and thus this step involves no query at all.

Step 2 prepares the recursive step by finding the hyperplanes containing q' . This can be implemented as a ray-shooting algorithm that performs a number of comparisons between projections of q on different hyperplanes of \mathcal{I} without explicitly computing them. In Appendix A, we prove that all such comparisons can be implemented using $O(|\mathcal{I}|)$ linear queries. Constructing \mathcal{E}' can be done by solving systems of linear equations that do not involve q .

In step 3, the input conditions are satisfied, that is, $q' \in [-1, 1]^n$, \mathcal{I}' is a set of hyperplanes not containing q' , \mathcal{E}' is a set of hyperplanes in general position containing q' , C' is a d -cell of C and is thus a bounded polytope. The value $\sigma(H, q')$ differs from $\sigma(H, q)$ only for hyperplanes that have been removed from \mathcal{I} , and for those $\sigma(H, q') = 0$, hence we know all necessary values $\sigma(H, q')$ in advance.

Since $|\mathcal{I}'| < |\mathcal{I}|$, $|\mathcal{E}'| > |\mathcal{E}|$, and $|\mathcal{I} \setminus \mathcal{I}'| - |\mathcal{E}' \setminus \mathcal{E}| \geq 0$, the complexity of the recursive call is no more than that of the parent call, and the maximal depth of the recursion is n . Thus, the total number of linear queries made to compute S is $O(n|\mathcal{I}|)$.


Hence given an input point $q \in [-1, 1]^n$, an arrangement of hyperplanes $\mathcal{A}(\mathcal{N})$, and the value of $\sigma(H, q)$ for all $H \in (\mathcal{N} \cup \mathcal{B})$, we can compute the desired simplex S by running Algorithm 1 on q , $\mathcal{I} = \{H \in (\mathcal{N} \cup \mathcal{B}): \sigma(H, q) \neq 0\}$, and $\mathcal{E} \subseteq (\mathcal{N} \cup \mathcal{B}) \setminus \mathcal{I}$. This uses $O(n^3 \log n)$ linear queries. Figure 1 illustrates a step of the algorithm.

3.3 Assembling the pieces

Let us summarize the algorithm

Algorithm 2.

Input: $q \in [-1, 1]^n$



■ **Figure 1** Illustration of a step of Algorithm 1.

1. Pick $O(n^2 \log n)$ hyperplanes of \mathcal{H} at random and locate q in this arrangement. Call C the cell containing q .
2. Construct the simplex S containing q and inscribed in C , using Algorithm 1.
3. For every hyperplane of \mathcal{H} containing S , output a solution.
4. Recurse on hyperplanes of \mathcal{H} intersecting the interior of S .

The query complexity of step 1 is $O(n^2 \log n)$, and that of step 2 is $O(n^3 \log n)$. Steps 3 and 4 do not involve any query at all. The recursion depth is $O(\log |\mathcal{H}|)$, with $|\mathcal{H}| = O(n^k)$, hence the total query complexity of this algorithm is $O(n^3 \log^2 n)$. This proves the first part of Theorem 3.

We can also consider the overall complexity of the algorithm in the RAM model, that is, taking into account the steps that do not require any query, but for which we still have to process the set \mathcal{H} . Note that the complexity bottleneck of the algorithm are steps 3-4, where we need to prune the list of hyperplanes according to their relative positions with respect to S . For this purpose, we simply maintain explicitly the list of all hyperplanes, starting with the initial set corresponding to all k -tuples. Then the pruning step can be performed by looking at the position of each vertex of S relative to each hyperplane of \mathcal{H} . Because in our case hyperplanes have only k nonzero coefficients, this uses a number of integer arithmetic operations on $\tilde{O}(n)$ bits integers that is proportional to the number of vertices times the number of hyperplanes. (For the justification of the bound on the number of bits needed to represent vertices of the arrangement see Appendix D.) Since we recurse on a fraction of the set, the overall complexity is $\tilde{O}(n^2 |\mathcal{H}|) = \tilde{O}(n^{k+2})$. The next section is devoted to improving this running time.

4 Time complexity

Proving the second part of Theorem 3 involves efficient implementations of the two most time-consuming steps of Algorithm 2. In order to efficiently implement the pruning step, we define an intermediate problem, that we call the *double k -SUM* problem.

► **Problem (double k -SUM).** *Given two vectors $\nu_1, \nu_2 \in [-1, 1]^n$, where the coordinates of ν_i can be written down as fractions whose numerator and denominator lie in the interval*

$[-M, M]$, enumerate all $i \in [n]^k$ such that

$$\left(\sum_{j=1}^k \nu_{1,i_j} \right) \left(\sum_{j=1}^k \nu_{2,i_j} \right) < 0.$$

In other words, we wish to list all hyperplanes of \mathcal{H} intersecting the open line segment $\nu_1\nu_2$. We give an efficient output-sensitive algorithm for this problem.

► **Lemma 5.** *The double k -SUM problem can be solved in time $O(n^{\lceil \frac{k}{2} \rceil} \log n \log M + Z)$, where Z is the size of the solution.*

Proof. If k is even, we consider all possible $\frac{k}{2}$ -tuples of numbers in ν_1 and ν_2 and sort their sums in increasing order. This takes time $O(n^{\frac{k}{2}} \log n)$ and yields two permutations π_1 and π_2 of $[n^{\frac{k}{2}}]$. If k is odd, then we sort both the $\lceil \frac{k}{2} \rceil$ -tuples and the $\lfloor \frac{k}{2} \rfloor$ -tuples. For simplicity, we will only consider the even case in what follows. The odd case carries through.

We let $N = n^{\frac{k}{2}}$. For $i \in [N]$ and $m \in \{1, 2\}$, let $\Sigma_{m,i}$ be the sum of the $\frac{k}{2}$ components of the i th $\frac{k}{2}$ -tuple in ν_m , in the order prescribed by π_m .

We now consider the two $N \times N$ matrices M_1 and M_2 giving all possible sums of two $\frac{k}{2}$ -tuples, for both ν_1 with the ordering π_1 and ν_2 with the ordering π_2 .

We first solve the k -SUM problem on ν_1 , by finding the sign of all pairs $\Sigma_{1,i} + \Sigma_{1,j}$, $i, j \in [N]$. This can be done in time $O(N)$ by parsing the matrix M_1 , just as in the standard k -SUM algorithm. We do the same with M_2 .

The set of all indices $i, j \in [N]$ such that $\Sigma_{1,i} + \Sigma_{1,j}$ is positive forms a staircase in M_1 . We sweep M_1 column by column in order of increasing $j \in [N]$, in such a way that the number of indices i such that $\Sigma_{1,i} + \Sigma_{1,j} > 0$ is growing. For each new such value i that is encountered during the sweep, we insert the corresponding $i' = \pi_2(\pi_1^{-1}(i))$ in a balanced binary search tree.

After each sweep step in M_1 — that is, after incrementing j and adding the set of new indices i' in the tree — we search the tree to identify all the indices i' such that $\Sigma_{2,i'} + \Sigma_{2,j'} < 0$, where $j' = \pi_2(\pi_1^{-1}(j))$. Since those indices form an interval in the ordering π_2 when restricted to the indices in the tree, we can search for the largest i'_0 such that $\Sigma_{2,i'_0} < -\Sigma_{2,j'}$ and retain all indices $i' \leq i'_0$ that are in the tree. If we denote by z the number of such indices, this can be done in $O(\log N + z) = O(\log n + z)$ time. Now all the pairs i', j' found in this way are such that $\Sigma_{1,i} + \Sigma_{1,j}$ is positive and $\Sigma_{2,i'} + \Sigma_{2,j'}$ is negative, hence we can output the corresponding k -tuples. To get all the pairs i', j' such that $\Sigma_{1,i} + \Sigma_{1,j}$ is negative and $\Sigma_{2,i'} + \Sigma_{2,j'}$ positive, we repeat the sweeping algorithm after swapping the roles of ν_1 and ν_2 .

Every matching k -tuple is output exactly once, and every $\frac{k}{2}$ -tuple is inserted at most once in the binary search tree. Hence the algorithm runs in the claimed time.

Note that we only manipulate rational numbers that are the sum of at most k rational numbers of size $O(\log M)$. ◀

Now observe that a hyperplane intersects the interior of a simplex if and only if it intersects the interior of one of its edges. Hence given a simplex S we can find all hyperplanes of \mathcal{H} intersecting its interior by running the above algorithm $\binom{n}{2}$ times, once for each pair of vertices (ν_1, ν_2) of S , and take the union of the solutions. The overall running time for this implementation will therefore be $\tilde{O}(n^2(n^{\lceil \frac{k}{2} \rceil} \log M + Z))$, where Z is at most the number of intersecting hyperplanes and M is to be determined later. This provides an implementation of the pruning step in Meiser’s algorithm, that is, step 4 of Algorithm 2.

25:10 Solving k -SUM Using Few Linear Queries

► **Corollary 6.** *Given a simplex S , we can compute all k -SUM hyperplanes intersecting its interior in $\tilde{O}(n^2(n^{\lceil \frac{k}{2} \rceil} \log M + Z))$ time, where $\log M$ is proportional to the number of bits necessary to represent S .*

In order to detect solutions in step **3** of Algorithm 2, we also need to be able to quickly solve the following problem.

► **Problem (multiple k -SUM).** *Given d points $\nu_1, \nu_2, \dots, \nu_d \in \mathbb{R}^n$, where the coordinates of ν_i can be written down as fractions whose numerator and denominator lie in the interval $[-M, M]$, decide whether there exists a hyperplane with equation of the form $x_{i_1} + x_{i_2} + \dots + x_{i_k} = 0$ containing all of them.*

Here the standard k -SUM algorithm can be applied, taking advantage of the fact that the coordinates lie in a small discrete set.

► **Lemma 7.** *k -SUM on n integers $\in [-V, V]$ can be solved in time $\tilde{O}(n^{\lceil \frac{k}{2} \rceil} \log V)$.*

► **Lemma 8.** *Multiple k -SUM can be solved in time $\tilde{O}(dn^{\lceil \frac{k}{2} \rceil + 2} \log M)$.*

Proof. Let $\mu_{i,j}$ and $\delta_{i,j}$ be the numerator and denominator of $\nu_{i,j}$ when written as an irreducible fraction. We define

$$\zeta_{i,j} = \nu_{i,j} \prod_{(i',j') \in [d] \times [n]} \delta_{i',j'} = \frac{\mu_{i,j} \prod_{(i',j') \in [d] \times [n]} \delta_{i',j'}}{\delta_{i,j}}.$$

By definition $\zeta_{i,j}$ is an integer and its absolute value is bounded by $U = M^{n^2}$, that is, it can be represented using $O(n^2 \log M)$ bits. Moreover, if one of the hyperplanes contains the point $(\zeta_{i,1}, \zeta_{i,2}, \dots, \zeta_{i,n})$, then it contains ν_i . Construct n integers of $O(dn^2 \log M)$ bits that can be written $\zeta_{1,j} + U, \zeta_{2,j} + U, \dots, \zeta_{d,j} + U$ in base $2Uk + 1$. The answer to our decision problem is “yes” if and only if there exists k of those numbers whose sum is kU, kU, \dots, kU . We simply subtract the number U, U, \dots, U to all n input numbers to obtain a standard k -SUM instance on n integers of $O(dn^2 \log M)$ bits. ◀

We now have efficient implementations of steps **3** and **4** of Algorithm 2 and can proceed to the proof of the second part of Theorem 3.

Proof. The main idea consists of modifying the first iteration of Algorithm 2, by letting $\varepsilon = \Theta(n^{-\frac{k}{2}})$. Hence we pick a random subset \mathcal{N} of $O(n^{k/2+2} \log n)$ hyperplanes in \mathcal{H} and use this as an ε -net. This can be done efficiently, as shown in Appendix C.

Next, we need to locate the input q in the arrangement induced by \mathcal{N} . This can be done by running Algorithm 2 on the set \mathcal{N} . From the previous considerations on Algorithm 2, the running time of this step is

$$O(n|\mathcal{N}|) = \tilde{O}(n^{k/2+4}),$$

and the number of queries is $O(n^3 \log^2 n)$.

Then, in order to prune the hyperplanes in \mathcal{H} , we have to compute a simplex S that does not intersect any hyperplane of \mathcal{N} . For this, we observe that the above call to Algorithm 2 involves computing a sequence of simplices for the successive pruning steps. We save the description of those simplices. Recall that there are $O(\log n)$ of them, all of them contain the input q and have vertices coinciding with vertices of the original arrangement $\mathcal{A}(\mathcal{H})$. In order to compute a simplex S avoiding all hyperplanes of \mathcal{N} , we can simply apply Algorithm 1 on

the set of hyperplanes bounding the intersection of these simplices. The running time and number of queries for this step are bounded respectively by $n^{O(1)}$ and $O(n^2 \log n)$.

Note that the vertices of S are not vertices of $\mathcal{A}(\mathcal{H})$ anymore. However, their coordinates lie in a finite set (see Appendix D)

► **Lemma 9.** *Vertices of S have rational coordinates whose fraction representations have their numerators and denominators absolute values bounded above by $C^{4n^5} n^{2n^5+n^3+\frac{n}{2}}$, where C is a constant.*

We now are in position to perform the pruning of the hyperplanes in \mathcal{H} with respect to S . The number of remaining hyperplanes after the pruning is at most $\varepsilon n^k = O(n^{k/2})$. Hence from Corollary 6, the pruning can be performed in time proportional $\tilde{O}(n^{\lceil k/2 \rceil + 7})$.

Similarly, we can detect any hyperplane of \mathcal{H} containing S using the result of Lemma 8 in time $\tilde{O}(n^{\lceil k/2 \rceil + 8})$. Note that those last two steps do not require any query.

Finally, it remains to detect any solution that may lie in the remaining set of hyperplanes of size $O(n^{k/2})$. We can again fall back on Algorithm 2, restricted to those hyperplanes. The running time is $\tilde{O}(n^{k/2+2})$, and the number of queries is still $O(n^3 \log^2 n)$.

Overall, the maximum running time of a step is $\tilde{O}(n^{\lceil \frac{k}{2} \rceil + 8})$, while the number of queries is always bounded by $O(n^3 \log^2 n)$. ◀

5 Query size

In this section, we consider a simple blocking scheme that allows us to explore a tradeoff between the number of queries and the size of the queries.

► **Lemma 10.** *For any integer $b > 0$, an instance of the k -SUM problem on $n > b$ numbers can be split into $O(b^{k-1})$ instances on at most $k \lceil \frac{n}{b} \rceil$ numbers, so that every k -tuple forming a solution is found in exactly one of the subproblems. The transformation can be carried out in time $O(n \log n + b^{k-1})$.*

Proof. Given an instance on n numbers, we can sort them in time $O(n \log n)$, then partition the sorted sequence into b consecutive blocks B_1, B_2, \dots, B_b of equal size. This partition can be associated with a partition of the real line into b intervals, say I_1, I_2, \dots, I_b . Now consider the partition of \mathbb{R}^k into grid cells defined by the k th power of the partition I_1, I_2, \dots, I_b . The hyperplane of equation $x_1 + x_2 + \dots + x_k = 0$ hits $O(b^{k-1})$ such grid cells. Each grid cell $I_{i_1} \times I_{i_2} \times \dots \times I_{i_k}$ corresponds to a k -SUM problem on the numbers in the set $B_{i_1} \cup B_{i_2} \cup \dots \cup B_{i_k}$ (note that the indices i_j need not be distinct). Hence each such instance has size at most $k \lceil \frac{n}{b} \rceil$. ◀

Combining Lemma 10 and Theorem 3 directly yields the following.

► **Theorem 11.** *For any integer $b > 0$, there exists a $k \lceil \frac{n}{b} \rceil$ -linear decision tree of depth $\tilde{O}(b^{k-4} n^3)$ solving the k -SUM problem. Moreover, this decision tree can be implemented as an $\tilde{O}(b^{\lfloor \frac{k}{2} \rfloor - 9} n^{\lceil \frac{k}{2} \rceil + 8})$ Las Vegas algorithm.*

The following two corollaries are obtained by taking $b = \Theta(\text{polylog}(n))$, and $b = \Theta(n^\alpha)$, respectively

► **Corollary 12.** *There exists an $o(n)$ -linear decision tree of depth $\tilde{O}(n^3)$ solving the k -SUM problem. Moreover, this decision tree can be implemented as an $\tilde{O}(n^{\lceil \frac{k}{2} \rceil + 8})$ Las Vegas algorithm.*

► **Corollary 13.** *For any α such that $0 < \alpha < 1$, there exists an $O(n^{1-\alpha})$ -linear decision tree of depth $\tilde{O}(n^{3+(k-4)\alpha})$ solving the k -SUM problem. Moreover, this decision tree can be implemented as an $\tilde{O}(n^{(1+\alpha)\frac{k}{2}+8.5})$ Las Vegas algorithm.*

Note that the latter query complexity improves on $\tilde{O}(n^{\frac{k}{2}})$ whenever $\alpha < \frac{k-6}{2k-8}$ and $k \geq 7$. By choosing $\alpha = \frac{k-6}{2k-8} - \frac{\beta}{k-4}$ we obtain $O(n^{1-\frac{k-6}{2k-8}+\frac{\beta}{k-4}})$ -linear decision trees of depth $\tilde{O}(n^{\frac{k}{2}-\beta})$ for any $k \geq 7$. Hence for instance, we obtain $O(n^{\frac{3}{4}+\frac{\beta}{4}})$ -linear decision trees of depth $\tilde{O}(n^{4-\beta})$ for the 8SUM problem.

References

- 1 Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In *European Symposium on Algorithms (ESA 2014)*, pages 1–12. Springer, 2014.
- 2 Amir Abboud, Virginia Vassilevska Williams, and Oren Weimann. Consequences of faster alignment of sequences. In *International Colloquium on Automata, Languages, and Programming (ICALP 2014)*, volume 8572 of *Lecture Notes in Computer Science*, pages 39–51. Springer, 2014.
- 3 Amir Abboud, Virginia Vassilevska Williams, and Huacheng Yu. Matching triangles and basing hardness on an extremely popular conjecture. In *Symposium on Theory of Computing (STOC 2015)*, pages 41–50, 2015.
- 4 Nir Ailon and Bernard Chazelle. Lower bounds for linear degeneracy testing. *J. ACM*, 52(2):157–171, 2005.
- 5 Amihoud Amir, Timothy M. Chan, Moshe Lewenstein, and Noa Lewenstein. On hardness of jumbled indexing. In *International Colloquium on Automata, Languages, and Programming (ICALP 2014)*, volume 8572 of *Lecture Notes in Computer Science*, pages 114–125. Springer, 2014.
- 6 Ilya Baran, Erik D. Demaine, and Mihai Patrascu. Subquadratic algorithms for 3SUM. *Algorithmica*, 50(4):584–596, 2008.
- 7 Gill Barequet and Sarel Har-Peled. Polygon-containment and translational min-hausdorff-distance between segment sets are 3SUM-hard. In *Symposium on Discrete Algorithms (SODA 1999)*, pages 862–863, 1999.
- 8 Anselm Blumer, Andrzej Ehrenfeucht, David Haussler, and Manfred K. Warmuth. Learnability and the vapnik-chervonenkis dimension. *J. ACM*, 36(4):929–965, 1989.
- 9 Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften*. Springer, 1997.
- 10 Jean Cardinal, Samuel Fiorini, Gwenaël Joret, Raphaël M Jungers, and J Ian Munro. An efficient algorithm for partial order production. *SIAM journal on computing*, 39(7):2927–2940, 2010.
- 11 Jean Cardinal, Samuel Fiorini, Gwenaël Joret, Raphaël M Jungers, and J Ian Munro. Sorting under partial information (without the ellipsoid algorithm). *Combinatorica*, 33(6):655–697, 2013.
- 12 Marco Carmosino, Jiawei Gao, Russell Impagliazzo, Ivan Mikhailin, Ramamohan Paturi, and Stefan Schneider. Nondeterministic extensions of the strong exponential time hypothesis and consequences for non-reducibility. *Electronic Colloquium on Computational Complexity (ECCC 2015)*, 22:148, 2015.
- 13 Timothy M. Chan and Moshe Lewenstein. Clustered integer 3SUM via additive combinatorics. In Rocco A. Servedio and Ronitt Rubinfeld, editors, *Symposium on Theory of Computing (STOC 2015)*, pages 31–40. ACM, 2015.
- 14 Kenneth L. Clarkson. A randomized algorithm for closest-point queries. *SIAM J. Comput.*, 17(4):830–847, 1988.

- 15 David P. Dobkin and Richard J. Lipton. On some generalizations of binary search. In *Symposium on Theory of Computing (STOC 1974)*, pages 310–316, 1974.
- 16 David P. Dobkin and Richard J. Lipton. A lower bound of the $\frac{1}{2}n^2$ on linear search programs for the knapsack problem. *J. Comput. Syst. Sci.*, 16(3):413–417, 1978.
- 17 Jeff Erickson. Lower bounds for linear satisfiability problems. *Chicago Journal of Theoretical Computer Science*, 1999.
- 18 Hervé Fournier. *Complexité et expressibilité sur les réels*. PhD thesis, École normale supérieure de Lyon, 2001.
- 19 Ari Freund. Improved subquadratic 3SUM. *Algorithmica*, 2015. To appear.
- 20 Anka Gajentaan and Mark H. Overmars. On a class of $O(n^2)$ problems in computational geometry. *Comput. Geom.*, 5:165–185, 1995.
- 21 O. Gold and M. Sharir. Improved bounds for 3SUM, k -SUM, and linear degeneracy. *ArXiv e-prints*, 2015. arXiv:1512.05279 [cs.DS].
- 22 Jacob E. Goodman and Joseph O’Rourke, editors. *Handbook of Discrete and Computational Geometry, Second Edition*. Chapman and Hall/CRC, 2004.
- 23 Allan Grønlund and Seth Pettie. Threesomes, degenerates, and love triangles. In *Foundations of Computer Science (FOCS 2014)*, pages 621–630. IEEE, 2014.
- 24 David Haussler and Emo Welzl. ϵ -nets and simplex range queries. *Discrete & Computational Geometry*, 2(1):127–151, 1987.
- 25 T. Kopelowitz, S. Pettie, and E. Porat. Higher lower bounds from the 3SUM conjecture. *ArXiv e-prints*, 2014. arXiv:1407.6756 [cs.DS].
- 26 S. Meiser. Point location in arrangements of hyperplanes. *Information and Computation*, 106(2):286–303, 1993.
- 27 Friedhelm Meyer auf der Heide. A polynomial linear search algorithm for the n -dimensional knapsack problem. *J. ACM*, 31(3):668–676, 1984.
- 28 Joseph S. B. Mitchell and Joseph O’Rourke. Computational geometry column 42. *Int. J. Comput. Geometry Appl.*, 11(5):573–582, 2001.
- 29 Mihai Patrascu. Towards polynomial lower bounds for dynamic problems. In *Symposium on Theory of Computing (STOC 2010)*, pages 603–610, 2010.
- 30 Mihai Patrascu and Ryan Williams. On the possibility of faster SAT algorithms. In *Symposium on Discrete Algorithms (SODA 2010)*, pages 1065–1075, 2010.
- 31 William L. Steiger and Ileana Streinu. A pseudo-algorithmic separation of lines from pseudolines. *Information Processing Letters*, 53(5):295–299, 1995.
- 32 Andrew Chi-Chih Yao. On parallel computation for the knapsack problem. *J. ACM*, 29(3):898–903, 1982.

A Keeping queries linear in Algorithm 1

In Algorithm 1, we want to ensure that the queries we make in step 2 are linear and that the queries we will make in the recursion step remain linear too.

► **Lemma 14.** *Algorithm 1 can be implemented so that it uses $O(n|I|)$ linear queries.*

Proof. Let us first analyze what the queries of step 2 look like. In addition to the input point q we are given a vertex ν and we want to find the projection q' of q in direction $\nu\vec{q}$ on the hyperplanes of \mathcal{I}_θ . Let the equation of H_i be $\Pi_i(x) = c_i + d_i \cdot x = 0$ where c_i is a scalar and d_i is a vector. The projection of q along $\nu\vec{q}$ on a hyperplane H_i can thus be written⁵

⁵ Note that we project from ν instead of q . We are allowed to do this since $\nu + \lambda_i\nu\vec{q} = q + (\lambda_i - 1)\nu\vec{q}$ and there is no hyperplane separating q from ν .

$\rho(q, \nu, H_i) = \nu + \lambda_i \vec{\nu}q$ such that $\Pi_i(\nu + \lambda_i \vec{\nu}q) = c_i + d_i \cdot \nu + \lambda_i d_i \cdot \vec{\nu}q = 0$. Computing the closest hyperplane amounts to finding $\lambda_\theta = \min_{\lambda_i > 0} \lambda_i$. Since $\lambda_i = -\frac{c_i + d_i \cdot \nu}{d_i \cdot \vec{\nu}q}$ we can test whether $\lambda_i > 0$ using the linear query⁶ $-\frac{d_i \cdot \vec{\nu}q}{c_i + d_i \cdot \nu} >? 0$. Moreover, if $\lambda_i > 0$ and $\lambda_j > 0$ we can test whether $\lambda_i < \lambda_j$ using the linear query $\frac{d_i \cdot \vec{\nu}q}{c_i + d_i \cdot \nu} <? \frac{d_j \cdot \vec{\nu}q}{c_j + d_j \cdot \nu}$. Step **2** can thus be achieved using $O(1)$ $(2k)$ -linear queries per hyperplane of \mathcal{N} .

In step **4**, the recursive step is carried out on $q' = \nu + \lambda_\theta \vec{\nu}q = \nu - \frac{c_\theta + d_\theta \cdot \nu}{d_\theta \cdot \vec{\nu}q} \vec{\nu}q$ hence comparing λ'_i to 0 amounts to performing the query $-\frac{d_i \cdot \vec{\nu}q'}{c_i + d_i \cdot \nu'} >? 0$, which is not linear in q . The same goes for comparing λ'_i to λ'_j with the query $\frac{d_i \cdot \vec{\nu}q'}{c_i + d_i \cdot \nu'} <? \frac{d_j \cdot \vec{\nu}q'}{c_j + d_j \cdot \nu'}$.

However, we can multiply both sides of the inequality test by $d_\theta \vec{\nu}q$ to keep the queries linear as shown below. We must be careful to take into account the sign of the expression $d_\theta \vec{\nu}q$, this costs us one additional linear query.

This trick can be used at each step of the recursion. Let $q^{(0)} = q$, then we have

$$q^{(s+1)} = \nu^{(s)} - \frac{c_{\theta_s} + d_{\theta_s} \cdot \nu^{(s)}}{d_{\theta_s} \cdot \vec{\nu}q^{(s)}} \vec{\nu}q^{(s)}$$

and $(d_{\theta_s} \cdot \vec{\nu}q^{(s)})q^{(s+1)}$ yields a vector whose components are linear in $q^{(s)}$. Hence, $(\prod_{k=0}^s d_{\theta_k} \cdot \vec{\nu}q^{(k)})q^{(s+1)}$ yields a vector whose components are linear in q , and for all pairs of vectors d_i and $\nu^{(s+1)}$ we have that $(\prod_{k=0}^s d_{\theta_k} \cdot \vec{\nu}q^{(k)})(d_i \cdot \vec{\nu}q^{(s+1)})$ is linear in q .

Hence at the s th recursive step of the algorithm, we will perform at most $|\mathcal{N}|$ linear queries of the type

$$-\left(\prod_{k=0}^{s-1} d_{\theta_k} \cdot \vec{\nu}q^{(k)}\right) \frac{d_i \cdot \vec{\nu}q^{(s)}}{c_i + d_i \cdot \nu^{(s)}} >? 0$$

$|\mathcal{N}| - 1$ linear queries of the type

$$\left(\prod_{k=0}^{s-1} d_{\theta_k} \cdot \vec{\nu}q^{(k)}\right) \frac{d_i \cdot \vec{\nu}q^{(s)}}{c_i + d_i \cdot \nu^{(s)}} <? \left(\prod_{k=0}^{s-1} d_{\theta_k} \cdot \vec{\nu}q^{(k)}\right) \frac{d_j \cdot \vec{\nu}q^{(s)}}{c_j + d_j \cdot \nu^{(s)}}$$

and a single linear query of the type

$$d_{\theta_{s-1}} \cdot \vec{\nu}q^{(s-1)} <? 0.$$

In order to detect all hyperplanes H_i such that $\lambda_i = \lambda_\theta$ we can afford to compute the query $f(q) > g(q)$ for all query $f(q) < g(q)$ that we issue, and vice versa.

Note that, without further analysis, the queries can become n -linear as soon as we enter the $\frac{n}{k}$ th recursive step. ◀

B Algebraic computation trees

We consider *algebraic computation trees*, whose internal nodes are labeled with arithmetic ($r \leftarrow o_1 \text{ op } o_2, \text{ op} \in \{+, -, \times, \div\}$) and branching ($z : 0$) operations. We say that an algebraic computation tree T *realizes* an algorithm A if the paths from the root to the leaves of T correspond to the execution paths of A on all possible inputs $q \in \mathbb{R}^n$, where n is fixed. A leaf is labeled with the output of the corresponding execution path of A . Such a tree is

⁶ Note that if $c_i + d_i \cdot \nu = 0$ then $\lambda_i = 0$, we can check this beforehand for free.

well-defined if any internal node labeled $r \leftarrow o_1 \text{ op } o_2$ has outdegree 1 and is such that either $o_k = q_i$ for some i or there exists an ancestor $o_k \leftarrow x \text{ op } y$ of this node, and any internal node labeled $z : 0$ has outdegree 3 and is such that either $z = q_i$ for some i or there exists an ancestor $z \leftarrow x \text{ op } y$ of this node. In the algebraic computation tree model, we define the complexity $f(n)$ of an algorithm A to be the minimum depth of a well-defined computation tree that realizes A for inputs of size n .

In the algebraic computation tree model, we only count the operations that involve the input, that is, members of the input or results of previous operations involving the input. The following theorem follows immediately from the analysis of the linearity of queries

► **Theorem 15.** *The algebraic computation tree complexity of k -LDT is $\tilde{O}(n^3)$.*

Proof. We go through each step of Algorithm 2. Indeed, each k -linear query of step 1 can be implemented as $O(k)$ arithmetic operations, so step 1 has complexity $O(|\mathcal{N}|)$. The construction of the simplex in step 2 must be handled carefully. What we need to show is that each n -linear query we use can be implemented using $O(k)$ arithmetic operations. It is not difficult to see from the expressions given in Appendix A that a constant number of arithmetic operations and dot products suffice to compute the queries. A dot product in this case involves a constant number of arithmetic operations because the d_i are such that they each have exactly k non-zero components. The only expression that involves a non-constant number of operations is the product $\prod_{k=0}^s d_{\theta_k} \cdot \vec{v}q^{(k)}$, but this is equivalent to $(\prod_{k=0}^{s-1} d_{\theta_k} \cdot \vec{v}q^{(k)})(d_{\theta_s} \cdot \vec{v}q^{(s)})$ where the first factor has already been computed during a previous step and the second factor is of constant complexity. Since each query costs a constant number of arithmetic operations and branching operations, step 2 has complexity $O(n|\mathcal{N}|)$. Finally, steps 3 and 4 are free since they do not involve the input. The complexity of Algorithm 2 in this model is thus also $O(n^3 \log n \log |\mathcal{H}|)$. ◀

C Uniform random sampling

Theorem 4 requires us to pick a sample of the hyperplanes uniformly at random. Actually the theorem is a little stronger; we can draw each element of \mathcal{N} uniformly at random, only keeping distinct elements. This is not too difficult to achieve for k -LDT when the $\alpha_i, i \in [k]$ are all distinct: to pick a hyperplane of the form $\alpha_0 + \alpha_1 x_{i_1} + \alpha_2 x_{i_2} + \dots + \alpha_k x_{i_k} = 0$ uniformly at random, we can draw each $i_j \in [n]$ independently and there are n^k possible outcomes. However, in the case of k -SUM, we only have $\binom{n}{k}$ distinct hyperplanes. A simple dynamic programming approach solves the problem for k -SUM. For k -LDT we can use the same approach, once for each class of equal α_i .

► **Lemma 16.** *Given $n \in \mathbb{N}$ and $(\alpha_0, \alpha_1, \dots, \alpha_k) \in \mathbb{R}^{k+1}$, m independent uniform random draws of hyperplanes in \mathbb{R}^n with equations of the form $\alpha_0 + \alpha_1 x_{i_1} + \alpha_2 x_{i_2} + \dots + \alpha_k x_{i_k} = 0$ can be computed in time $O(mk^2 \log n)$ and preprocessing time $O(k^2 n)$.*

Proof. We want to pick an assignment $a = \{(\alpha_1, x_{i_1}), (\alpha_2, x_{i_2}), \dots, (\alpha_k, x_{i_k})\}$ uniformly at random. Note that all x_i are distinct while the α_j can be equal.

Without loss of generality, suppose $\alpha_1 \leq \alpha_2 \leq \dots \leq \alpha_k$. There is a bijection between assignments and lexicographically sorted k -tuples $((\alpha_1, x_{i_1}), (\alpha_2, x_{i_2}), \dots, (\alpha_k, x_{i_k}))$.

Observe that x_{i_j} can be drawn independently of $x_{i_{j'}}$, whenever $\alpha_j \neq \alpha_{j'}$. Hence, it suffices to generate a lexicographically sorted $|\chi|$ -tuple of x_i for each class χ of equal α_i .

Let $\omega(m, l)$ denote the number of lexicographically sorted l -tuples, where each element

comes from a set of m distinct x_i . We have

$$\omega(m, l) = \begin{cases} 1 & \text{if } l = 0 \\ \sum_{i=1}^m \omega(i, l-1) & \text{otherwise.} \end{cases}$$

To pick such a tuple $(x_{i_1}, x_{i_2}, \dots, x_{i_l})$ uniformly at random we choose $x_{i_l} = x_o$ with probability

$$P(x_{i_l} = x_o) = \begin{cases} 0 & \text{if } o > m \\ \frac{\omega(o, l-1)}{\omega(m, l)} & \text{otherwise} \end{cases}$$

that we append to a prefix $(l-1)$ -tuple (apply the procedure recursively), whose elements come from a set of o symbols. If $l = 0$ we just return the empty tuple.

Obviously, the probability for a given l -tuple to be picked is equal to $\frac{1}{\omega(m, l)}$.

Let X denote the partition of the α_i into equivalence classes, then the number of assignments is equal to $\prod_{\chi \in X} \omega(n, |\chi|)$. (Note that for k -SUM this is simply $\omega(n, k)$ since there is only a single class of equivalence.) For each equivalence class χ we draw independently a lexicographically sorted $|\chi|$ -tuple on n symbols using the procedure above. This yields a given assignment with probability $\frac{1}{\prod_{\chi \in X} \omega(n, |\chi|)}$. Hence, this corresponds to a uniform random draw over the assignments.

It is a well known fact that $\omega(n, k) = \binom{n+k-1}{k-1}$, hence each number we manipulate fits in $O(k \log n)$ bits, that is, $O(k)$ words. Moreover $\omega(n, k) = \omega(n-1, k) + \omega(n-1, k-1)$ so each $\omega(m, l)$ can be computed using a single addition on numbers of $O(k)$ words.

For given n and k , there are at most nk values $\omega(m, l)$ to compute, and for a given k -LDT instance, it must be computed only once. One way to perform the random draws is to compute the cumulative distribution functions of the discrete distributions defined above, then to draw x_{i_l} , we use binary search to find a generated random integer of $O(k)$ words in the cumulative distribution function. Computing the values $\omega(m, l)$ and all cumulative distributions functions can be done as a preprocessing step in $O(k^2 n)$ time. Assuming the generation of a random sequence of words takes linear time, performing a random draw takes time $O(k^2 \log n)$. ◀

D Proof of Lemma 9

► **Theorem 17** (Cramer's rule). *If a system of n linear equations for n unknowns, represented in matrix multiplication form $Ax = b$, has a unique solution $x = (x_1, x_2, \dots, x_n)^T$ then, for all $i \in [n]$,*

$$x_i = \frac{\det(A_i)}{\det(A)}$$

where A_i is A with the i th column replaced by the column vector b .

► **Lemma 18** (Meyer auf der Heide[27]). *The absolute value of the determinant of an $n \times n$ matrix $M = M_{i=1 \dots n, j=1 \dots n}$ with integer entries is an integer that is at most $C^n n^{\frac{n}{2}}$, where C is the maximum absolute value in M .*

Proof. The determinant of M must be an integer and is the volume of the hyperparallelepiped spanned by the row vectors of M , hence

$$|\det(M)| \leq \prod_{i=1}^n \sqrt{\sum_{j=1}^n M_{i,j}^2} \leq (\sqrt{nC^2})^n \leq C^n n^{\frac{n}{2}}.$$

◀

► **Lemma 19.** *The determinant of an $n \times n$ matrix $M = M_{i=1\dots n, j=1\dots n}$ with rational entries can be represented as a fraction whose numerators and denominators absolute values are bounded above by $(ND^{n-1})^n n^{\frac{n}{2}}$ and D^{n^2} respectively, where N and D are respectively the maximum absolute value of a numerator and a denominator in M .*

Proof. Let $\delta_{i,j}$ denote the denominator of $M_{i,j}$. Multiply each row M_i of M by $\prod_j \delta_{i,j}$. Apply Lemma 18. ◀

We can now proceed to the proof of Lemma 9.

Proof. Coefficients of the hyperplanes of the arrangement are constant rational numbers, those can be changed to constant integers (because each hyperplane has at most k nonzero coefficients). Let C denote the maximum absolute value of those coefficients.

Because of Theorem 17 and Lemma 18, vertices of the arrangement have rational coordinates whose numerators and denominators absolute values are bounded above by $C^n n^{\frac{n}{2}}$.

Given simplices whose vertices are vertices of the arrangement, hyperplanes that define the faces of those simplices have rational coefficients whose numerators and denominators absolute values are bounded above by $C^{2n^3} n^{n^3 + \frac{n}{2}}$ by Theorem 17 and Lemma 19. (Note that some simplices might be not fully dimensional, but we can handle those by adding vertices with coordinates that are not much larger than that of already existing vertices).

By applying Theorem 17 and Lemma 19 again, we obtain that vertices of the arrangement of those new hyperplanes (and thus vertices of S) have rational coefficients whose numerators and denominators absolute values are bounded above by $C^{4n^5} n^{2n^5 + n^3 + \frac{n}{2}}$. ◀