Separable Partitions

Noga Alon [*]   Shmuel Onn [†]

Abstract

An ordered partition of a set of n points in the d dimensional Euclidean space is called a *separable partition* if the convex hulls of the parts are pairwise disjoint. For each fixed p and d we determine the maximum possible number $r_{p,d}(n)$ of separable partitions into p parts of n points in real d-space up to a constant factor. Of particular interest are the values $r_{p,d}(n) = \Theta(n^{d\binom{p}{2}})$ for every fixed p and $d \geq 3$, and $r_{p,2}(n) = \Theta(n^{6p-12})$ for every fixed $p \geq 3$. We establish similar results for spaces of finite Vapnik-Chervonenkis dimension and study the corresponding problem for points on the moment curve as well.

1   Introduction

A *separable p-partition* of a set of n points in the d-dimensional Euclidean space R^d is an ordered tuple $\pi = (\pi_1, \ldots, \pi_p)$ of p nonempty sets whose disjoint union is S, where the convex hulls of the sets π_j are pairwise disjoint. Let $r_{p,d}(n)$ denote the maximum possible number of separable p partitions of a set of n points in R^d. It is easy to see that $r_{p,1}(n) = p!\binom{n-1}{p-1} = \Theta(n^{p-1})$. The following theorem determines the asymptotic behavior of $r_{p,d}(n)$ for every fixed p, d, when n is large.

Theorem 1.1

- *For every fixed $p \geq 2$, $r_{p,1}(n) = \Theta\left(n^{p-1}\right)$.*

- *$r_{2,2}(n) = \Theta\left(n^2\right)$ and for every fixed $p \geq 3$, $r_{p,2}(n) = \Theta\left(n^{6p-12}\right)$.*

- *For every fixed $p \geq 2$ and $d \geq 3$,*

$$r_{p,d}(n) = \Theta\left(n^{d\binom{p}{2}}\right).$$

We also obtain similar estimates for the maximum possible number of separable partitions in spaces of finite Vapnik-Chervonenkis dimension [16]. Here are the relevant definitions. A *space* is a pair $(X, \mathcal{H})$ with $\mathcal{H}$ a collection of subsets of a set X. A p-partition $\pi = (\pi_1, \ldots, \pi_p)$ of a subset $S \subseteq X$ in a space $(X, \mathcal{H})$ is an ordered partition of S into p pairwise disjoint parts. It is *separable* if for each pair $1 \leq r < s \leq p$ there is a member $H_{r,s} \in \mathcal{H}$ such that $\pi_r \subseteq H_{r,s}$ and $\pi_s \subseteq X \setminus H_{r,s}$. A subset $S \subseteq X$ is *shattered* if all 2-partitions of S are separable. A space $(X, \mathcal{H})$ has finite *VC-dimension* d if X contains a shattered d-subset but not a shattered $(d+1)$-subset. An important example, which is of major concern in this article, is provided by real Euclidean d-space, which is the space $(X, \mathcal{H})$ with $X = \mathbb{R}^d$ and $\mathcal{H}$ the collection of closed halfspaces in $\mathbb{R}^d$. In this space, a partition is separable if and only if the convex hulls of the parts are pairwise disjoint, as defined earlier. The VC-dimension of this space is $d+1$.

Let $v_{p,d}(n)$ denote the maximum possible number of separable p-partitions of a set of n points in a space of finite VC-dimension d. We provide an upper bound on $v_{p,d}(n)$ which, together with the lower bound on $r_{p,d}(n)$ from Theorem 1.1, gives the following statement.

Theorem 1.2 *For every fixed $p \geq 2$ and $d \geq 4$,*

$$\Omega \left(n^{(d-1)\binom{p}{2}} \right) \leq v_{p,d}(n) \leq O \left(n^{d\binom{p}{2}} \right).$$

The special case $p = 2$ of 2-partitions in the real Euclidean space had been considered thirty years ago by Harding [7], who proved that $r_{2,d}(n) = 2 \sum_{i=1}^{d} \binom{n-1}{i} = \Theta(n^d)$ for all d. The case $p = d = 2$ of 2-partitions in the real plane under additional size constraints has been extensively studied ever since [9]. The upper bound in Theorem 1.1 for real Euclidean space with arbitrary p and $d \geq 3$ has been recently derived in [8] in the course of the study of a broad class of hard optimization problems over partitions. For $p = 2$ and an arbitrary space of finite VC-dimension d, an upper bound is available through the so-called Sauer's lemma [14].

The article is organized as follows. In the next Section we provide our lower bounds for the numbers $r_{p,d}(n)$. In Section 3 we describe the upper bounds in Theorems 1.1 and 1.2, and discuss the related problem for spaces of finite VC dimension. In Section 4 we study the problem of estimating the number of separable partitions of sets of points on the moment curve in $\mathbb{R}^d$ and show that it is closely related to the study of long Davenport-Schinzel sequences. The final Section 5 contains some concluding remarks.

2 The lower bounds

Throughout this section we restrict attention to the real Euclidean space. We assume some familiarity with the rudiments of convex polytopes theory (such as in [6, 17]), but include a compact description of all notions and facts that we use.

An *orientation* of a hyperplane H in $\mathbb{R}^d$ is a designation of the closed and open halfspaces $H^{\leq}, H^{<}$ *below* it and the closed and open halfspaces $H^{\geq}, H^{>}$ *above* it. A presentation of H as the zero set of a linear form $H = \{x \in \mathbb{R}^d : h_0 + h^T x = 0\}$ gives an orientation of

H in the obvious way. The hyperplane spanned by an *ordered* list $(v^1, \ldots, v^d)$ of d affinely independent points will be oriented by the linear form $\det[\bar{x}, \bar{v}^1, \ldots, \bar{v}^d]$, where $\bar{v}$ is used to denote the vector in $\mathbb{R}^{d+1}$ obtained by appending a first coordinate 1 to the vector $v \in \mathbb{R}^d$. We need the following building block. Let $H = \{x \in \mathbb{R}^d : h_0 + h^T x = 0\}$ be an oriented hyperplane in $\mathbb{R}^d$, let v be a point in H, let m be a positive integer and let $\epsilon > 0$. A set $S \subset \mathbb{R}^d$ will be called an (H, v, m, ϵ)-*set* if S is contained in the ball $B(v, \epsilon)$ of radius ϵ about v and is of the form

$$S = \cup_{i=1}^d \{v_i + j \cdot \delta \cdot h : j = 1, \ldots, m\}$$

where $(v_1, \ldots, v_d)$ is an ordered list of affinely independent points which span the oriented hyperplane H, and δ is a positive real. Thus, S consists of md points ϵ-close to v and above H, evenly spread on d parallel lines orthogonal to H. The *canonical partition of S associated with* $1 \le j_1, \ldots, j_d \le m$ is defined to be the 2-partition (π_1, π_2) of S with $\pi_1 = \cup_{i=1}^d \{v_i + j \cdot \delta \cdot h : 1 \le j \le j_i\}$ and $\pi_2 = S \setminus \pi_1$. The *canonical hyperplane of S associated with* $1 \le j_1, \ldots, j_d \le m$ is defined to be the oriented hyperplane spanned by the list $(v_1 + j_1 \cdot \delta \cdot h, \ldots, v_d + j_d \cdot \delta \cdot h)$. The verification of the following simple proposition is left to the reader.

Proposition 2.1 *Let S be an (H, v, m, ϵ)-set and let $\hat{H}$ and (π_1, π_2) be, respectively, the canonical hyperplane and the canonical partition of S associated with $1 \le j_1, \ldots, j_d \le m$. Then $\pi_1 \subset \hat{H}^{\le}$ and $\pi_2 \subset \hat{H}^{>}$. In particular, every canonical partition of S is separable.*

A *polytopal complex* is a nonempty finite collection $\mathcal{P}$ of convex polytopes in some $\mathbb{R}^d$ such that the face of any member of $\mathcal{P}$ is also in $\mathcal{P}$ and such that the intersection of any two members of $\mathcal{P}$ is a common face of both. A polytopal complex is *pure d-dimensional* if all (inclusion) maximal polytopes of $\mathcal{P}$ have the same dimension d. Two maximal polytopes in a pure d-dimensional polytopal complex are *adjacent* if their intersection is $(d-1)$-dimensional, i.e. a facet of both. The *graph $G(\mathcal{P})$* of a pure polytopal complex is the graph whose vertices are the maximal polytopes of $\mathcal{P}$ and whose edges are the pairs of adjacent maximal polytopes. We define a (p, d)-*complex* to be a pure d-dimensional polytopal complex embedded in $\mathbb{R}^d$ and containing p maximal polytopes.

Lemma 2.2 *For any fixed p, d, k, if the graph of some (p, d)-complex contains k edges then $r_{p,d}(n) = \Omega(n^{dk})$.*

Proof. Let $\mathcal{P}$ be a (p, d)-complex whose graph contains k edges. Let $P_1, \ldots, P_p$ be the maximal polytopes in $\mathcal{P}$. For each pair $1 \le r \ne s \le p$, let $F_{rs} = F_{sr} := P_r \cap P_s$ which is a common face of P_r and P_s, and let $H_{rs} = H_{sr}$ be a hyperplane such that $H_{rs} \cap P_r = H_{rs} \cap P_s = F_{rs}$ (note that if P_r, P_s are adjacent, then F_{rs} is a common facet of both and H_{rs} is uniquely defined). Let $H_{r,s}$ be the orientation of H_{rs} with P_r below $H_{r,s}$ and P_s above it and let $H_{s,r}$ be the opposite orientation.

For each r let $N_r := \{s : P_r, P_s \text{ adjacent}\}$ be the set of indices of neighbors of P_r. For each adjacent pair P_r, P_s, choose a point $v_{rs} = v_{sr}$ in the relative interior of F_{rs}. Then for all $1 \le r \ne s \le p$ the following hold: every point v_{ri} with $i \in N_r \setminus \{s\}$ lies strictly below $H_{r,s}$,

and every point v_{js} with $j \in N_s \setminus \{r\}$ lies strictly above $H_{r,s}$. For any given m it is therefore possible to choose an $\epsilon > 0$ sufficiently small, and an $(H_{r,s}, v_{rs}, m, \epsilon)$-set $S_{rs} = S_{sr}$ for each adjacent pair P_r and P_s with $r < s$, such that the following hold for all $1 \le r < s \le p$:

- Every set S_{ri} with $i \in N_r \setminus \{s\}$ lies strictly below $H_{r,s}$, and every set S_{js} with $j \in N_s \setminus \{r\}$ lies strictly above $H_{r,s}$.

- If P_r and P_s are adjacent and $\hat{H}$ is the canonical hyperplane of S_{rs} associated with any $1 \le j_1, \ldots, j_d \le m$ then every set S_{ri} with $i \in N_r \setminus \{s\}$ lies strictly below $\hat{H}$ and every set S_{js} with $j \in N_s \setminus \{r\}$ lies strictly above $\hat{H}$.

Now, let S be the union of all the S_{rs}. So S consists of mdk points. For each adjacent pair P_r, P_s with $r < s$ choose a canonical partition $\pi^{r,s} := (\pi_1^{r,s}, \pi_2^{r,s})$ of S_{rs} and let $\hat{H}_{r,s}$ be the corresponding canonical hyperplane which separates it. Now define a p-partition $\pi = (\pi_1, \ldots, \pi_p)$ of S as follows: for $i = 1, \ldots, p$ let

$$\pi_i := \left(\bigcup_{j \in N_i, \; j > i} \pi_1^{i,j} \right) \bigcup \left(\bigcup_{j \in N_i, \; j < i} \pi_2^{j,i} \right).$$

We claim that π is a separable partition of S. Indeed, it follows from the discussion above together with Proposition 2.1 that, for each pair $r < s$, we have that if ϵ is sufficiently small π_r and π_s are separated by $H_{r,s}$ if P_r and P_s are not adjacent, and by $\hat{H}_{r,s}$ if P_r and P_s are adjacent.

As there are m^d canonical partitions for each S_{rs}, we obtain this way m^{dk} separable p-partitions of the mdk-set S.

Now, given any positive integer n, let $m := \lfloor \frac{n}{dk} \rfloor$. Then, as claimed,

$$r_{p,d}(n) \ge r_{p,d}(mdk) \ge m^{dk} \ge \left(\frac{n}{2dk} \right)^{dk} = \Omega(n^{dk}). \quad \square$$

We proceed to describe a simple construction of (p, d)-complexes with dense graphs for all p and d. In particular, for every p and every $d \ge 3$ there is a (p, d)-complex whose graph is complete, a fact first established half a century ago by Rado [12].

In what follows, $G(P)$ denotes the graph of 0-faces and 1-faces of a convex polytope P.

Lemma 2.3 *Let P be a $(d+1)$-polytope with $p+1$ vertices and let v be an arbitrary vertex in its graph $G(P)$. Then there exists a (p, d)-complex $\mathcal{P}$ with $G(\mathcal{P})$ isomorphic to $G(P) - v$.*

Proof. Assume without loss of generality that P is full dimensional, let a be any interior point of P, and let Q be a *polar* of P,

$$Q := (P - a)^* = \{x \in \mathbb{R}^{d+1} : \; (y - a)^T x \le 1 \; \text{ for all } \; y \in P\}.$$

Then the face lattice of Q is the poset-dual of the face lattice of P. In particular, there is a bijection $v \mapsto F_v$ from the vertices of P to the facets of Q.

Let v be an arbitrary vertex of P and let F_v be the facet of Q corresponding to v. Define the polytopal complex $\mathcal{P}$ to be a *Schlegel* diagram of Q at F_v (see [6, 17] for details). Briefly, it is defined as follows. Let $H = \mathrm{aff}(F_v) = \{x : h_0 + h^T x = 0\}$ be the hyperplane supporting Q at F_v oriented so that $Q \subset H^{\leq}$. Choose any point b in the relative interior of F_v and let $u := b + \epsilon \cdot h \in H^{>}$, with $\epsilon > 0$ sufficiently small, so that for every point $x \in Q \setminus F_v$, the intersection point x' of the line segment $[u, x]$ with H is in the relative interior of F_v. The Schlegel diagram of Q at F_v is the (p, d)-complex $\mathcal{P}$ whose polytopes are the images of all proper faces of Q but F_v under the radial projection $x \mapsto x'$ (transformed by an affine map taking H onto $\mathbb{R}^d$). The face poset of $\mathcal{P}$ is then isomorphic to the poset of all proper faces of Q but F_v. Since the face lattice of P determines its graph $G(P)$ and the face poset of $\mathcal{P}$ determines its graph $G(\mathcal{P})$, it follows that $G(\mathcal{P})$ is isomorphic to $G(P) - v$ as desired. $\square$

We can now obtain our lower bounds.

Lemma 2.4 *For all $d \geq 3$ we have $r_{p,d}(n) = \Omega(n^{d\binom{p}{2}})$.*

Proof. Let $C(p, d) := \mathrm{conv}\{M_d(1), \ldots, M_d(p)\}$, with $M_d(i) := [i, i^2, \ldots, i^d]$, denote the *cyclic polytope* with p vertices in $\mathbb{R}^d$. It is well known (cf. [17]) that the graph of $C(p, d)$ is complete for all $d \geq 4$. Now, given any $p \geq 1$ and $d \geq 3$, let $P := C(p+1, d+1)$, and let v be any vertex in $G(P)$. By Lemma 2.3, there is a (p, d)-complex $\mathcal{P}$ whose graph $G(\mathcal{P})$ is isomorphic to $G(P) - v$ and hence is complete and has $\binom{p}{2}$ edges. The bound follows by Lemma 2.2. $\square$

Lemma 2.5 *For all $p \geq 3$ we have $r_{p,2}(n) = \Omega(n^{6p-12})$.*

Proof. First, we note that for every $p \geq 3$ there is a graph G_p with the following properties: it is planar; it is 3-connected; it is simplicial (i.e., all its faces are triangles in every planar embedding); it has $p + 1$ vertices and $3(p - 1)$ edges; and it has a vertex of degree 3. Clearly $G_3 := K_4$, the 4-clique, satisfies these properties. Proceeeding by induction, suppose G_p has been constructed and embedded in the plane. Choose any triangular face, insert a new vertex and connect it to each of the three vertices of that triangle. Clearly, this new graph G_{p+1} has again all desired properties. Now, each 3-connected planar graph is isomorphic to the graph of a 3-polytope by Steinitz' Theorem (see e.g. [10]). Given $p \geq 3$, let P be a 3-polytope with $p + 1$ vertices whose graph $G(P)$ is isomorphic to G_p, and choose a vertex v of degree 3 in that graph. By Lemma 2.3, there is a $(p, 2)$-complex $\mathcal{P}$ whose graph $G(\mathcal{P})$ is isomorphic to $G(P) - v$. Since G_p hence $G(P)$ have $3(p - 1)$ edges and v has degree 3 in $G(P)$, the graph of the $(p, 2)$-complex $\mathcal{P}$ has $3p - 6$ edges. The bound follows by Lemma 2.2. $\square$

3   The upper bounds

We now derive our upper bounds on $r_{p,d}(n)$ and $v_{p,d}(n)$. We start with real Euclidean space.

Lemma 3.1 *For any fixed p, d, k, if for every collection $\mathcal{S}$ of p compact, convex, pairwise disjoint sets in real d-space, there is a collection $\mathcal{H}$ of k hyperplanes such that any two members of $\mathcal{S}$ are separated by at least one hyperplane of $\mathcal{H}$, then $r_{p,d}(n) = O(n^{dk})$.*

Proof. Suppose that the hypothesis holds for p, d, k. Consider any set S of n points in $\mathbb{R}^d$, and consider any separable p-partition $\pi = (\pi_1, \ldots, \pi_p)$ of S. By the hypothesis, there is a collection of k hyperplanes that separate each pair among $\text{conv}(\pi_1), \ldots, \text{conv}(\pi_p)$. Clearly we may assume that each of these hyperplanes contains no point of S. Thus, by Harding's theorem mentioned above there are only $O(n^d)$ choices for each hyperplane. Suppose, now, that we are given S and the collection $\mathcal{H}$ of the k hyperplanes together with the information, for each of the convex hulls $\text{conv}(\pi_i)$ and each of the hyperplanes H in $\mathcal{H}$, if $conv(\pi_i)$ intersects H and in case it does not, in which side of H it lies. Then we can easily reconstruct the whole partition, as each π_i is simply the intersection of S with all the corresponding half spaces supported by members of $\mathcal{H}$ which contain it. It thus follows that the total number of separable p-partitions of S is bounded by 3^{pk} (which is fixed for fixed p, k) times the number of distinct collections $\mathcal{H}$ of k hyperplanes, which is bounded by $O(n^{dk})$, by Harding's Theorem. □

Since $\binom{p}{2}$ hyperplanes always suffice to pairwise separate p compact, convex, pairwise disjoint sets in any real space, an immediate consequence of this lemma is the following upper bound, which was first proved (by a slightly different argument) in [8, Lemma 4.1].

Lemma 3.2 *For all d and $p \geq 2$ we have $r_{p,d}(n) = O(n^{d\binom{p}{2}})$.*

We now turn to derive the following tighter upper bound for the real plane.

Lemma 3.3 *For all fixed $p \geq 3$ we have $r_{p,2}(n) = O(n^{6p-12})$.*

This lemma will follow at once from Lemma 3.1 and the following result, which is proved implicitly in [5] and [4] (see also [11]), and also follows from Lemma 8 of [2]. For the sake of completeness we sketch a proof.

Lemma 3.4 *For every collection $\mathcal{S}$ of p compact, convex, pairwise disjoint sets in the plane, there is a collection $\mathcal{H}$ of $3p - 6$ lines such that any two sets of $\mathcal{S}$ are separated by at least one line of $\mathcal{H}$.*

Proof. We begin by constructing pairwise disjoint circumscribing polygons around the sets in $\mathcal{S}$ and then we circumscribe a triangle T around all these polygons. Let A denote the set of these p polygons. It is convenient to assume that all directions of the sides of A and T have pairwise distinct slopes. Next we grow the polygons in A so as to maximize their area, thus obtaining polygons with overlapping boundaries but disjoint interiors. This is done by moving the sides of the polygons, one by one, until each polygon is of maximal area subject to the interiors of the polygons being disjoint, and subject to staying within the triangle T. The precise expansion process is done by choosing, arbitrarily, a side of a polygon and moving the corresponding half-plane in the direction perpendicular to the side

6

and away from the polygon's interior. The side stops moving further only when it touches another polygon's corner or it reaches the boundary of T or when it shrinks to a point and vanishes.

Once this process is finished, observe that the set of all lines containing all sides of the polygons without the lines containing the edges of T can serve as our separating set $\mathcal{H}$. Thus it suffices to show that the total number of such lines is at most $3p-6$. To this end we define, following [4], a graph whose vertices are all polygons, where two are adjacent iff a side of one of them intersects the boundary of another, where the side is considered here as a relatively open set (i.e., it does not contain its endpoints). As proved in Lemma 1 of [4] this graph is planar and hence its number of edges is at most $3p-6$. Moreover, as shown in Lemma 2 of [4] there is a simple way to embed this graph in the plane so that each line containing a side of a polygon which is not part of a side of T is crossed by at least one edge of this graph, and each edge of the graph crosses only one such line. This supplies the desired bound and completes the proof of the lemma. □

Another interesting consequence of Lemmas 2.4, 2.5, 3.1 and 3.4, is the following result.

Proposition 3.5 *Let $s(p,d)$ be the smallest number k such that for every collection $\mathcal{S}$ of p compact, convex, pairwise disjoint sets in real d-space, there is a collection $\mathcal{H}$ of k hyperplanes such that any two members of $\mathcal{S}$ are separated by at least one hyperplane of $\mathcal{H}$. Then:*

- *For every fixed $p \geq 2$, $s(p,1) = p-1$.*

- *$s(2,2) = 1$ and for every fixed $p \geq 3$, $s(p,2) = 3p-6$.*

- *For every fixed $p \geq 2$ and $d \geq 3$, $s(p,d) = \binom{p}{2}$.*

Note that the construction in Section 2 is not really needed in order to prove that for $p \geq 2$ and $d \geq 3$, $s(p,d) = \binom{p}{2}$. Indeed, it is easier to observe that a collection of p lines in general position in R^3 cannot be separated, in the sense of Lemma 3.1, by less then $\binom{p}{2}$ planes, and it is a simple matter to replace the lines with compact sets and to extend the result for higher dimensions as well. Note, also, that the assertion of Lemma 3.1 can be strengthened, with essentially the same proof, yielding the following result.

Proposition 3.6 *We say that a collection $\mathcal{H}$ of hyperplanes* separates *a collection $\mathcal{S}$ of compact, convex pairwise disjoint sets, if for any $S \in S$ the intersection of all half spaces bounded by an element of $\mathcal{H}$ which contain S contains no point of any other member of $\mathcal{S}$. If for some fixed p, d, k, every collection $\mathcal{S}$ of p compact, convex, pairwise disjoint sets in real d-space, is separated by a collection $\mathcal{H}$ of k hyperplanes, then $r_{p,d}(n) = O(n^{dk})$.*

Therefore, the construction in Section 2 provides examples of p compact, convex, pairwise disjoint sets in R^d which cannot be separated by less than $\binom{p}{2}$ hyperplanes even according to the separation as defined in the last proposition.

We proceed to derive an upper bound on $v_{p,d}(n)$. We use the following construction, which is similar to the one used in [8] for the real Euclidean space. Let S be any set and let $p \geq 2$. With each list $(\pi^{r,s} = (\pi_1^{r,s}, \pi_2^{r,s}) : 1 \leq r < s \leq p)$ of $\binom{p}{2}$ 2-partitions of S associate a p-tuple $\pi = (\pi_1, \ldots, \pi_p)$ of subsets of S as follows: for $i = 1, \ldots, p$ put

$$\pi_i := \left(\bigcap_{j=i+1}^{p} \pi_1^{i,j} \right) \bigcap \left(\bigcap_{j=1}^{i-1} \pi_2^{j,i} \right).$$

Since $\pi_r \subseteq \pi_1^{r,s}$ and $\pi_s \subseteq \pi_2^{r,s}$ for all $1 \leq r < s \leq p$, the π_i are pairwise disjoint. If moreover $\cup_{i=1}^{p} \pi_i = S$ then π will be called the p-partition of S associated with the given list. The following lemma extends the one provided in [8] to any space.

Lemma 3.7 *For any space $(X, \mathcal{H})$ and $p \geq 2$, the set of separable p-partitions of any subset S of X equals the set of p-partitions associated with lists of $\binom{p}{2}$ separable 2-partitions of S.*

Proof. Fix a subset S of X. Consider any p-partition π associated with a list of $\binom{p}{2}$ separable 2-partitions of S. For each pair $1 \leq r < s \leq p$, since $\pi^{r,s}$ is a separable 2-partition, there is an $H_{r,s} \in \mathcal{H}$ such that $\pi_r \subseteq \pi_1^{r,s} \subseteq H_{r,s}$ and $\pi_s \subseteq \pi_2^{r,s} \subseteq X \setminus H_{r,s}$. It follows that π is separable. Conversely, let $\pi = (\pi_1, \ldots, \pi_p)$ be any separable p-partition. Consider any pair $1 \leq r < s \leq p$. By definition, there exists an $H_{r,s} \in \mathcal{H}$ such that $\pi_r \subseteq H_{r,s}$ and $\pi_s \subseteq X \setminus H_{r,s}$. Let $\pi^{r,s} := (\pi_1^{r,s}, \pi_2^{r,s})$ be the separable 2-partition of S defined by $\pi_1^{r,s} := S \cap H_{r,s}$ and $\pi_2^{r,s} := S \setminus H_{r,s}$. Let π' be the p-tuple associated with the list of $\pi^{r,s}$ obtained that way. Then the π_i' are pairwise disjoint and for $i = 1, \ldots, p$ we have

$$\pi_i \subseteq \left(\bigcap_{j=i+1}^{p} \pi_1^{i,j} \right) \bigcap \left(\bigcap_{j=1}^{i-1} \pi_2^{j,i} \right) = \pi_i'.$$

Since

$$S = \cup_{i=1}^{p} \pi_i \subseteq \cup_{i=1}^{p} \pi_i' \subseteq S,$$

it follows that equality holds hence $\pi = \pi'$ is the p-tuple associated with the constructed list of $\binom{p}{2}$ separable 2-partitions. $\square$

This lemma allows to extend upper bounds on the number of 2-partitions in any space to upper bounds on the number of p-partitions in that space. For instance, it implies that any set in any real space has at most one separable p-partition per each list of $\binom{p}{2}$ separable 2-partitions, hence $r_{p,d}(n) \leq r_{2,d}(n)^{\binom{p}{2}}$. This together with the known bound $r_{2,d}(n) = O(n^d)$ gives a second proof (which is the one provided in [8]) of Lemma 3.2 above.

For spaces of finite VC-dimension, the so-called Sauer's lemma, suitably rephrased, provides the following upper bound on the number of separable 2-partitions.

Lemma 3.8 *[14]. Let $(X, \mathcal{H})$ be a space of VC-dimension d. Then the number of separable 2-partitions of any set of n points in X is at most $\sum_{i=0}^{d} \binom{n}{i} = O(n^d)$.*

We obtain the following upper bound.

Lemma 3.9 *For every p, d, the maximum number of separable p-partitions of any set of n points in any space of VC-dimension d satisfies*

$$v_{p,d}(n) = O\left(n^{d\binom{p}{2}}\right).$$

Proof. By Lemma 3.7, any n-subset of X has at most one separable p-partition per each list of $\binom{p}{2}$ separable 2-partitions. Since the number of separable 2-partitions of any n-subset is $O(n^d)$ by Lemma 3.8, the bound follows. ☐

We can now combine all the necessary ingredients and obtain Theorems 1.1 and 1.2.

Proof of Theorem 1.1. Follows from Lemmas 2.4, 2.5, 3.2 and 3.3. ☐

Proof of Theorem 1.2. Follows from Lemma 2.4, the fact that real Euclidean d-space has VC-dimension $d + 1$, and Lemma 3.9. ☐

4 Separable partitions on the moment curve

Recall that the *moment curve* in $\mathbb{R}^d$ is the image of the map

$$M_d : \mathbb{R} \longrightarrow \mathbb{R}^d : t \mapsto \begin{bmatrix} t \\ t^2 \\ \vdots \\ t^d \end{bmatrix}.$$

The moment curve is totally ordered in the obvious way. Therefore, we can and will identify any set S of n points on the moment curve with $[n] = \{1, \ldots, n\}$, and partitions of S with partitions of $[n]$. We will regard a p-partition π of $[n]$ also as the function $\pi : [n] \longrightarrow [p]$ defined by $\pi(j) = i$ for all $j \in \pi_i$ and as the sequence $\pi = [\pi(1), \ldots, \pi(n)]$. We shall move freely among these representations of π.

Let $m_{p,d}(n)$ denote the maximum number of separable p-partitions of a set of n points on the moment curve in R^d. In this section we first provide a combinatorial characterization of separable partitions of point sets on the moment curve, and then use it in estimating $m_{p,d}(n)$. Our characterization implies that, in fact, *any* set of n points on the moment curve admits the same number $m_{p,d}(n)$ of separable p-partitions.

A (p, d)-*sequence* is a sequence $\pi = [\pi(1), \ldots, \pi(n)]$ with $\pi(i) \in [p]$ for all $i \in [n]$, and with the property that for any pair $1 \leq r \neq s \leq p$, the length of any subsequence $[r, s, r, s, \ldots]$ of π whose elements alternate between r and s is at most $d + 1$. For instance,

$$[1, 2, 1, 2, 1, 3, 1, 3, 1, 3, 2, 3, 2, 3] \tag{1}$$

9

is a $(3,5)$-sequence with a (non unique) longest alternating subsequence $[1,2,1,2,1,2]$. The following lemma links the geometric property of a separable p-partition on the moment curve in $\mathbb{R}^d$ to the combinatorial property of a (p,d)-sequence.

Lemma 4.1 *A p-partition $\pi = (\pi_1, \ldots, \pi_p)$ of a set of n points on the moment curve in $\mathbb{R}^d$ is separable if and only if $[\pi(1), \ldots, \pi(n)]$ is a (p,d)-sequence.*

Proof. Let $\pi = (\pi_1, \ldots, \pi_p)$ be a p-partition of a set of n points on the moment curve in $\mathbb{R}^d$. The claim being trivial for $p = 1$ assume that $p \geq 2$. Now, π is separable if and only if for each pair $1 \leq r \neq s \leq p$ the 2-partition (π_r, π_s) is separable. It suffices then to show that for each such pair, (π_r, π_s) is separable if and only if the subsequence of π consisting of all occurrences of r, s does not contain an alternating subsequence $[r, s, r, s, \ldots]$ of length $d + 2$. This is essentially known, see, e.g., [1]. For completeness, we include the short proof.

Consider any such pair $1 \leq r \neq s \leq p$, and let $V := \pi_r \cup \pi_s$. Then $V = \{v_1, \ldots, v_m\}$ with $v_i = M_d(t_i)$ for some $t_1 < \ldots < t_m$. Let $\mu := [\mu(1), \ldots, \mu(m)]$ denote the sequence which is the restriction of π to V, so that $\mu(i) = r$ if $v_i \in \pi_r$ and $\mu(i) = s$ if $v_i \in \pi_s$. If μ contains an alternating subsequence of length $d + 2$, then it is obvious that π_r and π_s cannot be separated, since any hyperplane intersects the moment curve in at most d points. On the other hand, if there is no such subsequence, then there are d real numbers $y_1 < y_2 < y_3 \ldots < y_d$ whose images on the moment curve split it into $d + 1$ intervals, so that each element of π_r lies in one of the even intervals and each element of π_s lies in one of the odd intervals. Put $\prod_{i=1}^d (t - y_i) = c_0 + c_1 t + \ldots + c_{d-1} t^{d-1} + t^d$. It is easy to check that the hyperplane $H = \{(x_1, x_2, \ldots, x_d) : c_0 + c_1 x_1 + c_2 x_2 + \ldots + c_{d-1} x_{d-1} + x_d = 0\}$ separates π_r and π_s.

Concluding, we see that for each pair $1 \leq r \neq s \leq p$, the 2-partition (π_r, π_s) is separable if and only if the length of any alternating subsequence $[r, s, r, s, \ldots]$ of π is at most $d + 1$, and so $\pi = (\pi_1, \ldots, \pi_p)$ is separable if and only if $[\pi(1), \ldots, \pi(n)]$ is a (p,d)-sequence. $\square$

Lemma 4.1 implies that the number of (p,d)-sequences of length n is equal to $m_{p,d}(n)$. We proceed to estimate the asymptotics of this number.

A *$DS(p,d)$-sequence*, or *Davenport-Schinzel (p,d)-sequence* (termed so in [13] and introduced in [3]), is a (p,d)-sequence $\pi = [\pi(1), \ldots, \pi(n)]$ in which $\pi(i+1) \neq \pi(i)$ for all i. For example, the sequence in equation (1) above is a $DS(3,5)$-sequence of maximum possible length 14. A *k-composition* of n is a k-tuple $\alpha = (\alpha_1, \ldots, \alpha_k)$ of positive integers summing up to n. Clearly, each (p,d)-sequence of length n is uniquely obtainable as the replication $[\mu(1)^{\alpha_1}, \ldots, \mu(k)^{\alpha_k}]$ of a $DS(p,d)$-sequence $[\mu(1), \ldots, \mu(k)]$ by a k-composition $(\alpha_1, \ldots, \alpha_k)$ of n. Therefore, an estimate on the total number of $DS(p,d)$-sequences of each length will lead to an estimate on the number $m_{p,d}(n)$ of (p,d)-sequences of length n. But for our purpose here it suffices to consider the replications of a single long $DS(p,d)$-sequence, since the number of such sequences is a function of p and d (and does not depend on n). Let $\lambda_d(p)$ denote the maximum length of any $DS(p,d)$-sequence. It is easy to see that $\lambda_d(p)$ is finite and satisfies $\lambda_d(p) \leq d\binom{p}{2} + 1$. We have the following simple estimate on $m_{p,d}(n)$, where p and d are regarded as fixed parameters.

10

Proposition 4.2 *For every fixed p, d, the number of (p, d)-sequences of length n satisfies*

$$m_{p,d}(n) = \Theta\left(n^{\lambda_d(p)-1}\right).$$

Proof. Let μ be any $DS(p, d)$-sequence of length $k = \lambda_d(p)$. This sequence gives rise to $\binom{n-1}{k-1}$ replications, one for each k-compositions of n. This way, $\binom{n-1}{k-1}$ distinct (p, d)-sequences of length n are obtained. The upper bound follows in a similar manner, using the fact that the number of $DS(p, d)$-sequences is a function of p and d only. □

A (p, d)-sequence is *normal* if the first occurrence of i precedes that of $i+1$ for all i. In [13], the following recursive construction of a long normal $DS(p, d)$-sequence $\mu(p, d)$ was described for all $d \geq p$. For $p = 1$ and any d set $\mu(p, d) = [1]$. For $d \geq p \geq 2$ construct $\mu(p, d)$ as follows: start with a sequence of length $(p-1)\lfloor\frac{d-1}{2}\rfloor + 1$ of 1's; next, insert $(p-1)\lfloor\frac{d-1}{2}\rfloor$ symbols so as to separate each pair of 1's, where the first $\lfloor\frac{d-1}{2}\rfloor$ of the new symbols are p, the next $\lfloor\frac{d-1}{2}\rfloor$ are $p - 1$, and so on; then, append to the right the normal sequence $\mu(p - 1, d - 1)$ on the symbols $\{2, 3, \ldots, p\}$, and, if d is even, an additional last symbol 1 to its right; finally, apply the necessary permutation of $[p]$ to the elements of the sequence so as to make it normal. For example, the sequence in equation (1) above is precisely the normal $DS(3, 5)$-sequence $\mu(3, 5)$ obtained this way. With some care, the results of [13] can be shown to imply the following bound on the length of $\mu(p, d)$ and hence on the value of $\lambda_d(p)$.

Proposition 4.3 *The bound $\lambda_d(p) \geq (d - \frac{1}{2}p)\binom{p}{2} + 1$ holds for all $d \geq p$.*

Propositions 4.2 and 4.3 give the following lower bound on $m_{p,d}(n)$.

Proposition 4.4 *For every fixed p, d with $d \geq p$,*

$$m_{p,d}(n) = \Omega\left(n^{(d-\frac{1}{2}p)\binom{p}{2}}\right).$$

Theorem 4.5 *Let $0 < \epsilon \leq \frac{1}{2}$ be any constant. For every p, d with $d \geq \frac{1}{2\epsilon}p$, the number $m_{p,d}(n)$ of separable p-partitions of any set of n points on the moment curve in real d-space satisfies*

$$\Omega\left(n^{(1-\epsilon)d\binom{p}{2}}\right) \leq m_{p,d}(n) \leq O\left(n^{d\binom{p}{2}}\right).$$

Proof. If $0 < \epsilon \leq \frac{1}{2}$ and $d \geq \frac{1}{2\epsilon}p$ then $d \geq p$ and $d - \frac{1}{2}p \geq (1 - \epsilon)d$. The lower bound then follows from Proposition 4.4. The upper bound follows from Lemma 3.2 and the inequality $m_{p,d}(n) \leq r_{p,d}(n)$. □

5 Conclusions and remarks

Theorem 1.1 determines the asymptotic behavior of $r_{p,d}(n)$ for all fixed admissible p and d up to a constant factor. However, the bounds for spaces of finite VC-dimension d given in

Theorem 1.2 are not that tight. It might be interesting to close the gap between the lower and upper bounds here.

For points on the moment curve, the asymptotic behavior of $m_{p,d}(n)$ is reduced, in Proposition 4.2, to the well studied problem of determining or estimating the maximum possible length $\lambda_d(p)$ of a $DS(p,d)$-sequence. The known bounds for this function can be found in [15]. In particular, $\lambda_2(p) = 2p - 2$, $\lambda_3(p) = \Theta(p\alpha(p))$ and $\lambda_4(p) = \Theta(p2^{\alpha(p)})$, where $\alpha(p)$ is the inverse of Ackermann's function. Thus $\lambda_3(p)$ is already a superlinear function of p.

References

[1] N. Alon, P. Frankl and V. Rödl, Geometrical realization of set systems and probabilistic communication complexity, *Proc. 26th Annual Symp. on Foundations of Computer Science (FOCS)*, Portland, Oregon, IEEE, 277–280, 1985.

[2] V. Capoyleas, G. Rote and G. Woeginger, Geometric clustering, *J. Algorithms* 12:341–356, 1991.

[3] H. Davenport and A. Schinzel, A combinatorial problem connected with differential equations, *American Journal of Mathematics*, 87:684–690, 1965.

[4] H. Edelsbrunner, A. D. Robison and X. J. Shen, Covering convex sets with non-overlapping polygons, *Discrete Math.* 81:153–164, 1990.

[5] L. Fejes Tóth, Some packing and covering theorems, *Acta Sci. Math.* 12:62–67, 1950.

[6] B. Grünbaum, *Convex Polytopes*, John Wiley, New York, 1967.

[7] E.F. Harding, The number of partitions of a set of n points in k dimensions induced by hyperplanes, *Proc. Edinburgh. Math. Soc.* 15:285–289, 1967.

[8] F.K. Hwang, S. Onn and U.G. Rothblum, A polynomial time algorithm for vertex enumeration and optimization over shaped partition polytopes, MSRI Preprint Series, Mathematical Sciences Research Institute, Berkeley, CA, 1997-041:1–14, 1997.

[9] L. Lovász, On the number of halving lines, *Annales Universitatis Scientiarum Budapestinensis de Rolando Eötvös Nominatae, Sectio Mathematica* 14:107–108, 1970.

[10] S. Onn and B. Sturmfels, A quantitative Steinitz' theorem, *Beiträge zur Algebra und Geometrie*, 35:125–129, 1994.

[11] J. Pach and P. K. Agarwal, *Combinatorial Geometry*, Wiley, New York, 1995.

[12] R. Rado, A sequence of polyhedra having intersections of specified dimensions, *J. London Math. Soc.* 22:287–289, 1947.

[13] D. Roselle and R. Stanton, Some properties of Davenport-Schinzel sequences, *Acta Arithmetica*, 17:355–362, 1971.

[14] N. Sauer, On the density of families of sets, *J. Combinatorial Theory (A)*, 13:145–147, 1972.

[15] M. Sharir and P.K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*, Cambridge University Press, New York, 1995.

[16] V. Vapnik and A. Chervonenkis, On the uniform convergence of relative frequencies of events to their probabilities, *Theory of Probability and its Applications*, 16:264–280, 1971.

[17] G.M. Ziegler, *Lectures on Polytopes*, Springer-Verlag, 1995.