

Geometric clustering: fixed-parameter tractability and lower bounds with respect to the dimension

SERGIO CABELLO, University of Ljubljana and IMFM

PANOS GIANNOPOULOS, Humboldt-Universität zu Berlin

CHRISTIAN KNAUER, Freie Universität Berlin

DÁNIEL MARX, Budapest University of Technology and Economics

GÜNTER ROTE, Freie Universität Berlin

We study the parameterized complexity of the k -center problem on a given n -point set P in \mathbb{R}^d , with the dimension d as the parameter. We show that the rectilinear 3-center problem is fixed-parameter tractable, by giving an algorithm that runs in $O(n \log n)$ time for any fixed dimension d . On the other hand, we show that this is unlikely to be the case with both the Euclidean and rectilinear k -center problems for any $k \geq 2$ and $k \geq 4$ respectively. In particular, we prove that deciding whether P can be covered by the union of 2 balls of given radius or by the union of 4 cubes of given side length is $W[1]$ -hard with respect to d , and thus not fixed-parameter tractable unless $FPT=W[1]$. For the Euclidean case we also show that even an $n^{o(d)}$ -time algorithm does not exist, unless there is a $2^{o(n)}$ -time algorithm for n -variable 3SAT, i.e., the Exponential Time Hypothesis fails.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems

1. INTRODUCTION

A common type of facility location or clustering problem is the k -center problem, which is defined as follows: Given a set P of n points in a metric space and a positive integer k , find a set of k supply points such that the maximum distance between a point in P and its nearest supply point is minimized. For the cases of the (\mathbb{R}^d, L_2) and (\mathbb{R}^d, L_∞) -metric the problem is usually referred to as the *Euclidean* and *rectilinear* k -center respectively. Drezner [1995] describes many variations of the facility location problem and their numerous applications.

Research partially supported by the Slovenian Research Agency, project J1-7218 and program P1-0297, by a Magyary Zoltán post-doctoral research fellowship, and by the Hungarian National Research Fund (Grant Number OTKA 67651).

Part of these results (excluding Section 3) were presented at the 19th ACM-SIAM Symposium on Discrete Algorithms (SODA), San Francisco, in January 2008 [Cabello et al. 2008]. Part of this research was done during the Dagstuhl seminar ‘Geometric Networks and Metric Space Embeddings’, no. 06481, in November 2006.

Authors’ address: Sergio Cabello, IMFM & FMF, University of Ljubljana, Jadranska 19, SI-1000 Ljubljana, Slovenia, sergio.cabello@fmf.uni-lj.si. Panos Giannopoulos, Humboldt-Universität zu Berlin, Institut für Informatik, Unter den Linden 6, D-10099 Berlin, Germany, panos@mi.fu-berlin.de. Christian Knauer and Günter Rote, Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany, {knauer,rote}@inf.fu-berlin.de. Dániel Marx, Department of Computer Science and Information Theory, Budapest University of Technology and Economics, Budapest H-1521, Hungary, dmarx@cs.bme.hu.

k -center problems as well as other clustering problems can be formulated as geometric optimization problems and, as such, they have been studied extensively in the field of computational geometry; see, for example, the survey by Agarwal and Sharir [1998] and the references therein.

For solving a k -center problem, one usually looks at its corresponding *decision* problem: In the Euclidean k -center problem, one wants to decide whether P can be covered by the union of k balls of given radius, and if so, return such a covering; in the rectilinear case, a covering with k axis-aligned cubes of given size is sought. Once an algorithm for the decision problem is available, a solution for the k -center problem can be found using search techniques, e.g., binary search, parametric search, on a finite set of candidate values for the optimal size of the balls or cubes [Agarwal and Sharir 1998; Chan 1999a].

Efficient polynomial-time algorithms have been found for the planar k -center problem when k is a small constant [Chan 1999a; 1999b; Eppstein 1997; Nussbaum 1997; Segal 1999; Sharir and Welzl 1996]. Also, the Euclidean 1-center and rectilinear 1- and 2-center problems can be solved in polynomial time when d is part of the input [Megiddo 1990]. However, only $O(n^{O(kd)})$ -time algorithms are known when both k and d are part of the input, in particular, for $k \geq 2$ and $d > 2$ for the Euclidean case [Agarwal and Sharir 1998], and $k \geq 3$ and $d \geq 6$ for the rectilinear case [Assa and Katz 1999]. The fastest previously known algorithm for the rectilinear 3-center problem, due to Assa and Katz [1999], runs in $O(n^{\lfloor d/3 \rfloor} \log n)$ time.

As for lower bounds, both the Euclidean and rectilinear (decision) problems are NP-hard, for $d = 2$ when k is part of the input [Fowler et al. 1981; Megiddo and Supowit 1984], while, the Euclidean 2-center and rectilinear 3-center are NP-hard when d is part of the input [Megiddo 1990]. Moreover, a lower bound of $\Omega(n \log n)$ in the algebraic computation tree model has been shown for both the rectilinear 4-center problem in the plane [Sharir and Welzl 1996] and the rectilinear 3-center problem in 3-dimensional space [Hoffmann 2001; 2002].

The above NP-hardness results do not exclude the possibility of algorithms in which the exponent of n in the running time is independent of the parameter k or d or both. In terms of parameterized complexity theory (see below), the question is whether the problem is fixed-parameter tractable with respect to any of these two parameters. When k is considered as the parameter, Marx [2005] showed that this is most probably not the case for the rectilinear k -center problem, for any $d \geq 2$, by proving the respective decision problem to be W[1]-hard. We note here that Demaine et al. [2005] showed that the related (unweighted) (k, r) -center problem in planar graphs and map graphs, where one wants to find k vertices (centers) of the input graph such that every vertex of the graph is within distance at most r from some center, is fixed-parameter tractable when parameterized by k and r .

Parameterized Complexity. We review some basic definitions of parameterized complexity theory; for an introduction to the field, the reader is referred to the textbooks by Downey and Fellows [1999], and Flum and Grohe [2006]. Formally, a parameterized problem Π is a subset of $\Sigma^* \times \mathbb{N}$, where Σ is a fixed alphabet. An instance of Π is a pair (x, k) , where the second component k is called *the*

k	L_2	L_∞
1	P (weakly polynomial)	P
2	NP-hard [Megiddo 1990], W[1]-hard	P [Megiddo 1990]
3	NP-hard [Megiddo 1990], W[1]-hard	NP-hard [Megiddo 1990], FPT
≥ 4	NP-hard [Megiddo 1990], W[1]-hard	NP-hard [Megiddo 1990], W[1]-hard

Table I. Complexity classification of k -center problems for unbounded dimension.

parameter. A parameterized problem Π is *fixed-parameter tractable* if there is an algorithm that decides whether or not $(x, k) \in \Pi$ in $f(k) \cdot |(x, k)|^c$ time, where f is a computable function depending only on k , and c is a constant independent of k ; such an algorithm is (informally) said to run in *fpt-time*. The class of all fixed-parameter tractable problems is denoted by FPT. An infinite hierarchy of classes, the W-hierarchy, has been introduced for establishing fixed-parameter intractability. Its first level, W[1], can be thought of as the parameterized analog of NP: a parameterized problem that is hard for W[1] is not in FPT unless $\text{FPT} = \text{W}[1]$, which is considered highly unlikely under standard complexity theoretic assumptions. Hardness is sought via *fpt-reductions*: an *fpt-reduction* is an *fpt-time* Turing reduction from a problem Π , parameterized with k , to a problem Π' , parameterized with k' , such that $k' \leq g(k)$ for some computable function g .

Our Results. In this paper, we give a fine classification of the complexity of the k -center problem parameterized by the dimension d in the L_2 and L_∞ metric; see Table I. We show that the rectilinear 3-center problem can be solved in $O(6^d dn \log(dn))$ time, which is a considerable improvement over the $O(n^{\lfloor d/3 \rfloor} \log n)$ -time algorithm by Assa and Katz [1999]. Thus, the rectilinear 3-center problem is fixed-parameter tractable with respect to d . Our algorithm is based on two ingredients. First, we solve the corresponding *decision problem* in $O(6^d dn + dn \log n)$ time by a quite simple reduction to 2-satisfiability (2SAT). Second, we use the technique by Frederickson and Johnson [1984] to efficiently search among the candidate values for the optimal side length of the cubes. In view of the $O(n \log n)$ lower bound for $d = 3$ [Hoffmann 2001; 2002], our algorithm is asymptotically optimal for any $d \geq 3$.

On the negative side, we prove that both the Euclidean and rectilinear k -center (decision) problems are W[1]-hard with respect to d , for $k \geq 2$ and $k \geq 4$ respectively. For the Euclidean case, our reduction also implies that the problem cannot be solved in $O(n^{o(d)})$ time unless the Exponential Time Hypothesis fails.

2. THE RECTILINEAR 3-CENTER PROBLEM

In this section we show that the rectilinear 3-center problem is fixed-parameter tractable with respect to the dimension of the input point set.

- THEOREM 2.1.** a) *Given n points in d dimensions, we can decide whether they can be covered by three axis-aligned cubes of given side length m in $O(6^d \cdot dn + dn \log n)$ time.*
- b) *The smallest side length m for which the given points can be covered can be determined in $O(6^d \cdot dn \log(dn))$ time.*

PROOF. (a) We can assume w.l.o.g. that $m = 1$. Let $P = \{p_1, \dots, p_n\}$ be the

input point set. We denote the three cubes by A , B , and C . Each cube is the Cartesian product of d unit intervals.

The function x_j denotes the projection onto the j -coordinate axis. Therefore, $x_j(p)$ is the j -th coordinate of a point p and $x_j(A)$ is a unit interval for any cube A . Projecting the n points and the cubes on the j -th coordinate axis, we get n real numbers $x_j(p_u)$ and 3 unit intervals $x_j(A)$, $x_j(B)$, and $x_j(C)$ (whose positions are to be determined). We sort the coordinates of the points in each of the coordinate directions in $O(dn \log n)$ time.

We have a covering if we can assign every point p_u to one of the cubes (A , B , or C) such that, in each coordinate, this point is covered by the interval corresponding to the assigned cube.

In the following, we will consider the dimensions separately. We will look at the projection on each coordinate j and try to see by which interval a point can be covered in this coordinate. Let the minimum and maximum coordinate values be l_j and r_j .

If the diameter $r_j - l_j$ is at most one, we can, for example, align the three left interval endpoints with the leftmost point l_j . Then, in this coordinate, all points are covered by all intervals. This means that we can eliminate this coordinate from consideration. From now on, we will assume that all these irrelevant coordinates have been eliminated, and thus, the diameter in coordinate j is bigger than one. Then we can assume, w.l.o.g., that no interval sticks out to the left of l_j or to the right of r_j . On the other hand, these points must be covered by *some* interval. Thus we can make the following assumption:

In dimension j , one of the intervals $x_j(A)$, $x_j(B)$, $x_j(C)$ has its left endpoint aligned with the leftmost point l_j . Another interval has its right endpoint aligned with the rightmost point r_j . The third interval (the “middle” interval) lies between these two positions. Intuitively we can see the middle interval “floating” between l_j and r_j because its position is not yet determined. The boundary cases, where the middle interval coincides with the left or right interval, are permitted.

We can thus classify the solutions into 6^d *patterns*, according to the intervals ($x_j(A)$, $x_j(B)$, or $x_j(C)$) that are the left, middle, and right intervals in each coordinate direction. Formally, a pattern is represented as a sequence

$$(L_1, M_1, R_1), \dots, (L_d, M_d, R_d),$$

where each triplet (L_j, M_j, R_j) is a permutation of the three symbols A, B, C .

Let us restrict our attention to one fixed pattern. We now describe how to model this restricted covering problem as a logical satisfiability problem in conjunctive normal form, and decide whether such restricted covering exists in $O(dn)$ time.

We have $3n$ Boolean variables y_{Au} , y_{Bu} , y_{Cu} . The variable y_{Xu} represents the fact that point p_u is covered by box X , for $X = A, B, C$.

We have the n *covering clauses*

$$(y_{Au} \vee y_{Bu} \vee y_{Cu}), \tag{1}$$

for every $u \in \{1, \dots, n\}$, expressing the fact that every point is covered (by at least one box).

Let us now look at some dimension j , where $x_j(L_j)$, $x_j(M_j)$, $x_j(R_j)$ are the

left, middle, and right interval in dimension j according to the chosen pattern. (L_j, M_j, R_j is a permutation of A, B, C .)

The positions of the intervals $x_j(L_j)$ and $x_j(R_j)$ are fixed, and we only have to decide the position of the middle interval $x_j(M_j)$, that floats between l_j and r_j .

When $x_j(p_u) > l_j + 1$, the point p_u cannot be covered by the box L_j and we can put the following set of clauses with one literal:

$$(\neg y_{L_j u}), \quad (2)$$

for all u with $x_j(p_u) > l_j + 1$. A similar argument applies to the box R_j , and we can put the following set of clauses:

$$(\neg y_{R_j u}) \quad (3)$$

for all u with $x_j(p_u) < r_j - 1$. We can cover two points p_u and p_v with the box M_j only if the distance between $x_j(p_u)$ and $x_j(p_v)$ is at most one. Thus we add the following set of clauses:

$$(\neg y_{M_j u} \vee \neg y_{M_j v}), \quad (4)$$

for all $u, v \in \{1, \dots, n\}$ with $x_j(p_u) - x_j(p_v) > 1$.

The above model is a somewhat abbreviated story of the situation. For example, we have no variable to express explicitly the fact that box X covers point u in dimension j (in the projection). Nevertheless, the model captures the problem correctly:

LEMMA 2.2. *There is a covering conforming to the chosen pattern if and only if the clauses (1–4) are satisfiable.*

PROOF. Suppose we have a covering conforming to the chosen pattern. Set y_{Xu} to true if and only if point p_u is covered by box X . Then it is easy to check that all clauses are satisfied.

Conversely, assume that we have a Boolean assignment that satisfies all clauses. In each dimension j , the intervals $x_j(L_j)$ and $x_j(R_j)$ are already fixed, and we place the interval $x_j(M_j)$ as follows: we align its left endpoint with the leftmost point $x_j(p_u)$ (in dimension j) for which $y_{M_j u}$ is true. This defines the position of the boxes A, B, C .

For a point p_u the clauses (1) imply that at least one of y_{Au}, y_{Bu}, y_{Cu} is true. We have to show that, if y_{Xu} is true, then these chosen unit intervals for box X cover point p_u in every dimension.

If $X = L_j$ or $X = R_j$ in dimension j , the clauses (2) or (3) ensure that point p_u is covered in dimension j . Thus, suppose finally that $X = M_j$. The interval for M_j was chosen such that $x_j(p_u)$ does not lie to the left of $x_j(M_j)$. If $x_j(p_u)$ lies to the right of $x_j(M_j)$ it means that some point p_v , whose distance $x_j(p_u) - x_j(p_v)$ from p_u is bigger than 1, has also $y_{M_j v}$ true. This contradicts the clause (4). \square

All clauses except the clauses (1) contain at most two literals. We will now show that the clauses (1) can be eliminated, turning the problem into a 2-satisfiability problem, which can be solved in linear time.

Any of the clauses (2) or (3) effectively sets a variable to false, and it can be immediately used to eliminate a literal from one of the clauses (1). If we perform




Fig. 1. The points in the indicated region do not appear in a clause of the form (2–4).

this elimination for all literals, we end up with n modified covering clauses, each of which contains a proper subset of $\{y_{Au}, y_{Bu}, y_{Cu}\}$. (If we obtain an empty clause, we know that the problem is not satisfiable.) We refer to the resulting clauses, which contain at most two literals each, as the *reduced covering clauses*, and we denote them by (1').

LEMMA 2.3. *There is a covering conforming to the chosen pattern if and only if the clauses (1') and (2–4) are satisfiable.*

PROOF. The new set of clauses is weaker than the old one: It is derived by drawing logical conclusions (actually, some form of resolution), and omitting the clauses with three literals. Therefore, when the clauses (1–4) are satisfiable, the new set of clauses is satisfiable too.

Thus we only have to show that the clauses (1–4) are satisfiable whenever the reduced system of clauses is satisfiable.

A reduced clause (1') implies that the corresponding original clause is also satisfied. Consider now a clause (1) for a point p_u which remains intact during the reduction process. None of y_{Au} , y_{Bu} , and y_{Cu} , ever appears in a clause (2) or (3). In other words, in each dimension j , point p_u lies within distance 1 both of the leftmost point l_j and of the rightmost point r_j ; see Fig. 1. This means that point p_u is covered by all three intervals, no matter where the interval $x_j(M_j)$ is. (Here we are using the fact that the three cubes have equal size.)

On the logical level, none of y_{Au} , y_{Bu} , and y_{Cu} appears in the clauses (4), and thus they do not appear in negated form at all. We can thus satisfy the clause $(y_{Au} \vee y_{Bu} \vee y_{Cu})$ simply by setting all three variables to true. \square

Thus we have reduced the covering problem for a fixed pattern to an equivalent 2-SAT instance. There are $O(n)$ clauses of type (1'), $O(dn)$ clauses of types (2) and (3), but $O(dn^2)$ clauses of type (4).

The clauses of the last type can be replaced by $O(dn)$ clauses by introducing auxiliary variables, as follows: Let us look at a fixed dimension j . The $O(n^2)$ clauses of the form (4) involve the n variables $y_{M_j u}$, which we abbreviate by w_u , and we assume for simplicity of notation that the points are ordered by the j -th coordinate: $x_j(p_1) \leq x_j(p_2) \leq \dots \leq x_j(p_n)$.

The $O(n^2)$ clauses of the form (4) can be equivalently written as implications:

$$w_u \Rightarrow \neg w_v, \quad (5)$$

whenever $x_j(p_v) - x_j(p_u) > 1$.

We introduce auxiliary variables z_u that are intended to represent the fact that the interval for M_j starts left of $x_j(p_u)$ or at $x_j(p_u)$. Then we have the implications

$$w_u \Rightarrow z_u, \quad (6)$$

for $u = 1, \dots, n$, and

$$z_u \Rightarrow z_{u+1}, \quad (7)$$

for $u = 1, \dots, n - 1$.

Finally, for a given point p_v with $x_j(p_v) > l_j + 1$, let $\bar{u}(v)$ denote the largest index u such that $x_j(p_u) < x_j(p_v) - 1$, (i. e., $p_{\bar{u}(v)}$ is the right-most point with this property). Then we add the $O(n)$ clauses

$$z_{\bar{u}(v)} \Rightarrow \neg w_v, \quad (8)$$

for all $v = 1, \dots, n$ with $x_j(p_v) > l_j + 1$. We have omitted the reference to j for the variables w and z , but it should be kept in mind that this procedure has to be carried out for each dimension j separately.

LEMMA 2.4. *For given values of the variables w_1, \dots, w_n , the clauses (5) are satisfied if and only if there is a truth assignment for the variables z_1, \dots, z_n , that satisfies (6–8).*

PROOF. If we have a truth assignment w_1, \dots, w_n satisfying the clauses (5), we set $z_u := w_1 \vee w_2 \vee \dots \vee w_u$. Then (6) and (7) are satisfied by construction. To prove (8), assume for contradiction that w_v and $z_{\bar{u}(v)}$ are true, for some v . By the definition of $z_{\bar{u}(v)}$, there is some true w_u with $u \leq \bar{u}(v)$. Since we have $x_j(p_u) \leq x_j(p_{\bar{u}(v)}) < x_j(p_v) - 1$ and w_u, w_v are true, then w_u, w_v violate (5).

Conversely, assume that (6–8) is fulfilled, and let us prove (5) for each pair u, v with $x_j(p_u) < x_j(p_v) - 1$. The clauses (8) include the clause $z_{\bar{u}(v)} \Rightarrow \neg w_v$, and from the definition of $\bar{u}(v)$ we have $u \leq \bar{u}(v) < v$. Thus, the chain of implications $w_u \Rightarrow z_u \Rightarrow z_{u+1} \Rightarrow \dots \Rightarrow z_{\bar{u}(v)} \Rightarrow \neg w_v$ proves (5). \square

We have reduced the number of clauses to $O(dn)$, and each clause has at most two literals. The clauses can be generated in $O(dn)$ time if the input coordinates are sorted in each dimension, and the satisfiability of these clauses can be tested in $O(dn)$ time as well. This procedure has to be repeated for each of the 6^d patterns. This concludes the proof of part (a) of Theorem 2.1.

(b) The minimum side length m for which the given points are covered is one of the $O(dn^2)$ pairwise distances $|x_j(p_u) - x_j(p_v)|$. We initially sort in $O(dn \log n)$ time the input coordinates in each dimension. For each dimension j , assuming for simplicity of notation that the points are indexed such that $x_j(p_1) \leq x_j(p_2) \leq \dots \leq x_j(p_n)$, we define an $n \times n$ matrix $\Delta^j = \{\delta_{uv}^j\}$ with entries $\delta_{uv}^j = x_j(p_u) - x_j(p_v)$. Each matrix Δ^j is a sorted matrix: each column has nondecreasing values and each row has non-increasing values. The matrices $\Delta^1, \dots, \Delta^d$ are not constructed explicitly, but only some of their entries will be evaluated. Let Δ denote the multiset of dn^2 entries in $\Delta^1, \dots, \Delta^d$. Clearly, the sought value m is one of the values in Δ .

Frederickson and Johnson [1984] showed how to select for any $1 \leq k \leq dn^2$ the k -th largest entry in the collection of sorted matrices $\Delta^1, \dots, \Delta^d$ evaluating $O(dn)$

entries. In our scenario, any desired entry δ_{uv}^j can be obtained in $O(1)$ time, after the initial sorting of the coordinates. Thus, we can find the k -th largest value of Δ in $O(dn)$ time.

We can now perform a binary search for m on the entries of Δ . Since Δ has dn^2 values, the binary search requires $O(\log(dn))$ calls to the selection procedure and applications of the decision algorithm from part (a). Therefore, each of the $O(\log(dn))$ steps of the binary search requires $O(6^d \cdot dn)$ time, after the initial sorting of the coordinates.

3. THE RECTILINEAR 4-CENTER PROBLEM

In this section we show that the rectilinear 4-center decision problem, parameterized with the dimension d , is W[1]-hard. The problem asks whether n given points can be covered by 4 axis-aligned cubes of a given side length. Without loss of generality, we will only consider unit cubes throughout.

Our reduction builds point sets from basic building blocks with the *join* operation: Let P_1 and P_2 be two sets of points in d_1 and d_2 dimensions, respectively. The *join* of P_1 and P_2 is a set P of $|P_1| + |P_2|$ points in $d_1 + d_2$ dimensions that are obtained by padding the points in P_1 with d_2 zero coordinates at the end and by padding the points in P_2 with d_1 zero coordinates at the beginning. The join of multiple point sets is defined similarly.

We say that a cube is a *0-covering* cube if it contains the origin. A set C of cubes is a *0-covering* set if every cube in C contains the origin. In order to make the gadget construction in the hardness proof easier, we consider a variant of the problem where the 4 cubes have to form a 0-covering set. In general, it is not true that if P_1 can be covered by 4 cubes and P_2 can be covered by 4 cubes, then their join can be also covered by 4 cubes. However, this is true if P_1 and P_2 can be covered by 0-covering cubes:

PROPOSITION 3.1. *The join P of point sets P_1, \dots, P_ℓ can be covered by a 0-covering set of k unit cubes if and only if each set P_j can be covered by a 0-covering set of k unit cubes.*

PROOF. Let d_i be the dimension of the point set P_i and let $d = \sum_{i=1}^{\ell} d_i$. For each $i = 1, \dots, \ell$, let $C_i = \{c_{i,1}, \dots, c_{i,k}\}$ be a set of d_i -dimensional 0-covering unit cubes that cover P_i . We claim that the set $C = \{c_1, \dots, c_k\}$ of d -dimensional 0-covering unit cubes $c_j = c_{1,j} \times c_{2,j} \times \dots \times c_{\ell,j}$ covers P . Let x be a point in P_i that is covered by $c_{i,j}$ and let x' be the corresponding point in the join P . It is easy to see that c_j covers x' : this follows from the facts that $c_{i,j}$ covers x and $c_{i',j}$ is a 0-covering cube for $i' \neq i$.

The other direction is obvious. \square

To simplify the gadget construction further, we consider a generalization of the problem, where each point has a prescribed list of cubes and a point has to be covered with one of the cubes on its list. For a set P of n points in d dimensions and a subset $L(p) \subseteq \{1, 2, 3, 4\}$ for every point $p \in P$, the *constrained* rectilinear 4-center decision problem asks for a 4-tuple (c_1, c_2, c_3, c_4) of d -dimensional 0-covering unit cubes such that every point $p \in P$ is contained in c_r for some $r \in L(p)$. If a quadruple (c_1, c_2, c_3, c_4) of 0-covering unit cubes covers every

point in the sense defined above, then we say that (c_1, c_2, c_3, c_4) is a *proper cover* of P . The join operation extends to point sets with constraint lists $L(p)$, and Proposition 3.1 remains true for this constrained version of the problem. (In contrast to Proposition 3.1, we restrict the following statement to coverings by 4 cubes; the proof carries over almost verbatim.)

PROPOSITION 3.2. *Let P_1, \dots, P_ℓ be point sets with constraint lists, and let P be their join. Then there is a proper cover for P if and only if each set P_j has a proper cover. \square*

We will prove hardness for the constrained rectilinear 4-center decision problem parameterized with d , which, as shown below, is not harder than the original version.

LEMMA 3.3. *There is an fpt-reduction from the the constrained rectilinear 4-center problem to the rectilinear 4-center problem, with respect to the dimension d .*

PROOF. Let P be a set of points in d dimensions with constraint lists $L(p)$ for each $p \in P$. Let us augment each point with 4 new coordinates; denote these extra coordinates of a point p with $x_i(p)$ for $i = 1, 2, 3, 4$. For every point $p \in P$, we define

$$x_i(p) = \begin{cases} 0 & \text{if } i \notin L(p), \\ 1 & \text{if } i \in L(p) \end{cases}$$

for $i = 1, 2, 3, 4$. (We will have no need to refer to the d original coordinates; therefore we number the additional coordinates starting from 1.) Let us add 4 new points p_1, p_2, p_3, p_4 such that

$$x_i(p_j) = \begin{cases} 2 & \text{if } i = j, \\ 0 & \text{if } i \neq j \end{cases}$$

for $1 \leq i, j \leq 4$ and every other coordinate is 0. Denote by P' this set of $|P| + 4$ points in $d + 4$ dimensions.

We claim that P' can be covered with 4 unit cubes if and only if the constrained rectilinear 4-center decision problem on P has a solution.

Let (c_1, c_2, c_3, c_4) be a proper cover for P . Augment c_r to the 4 extra coordinates by defining the projection of c_r to the i -th extra coordinate to be

$$x_i(c_r) = \begin{cases} [1, 2] & \text{if } i = r, \\ [0, 1] & \text{if } i \neq r. \end{cases}$$

It is clear that the resulting unit cubes c'_1, c'_2, c'_3, c'_4 cover the $|P|$ points in P' that were created from P . Furthermore, c'_r covers p_r : this follows from the definitions of $x_i(p_r), x_i(c_r)$ above and from the fact that c_r is 0-covering.

Conversely, let us assume that there are 4 unit cubes that cover the points in P' . Observe that two extra points p_r and p_s must be covered by different cubes since $|x_r(p_r) - x_r(p_s)| = 2 > 1$ if $r \neq s$, $1 \leq r, s \leq 4$. Let us denote the cube that covers p_r by c'_r . Since $x_r(p_r) = 2$, we have $x_r(c'_r) = [a, a + 1]$ with $a \geq 1$. Let p

be a point in P and let p' be the corresponding point in P' . If c'_r covers p' , then $r \in L(p)$: otherwise $x_r(p) = 0$ would not be contained in $x_r(c'_r)$. Let c_r be the d -dimensional cube obtained from c'_r by projecting out the 4 extra dimensions. Clearly, c_1, c_2, c_3, c_4 cover every point in P ; more precisely, for every $p \in P$, there is a $r \in L(p)$ such that c_r covers p . Furthermore, c_r is 0-covering: c'_r covers p_r and p_r is nonzero only in the 4 extra coordinates. \square

We now come to the main lemma of this section: proving that the constrained rectilinear 4-center problem is $W[1]$ -hard. Let us give a short intuitive overview of the proof. The reduction is from the k -clique problem. The first ingredient of the reduction is the construction of the selection gadgets. The vertex selection gadget ensures that among a set of n points there is at least one which is covered only by cube c_4 in a solution. The n possible choices of this point will correspond to the n possible way of choosing a vertex of the clique. In the reduction, we build k vertex selection gadgets, for selecting k vertices that are supposed to form a clique. Analogously, the edge selection gadget constrains a set of m points in such a way that at least one of the points is covered only by cube c_1 ; the choice of this point corresponds to selecting an edge of the clique. There are $\binom{k}{2}$ edge selection gadgets, for selecting the edges of the clique. The second ingredient is the construction of the incidence testing gadgets. The role of these gadgets is to ensure that the selected vertices are compatible with the selected edges: the endpoints of the selected edges are exactly the selected vertices. More precisely, the incident testing gadget on a set A of n points and a set B of m points constrains these points in such a way that if the i -th point of A is covered by cube c_4 and the j -th point of B is covered by cube c_1 , then vertex i is an endpoint of the j -th edge (for some given numbering of the vertices and edges of the graph.) The third ingredient of the proof is finding a way of joining the gadgets such that “points with the same meaning” in the different gadgets are covered by the same cubes. We will use here the join operation defined above and the fact that the cubes are 0-covering (recall that the definition of the constrained problem requires that the cubes in the solution are 0-covering).

LEMMA 3.4. *The constrained rectilinear 4-center decision problem is $W[1]$ -hard with respect to the dimension d .*

PROOF. The proof is by an fpt-reduction from the parameterized k -clique problem, which is $W[1]$ -complete with respect to k [Flum and Grohe 2006]. Let $[n] = \{1, \dots, n\}$. We look for a clique of size k in a graph $G([n], E)$, with $|E| = m$; for convenience we identify E with the set $[m]$. Each edge t has an endpoint $\text{low}(t)$ with smaller value and an endpoint $\text{up}(t)$ with larger value; we will call them the *lower endpoint* and the *upper endpoint*. We also define a value $\epsilon := 1/(10n + 10m)$.

A k -clique in G will be represented as a mapping from the vertices of a complete graph of order k (i. e., from the integers $p = 1, \dots, k$) to the vertices of G , and from the edges of the complete graph (i. e., from the pairs (p, q) with $1 \leq p < q \leq k$) to E . In the reduction we construct a point set P that consists of several *gadgets*: each gadget is a point set having certain properties and P is obtained by combining the gadgets with the join operation. We have k gadgets that select the k vertices of the clique, $\binom{k}{2}$ gadgets that select the edges of the clique, and $2\binom{k}{2}$ incidence

testing gadgets that ensure the consistency of these selections.

3.1 Vertex selection gadget

The vertices of the clique will be represented by k three-dimensional vertex selection gadgets. Each gadget consists of the following $4n + 6$ points:

- $x_i = ((i + 1)\epsilon, (i - 1)\epsilon - 1, 0)$ and $L(x_i) = \{1, 2, 4\}$ for $1 \leq i \leq n$.
- $z_i^1 = (i\epsilon, (i - 1)\epsilon - 1, 0)$ and $L(z_i^1) = \{1, 2\}$ for $1 \leq i \leq n$.
- $z_i^2 = (0, i\epsilon, (i - 1)\epsilon - 1)$ and $L(z_i^2) = \{2, 3\}$ for $1 \leq i \leq n$.
- $z_i^3 = ((i - 1)\epsilon - 1, 0, i\epsilon)$ and $L(z_i^3) = \{1, 3\}$ for $1 \leq i \leq n$.
- $u_1 = (n\epsilon - 1, 0, 0)$ and $v_1 = (\epsilon, 0, 0)$ with $L(u_1) = L(v_1) = \{1\}$.
- $u_2 = (0, n\epsilon - 1, 0)$ and $v_2 = (0, \epsilon, 0)$ with $L(u_2) = L(v_2) = \{2\}$.
- $u_3 = (0, 0, n\epsilon - 1)$ and $v_3 = (0, 0, \epsilon)$ with $L(u_3) = L(v_3) = \{3\}$.

Each gadget has n possible “states”: there is some point x_i that is covered only by cube c_4 ; the choice of this point corresponds to the vertex that is selected by this gadget. To achieve this effect, the points z_i^1, z_i^2, z_i^3 set up a “cycle of implications” in the following sense. If c_1 covers many of the z_i^1 vertices, then it can cover only a few of z_i^3 vertices, which means that c_3 has to cover many of these vertices. But this implies that c_3 covers only a few of the z_i^2 vertices, which in turn implies that c_2 covers many of the z_i^2 vertices, but only few of the z_i^1 . Thus all three cubes have to be perfectly aligned if c_1 and c_2 cover all the z_i^1 vertices. Finally, looking at the x_i ’s (which are almost the same as the z_i^1 vertices), it turns out that even if c_1 and c_2 completely cover the z_i^1 vertices, there is one x_i that is covered by neither c_1 nor c_2 . The role of the u_1, v_1 , etc. vertices is only to give some simple bounds on the possible locations of the cubes.

LEMMA 3.5. *The vertex gadget has the following properties:*

- i) *For every $1 \leq s \leq n$, there is a proper cover (c_1, c_2, c_3, c_4) of P such that $x_s \in c_4$, and $x_i \in c_1$ or $x_i \in c_2$ for every $i \neq s$.*
- ii) *In every proper cover (c_1, c_2, c_3, c_4) , there is a $1 \leq s \leq n$ such that x_s is only covered by c_4 , among the cubes c_1, c_2, c_4 to which x_s is constrained.*

PROOF. (ii) We begin with the proof of the second statement, which shows the gadget “at work”. Assume that we have a proper cover (c_1, c_2, c_3, c_4) . We can assume that the coordinates of the cubes are integer multiples of ϵ . For $d = 1, 2, 3$, let $[p_d\epsilon - 1, p_d\epsilon]$ be the projection of c_d to the d -th coordinate. Since c_j covers both u_j and v_j , we have $1 \leq p_d \leq n$. If $p_2 > p_1$, then $z_{p_2}^1$ is covered by neither c_1 nor c_2 : cube c_1 does not cover $z_{p_2}^1$ in the first coordinate and cube c_2 does not cover $z_{p_2}^1$ in the second coordinate. Thus $p_1 \geq p_2$. Similar arguments show that $p_1 \geq p_2 \geq p_3 \geq p_1$, thus there is equality throughout; let $s := p_1 = p_2 = p_3$. Now point x_s is covered by neither c_1 nor c_2 .

(i) For a given s , we define the 4 cubes as follows:

$$\begin{aligned} c_1 &= [s\epsilon - 1, s\epsilon] \times [-1, 0] \times [0, 1] & c_2 &= [0, 1] \times [s\epsilon - 1, s\epsilon] \times [-1, 0] \\ c_3 &= [-1, 0] \times [0, 1] \times [s\epsilon - 1, s\epsilon] & c_4 &= [0, 1] \times [-1, 0] \times [0, 1] \end{aligned}$$

It can be verified that these 4 cubes cover every point. For example, z_i^1 is covered by c_1 for $i \leq s$ and by c_2 for $i > s$. Every x_i is covered by c_4 , but point x_i is also covered by c_1 for $i < s$ and by c_2 for $i > s$. Thus we have a proper cover where $x_i \in c_1$ or $x_i \in c_2$ for every $i \neq s$. Furthermore, for $d = 1, 2, 3$, cube c_d covers both u_d and v_d . \square

3.2 Edge selection gadget

The edge selection gadget is very similar to the vertex selection gadget, but it has m states and the cubes 1 and 4 are exchanged in the lists. That is, the gadget consists of the following $4m + 6$ points:

- $y_j = ((j + 1)\epsilon, (j - 1)\epsilon - 1, 0)$ and $L(y_j) = \{1, 2, 4\}$ for $1 \leq j \leq m$.
- $z_j^1 = (j\epsilon, (j - 1)\epsilon - 1, 0)$ and $L(z_j^1) = \{2, 4\}$ for $1 \leq j \leq m$.
- $z_j^2 = (0, j\epsilon, (j - 1)\epsilon - 1)$ and $L(z_j^2) = \{2, 3\}$ for $1 \leq j \leq m$.
- $z_j^3 = ((j - 1)\epsilon - 1, 0, j\epsilon)$ and $L(z_j^3) = \{3, 4\}$ for $1 \leq j \leq m$.
- $u_1 = (m\epsilon - 1, 0, 0)$ and $v_1 = (\epsilon, 0, 0)$ with $L(u_1) = L(v_1) = \{4\}$.
- $u_2 = (0, m\epsilon - 1, 0)$ and $v_2 = (0, \epsilon, 0)$ with $L(u_2) = L(v_2) = \{2\}$.
- $u_3 = (0, 0, m\epsilon - 1)$ and $v_3 = (0, 0, \epsilon)$ with $L(u_3) = L(v_3) = \{3\}$.

The gadget selects one edge: there is one vertex y_i that is covered only by cube c_1 .

LEMMA 3.6. *The edge gadget has the following properties:*

- i) *For every $1 \leq t \leq m$, there is a proper cover (c_1, c_2, c_3, c_4) of P such that $y_t \in c_1$, and $y_j \in c_2$ or $y_j \in c_4$ for every $j \neq t$.*
- ii) *In every proper cover (c_1, c_2, c_3, c_4) , there is a $1 \leq t \leq m$ such that y_t is only covered by c_1 , among the cubes c_1, c_2, c_4 to which y_t is constrained. \square*

3.3 Incidence testing gadget

The role of the incidence testing gadget is to ensure that the edge connecting two vertices of the clique is incident to these vertices: More precisely, the edge j that is chosen in the edge selection gadget for the edge between the p -th and the q -th vertex of the clique ($1 \leq p, q \leq k$) must be incident to the vertices i and i' that are chosen in the p -th and the q -th vertex selection gadget. There are two types of incident testing gadgets: the first one tests whether the lower endpoint of a selected edge is the same as a selected vertex, while the second tests the upper endpoint. The lower incidence testing gadget consists of the following $5n + m + 6$ points in six dimensions:

- $x_i = (i\epsilon, i\epsilon - 1, 0, 0, 0, 0)$ for $1 \leq i \leq n$ with $L(x_i) = \{1, 2, 4\}$,
- $y_j = (0, 0, 0, \text{low}(j)\epsilon, \text{low}(j)\epsilon - 1, 0)$ for $1 \leq j \leq m$ with $L(y_j) = \{1, 2, 4\}$,
- $z_i^1 = (0, (i + 1)\epsilon, i\epsilon - 1, 0, 0, 0)$ for $1 \leq i \leq n$ with $L(z_i^1) = \{3, 4\}$,
- $z_i^2 = (0, 0, (i + 1)\epsilon, i\epsilon - 1, 0, 0)$ for $1 \leq i \leq n$ with $L(z_i^2) = \{1, 3\}$,
- $z_i^3 = (0, 0, 0, 0, (i + 1)\epsilon, i\epsilon - 1)$ for $1 \leq i \leq n$ with $L(z_i^3) = \{1, 3\}$,
- $z_i^4 = (i\epsilon - 1, 0, 0, 0, 0, (i + 1)\epsilon)$ for $1 \leq i \leq n$ with $L(z_i^4) = \{3, 4\}$,
- $u_1 = (0, 0, 0, n\epsilon - 1, n\epsilon - 1, 0)$ and $v_1 = (0, 0, 0, \epsilon, \epsilon, 0)$ with $L(u_1) = L(v_1) = \{1\}$.

— $u_3 = (0, 0, n\epsilon - 1, 0, 0, n\epsilon - 1)$ and $v_3 = (0, 0, \epsilon, 0, 0, \epsilon)$ with $L(u_3) = L(v_3) = \{3\}$.
 — $u_4 = (n\epsilon - 1, n\epsilon - 1, 0, 0, 0, 0)$ and $v_4 = (\epsilon, \epsilon, 0, 0, 0, 0)$ with $L(u_4) = L(v_4) = \{4\}$.

The points x_i and y_j will be associated to corresponding points in vertex and edge selection gadgets, as described later in Section 3.4.

Similarly to the vertex/edge selection gadgets, the points form a cycle of implications, but this time in a somewhat more complicated way. The points z_i^1, z_i^2 and coordinates 2 and 3 ensure that if c_4 cover an x_i with small i , then c_1 can cover only an y_j with small $\text{low}(j)$. The points z_i^3, z_i^4 and coordinates 5 and 6 ensure that if c_4 cover an x_i with large i , then c_1 can cover only a y_j with large $\text{low}(j)$, making the connection tight.

LEMMA 3.7. *The lower incidence testing gadget has the following properties:*

- i) *For every $1 \leq t \leq m$ and $s = \text{low}(t)$, there is a proper cover (c_1, c_2, c_3, c_4) of P such that $x_i \in c_1 \cap c_2$ for every $1 \leq i \leq n$, $y_j \in c_2 \cap c_4$ for every $1 \leq j \leq m$, $x_s \in c_4$, and $y_t \in c_1$.*
- ii) *In every proper cover (c_1, c_2, c_3, c_4) , if $x_s \in c_4$ for some $1 \leq s \leq n$ and $y_t \in c_1$ for some $1 \leq t \leq m$, then $s = \text{low}(t)$.*

PROOF. (ii) Again we start with the second property. Assume that (c_1, c_2, c_3, c_4) is a proper cover. Without loss of generality, it can be assumed that the coordinates of the cubes are integer multiples of ϵ ; let $c_r = \prod_{d=1}^6 [a_{rd}\epsilon - 1, a_{rd}\epsilon]$ for $r = 1, 2, 3, 4$. We are interested in $a_{14}, a_{15}, a_{33}, a_{36}, a_{41}, a_{42}$, so the cubes have the structure

$$\begin{aligned} c_1 &= [*] \times [*] \times [*] \times [a_{14}\epsilon - 1, a_{14}\epsilon] \times [a_{15}\epsilon - 1, a_{15}\epsilon] \times [*] \\ c_3 &= [*] \times [*] \times [a_{33}\epsilon - 1, a_{33}\epsilon] \times [*] \times [*] \times [a_{36}\epsilon - 1, a_{36}\epsilon] \\ c_4 &= [a_{41}\epsilon - 1, a_{41}\epsilon] \times [a_{42}\epsilon - 1, a_{42}\epsilon] \times [*] \times [*] \times [*] \times [*], \end{aligned}$$

where $[*]$ denotes an arbitrary interval. From the fact that c_r covers u_r and v_r for $r = 1, 3, 4$, it follows that $1 \leq a_{14}, a_{15}, a_{33}, a_{36}, a_{41}, a_{42} \leq n$. Observe that $a_{15} \geq a_{36}$: otherwise $z_{a_{15}}^3$ would be covered by neither c_1 (because of the fifth coordinate) nor c_3 (because of the sixth coordinate). Also, $a_{42} \geq a_{33}$, because otherwise $z_{a_{42}}^1$ would be covered by neither c_3 (because of the third coordinate) nor by c_4 (because of the second coordinate). Similar arguments show that $a_{36} \geq a_{41}$ (because of $z_{a_{36}}^4$) and $a_{33} \geq a_{14}$ (because of $z_{a_{33}}^2$). Assume now that $x_s \in c_4$ and $y_t \in c_1$ for some s and t . We have that $a_{14} \geq \text{low}(t) \geq a_{15}$ (because of the fourth and fifth coordinates) and $a_{41} \geq s \geq a_{42}$ (because of the first two coordinates). Thus we have the following chain of inequalities

$$\text{low}(t) \geq a_{15} \geq a_{36} \geq a_{41} \geq s \geq a_{42} \geq a_{33} \geq a_{14} \geq \text{low}(t),$$

which implies that $\text{low}(t) = s$, as required.

(i) For some $1 \leq t \leq m$, let $s = \text{low}(t)$ be the lower endpoint of t . Consider the following cubes:

$$\begin{aligned} c_1 &= [0, 1] \times [-1, 0] \times [0, 1] \times [s\epsilon - 1, s\epsilon] \times [s\epsilon - 1, s\epsilon] \times [-1, 0] \\ c_2 &= [0, 1] \times [-1, 0] \times [0, 1] \times [0, 1] \times [-1, 0] \times [0, 1] \\ c_3 &= [-1, 0] \times [0, 1] \times [s\epsilon - 1, s\epsilon] \times [-1, 0] \times [0, 1] \times [s\epsilon - 1, s\epsilon] \\ c_4 &= [s\epsilon - 1, s\epsilon] \times [s\epsilon - 1, s\epsilon] \times [-1, 0] \times [0, 1] \times [-1, 0] \times [0, 1] \end{aligned}$$

It is clear that both cubes c_1, c_2 cover every x_i ; both cubes c_2, c_4 cover every y_j ; cube c_4 covers x_s ; and cube c_1 covers y_t . For $i < s$, z_i^3 is covered by c_1 ; z_i^4 and z_i^2 are covered by c_3 ; and z_i^1 is covered by c_4 . For $i \geq s$, z_i^3 is covered by c_3 ; z_i^4 is covered by c_4 ; z_i^2 is covered by c_1 ; and z_i^1 is covered by c_3 . Finally, for $r = 1, 3, 4$, points u_r and v_r are covered by c_r . \square

By replacing “lower” with “upper” everywhere, we get the upper incidence testing gadget in a symmetrical way:

LEMMA 3.8. *The upper incidence testing gadget has the following properties:*

- i) *For every $1 \leq t \leq m$ and $s = \text{up}(t)$, there is a proper cover (c_1, c_2, c_3, c_4) of P such that $x_i \in c_1 \cap c_2$ for every $1 \leq i \leq n$, $y_j \in c_2 \cap c_4$ for every $1 \leq j \leq m$, $x_s \in c_4$, and $y_t \in c_1$.*
- ii) *In every proper cover (c_1, c_2, c_3, c_4) , if $x_s \in c_4$ for some $1 \leq s \leq n$ and $y_t \in c_1$ for some $1 \leq t \leq m$, then $s = \text{up}(t)$.* \square

3.4 The construction

Given a graph G with n vertices and m edges and an integer k , we construct a point set P by taking the join of the following point sets:

- k copies of the vertex selection gadget (denote these sets by VS_p for $1 \leq p \leq k$),
- $\binom{k}{2}$ copies of the edge selection gadget (denote these sets by $ES_{p,q}$ for $1 \leq p < q \leq k$),
- $\binom{k}{2}$ copies of the lower incidence testing gadget (denote these sets by $TL_{p,q}$ for $1 \leq p < q \leq k$), and
- $\binom{k}{2}$ copies of the upper incidence testing gadget (denote these sets by $TU_{p,q}$ for $1 \leq p < q \leq k$).

Thus P contains $k(4n + 6) + \binom{k}{2}(4m + 6 + 2m + 10n + 12)$ points in $3k + 3\binom{k}{2} + 6\binom{k}{2} + 6\binom{k}{2}$ dimensions. By Proposition 3.2, Lemmas 3.5–3.8 hold for each gadget in the combined set P , but so far, the gadgets are completely independent of each other. (Every coordinate is nonzero only in one gadget.) We connect the gadgets by requiring that certain pairs of points are covered by the same cube.

For $1 \leq i \leq n$ and $1 \leq p \leq k$, let set V_{pi} contain the following k points:

- point x_i of VS_p ,
- point x_i in $TU_{q,p}$ for $1 \leq q < p$, and
- point x_i in $TL_{p,q}$ for $p < q \leq k$.

For $1 \leq j \leq m$ and $1 \leq p < q \leq k$, let set E_{pqj} contain the following 3 points:

- point y_j of $ES_{p,q}$,
- point y_j in $TU_{p,q}$, and
- point y_j in $TL_{p,q}$.

If vertex i is chosen as the p -th vertex of the clique, this will be represented by the fact that cube c_4 covers V_{pi} . Similarly, if the edge between the p -th and q -th vertices of the clique is edge j , then this is represented by the fact that c_1 covers E_{pqj} . In order to ensure that the correspondence works in both directions, we make the following additional requirement on the way the cubes cover the points:

DEFINITION 3.9. *Given a proper cover (c_1, c_2, c_3, c_4) , we say that a point set J is simultaneously covered if there is an $r \in \bigcap_{x \in J} L(x)$ such that c_r covers every $x \in J$. A proper cover of P is consistent if for every $1 \leq i \leq n$ and $1 \leq p \leq k$, the set V_{pi} is simultaneously covered, and for every $1 \leq j \leq m$ and $1 \leq p < q \leq k$, the set E_{pqj} is simultaneously covered.*

As we shall see, by combining points that must be simultaneously covered to new points we can enforce the requirement that the cover is consistent. Before this, we prove the correspondence between consistent covers and the cliques of G .

LEMMA 3.10. *The point set P has a consistent proper cover (c_1, c_2, c_3, c_4) if and only if G has a clique of size k .*

PROOF. Assume that (c_1, c_2, c_3, c_4) is consistent. Since it is a proper cover of VS_p , for every $1 \leq p \leq k$ there is a value v_p such that point x_{v_p} of VS_p is covered only by c_4 (Lemma 3.5). Furthermore, for every $1 \leq p < q \leq k$, there is a value e_{pq} such that point $y_{e_{pq}}$ of ES_{pq} is covered only by c_1 . We claim that v_1, \dots, v_k is a clique of G with e_{pq} being the edge connecting v_p and v_q . Let us show that v_p is the lower endpoint of e_{pq} . The definition of consistency implies that point x_{v_p} of $TL_{p,q}$ is also covered by c_4 (since this point and point x_{v_p} of VS_p are in the simultaneously covered set V_{p,v_q} and the latter point is covered only by c_4). Similarly, point $y_{e_{pq}}$ of $TL_{p,q}$ is covered by c_1 . Lemma 3.7 implies that v_p is the lower endpoint of e_{pq} . Similarly, $TU_{p,q}$ ensures that v_q is the upper endpoint of e_{pq} .

Assume now that $v_1 < v_2 < \dots < v_k$ is a clique in G ; let e_{pq} be the edge connecting v_p and v_q . By Lemma 3.5, for every $1 \leq p \leq k$, point set VS_p has a proper cover where $x_{v_p} \in c_4$, and $x_i \in c_1$ or $x_i \in c_2$ for every $i \neq v_p$. By Lemma 3.6, for every $1 \leq p < q \leq k$, $ES_{p,q}$ has a proper cover where $y_{e_{pq}} \in c_1$, and $y_j \in c_2$ or $y_j \in c_4$ for every $j \neq e_{pq}$. By Lemmas 3.7 and 3.8, for every $1 \leq p < q \leq k$, TL_{pq} (resp. TU_{pq}) has a proper cover where both c_1 and c_2 cover all x_i 's, both c_2 and c_4 cover all y_j 's, c_4 covers x_{v_p} (resp., x_{v_q}), and c_1 covers $y_{e_{pq}}$. By Proposition 3.2, these covers for the gadgets can be joined together to obtain a proper cover of the whole point set P . Let us verify that this cover is consistent. The set V_{p,v_p} is covered by c_4 and if $i \neq v_p$, then V_{pi} is covered by either c_1 or c_2 . The set $E_{p,q,e_{pq}}$ is covered by c_1 and if $j \neq e_{pq}$, then $E_{p,q,j}$ is covered by c_2 or c_4 . \square

Finally, we modify P to obtain a point set P' such that P has a consistent proper cover if and only if P' has a proper cover. Each set V_{pi} and E_{pqj} that is required to be simultaneously covered in a consistent cover is replaced by a single point that is the coordinatewise sum of the points in the set. Observe that in each coordinate at most one point in the set is nonzero, since they belong to different gadgets. Moreover, all sets V_{pi} and E_{pqj} are disjoint, and all points in the same set V_{pi} and E_{pqj} have equal constraint lists, and so there is no conflict in defining the constraint list of the new point.

This decreases the number of points by $nk \cdot (k-1) + m \binom{k}{2} \cdot 2 = \binom{k}{2} (2n + 2m)$, hence $|P'| = k(4n + 6) + \binom{k}{2} (4m + 8n + 18)$.

Assume that P has a consistent proper cover. Suppose that some c_r covers a particular V_{pi} . Then c_r covers the point that replaces V_{pi} : the point is covered in

every coordinate, since every point in V_{pi} is covered by c_r . Assume now that P' has a proper cover; let c_r be a cube covering the point replacing V_{pi} . Since c_r is 0-covering, if we set a coordinate of the point to 0, then the point remains in c_r . This means that c_r covers every point of V_{pi} , hence V_{pi} is simultaneously covered. The same argument applies to E_{pqj} .

We have shown that P' has a proper cover if and only if G has a clique of size k , proving the correctness of the reduction. The reduction can be done in polynomial time, hence the constrained rectilinear 4-center decision problem parameterized with the dimension d is W[1]-hard. This completes the proof of Lemma 3.4.

Lemma 3.3 and Lemma 3.4 imply the following:

THEOREM 3.11. *The rectilinear 4-center decision problem is W[1]-hard with respect to the dimension d . □*

4. THE EUCLIDEAN 2-CENTER PROBLEM

In this section we give an fpt-reduction from the parameterized k -independent set problem in general graphs, which is known to be W[1]-complete with respect to k [Flum and Grohe 2006], to the Euclidean 2-center decision problem, parameterized with the dimension d . Let $[n] = \{1, \dots, n\}$. We look for an independent set of size k in a graph $G([n], E)$. We assume $k \geq 4$, and we assume that $n \geq 4$ and n is even, by adding an additional vertex to G if necessary and connecting it to all other vertices. Using G , we will construct a point set P in \mathbb{R}^{2k+1} with the property that P can be covered by 2 unit balls if and only if G has a independent set of size k .

We first give a high-level overview of our reduction at the logical level. We start with a *scaffolding point set* P^0 of $nk + 2$ points. For an appropriate radius ρ , the set P^0 has the property that there are n^k ways to cover it with two balls of radius ρ , in one-to-one correspondence with all k -tuples (u_1, \dots, u_k) with $1 \leq u_i \leq n$. These coverings allow us to represent the potential independent sets of vertices in the graph. More precisely, they represent *ordered* selections of k (not necessarily distinct) vertices of the graph.

Geometrically, the scaffolding set will consist of a set P_i ($i = 1, \dots, k$) of n equally spaced points on a circle in each of k orthogonal 2-dimensional planes, plus a “top” and a “bottom” *anchor point* on the remaining orthogonal axis. The ball containing the top point can cover $n/2$ consecutive points of each P_i , but not more. So there are n choices for these points. Since the sets P_1, \dots, P_k lie in orthogonal planes, the $n/2$ covered points can be chosen independently in each plane, giving n^k choices altogether, see Lemma 4.4 below. The complementary half of the point set, together with the bottom point, can be covered by the other ball.

The structure of the input graph is represented using additional *constraint points*: for each pair of distinct indices $i \neq j$ ($1 \leq i, j \leq k$) and for each pair of (possibly equal) vertices $u, v \in [n]$, we define a constraint point q_{ij}^{uv} which is covered by all solutions (u_1, \dots, u_k) with the exception of those with $u_i = u$ and $u_j = v$. In particular, we add to P^0 $\binom{k}{2}n$ constraint points $Q_V = \{q_{ij}^{uu} \mid 1 \leq u \leq n, 1 \leq i < j \leq k\}$ to ensure that all components u_i in a solution

must be distinct. Also, for each edge $uv \in E$ we add all $k(k-1)$ points q_{ij}^{uv} with $i \neq j$. In this way, we ensure that the remaining coverings (u_1, \dots, u_k) represent independent sets of size k . In total the edges are represented by the $k(k-1)|E|$ points $Q_E = \{q_{ij}^{uv} \mid uv \in E, 1 \leq i, j \leq k, i \neq j\}$. The resulting set $P = P^0 \cup Q_V \cup Q_E$ will have in total $nk + 2 + \binom{k}{2}(n+2|E|)$ points. A covering of P exists if and only if the graph has an independent set of size k . (Each independent set of size k is represented by $k!$ coverings.)

We will first describe the geometry of the point sets exactly, as if exact square roots and expressions of the form $\sin \frac{\pi}{n}$ were available. We will later show that the essential features of our construction are preserved when the data are perturbed within some tolerance. This allows us to work with fixed-precision roundings of the exact construction, making the reduction suitable for the Turing machine model.

Notation. For our construction it is convenient to view \mathbb{R}^{2k+1} as the product of k orthogonal planes E_1, \dots, E_k plus one extra axis. Each E_i has coordinate axes X_i, Y_i and the extra axis is denoted by Z . For giving coordinates, the axes are considered in the order $X_1, Y_1, \dots, X_k, Y_k, Z$. The coordinate on X_i, Y_i , and Z of a point p is denoted by $x_i(p)$, $y_i(p)$, and $z(p)$, respectively.

4.1 The Scaffolding Point Set

On each plane E_i we define a set P_i consisting of n points regularly spaced on the unit circle C_i centered at the origin o :

$$P_i = \{p_{iu} \in E_i \mid x_i(p_{iu}) = \cos(2u-3)\frac{\pi}{n}, y_i(p_{iu}) = \sin(2u-3)\frac{\pi}{n}, u = 1, \dots, n\}$$

We also use two *anchor* points $p_z, -p_z$ on the Z -axis with $z(p_z) = 2$. The scaffolding point set P^0 is defined as $P^0 = P_1 \cup \dots \cup P_k \cup \{p_z, -p_z\}$. We have $|P^0| = nk + 2$. This point set is highly symmetric. In particular, since the planes E_i are orthogonal, we can independently rotate each plane E_i by multiples of $2\pi/n$.

Two points p and $-p$ of P_i are called *antipodal*. For any u with $1 \leq u \leq n$ we define the index of its *almost antipodal* partner as $\bar{u} = ((u + \frac{n}{2}) \bmod n) + 1$. The pair a_{iu} of *almost antipodal points* is defined as $a_{iu} = \{p_{iu}, p_{i\bar{u}}\}$, for $i = 1, \dots, k$ and $u = 1, \dots, n$. See Fig. 2 for an illustration. Pair a_{iu} results from a counter-clockwise rotation, on the plane E_i , of the pair a_{i1} about o by $(u-1)\frac{2\pi}{n}$.

The radius ρ of the two balls with which we would like to cover P^0 will be slightly smaller than $5/4$. Thus the two anchor points must be covered by two different balls. The ball containing p_z will be called the *top ball*, while the ball containing $-p_z$ is called the *bottom ball*.

Two antipodal points of P_i and the top anchor form an isosceles triangle whose circumradius is $5/4$. Therefore, if $\rho < 5/4$, the top ball (or the bottom ball) cannot contain two antipodal points. (With a radius of $5/4$, the top ball could be centered on the Z -axis at height $3/4$ and cover *all* points P^0 except the bottom anchor.) By choosing $\rho < 5/4$, we ensure that each ball can cover at most half of the points from every P_i . We define the radius ρ as the smallest radius such that the top ball can cover precisely $n/2$ consecutive points of each subset P_i , besides the anchor p_z . The precise value of ρ is given below.




Fig. 2. Point set P_i and the two pairs $a_{i1} = \{p_{i1}, p_{i, \frac{n}{2}+2}\}$ and $a_{i2} = \{p_{i2}, p_{i, \frac{n}{2}+3}\}$ for $n = 12$.

Let $A(u_1, \dots, u_k)$ be the set of $2k$ points

$$A(u_1, \dots, u_k) = a_{1u_1} \cup a_{2u_2} \cup \dots \cup a_{ku_k}.$$

We denote by $B(u_1, \dots, u_k)$ the smallest enclosing ball of $A(u_1, \dots, u_k) \cup \{p_z\}$. Note that all the $2k + 1$ points in this set are affinely independent.

LEMMA 4.1. *All balls $B(u_1, \dots, u_k)$ have the same radius*

$$\rho = \frac{5}{4} \cdot \sqrt{\frac{k}{k + \sigma^2/4}}, \text{ with } \sigma := -y_i(p_{i1}) = \sin \frac{\pi}{n}.$$

The center c of $B(u_1, \dots, u_k)$ has coordinates

$$x_i(c) = +w \sin \frac{(u_i-1)2\pi}{n}, \quad y_i(c) = -w \cos \frac{(u_i-1)2\pi}{n}, \text{ for } i = 1, \dots, k, \text{ and } z(c) = h,$$

with

$$w = \frac{5\sigma}{2(4k + \sigma^2)}, \quad h = \frac{3k + 2\sigma^2}{4k + \sigma^2}. \tag{9}$$

All points in $A(u_1, \dots, u_k) \cup \{p_z\}$ lie on the boundary of $B(u_1, \dots, u_k)$.

PROOF. By symmetry, it is sufficient to show this for the ball $B(1, \dots, 1)$, whose center we claim to be $c = (0, -w, 0, -w, \dots, 0, -w, h)$. We use the following well-known characterization of the smallest enclosing ball:

PROPOSITION 4.2. *A ball B containing a finite set of points A is the smallest enclosing ball for A if and only if its center lies in the convex hull of the points of A that lie on the boundary of B .*

The set $A(1, \dots, 1)$ consists of the $2k$ points of the form

$$(0, \dots, 0, \pm \cos \frac{\pi}{n}, -\sigma, 0, \dots, 0, 0),$$

where the pair $(\pm \cos \frac{\pi}{n}, -\sigma)$ cycles through all k planes E_i . It is straightforward to check that all these points and the point p_z have the same distance ρ to c . Moreover, the center c lies on the line segment between p_z and the center of gravity of the $2k$ points of $A(1, \dots, 1)$, which is the point $(0, -\sigma/k, 0, -\sigma/k, \dots, 0, -\sigma/k, 0)$. Thus, it lies in the convex hull of $A(1, \dots, 1) \cup \{p_z\}$. \square

Note that by symmetry, any bottom ball will have its center c' with $z(c') = -h$.

PROPOSITION 4.3. i) *The height $h = z(c)$ of the center lies in the range $3/4 < h < 1$.*

ii) *The radius ρ lies in the range $1 < \rho < 5/4$, with $\rho = 5/4 - \Theta(1/(n^2k))$.*

PROOF. (i) This follows easily from $\sigma^2 < k$ and $h = 1 - (k - \sigma^2)/(4k + \sigma^2) > 1 - k/(4k) = 3/4$.

(ii) The upper bound on ρ is obvious, and the lower bound follows from part (i), since ρ is the distance between p_z and the point c with z -coordinate $h < 1$. We have $\rho = \frac{5}{4} \cdot (1 + \frac{\sigma^2}{4k})^{-1/2} = \frac{5}{4} \cdot (1 - \frac{1}{2} \cdot \frac{\sigma^2}{4k} + O(\frac{\sigma^2}{4k})^2) = \frac{5}{4} - \frac{5}{32} \cdot \frac{\sigma^2}{k} + O(\frac{\sigma^4}{k^2}) = \frac{5}{4} - \frac{5}{32} \cdot \frac{1}{k} \cdot (\frac{\pi}{n} + O(\frac{\pi}{n}))^2 + O(\frac{1}{n^4k^2}) = \frac{5}{4} - \frac{5\pi^2}{32} \cdot \frac{1}{n^2k} + O(\frac{1}{n^3k})$. \square

The intersection of $B(u_1, \dots, u_k)$ with E_i is a disk of radius

$$\rho_i = \sqrt{\rho^2 - \sum_{\substack{1 \leq r \leq k \\ r \neq i}} (x_r^2(c) + y_r^2(c)) - z^2(c)} < \sqrt{\rho^2 - z^2(c)} < \sqrt{25/16 - 9/16} = 1$$

since $\rho < 5/4$ and $z(c) = h > 3/4$. Since the disk has radius smaller than 1, it cannot contain any two antipodal points of P_i . And, since it contains the pair a_{iu_i} on its boundary, it follows that $B(u_1, \dots, u_k)$ also contains the $n/2 - 2$ consecutive points of P_i between the points of the pair a_{iu_i} . Since the planes E_1, \dots, E_k are orthogonal, each u_i independently defines which of the $n/2$ consecutive points of set P_i is covered by $B(u_1, \dots, u_k)$. The complementary halves can then be covered by the bottom ball. In total, we have n^k possible partitions of P^0 into two groups covered by the two balls, which correspond to the n^k possible tuples $(u_1, \dots, u_k) \in [n]^k$.

We conclude with the following characterization of the possible coverings of P^0 with two balls of radius ρ .

LEMMA 4.4. *Assume that two balls B, B' of radius ρ cover P^0 , and that $p_z \in B$. Then there is a tuple $(u_1, \dots, u_k) \in [n]^k$ such that $B = B(u_1, \dots, u_k)$ and $B' = -B(u_1, \dots, u_k)$.*

PROOF. As discussed before, B or B' can contain at most $n/2$ consecutive points of P_i , and therefore, B and B' must cover complementary halves of each P_i . If B covers the halves between the pairs $a_{1u_1}, \dots, a_{ku_k}$, it follows from the uniqueness of the minimum enclosing ball that $B = B(u_1, \dots, u_k)$. Since B' covers the complementary halves of each P_i and $-p_z$, it follows that $B' = -B(u_1, \dots, u_k)$. \square

From this characterization, the bijection between the possible coverings of P^0 and $[n]^k$ is clear. Every covering consists of a symmetric pair of balls $B = B(u_1, \dots, u_k)$ and $B' = -B(u_1, \dots, u_k)$.

4.2 Constraint Points

We continue now the construction of point set P , by showing how we encode the structure of G . For each pair of distinct indices $i \neq j$ ($1 \leq i, j \leq k$) and for each pair of (possibly equal) vertices $u, v \in [n]$, we define a *constraint point* q_{ij}^{uv} . All constraint points lie on the hyperplane $H = \{p \in \mathbb{R}^{2k+1} \mid z(p) = h\}$, which also contains the center of the top ball. (Recall that h is the coordinate on the Z -axis (“height”) of the center of the top ball, see (9).) This breaks the symmetry between top and bottom balls that the construction had until now. Since the center of the bottom ball lies on the hyperplane $-H$ and $\rho < 5/4 < 2h$, none of the constraint points can be covered by a bottom ball. Therefore, our discussion will only consider top balls. The constraint point q_{ij}^{uv} will lie in all top balls $B(u_1, \dots, u_k)$ except in those with $u_i = u$ and $u_j = v$.

We choose q_{ij}^{uv} in the four-dimensional affine subspace

$$F_{ij} = \{p \in \mathbb{R}^{2k+1} \mid x_r(p) = y_r(p) = 0 \text{ for } r \neq i, j, \text{ and } z(p) = h\} = o' + E_i \times E_j,$$

where $o' = (0, \dots, 0, h)$. We look at the intersections of the balls $B(u_1, \dots, u_k)$ with F_{ij} . Let

$$\mathcal{D} = \{B(u_1, \dots, u_k) \cap F_{ij} \mid (u_1, \dots, u_k) \in [n]^k\}.$$

The intersection of any ball $B = B(u_1, \dots, u_k)$ with F_{ij} is a 4-dimensional ball D , whose center c is the orthogonal projection of the center of B on F_{ij} . From Lemma 4.1, we have

$$x_i(c) = +w \sin \theta_i, \quad y_i(c) = -w \cos \theta_i, \quad x_j(c) = +w \sin \theta_j, \quad y_j(c) = -w \cos \theta_j, \quad (10)$$

where $\theta_i = (u_i - 1) \frac{2\pi}{n}$ and $\theta_j = (u_j - 1) \frac{2\pi}{n}$. The location of the center c thus depends only on u_i and u_j . We denote this center by $c_{ij}^{u_i u_j}$.

Looking at the distance between the centers of B , D , and o' , we get the following properties:

- a) every ball $D \in \mathcal{D}$ has radius $\rho_* = \sqrt{\rho^2 - (k-2)w^2}$;
- b) a point $q \in F_{ij}$ lies in the ball $B(u_1, \dots, u_k)$ if and only if $\|q - c_{ij}^{u_i u_j}\| \leq \rho_*$;
- c) the center of D lies on the three-dimensional sphere $C = \{p \in F_{ij} \mid \|p - o'\| = w\sqrt{2}\}$;
- d) D contains the sphere C in its interior (as $\rho_* > 2\sqrt{2}w$.)

Let D_{ij}^{uv} denote the ball in F_{ij} with center c_{ij}^{uv} and radius ρ_* . For each u, v , we want to find a point $q_{ij}^{uv} \in F_{ij}$ that lies outside the ball D_{ij}^{uv} but in all other balls of \mathcal{D} .

Since the centers $c_{ij}^{uv} \in F_{ij}$ form a completely symmetric set and all balls have the same radius, we can find this point as follows. (See Fig. 3a for a two-dimensional analog of this situation.) Start at c_{ij}^{uv} and move along the ray $L^+ = \{c_{ij}^{uv} + \lambda(o' - c_{ij}^{uv}) \mid \lambda \geq 0\}$ through o' . By properties (c) and (d), we are initially inside all balls $D \in \mathcal{D}$. At some point l_1 , we hit the boundary of some ball. We prove below that this ball is D_{ij}^{uv} . Thus, after passing l_1 , we are outside D_{ij}^{uv} but still inside any ball $D \neq D_{ij}^{uv}$. We place q_{ij}^{uv} at the point l_2 where L^+ intersects the boundary of the next ball.

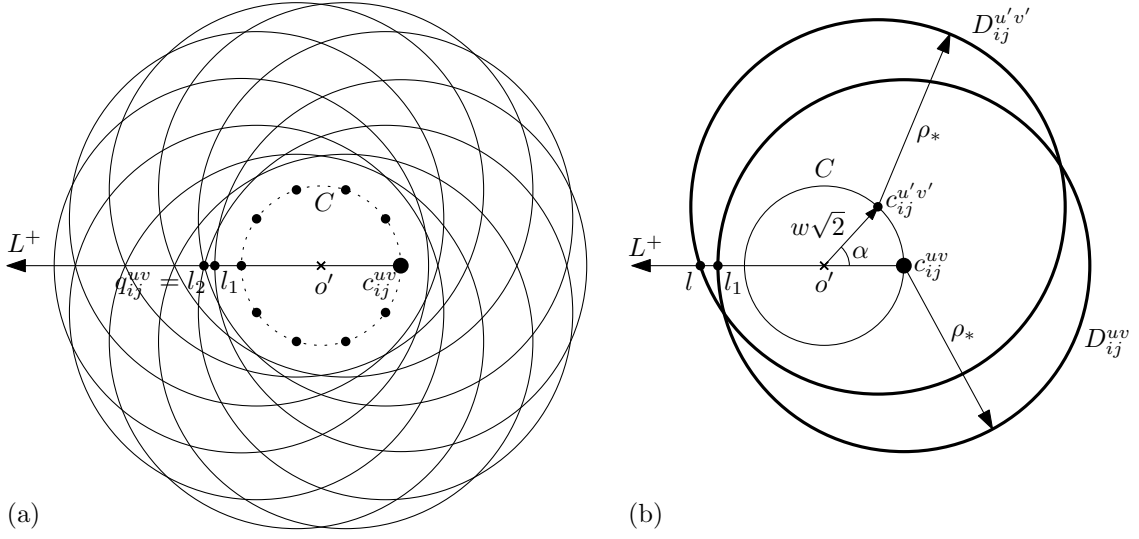


Fig. 3. (a) Finding the point $l_2 = q_{ij}^{uv}$. (b) The situation of Lemma 4.5, in a two-dimensional intersection with the plane through o' , c_{ij}^{uv} , and $c_{ij}^{u'v'}$.

LEMMA 4.5. *The ray L^+ , after having visited o' , hits the boundary of the ball D_{ij}^{uv} before the boundary of any other ball $D \in \mathcal{D}$.*

PROOF. Let l_1 be the point where L^+ intersects the boundary of D_{ij}^{uv} , and let $c_{ij}^{u'v'}$ be the center of any other ball $D_{ij}^{u'v'} \in \mathcal{D}$, $D_{ij}^{u'v'} \neq D_{ij}^{uv}$ (see Fig. 3b). Using the triangle inequality, we have

$$\|c_{ij}^{u'v'} - l_1\| < \|c_{ij}^{u'v'} - o'\| + \|o' - l_1\| = \|c_{ij}^{uv} - o'\| + \|o' - l_1\| = \|c_{ij}^{uv} - l_1\| = \rho_*.$$

This implies that the boundary of D_{ij}^{uv} is the first boundary intersected by L^+ . \square

A quantitative version of this lemma is given below in Lemma 4.7.

4.3 The Reduction

As mentioned in the beginning of this section, we add $\binom{k}{2}(n + 2|E|)$ constraint points Q_V and Q_E to the scaffolding set P^0 to represent the structure of the input graph $G([n], E)$.

LEMMA 4.6. *The set $P = P^0 \cup Q_V \cup Q_E$ can be covered with two balls of radius ρ if and only if G has an independent set of size k .*

PROOF. Any covering of P with two balls B, B' of radius ρ must consist of the two balls $B = B(u_1, \dots, u_k)$ and $B' = -B(u_1, \dots, u_k)$ for some tuple (u_1, \dots, u_k) , by Lemma 4.4. Since the constraint points exclude the tuples with two equal indices $u_i = u_j$, or with indices u_i and u_j when (u_i, u_j) is an edge of G , the coverings represent precisely the independent sets of G . \square

Rounding coordinates. To make the reduction suitable for a Turing machine, we round all data to multiples of a small “unit” U . Scaling by U will then convert the input to integral data. We will show that choosing $U = \Theta(1/(n^6 k^2))$ will preserve all important characteristics of our point set. Since it is easy to evaluate $\sin \frac{u\pi}{n}$ or $\sqrt{\cdot}$ to this precision of $O(\log(nk))$ bits, the reduction can be carried out

in polynomial time. Intuitively, it is quite clear that rounding to a grid of size polynomial in n and k should be sufficient. In this section, we only sketch the main flow of ideas that leads to our claimed bound of U . (The elementary but tedious calculations that prove the following lemmas are deferred to the appendix.)

More precisely, we replace each nonzero input coordinate x by a multiple \hat{x} of U in the range $x - U < \hat{x} < x + U$. This ensures that each input point is moved at most $\sqrt{5} \cdot U$ from its original position. (Recall that most coordinates of our input points are 0.) We replace ρ by a multiple $\hat{\rho}$ of U in the range $\rho + \sqrt{5} \cdot U \leq \hat{\rho} \leq \rho + \sqrt{5} \cdot U + 2U$. In this way we ensure that for each ball that covers some set of ideal input points, there exists an enlarged ball with radius $\hat{\rho}$ that covers the same input points after rounding. We want to exclude the possibility that the enlarged ball covers additional points (possibly after moving its center).

LEMMA 4.7. *Every input point that is not in a ball $B(u_1, \dots, u_k)$ is at least $\varepsilon_1 = \frac{1}{2n^3k}$ away from this ball. \square*

(This bound is asymptotically tight for the constraint points: the distance between l_1 and l_2 is $\Theta(w/n^2) = \Theta(1/(n^3k))$ (see Fig. 3a). The scaffolding points have a much larger distance: they are at least $1/(nk)$ away from the ball.) Since we have introduced some slack by enlarging the radius and moving the points, the center of the ball may move away from the original center. The following lemma bounds this movement.

LEMMA 4.8. *If we move the center of the ball $B(u_1, \dots, u_k)$ by ε_2 , the distance to some point from the scaffolding set on its boundary increases by at least $\varepsilon_3 = \Omega(\min(\varepsilon_2^2, \varepsilon_2/(nk)))$. \square*

(The first bound comes from moving the center perpendicular to the hyperplane through the boundary points of the ball; the second one comes from any movement in this hyperplane.)

We also maintain the global structure of the potential coverings by ensuring that there must be a top ball and a bottom ball.

LEMMA 4.9. *If we chose $U = \text{const}/(n^2k)$, for a sufficiently small constant, then no ball can cover both anchor points p_z and $-p_z$, and the top ball or the bottom ball cannot cover two (perturbed) antipodal points.*

PROOF. To cover an anchor point and two (precise) antipodal points (or even both anchor points), the radius must be at least $5/4$. Thus, to cover the perturbed points, it must grow from $\rho = 5/4 - \Theta(1/(n^2k))$ (by Proposition 4.3(ii)), to at least $5/4 - \sqrt{5}U$. If U is small enough, this distance is larger than $\sqrt{5}U + 2U$, the maximum possible radius increase. \square

From the above lemmas, we obtain the following

THEOREM 4.10. *If we chose $U = \text{const}/(n^6k^2)$, for a sufficiently small constant, all possible coverings by two balls of radius $\hat{\rho}$ partition the rounded point set in the same way as two balls $B(u_1, \dots, u_k)$ and $-B(u_1, \dots, u_k)$ for the original data.*

PROOF. Recall that each point moves by at most $\sqrt{5} \cdot U$, and the radius is increased by at most $\sqrt{5} \cdot U + 2U$. Lemma 4.9 implies that a 2-covering has the same structure as for the ideal point set as far as the scaffolding points are concerned. We can choose U small enough to ensure that the center \hat{c} of a ball with the modified data can move only $\varepsilon_2 = 1/(5n^3k)$ from its original position: Otherwise, by Lemma 4.8, its distance from some boundary point of the scaffolding point set would increase by more than $\varepsilon_3 = \Omega(\min(\varepsilon_2^2, \varepsilon_2/(nk))) = \Omega(1/(n^6k^2))$. If U is chosen such that $\varepsilon_3 > \sqrt{5} \cdot U + (\sqrt{5} \cdot U + 2U)$, this would mean that the ball can no longer contain this boundary point.

By Lemma 4.7 we know that each point that is not covered is at least $\varepsilon_1 = 1/(2n^3k)$ away from the original ball. Now, choosing U in such a way that $\varepsilon_2 + (\sqrt{5} \cdot U + 2U) < \varepsilon_1 - \sqrt{5} \cdot U$, the ball cannot swallow any additional points. \square

Using the rounded coordinates for the points of P , and since $|P| = nk + 2 + \binom{k}{2}(n + 2|E|)$ and $d = 2k + 1$, we see that this is an fpt-reduction. Thus, we have the following:

THEOREM 4.11. *The Euclidean 2-center decision problem is W[1]-hard with respect to the dimension d . \square*

For a parameterized complexity upper bound, we mention that, trivially, the (integral) Euclidean 2-center decision problem, parameterized with d , is in W[P]; see Downey and Fellows [1999].

Since $d = 2k + 1$ in the above fpt-reduction, an $n^{o(d)}$ -time algorithm for the Euclidean 2-center decision problem implies an $n^{o(k)}$ -time algorithm for the parameterized k -independent set problem, which in turn implies that n -variable 3SAT can be solved in time $2^{o(n)}$ [Chen et al. 2004; Chen et al. 2005]. The conjecture that there is no such algorithm is called the Exponential Time Hypothesis (ETH), which was formulated and investigated in [Impagliazzo et al. 2001]. Thus we have the following:

COROLLARY 4.12. *The Euclidean k -center problem, for any $k \geq 2$, cannot be solved in $n^{o(d)}$ time, unless ETH fails. \square*

5. OPEN PROBLEMS

We have given a fine classification of the complexity of the k -center problem parameterized by the dimension. Our fpt-reduction from the parameterized k -clique problem to the rectilinear 4-center problem in $\Theta(k^2)$ dimensions implies that an $n^{o(\sqrt{d})}$ -time algorithm for the latter does not exist, unless the Exponential Time Hypothesis fails. However, it does not exclude the existence of algorithms taking $n^{o(d)}$ time. It would be interesting to find the appropriate order of magnitude in the exponent for the rectilinear 4-center problem. More generally, it would be interesting to study the order of magnitude for the k -center problem when parameterized by both k and d . Finally, one could also investigate k -center problems for the (\mathbb{R}^d, L_1) -metric or for variations where the centers are not points but higher-dimensional objects (e. g., lines) that should lie close to the input points.

REFERENCES

- AGARWAL, P. K. AND SHARIR, M. 1998. Efficient algorithms for geometric optimization. *ACM Comput. Surv.* 30, 412–458.
- ASSA, E. AND KATZ, M. J. 1999. 3-piercing of d -dimensional boxes and homothetic triangles. *Int. J. Comput. Geometry Appl.* 9, 3, 249–260.
- CABELLO, S., GIANNOPOULOS, P., KNAUER, C., AND ROTE, G. 2008. Geometric clustering: fixed-parameter tractability and lower bounds with respect to the dimension. In *Proceedings of the 19th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, San Francisco. 836–843.
- CHAN, T. M. 1999a. Geometric applications of a randomized optimization technique. *Discrete & Computational Geometry* 22, 4, 547–567.
- CHAN, T. M. 1999b. More planar two-center algorithms. *Comput. Geom. Theory Appl.* 13, 3, 189–198.
- CHEN, J., CHOR, B., FELLOWS, M., HUANG, X., JUEDES, D., KANJ, I. A., AND XIA, G. 2005. Tight lower bounds for certain parameterized NP-hard problems. *Information and Computation* 201, 2, 216–231.
- CHEN, J., HUANG, X., KANJ, I. A., AND XIA, G. 2004. Linear FPT reductions and computational lower bounds. In *Proceedings of the 36th Annual ACM Symposium on Theory of Computing*. ACM, New York, 212–221.
- DEMAINE, E. D., FOMIN, F. V., HAJIAGHAYI, M., AND THILIKOS, D. M. 2005. Fixed-parameter algorithms for (k, r) -center in planar graphs and map graphs. *ACM Trans. Algorithms* 1, 1, 33–47.
- DOWNEY, R. AND FELLOWS, M. 1999. *Parameterized Complexity*. Springer-Verlag.
- DREZNER, Z., Ed. 1995. *Facility Location: A Survey of Applications and Methods*. Springer-Verlag.
- EPPSTEIN, D. 1997. Fast construction of planar two-centers. In *Proc. 8th Ann. ACM-SIAM Sympos. Discrete Algorithms (SODA)*. 131–138.
- FLUM, J. AND GROHE, M. 2006. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science, vol. XIV. Springer-Verlag.
- FOWLER, R. J., PATERSON, M. S., AND TANIMOTO, S. L. 1981. Optimal packing and covering in the plane are NP-complete. *Inf. Proc. Lett.* 12, 3, 133–137.
- FREDERICKSON, G. N. AND JOHNSON, D. B. 1984. Generalized selection and ranking: sorted matrices. *SIAM J. Comput.* 13, 14–30.
- HOFFMANN, M. 2001. Covering polygons with few rectangles. In *Abstracts of the 17th European Workshop on Computational Geometry (EuroCG '01)*. 39–42.
- HOFFMANN, M. 2002. A matrix search algorithm with applications to rectilinear 3-centers of polygons. Manuscript.
- IMPAGLIAZZO, R., PATURI, R., AND ZANE, F. 2001. Which problems have strongly exponential complexity? *J. Comput. System Sci.* 63, 4, 512–530. Special issue on FOCS 98 (Palo Alto, CA).
- MARX, D. 2005. Efficient approximation schemes for geometric problems? In *Proc. 13th Ann. European Sympos. Algorithms (ESA)*. Lect. Notes Comput. Sci., vol. 3669. Springer-Verlag, 448–459.
- MEGIDDO, N. 1990. On the complexity of some geometric problems in unbounded dimension. *J. Symb. Comput.* 10, 3-4, 327–334.
- MEGIDDO, N. AND SUPOWIT, K. J. 1984. On the complexity of some common geometric location problems. *SIAM J. Comput.* 13, 182–196.
- NUSSBAUM, D. 1997. Rectilinear p -piercing problems. In *Proc. of the 1997 Annual International Symposium on Symbolic and Algebraic Computation*. ACM, 316–323.
- SEGAL, M. 1999. On piercing sets of axis-parallel rectangles and rings. *Int. J. Comput. Geometry Appl.* 9, 3, 219–234.
- SHARIR, M. AND WELZL, E. 1996. Rectilinear and polygonal p -piercing and p -center problems. In *Proc. of the 12th Ann. Sympos. Computational Geometry*. ACM Press, 122–132.

A. GEOMETRIC ESTIMATES

First, we collect some useful inequalities.

- PROPOSITION A.1. (a) $2/n \leq \sigma = \sin \frac{\pi}{n} \leq \pi/n$, for $n \geq 4$.
 (b) $1 - \cos \alpha \geq \alpha^2/4$, for $|\alpha| \leq \pi/3$ and for $\alpha = \pi/2$.
 (c) $\frac{1}{kn} < w < \frac{2}{kn}$, for $k, n \geq 4$.
 (d) $\rho_* > 3/4$, for $k, n \geq 4$.

PROOF. (a) We have

$$\alpha \cdot \frac{2}{\pi} \leq \sin \alpha \leq \alpha$$

for $0 \leq \alpha \leq \pi/4$. The right inequality holds for all $\alpha \geq 0$, and the left inequality is established by comparing the concave sine function with the linear function at the endpoints of the range.

(b) The function $1 - \cos \alpha - \alpha^2/4$ is concave and nonnegative at the boundaries of the range.

(c) We have

$$w = \frac{5\sigma}{2(4k + \sigma^2)} > \frac{5(2/n)}{2(4k + (\pi/n)^2)} = \frac{10n}{8kn^2 + 2\pi^2} > \frac{1}{kn},$$

for $k, n \geq 4$, and,

$$w = \frac{5\sigma}{2(4k + \sigma^2)} < \frac{5(\pi/n)}{2(4k + (2/n)^2)} = \frac{5\pi n}{8kn^2 + 8} < \frac{5\pi}{8kn} < \frac{2}{kn}.$$

(d) We have

$$\begin{aligned} \rho_* &= \sqrt{\rho^2 - (k-2)w^2} > \sqrt{1 - kw^2} > \sqrt{1 - k(2/(kn))^2} \\ &= \sqrt{1 - 4/(kn^2)} \geq \frac{\sqrt{15}}{4} > \frac{3}{4}, \end{aligned}$$

for $k, n \geq 4$. \square

Now we are ready to prove the lemmas of Section 4.3.

PROOF OF LEMMA 4.7. There are three types of points outside $B(u_1, \dots, u_k)$: (a) the opposite anchor point; (b) scaffolding points on the circles P_i ; (c) constraint points. We consider each type separately.

(a) The opposite anchor point, $-p^z$, is at least $4 - 2\rho > 4 - 2 \cdot \frac{5}{4} = \frac{3}{4}$ away from the ball.

(b) By symmetry, for the scaffolding points it is sufficient to look at the ball $B(1, \dots, 1)$. Consider some fixed plane E_i . Fig. 4 shows the projection \hat{c} of the ball center c onto this plane, which forms the center of the disk $\hat{B} = B(1, \dots, 1) \cap E_i$, whose radius is denoted by $\hat{\rho}$. The point p_{i1} is on the boundary of this disk. For any point $p_{ij} \in P_i$, we have $\|p_{ij} - c\|^2 = \|p_{ij} - \hat{c}\|^2 + \|\hat{c} - c\|^2$. Thus, for any point p_{ij} in E_i , we can calculate the distance from $B(1, \dots, 1)$ by using the equation

$$\begin{aligned} \|p_{ij} - c\|^2 - \rho^2 &= \|p_{ij} - c\|^2 - \|p_{i1} - c\|^2 \\ &= (\|p_{ij} - \hat{c}\|^2 - \|\hat{c} - c\|^2) - (\|p_{i1} - \hat{c}\|^2 - \|\hat{c} - c\|^2) \\ &= \|p_{ij} - \hat{c}\|^2 - \|p_{i1} - \hat{c}\|^2 \end{aligned}$$




Fig. 4. The intersection of the ball $B(1, \dots, 1)$ with the $X_i Y_i$ -plane. The drawing is only schematic. The center \hat{c} lies much closer to the center o than shown.

Among the points p_{ij} in $E_i \setminus \hat{B}$, a point that minimizes $\|p_{ij} - \hat{c}\|$ is p_{i2} . We use the abbreviations

$$\eta := x_i(p_{i1}) = x_i(p_{i2}) = \cos \frac{\pi}{n} \quad \text{and} \quad \sigma := y_i(p_{i2}) = -y_i(p_{i1}) = \sin \frac{\pi}{n}.$$

We can thus estimate the above difference as follows:

$$\begin{aligned} \|p_{ij} - c\|^2 - \rho^2 &\geq \|p_{i2} - \hat{c}\|^2 - \|p_{i1} - \hat{c}\|^2 \\ &= (\eta^2 + (\sigma + w)^2) - (\eta^2 + (-\sigma + w)^2) \\ &= 4\sigma w \\ &\geq 4 \cdot \frac{2}{n} \cdot \frac{1}{kn} \\ &\geq \frac{8}{kn^2}. \end{aligned}$$

If $\|p_{ij} - c\| > 7/4 \geq \rho + 1/2$, there is nothing left to show. Let us therefore assume that $\rho < \|p_{ij} - c\| \leq 7/4$. Then,

$$\|p_{ij} - c\| - \rho = \frac{\|p_{ij} - c\|^2 - \rho^2}{\|p_{ij} - c\| + \rho} > \frac{\|p_{ij} - c\|^2 - \rho^2}{7/4 + 5/4} \geq \frac{8}{kn^2} \cdot \frac{1}{3} > \frac{1}{kn^2} > \frac{1}{2kn^3}.$$

(c) Finally, let us look at the constraint points, see Fig. 3b. We will first give a lower bound on the distance from l_1 to the intersection l of any disk $D_{ij}^{u'v'} \in (\mathcal{D} \setminus \{D_{ij}^{uv}\})$ with the ray L^+ . Let $\alpha = \angle c_{ij}^{u'v'} o' c_{ij}^{uv} \geq 2\pi/n$. To see this, observe that, for two centers c_{ij}^{uv} and $c_{ij}^{u'v'}$, the orthogonal projection of segment $c_{ij}^{uv} c_{ij}^{u'v'}$ on the plane E_i is the base of an isosceles triangle with legs of length w and apex angle $(u - u')\frac{2\pi}{n}$ and has length $2w \cdot \sin \frac{(u - u')\pi}{n}$. So, the

Euclidean distance between the two centers is

$$w \cdot \sqrt{(2 \sin \frac{(u-u')\pi}{n})^2 + (2 \sin \frac{(v-v')\pi}{n})^2}.$$

This expression is minimized when $u' \equiv u \pm 1 \pmod{n}$ and $v = v'$ or when $v' \equiv v \pm 1 \pmod{n}$ and $u = u'$. Since all centers lie on a 3-sphere, the central angle α is minimized precisely if the Euclidean distance is minimized. The angle that corresponds to these minimizing cases is clearly $2\pi/n$.

The length of the projection of the vector $c_{ij}^{u'v'} l$ on L^+ is $\sqrt{\rho_*^2 - 2w^2 \sin^2 \alpha}$ since $c_{ij}^{u'v'} l$ has length ρ_* and the distance from $c_{ij}^{u'v'} l$ to L^+ is $\sqrt{2}w \sin \alpha$. Hence

$$\|o' - l\| = \sqrt{\rho_*^2 - 2w^2 \sin^2 \alpha} - \sqrt{2}w \cos \alpha.$$

The distance $\|o' - l_1\|$ is $\rho_* - \sqrt{2}w$. We estimate the distance $\|l_1 - l\|$ as follows, using Proposition A.1 where appropriate.

$$\begin{aligned} \|l_1 - l\| &= \|o' - l\| - \|o' - l_1\| \\ &= \sqrt{\rho_*^2 - 2w^2 \sin^2 \alpha} - \sqrt{2}w \cos \alpha - (\rho_* - \sqrt{2}w) \\ &= \sqrt{2}w(1 - \cos \alpha) + \frac{\rho_*^2 - 2w^2 \sin^2 \alpha - \rho_*^2}{\sqrt{\rho_*^2 - 2w^2 \sin^2 \alpha} + \rho_*} \\ &= \sqrt{2}w(1 - \cos \alpha) - \frac{2w^2 \sin^2 \alpha}{\sqrt{\rho_*^2 - 2w^2 \sin^2 \alpha} + \rho_*} \\ &> w(1 - \cos \alpha) - \frac{2w^2 \sin^2 \alpha}{\rho_*} \\ &= w \left(1 - \cos \alpha - \frac{2w}{\rho_*} \cdot \sin^2 \alpha \right) \\ &= w \left(1 - \cos \alpha - \frac{2w}{\rho_*} \cdot (1 - \cos^2 \alpha) \right) \\ &= w \left((1 - \cos \alpha) \cdot \left(1 - \frac{4w}{\rho_*} \right) + (1 - \cos \alpha)^2 \cdot \frac{2w}{\rho_*} \right) \\ &> w \cdot (1 - \cos \alpha) \cdot \left(1 - \frac{4w}{\rho_*} \right) > w \cdot (1 - \cos \frac{2\pi}{n}) \cdot \left(1 - \frac{32}{3kn} \right) \\ &\geq w \cdot (1 - \cos \frac{2\pi}{n}) \cdot \frac{1}{3} > w \cdot \frac{1}{4} \left(\frac{2\pi}{n} \right)^2 \cdot \frac{1}{\pi^2} \\ &= \frac{w}{n^2} \\ &> \frac{1}{kn^3}, \end{aligned}$$

for $k, n \geq 4$. Thus, for the constraint point $q_{ij}^{uv} = l_2$ we also have that $\|l_2 - l_1\| > 1/(kn^3)$.

We can now estimate the distance from l_2 to the ball $B(u_1, \dots, u_n)$ with $u_i = u$ and $u_j = v$. Let c be the center of the ball in \mathbb{R}^{2k+1} , and as before, let c_{ij}^{uv} be its projection on $E_i \times E_j$. If $\|l_2 - c\| > 7/4 \geq \rho + 1/2$, there is nothing left to show.

Let us therefore assume that $\rho < \|l_2 - c\| \leq 7/4$. Then,

$$\begin{aligned}
\|l_2 - c\| - \rho &= \|l_2 - c\| - \|l_1 - c\| \\
&= \frac{\|l_2 - c\|^2 - \|l_1 - c\|^2}{\|l_2 - c\| + \|l_1 - c\|} \\
&> \frac{\|l_2 - c\|^2 - \|l_1 - c\|^2}{7/4 + 5/4} \\
&= \frac{1}{3} \cdot ((\|l_2 - c_{ij}^{uv}\|^2 + \|c_{ij}^{uv} - c\|^2) - (\|l_1 - c_{ij}^{uv}\|^2 + \|c_{ij}^{uv} - c\|^2)) \\
&= \frac{1}{3} \cdot (\|l_2 - c_{ij}^{uv}\|^2 - \|l_1 - c_{ij}^{uv}\|^2) \\
&= \frac{1}{3} \cdot (\|l_2 - c_{ij}^{uv}\| - \|l_1 - c_{ij}^{uv}\|) \cdot (\|l_2 - c_{ij}^{uv}\| + \|l_1 - c_{ij}^{uv}\|) \\
&\geq \frac{1}{3} \cdot \|l_2 - l_1\| \cdot 2\rho_* \\
&> \frac{1}{3} \cdot \frac{1}{kn^3} \cdot 2 \cdot \frac{3}{4} = \frac{1}{2n^3k} \quad \square
\end{aligned}$$

The following lemma is a more precise version of Lemma 4.8.

LEMMA A.2. *If we move the center of the ball $B(u_1, \dots, u_k)$ by more than ε_2 , for $\varepsilon_2 \leq 1$, the distance to some point on its boundary increases by at least*

$$\varepsilon_3 := \min\left(\frac{\varepsilon_2^2}{6}, \frac{\varepsilon_2}{15nk}\right) = \Omega(\min(\varepsilon_2^2, \varepsilon_2/(nk))).$$

PROOF. By symmetry, it suffices to study the case of the ball $B(1, \dots, 1)$. Let us assume that the center moves from c to $c + \Delta c$, with $\|\Delta c\| > \varepsilon_2$. The motion of the center can be decomposed into two orthogonal components: a component Δ_1 parallel to the hyperplane H through the boundary points and the center of the ball, and a second component Δ_2 perpendicular to this hyperplane.

$$\Delta c = \Delta_1 + \Delta_2$$

with $\|\Delta c\|^2 = \|\Delta_1\|^2 + \|\Delta_2\|^2$.

The motion Δ_2 increases the distance d from *any* point in the hyperplane H to $\sqrt{d^2 + \|\Delta_2\|^2}$, and is therefore easy to analyze. Suppose $\|\Delta_2\|^2 \geq \frac{1}{2}\varepsilon_2^2$. Since $B(1, \dots, 1)$ is the smallest enclosing ball, the distance from c to some boundary point is at least $\sqrt{\rho^2 + \|\Delta_2\|^2}$. Thus, the distance increases by at least

$$\sqrt{\rho^2 + \|\Delta_2\|^2} - \rho \geq \sqrt{\rho^2 + \frac{\varepsilon_2^2}{2}} - \rho = \frac{\rho^2 + \varepsilon_2^2/2 - \rho^2}{\sqrt{\rho^2 + \varepsilon_2^2/2} + \rho} \geq \frac{\varepsilon_2^2/2}{3} = \frac{\varepsilon_2^2}{6}.$$

On the other hand, let us assume that $\|\Delta_2\|^2 < \frac{1}{2}\varepsilon_2^2$, and therefore $\|\Delta_1\| \geq \frac{1}{2}\varepsilon_2$. We take the alternative viewpoint and fix the maximum distance increase ε_3 , and we bound the amount by which the center c can move. The hyperplane H in which c moves is given by the equation $\sum_{i=1}^k y_i - \frac{\sigma}{2} \cdot z = -\sigma$. Consider the set of points $c + \Delta_1$ that lie in H and whose distance from each point of $q \in A(1, \dots, 1) \cup \{p^z\}$ is increased by at most ε_3 :

$$\|(c + \Delta_1) - q\| \leq \rho + \varepsilon_3, \text{ for all } q \in A(1, \dots, 1) \cup \{p^z\}$$

We relax this constraint to a linear inequality with the help of the Cauchy-
to appear in ACM Transactions on Algorithms

Schwarz-Bunyakovski inequality:

$$\begin{aligned} \|(c + \Delta_1) - q\| \cdot \|c - q\| &\geq \langle (c + \Delta_1) - q, c - q \rangle \\ &= \langle \Delta_1, c - q \rangle + \langle c - q, c - q \rangle \\ &= \langle \Delta_1, c - q \rangle + \rho^2 \end{aligned}$$

and therefore

$$\langle \Delta_1, c - q \rangle \leq \|(c + \Delta_1) - q\| \cdot \rho - \rho^2 = \rho (\|(c + \Delta_1) - q\| - \rho) \leq \rho \varepsilon_3.$$

For bounding the length of the vectors Δ_1 that satisfy $\langle \Delta_1, c - q \rangle \leq \rho \varepsilon_3$, we use the following lemma, whose proof is given afterwards.

LEMMA A.3. *Let c be the center of the ball $B(1, \dots, 1)$. If a vector*

$$p = (x_1, y_1, \dots, x_k, y_k, z)$$

lies in the hyperplane $\sum_{i=1}^k y_i - \frac{\sigma}{2} \cdot z = 0$ and satisfies the conditions

$$\langle p, c - q \rangle \leq a \text{ for all } q \in A(1, \dots, 1) \cup \{p_z\}, \quad (11)$$

for some $a \geq 0$, then $\|p\| \leq a \cdot 6nk$.

It follows that $\frac{1}{2}\varepsilon_2 \leq \|\Delta_1\| \leq 6nk \cdot \rho \varepsilon_3 \leq nk\varepsilon_3 \cdot \frac{30}{4}$ and thus, $\varepsilon_3 \geq \varepsilon_2/(15nk)$. \square

PROOF OF LEMMA A.3. Let $\eta = x_i(p_{i1}) = \cos(\pi/n)$. The vectors $c - q$ are the $2k$ vectors of the form

$$(\pm\eta, \sigma - w, 0, -w, 0, -w, \dots, 0, -w, h), \quad (12)$$

where the pair $(\pm\eta, \sigma - w)$ cycles through all k planes E_i , and the vector

$$(0, -w, 0, -w, \dots, 0, -w, h - 2).$$

The $2k + 1$ constraints (11) define a $2k$ -dimensional simplex in the hyperplane in which p is constrained to lie. The longest vector in this simplex must be a vertex of the simplex. Each of the $2k + 1$ vertices can be obtained by dropping one constraint from (11) and setting the remaining constraints to equations. The statement of the lemma is invariant under scaling of a : Scaling a will multiply all right-hand sides of the inequalities (11) by the same constant, and will therefore simply scale the polytope of solutions p of the system by the same factor. (The additional hyperplane constraint is a homogeneous equation, and thus unaffected by scaling.) But the same factor appears in the claimed bound on $|p|$. Therefore, it is sufficient to calculate the vertices for any convenient value of a and prove that $\|p\|/a \leq 6nk$. This implies that the same relation will hold for all $a \geq 0$.

If we omit any of the first $2k$ constraints that correspond to the vectors (12), this leads to a symmetric set of $2k$ vertices p of the form

$$(\pm x_1, y_1, 0, y, \dots, 0, y, z), (0, y, \pm x_1, y_1, \dots, 0, y, z), \dots, (0, y, 0, y, \dots, \pm x_1, y_1, z),$$

whose coordinates and associated value of a can be calculated as

$$\begin{aligned}x_1 &= \sigma(4k + \sigma^2)/\eta \\y_1 &= -(4(k-1) + \sigma^2) \\y &= 4 \\z &= -2\sigma \\a &= \sigma(4 - 2h + w\sigma)\end{aligned}$$

To bound $\|p\|$, we individually bound each coordinate. Since there is no point in optimizing the constant factors for the reduction, we will mostly be very generous and use only crude bounds like $|\sigma| \leq 1$ and $|w| \leq 1/2$ and $\eta \geq 1/2$.

$$\begin{aligned}|x_1| &= |\sigma(4k + \sigma^2)/\eta| \leq (4k + 1)/\frac{1}{2} = 8k + 2 \leq 10k \\|y_1| &= |4(k-1) + \sigma^2| \leq 4(k-1) + 1 \leq 4k \\|z| &= |2\sigma| \leq 2 \\a &= \sigma(4 - 2h + w\sigma) \geq \sigma(4 - 2 \cdot 1) = 2\sigma \geq \frac{2}{n}\end{aligned}$$

We get

$$\|p\|^2 = x_1^2 + y_1^2 + (k-1)y^2 + z^2 \leq 100k^2 + 16k^2 + 16(k-1) + 4 < 132k^2 < (12k)^2$$

and thus $\|p\| \leq 12k$. The ratio $\|p\|/a$ is bounded by $6nk$.

The last vertex p , for which all constraints corresponding to (12) are fulfilled, has coordinates

$$p = (0, \frac{\sigma}{k}, 0, \frac{\sigma}{k}, \dots, 0, \frac{\sigma}{k}, 2).$$

(This vector is the vector from the centroid of $A(1, \dots, 1)$ to p^z .) Its associated value of a is $a = 2h + \sigma^2/k - \sigma w \geq 2 \cdot \frac{3}{4} - 1 = 1/2$.

$$\|p\| = \sqrt{k(\frac{\sigma}{k})^2 + 4} = \sqrt{\sigma^2/k + 4} \leq \sqrt{5} < 3$$

Thus, for this vertex, the ratio $\|p\|/a$ is bounded by $6 < 6nk$. \square