J. Eur. Math. Soc. 18, 2469–2482

DOI 10.4171/JEMS/645

Imre Bárány · Jiří Matoušek · Attila Pór

Curves in $\mathbb{R}^d$ intersecting every hyperplane at most $d+1$ times

Received October 17, 2013 and in revised form July 8, 2014

Abstract. By a curve in $\mathbb{R}^d$ we mean a continuous map $\gamma\colon I \to \mathbb{R}^d$, where $I \subset \mathbb{R}$ is a closed interval. We call a curve γ in $\mathbb{R}^d$ $(\le k)$-crossing if it intersects every hyperplane at most k times (counted with multiplicity). The $(\le d)$-crossing curves in $\mathbb{R}^d$ are often called *convex curves* and they form an important class; a primary example is the *moment curve* $\{(t, t^2, \dots, t^d) : t \in [0, 1]\}$. They are also closely related to *Chebyshev systems*, which is a notion of considerable importance, e.g., in approximation theory. Our main result is that for every d there is $M = M(d)$ such that every $(\le d+1)$-crossing curve in $\mathbb{R}^d$ can be subdivided into at most M $(\le d)$-crossing curve segments. As a consequence, based on the work of Eliáš, Roldán, Safernová, and the second author, we obtain an essentially tight lower bound for a geometric Ramsey-type problem in $\mathbb{R}^d$ concerning order-type homogeneous sequences of points, investigated in several previous papers.

Keywords. Ramsey function, order type, convex curve, moment curve, Chebyshev system

1. Introduction

The most intuitive statement of the problem investigated in this paper involves curves in $\mathbb{R}^d$. By a curve we mean an arbitrary continuous mapping $\gamma\colon I \to \mathbb{R}^d$, where $I \subset \mathbb{R}$ is a closed interval (we could admit an open interval as well, but this would add unnecessary technical complications). Let us say that a curve γ in $\mathbb{R}^d$ is $(\le k)$-*crossing* if it intersects every hyperplane h at most k times.[1] Here the intersections are counted with multiplicity; that is, the condition of $(\le k)$-crossing reads $|\{t \in I : \gamma(t) \in h\}| \le k$.

I. Bárány: Rényi Institute of Mathematics, Hungarian Academy of Sciences, PO Box 127, 1364 Budapest, Hungary, and Department of Mathematics, University College London, Gower Street, London WC1E 6BT, England; e-mail: barany@math-inst.hu

J. Matoušek: Department of Applied Mathematics, Charles University, Malostranské nám. 25, 118 00 Praha 1, Czech Republic, and Department of Computer Science, ETH Zürich, 8092 Zürich, Switzerland; e-mail: matousek@kam.mff.cuni.cz

A. Pór: Department of Mathematics, Western Kentucky University, 1906 College Heights Blvd. #11078, Bowling Green, KY 42101, USA; e-mail: attila.por@wku.edu

Mathematics Subject Classification (2010): Primary 05D10; Secondary 52A20, 52A37

[1] For algebraic curves in the complex projective space, the number of intersections with a generic hyperplane is the *degree*, but we prefer using a different term, since we deal with much more general curves, which are typically not algebraic.

It will be useful to observe that a $(\leq k)$-crossing curve is not constant on any nonempty open interval, and its image contains no segment.

$(\leq d)$-crossing (= convex) curves. The $(\leq d)$-crossing curves in $\mathbb{R}^d$ are called *convex curves* in a significant part of the literature (e.g., [Arn04, Živ04, SS00, SS05, Mus98]), and they are of considerable interest in several areas. In the plane, a convex curve in this sense is a connected piece of the boundary of a convex set. A primary example of a higher-dimensional convex curve is the *moment curve* $\{(t, t^2, \ldots, t^d) : t \in [0, 1]\}$. The convex hull of $n \geq d + 1$ points on a convex curve in $\mathbb{R}^d$ is a *cyclic polytope*, one of the most important examples in the theory of convex polytopes and in discrete geometry in general.

If we regard a convex curve $\gamma : I \to \mathbb{R}^d$ as a d-tuple $(\gamma_1, \ldots, \gamma_d)$ of functions $I \to \mathbb{R}$, and define $\gamma_0 \equiv 1$, then the $(d+1)$-tuple $(\gamma_0, \gamma_1, \ldots, \gamma_d)$ (or possibly $(-\gamma_0, \gamma_1, \ldots, \gamma_d)$) forms a *Chebyshev system*,[2] which is an important notion in approximation theory, theory of finite moments, and other areas—see, e.g., [KS66, CPZ98]. Conversely, every Chebyshev system $(\gamma_0, \ldots, \gamma_d)$ on an interval I with $\gamma_0 \equiv 1$ (or more generally, γ_0 strictly monotone) gives rise to a convex curve in $\mathbb{R}^d$.

Subdividing $(\leq d{+}1)$-crossing curves. The following question is quite natural and interesting in its own right and it has been motivated by the work [EMRS14] in geometric Ramsey theory, as will be explained below. Given an integer $d \geq 2$, does there exist $M = M(d)$ such that every $(\leq d{+}1)$-crossing curve γ in $\mathbb{R}^d$ can be subdivided into at most M convex curves? In more detail, if γ is a map $I \to \mathbb{R}^d$, we want to subdivide I into subintervals $I_1, \ldots, I_k$, $k \leq M$, so that the restriction of γ to each I_i is convex (i.e., $(\leq d)$-crossing). Our main result answers this question in the affirmative.

Theorem 1.1. *For every integer $d \geq 2$ there exists $M = M(d)$ such that every $(\leq d{+}1)$-crossing curve γ in $\mathbb{R}^d$ can be subdivided into at most M convex curves.*

We note that the value $d+1$ is important, since a $(\leq d+2)$-crossing curve in $\mathbb{R}^d$ in general cannot be subdivided into a bounded number of convex curves. An example for $d = 2$ can be obtained, e.g., by starting with a circular arc and making many very small and flat inward dents in it.

The case $d = 2$ is already nontrivial, but to our surprise, we have not found it mentioned in the literature. The following picture shows a planar curve, namely, the graph of $x(1 - x^2)^2$ on $[-1, 1]$, which can be checked to be (≤ 3)-crossing, but obviously cannot be subdivided into fewer than four convex arcs:

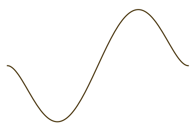

Hence $M(2) \geq 4$. We can prove that $M(2)$ actually equals 4, and that $M(3) \leq 22$. The proofs can be found in an earlier version of this paper [BM13] by the first two authors.

[2] Let A be a linearly ordered set of at least $k + 1$ elements. A (real) *Chebyshev system* on A is a system of continuous real functions $f_0, f_1, \ldots, f_k : A \to \mathbb{R}$ such that for every choice of elements $t_0 < t_1 < \cdots < t_k$ in A, the matrix $(f_i(t_j))_{i,j=0}^k$ has a (strictly) positive determinant.

Theorem 1.1 for polygonal paths. For technical reasons, and also from the point of view of our motivation in geometric Ramsey theory, it is more convenient to work with polygonal paths. A *polygonal path* is a curve made of finitely many straight segments; we call these segments the *edges* of the polygonal path, and their endpoints are the *vertices*. For a point sequence $(p_1, \ldots, p_n)$, we write $p_1 \cdots p_n$ for the polygonal path consisting of the segments $p_1 p_2, \ldots, p_{n-1} p_n$.

The definition of $(\leq k)$-crossing needs to be modified: we call a polygonal path π *$(\leq k)$-crossing* if it intersects every hyperplane in at most k points, *with the exception of the hyperplanes that contain an edge of* π. Moreover, we will also consider only *polygonal paths in general position*, meaning that any $k \leq d+1$ vertices of the polygonal path are affinely independent. The polygonal path version of Theorem 1.1 says the following.

Theorem 1.2. *For every integer $d \geq 2$ there exists $M = M(d)$ such that every $(\leq d+1)$-crossing polygonal path π in $\mathbb{R}^d$ can be subdivided into at most M convex (i.e., $(\leq d)$-crossing) polygonal paths.*

In Section 6 we prove by a limit argument that Theorem 1.2 implies Theorem 1.1.

Order-type homogeneous subsequences. Now we come to the geometric Ramsey-type problem motivating our work.

Let $T = (p_1, \ldots, p_{d+1})$ be an ordered $(d+1)$-tuple of points in $\mathbb{R}^d$. We recall that the *sign* (or *orientation*) of T is defined as $\operatorname{sgn}\det X$, where the jth column of the $(d+1) \times (d+1)$ matrix X is $(1, p_{j,1}, \ldots, p_{j,d})$, with $p_{j,i}$ denoting the ith coordinate of p_j. Geometrically, the sign is $+1$ if the d-tuple of vectors $p_1 - p_{d+1}, \ldots, p_d - p_{d+1}$ forms a positively oriented basis of $\mathbb{R}^d$, it is -1 if it forms a negatively oriented basis, and it is 0 if these vectors are linearly dependent.

We call a sequence $(p_1, \ldots, p_n)$ of points in $\mathbb{R}^d$ in general position *order-type homogeneous* if all $(d+1)$-tuples $(p_{i_1}, \ldots, p_{i_{d+1}})$, $i_1 < \cdots < i_{d+1}$, have the same sign (which is nonzero, by the general position assumption).

Let $\mathrm{OT}_d(n)$ be the smallest N such that every sequence of N points in general position in $\mathbb{R}^d$ contains an order-type homogeneous subsequence of length n. The existence of $\mathrm{OT}_d(n)$ for all d and n follows immediately from Ramsey's theorem, but several recent papers [EM13, CFP$^+$14, Suk14, EMRS14] considered the order of magnitude of $\mathrm{OT}_d(n)$, for d fixed and n large.

For $d = 2$, the classical paper of Erdős and Szekeres [ES35] implies that $\mathrm{OT}_2(n) = 2^{\Theta(n)}$.[3] Suk [Suk14], improving on a somewhat weaker bound by Conlon et al. [CFP$^+$14], proved the upper bound $\mathrm{OT}_d(n) \leq \mathrm{twr}_d(O(n))$ for every fixed d, where the *tower function* $\mathrm{twr}_k(x)$ is defined by $\mathrm{twr}_1(x) = x$ and $\mathrm{twr}_{i+1}(x) = 2^{\mathrm{twr}_i(x)}$. He conjectured this to be optimal, but so far matching lower bounds have been known only for $d = 2$ (by [ES35]) and $d = 3$ [EM13].

By combining the results of [EMRS14] with Theorem 1.2, we obtain a matching lower bound for all $d \geq 2$:

[3] We employ the usual asymptotic notation for comparing functions: $f(n) = O(g(n))$ means that $|f(n)| \leq C|g(n)|$ for some C and all n, where C may depend on parameters declared as constants (in our case on d); $f(n) = \Omega(g(n))$ is equivalent to $g(n) = O(f(n))$; and $f(n) = \Theta(g(n))$ means that both $f(n) = O(g(n))$ and $f(n) = \Omega(g(n))$.

Theorem 1.3. *We have* $\mathrm{OT}_d(n) \geq \mathrm{twr}_d(\Omega(n))$.

The argument is given in Section 7.

The question of estimating $M(d)$**.** In our proof of Theorem 1.2, we will bound $M(d)$ from above by another, combinatorially defined quantity $c(d)$, introduced in Theorem 4.2 below. We will show that $c(1) = 3$, and since $M(1) = 3$ trivially, we have $c(1) = M(1) = 3$. In general, our proof yields $M(d) \leq c(d) \leq \exp(O(d))$, where the constant in the $O(\cdot)$ notation could easily be made explicit. In particular, for $d = 2$ the general proof provides the bound $M(2) \leq c(2) \leq 28$, while, as was mentioned earlier, with a careful argument it is possible to show that $M(2) = 4$. It would be interesting to find the correct order of magnitude of $M(d)$, or at least reasonable upper and lower bounds.

2. Order-type homogeneity and path convexity

We need the following fact.

Lemma 2.1. *A sequence* $P = (p_1, \ldots, p_n)$ *in general position in* $\mathbb{R}^d$ *is order-type homogeneous iff the polygonal path* $\pi = p_1 \cdots p_n$ *is convex.*

Proof. First we assume that P is not order-type homogeneous. Then it has two $(d+1)$-tuples, of the form $Q = (q_1, \ldots, q_{d+1})$ and $R = (r_1, \ldots, r_{d+1})$, with opposite signs (both Q and R are subsequences of P, i.e., the q_i and the r_j appear in P in this order).

It is easy to check that we can also find Q and R with opposite signs that differ in a single point; more precisely, there is an index k such that $q_i = r_i$ for all $i \neq k$. Indeed, given arbitrary Q and R with opposite signs, we can convert Q into R by a sequence of moves, each of them changing a single element: we always move the first element in which the current Q differs from R to the correct position. Then at least one of the moves involves two $(d+1)$-tuples with opposite signs.

Having Q and R as above with $q_i = r_i$ for all $i \neq k$, we consider the hyperplane h spanned by the points of $Q' := \{q_i : i \neq k\}$. Then q_k and r_k lie on opposite sides of h, and hence π intersects h between q_k and r_k. Together with the d points Q', we have $d+1$ intersections of π with h.

This h may still contain edges of π, so we may need to move it slightly. For simpler description, we think of h as horizontal, and say that q_k is below h, r_k is above h, and q_k precedes r_k in P.

If we choose, for every point $q_i \in Q'$, a point $\tilde{q}_i$ sufficiently close to q_i, then these $\tilde{q}_i$ span a hyperplane h', since Q' is affinely independent. Moreover, as the $\tilde{q}_i$ get closer and closer to the corresponding q_i, the hyperplane h' gets arbitrarily close to h within every bounded region of $\mathbb{R}^d$. It follows that if the points $\tilde{q}_i$ are chosen suitably (and in particular on appropriate sides of h), then the points in the sequence $(q_1, \ldots, q_k, r_k, q_{k+1}, \ldots, q_{d+1})$ are alternately above and below h'. This implies that π intersects h' at least $d+1$ times, and since the move of h was generic, we may assume that h' contains no edges of π.

For the reverse implication, we need the following claim: *If* $P = (p_1, \ldots, p_n)$ *is an order-type homogeneous sequence and* q *is an interior point of the segment* $p_i p_{i+1}$*, then* *the sequence* $P' = (p_1, \ldots, p_i, q, p_{i+2}, \ldots, p_n)$ *(*p_{i+1} *replaced with* q*) is order-type homogeneous as well.*

To verify this claim, we suppose without loss of generality that all $(d + 1)$-tuples of P are positive, and we consider an arbitrary $(d + 1)$-tuple in P' involving q, of the form

$$T = (p_{j_1}, \ldots, p_{j_{k-1}}, q, p_{j_{k+1}}, \ldots, p_{j_{d+1}}),$$
$$1 \le j_1 < \cdots < j_{k-1} < i + 1 < j_{k+1} < \cdots < j_{d+1} \le n.$$

We think of q moving from p_i to p_{i+1} along the segment $p_i p_{i+1}$. The determinant whose sign defines the sign of T is an affine function of q (considering the remaining points of T fixed). For $q = p_i$ it is either 0 (if $j_{k-1} = i$) or strictly positive, and for $q = p_{i+1}$ it is strictly positive. Therefore, for q in between, it is strictly positive too, which proves the claim.

Now we assume for contradiction that the sequence $P = (p_1, \ldots, p_n)$ is order-type homogeneous, but the corresponding polygonal path π is not convex, and so it has at least $d + 1$ intersections with some hyperplane h not containing any edge of π. Let us fix intersections $q_1, \ldots, q_{d+1}$; at least one of them, call it q_ℓ, is an interior point of an edge $p_j p_{j+1}$ of π (since the p_i are in general position).

Using the claim above, we now want to replace π by another polygonal path π', whose vertex sequence is still order-type homogeneous and includes all q_i with $i \neq \ell$, as well as p_j and p_{j+1}. To this end, we first observe that no two q_i share a segment of π (since h contains no such segment).

When producing π', first, if there is a q_i with $i > \ell$ that is not a vertex of the current polygonal path, we take the last such q_i. We replace the vertex of the current polygonal path immediately following q_i with q_i. By the claim, the new vertex sequence is still order-type homogeneous. We repeat this step until all q_i with $i > \ell$ become vertices.

Then we proceed analogously with the q_i, $i < \ell$, that are not vertices. This time we start with the smallest i, and q_i always replaces the vertex immediately preceding it (and we apply the claim to the reversal of the sequences under consideration). Here is an illustration, with q_ℓ, p_j, and p_{j+1} marked white:

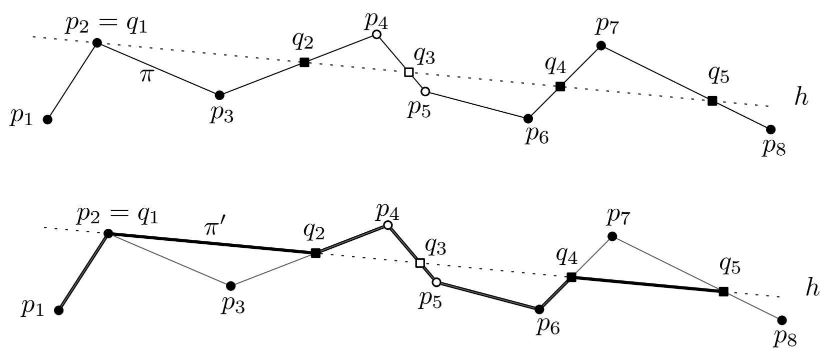

In this way, we obtain the polygonal path π' with order-type homogeneous vertex sequence that is intersected by the hyperplane h in the d vertices q_i, $i \neq \ell$, and in q_ℓ, which is an interior point of the segment $p_j p_{j+1}$ (neither p_j nor p_{j+1} has been replaced). But then the $(d+1)$-tuples $(q_1, \ldots, q_{\ell-1}, p_j, q_{\ell+1}, \ldots, q_{d+1})$ and $(q_1, \ldots, q_{\ell-1}, p_{j+1}, q_{\ell+1}, \ldots, q_{d+1})$ have opposite signs—a contradiction. □

3. A combinatorial property of $(\leq d+1)$-crossing paths

Here we prove a combinatorial property of point sequences in $\mathbb{R}^d$ for which the corresponding polygonal path is $(\leq d+1)$-crossing. In the two subsequent sections we will derive Theorem 1.2 from this property in a purely combinatorial way.

Let $P = (p_1, \ldots, p_n)$ be a sequence in general position in $\mathbb{R}^d$ and let $\pi = p_1 \cdots p_n$ be the corresponding polygonal path. For notational convenience, for $Q \subset P$ with $|Q| = d+1$, we define $\operatorname{sgn} Q$ as the sign of the sequence $(p_{i_1}, \ldots, p_{i_{d+1}})$, where $Q = \{p_{i_1}, \ldots, p_{i_{d+1}}\}$ with $i_1 < \cdots < i_{d+1}$. For a fixed subset $R \subset P$ with $|R| = d$, we consider the following sequence, which we call the *sign sequence of* R:

$$\big(\operatorname{sgn}(\{p_i\} \cup R) : i = 1, \ldots, n,\ p_i \notin R\big) \in \{-1, +1\}^{n-d}. \tag{3.1}$$

Lemma 3.1. *If π is $(\leq d+1)$-crossing, then for every R as above, the sign sequence (3.1) of R has at most one sign change.*

A simple case. To prove the lemma, we first consider a simple special case. Letting H be the hyperplane spanned by R, we assume that R contains no consecutive elements of P, and moreover that H separates p_{i-1} from p_{i+1} whenever $p_i \in R$.

Because of the $(\leq d+1)$-crossing condition, $(\pi \cap H) \setminus R$ is either the empty set or a single point, which we call q. Then for $x \in \pi$, we have $\operatorname{sgn}(\{x\} \cup R) = 0$ iff $x \in R$ or $x = q$.

Let us think of x moving along π. When it passes through a point $p \in R$, $\operatorname{sgn}(\{x\} \cup R)$ does not change because x moves from one side of H to the other, while x changes places with p in the order on π. The same argument shows that $\operatorname{sgn}(\{x\} \cup R)$ changes only if x passes through q.

Auxiliary claims. Next, we make preparations for proving the lemma in general.

The set $P \setminus R$ is non-empty, so we fix one of its elements and call it p_α. We define $\mathcal{R}_\delta$ as the set of all sequences $(q_i \in \pi : p_i \in R)$ such that $|q_i - p_i| < \delta$, and for $i > \alpha$, q_i lies on the open segment (p_{i-1}, p_i), while for $i < \alpha$ it lies on (p_i, p_{i+1}). Here is a schematic illustration:

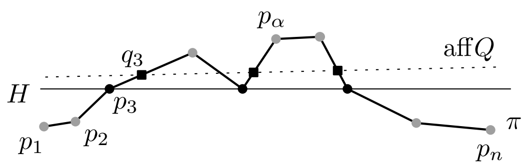

Since R spans the hyperplane H, every set $Q \in \mathcal{R}_\delta$ for sufficiently small δ spans a hyperplane as well. By general position, we have $\varepsilon_0 := \operatorname{dist}(P \setminus R, H) > 0$. By continuity, we also get the next claim:

Claim 3.2. *There is* $\delta_1 > 0$ *such that* $\operatorname{dist}(P \setminus R, \operatorname{aff} Q) > \frac{1}{2}\varepsilon_0$ *for all* $Q \in \mathcal{R}_{\delta_1}$.

This has the following consequence:

Corollary 3.3. *If* $p_h, p_{h+1} \notin R$ *and* $H \cap p_h p_{h+1} \neq \emptyset$, *then* $\operatorname{aff} Q \cap p_h p_{h+1} \neq \emptyset$ *for all* $Q \in \mathcal{R}_{\delta_1}$.

Claim 3.4. *There is a* $\delta_2 \in (0, \delta_1)$ *such that* $P \cap \operatorname{aff} Q = \emptyset$ *for all* $Q \in \mathcal{R}_{\delta_2}$.

Proof. If not, then there is a sequence $\delta_m \to 0$ and $Q_m \in \mathcal{R}_{\delta_m}$ with $P \cap \operatorname{aff} Q_m \neq \emptyset$. Then, for a suitable subsequence, $P \cap \operatorname{aff} Q_m$ contains a fixed element $p_h \in P$. We have $p_h \in R$ because the Q_m have distance at least $\varepsilon_0/2$ to $P \setminus R$.

Let $(p_i, p_{i+1}, \dots, p_j)$ be the *string* of R containing p_h, i.e., a maximal contiguous subsequence of P whose points all lie in R (i.e., $p_{i-1}, p_{j+1} \notin R$; we also admit $i = 1$ and $j = n$, as well as $i = j$). Thus $i \le h \le j$ and the polygonal path $p_i \dots p_j$ is contained in H.

Let us assume $h > \alpha$; then $i > \alpha$ as well. Since $p_h \in \operatorname{aff} Q_m$ and $q_h \in Q_m$, the whole line $\operatorname{aff}\{p_h, q_h\}$ is contained in $\operatorname{aff} Q_m$. Since p_{h-1} is on this line, it is in $\operatorname{aff} Q_m$ as well. This shows (by induction) that $p_h, p_{h-1}, \dots, p_i, p_{i-1} \in \operatorname{aff} Q_m$. Thus $p_{i-1} \in \operatorname{aff} Q_m$, which contradicts Claim 3.2. The argument for $h < \alpha$ is symmetric. □

Proof of Lemma 3.1. We fix some $\delta \in (0, \delta_2)$ and $Q \in \mathcal{R}_\delta$, and set $H^* = \operatorname{aff} Q$. We observe that H and H^* separate the points of $P \setminus R$ the same way. Moreover, if $(p_i, \dots, p_j)$ is a string of R and $i > \alpha$, then the points $p_{i-1}, p_i, \dots, p_j$ lie alternately on the two sides of H^*. This follows from the fact that the path $p_{i-1} p_i \dots p_j$ intersects H^* in the points $q_i, \dots, q_j$. Similarly, for $i < \alpha$, the points $p_i, \dots, p_j, p_{j+1}$ lie alternately on the two sides of H^*.

We again let x move along π. With $R = (p_{i_1}, \dots, p_{i_{d+1}})$, we have

$$\operatorname{sgn}(\{x\} \cup R) = \operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & 1 & 1 & \dots & 1 \\ p_{i_1} & \dots & p_{i_{j-1}} & x & p_{i_j} & \dots & p_{i_{d+1}} \end{pmatrix}$$

where the position of the column with x is determined by x lying between $p_{i_{j-1}}$ and p_{i_j}. Then

$$\operatorname{sgn}(\{x\} \cup Q) = \operatorname{sgn}\det\begin{pmatrix} 1 & \dots & 1 & 1 & 1 & \dots & 1 \\ q_{i_1} & \dots & q_{i_{j-1}} & x & q_{i_j} & \dots & q_{i_{d+1}} \end{pmatrix}$$

where $Q = (q_{i_1}, \dots, q_{i_{d+1}})$ and the same remark applies to the position of the x column.

Clearly $\operatorname{sgn}(\{x\} \cup R) = \operatorname{sgn}(\{x\} \cup Q)$ when $x \in P \setminus R$. Thus, it suffices to check how $\operatorname{sgn}(\{x\} \cup Q)$ changes when x moves through $q_i, \dots, q_j$ for the string $p_i, \dots, p_j$. Note that $\operatorname{sgn}(\{x\} \cup Q)$ changes only when x passes some point in $Q \cap \pi$.

Just as in the basic case, $\operatorname{sgn}(\{x\} \cup Q)$ does not change when x passes q_h because then x moves from one side of H^* to the other and it also changes places with q_h. Thus, $\operatorname{sgn}(\{p_{i-1}\} \cup Q) = \operatorname{sgn}(\{x\} \cup Q)$ when x just passed q_j.

Now we assume that $\alpha < i$; the other option $\alpha > i$ is symmetric and follows the same way. There are two cases.

Case 1: p_j and p_{j+1} are on the same side of H^*. Then $\mathrm{sgn}(\{p_j\}\cup Q) = \mathrm{sgn}(\{p_{j+1}\}\cup Q)$, and so $\mathrm{sgn}(\{p_{i-1}\}\cup Q) = \mathrm{sgn}(\{p_{j+1}\}\cup Q)$, implying $\mathrm{sgn}(\{p_{i-1}\}\cup R) = \mathrm{sgn}(\{p_{j+1}\}\cup R)$. So there is no sign change between p_{i-1} and p_{j+1} in the sign sequence of R.

Case 2: p_j and p_{j+1} are on opposite sides of H^*. Then $H^* \cap p_j p_{j+1}$ is a point q, and $\mathrm{sgn}(\{x\} \cup Q)$ changes sign when x moves through q. Consequently, $\mathrm{sgn}(\{p_{j+1}\} \cup R) = -\,\mathrm{sgn}(\{p_{i-1}\} \cup R)$, and there is a sign change in the sign sequence of R here.

But since $H^* \cap \pi$ contains already $d + 1$ points, Case 2 cannot occur anywhere else. Also, the case in Corollary 3.3 cannot come up either, since that would mean $H^* \cap \pi$ contains $d+2$ points. Thus, the only sign change in the sign sequence of R occurs between p_{i-1} and p_{j+1}. □

4. k-sequences and flip k-sequences

Now we will define a combinatorial abstraction of point sequences in $\mathbb{R}^k$. A *k-sequence* is a sequence $S = (a_1, \ldots, a_n)$, where $a_1, \ldots, a_n$ are distinct (abstract) elements, together with a mapping sgn that assigns either $+1$ or -1 to every $(k + 1)$-element subset $A \subseteq \{a_1, \ldots, a_n\}$ (sometimes we will regard A as a subsequence, with the elements in the same order as in S). We will also say that A is *positive* or *negative* if $\mathrm{sgn}\, A = 1$ or $\mathrm{sgn}\, A = -1$, respectively.

We subdivide the sequence S into contiguous blocks with one-point overlaps: The first block is $B_1 = (a_1, \ldots, a_{i_1})$ with i_1 maximal such that all $(k+1)$-point subsequences in B_1 have the same sign σ_1. The next one is $B_2 = (a_{i_1}, \ldots, a_{i_2})$ with i_2 maximal such that all $(k + 1)$-point subsequences in B_2 have the same sign σ_2, and so on, up until some $B_m = (a_{i_{m-1}}, \ldots, a_n)$, where B_m either has at most k elements, or it has more than k elements and every $(k + 1)$-tuple in it has the same sign σ_m.

We call this partition the *greedy partition* of S; here both $m = m(S)$ and the blocks B_j are uniquely determined. Note that each B_j, $j < m$, contains a subset D_j of size k such that $\mathrm{sgn}(\{a_{i_j+1}\} \cup D_j) \neq \sigma_j$.

The following lemma shows that S has a short subsequence S^* whose greedy partition is similar to that of S.

Lemma 4.1. *There is a subsequence S^* of S, which we call the* reduced version *of S, such that $m(S^*) = m(S)$, every block of the greedy partition of S^* contains at most $k + 3$ elements, and the last one exactly two. Moreover, every string of $2k + 5$ consecutive elements of S^* contains both a positive $(k + 1)$-tuple and a negative one.*

Proof. Let $B_j = (a_{i_{j-1}}, \ldots, a_{i_j})$ be a block of the greedy partition of S with $j < m$. Let us fix a d-element subset D_j of B_j as above, i.e., with $\mathrm{sgn}(\{a_{i_j+1}\} \cup D_j) \neq \sigma_j$.

The subsequence S^* contains the following elements of B_j: $a_{i_{j-1}}, a_{i_{j-1}+1}, a_{i_j}$, the elements of D_j, and one more (arbitrarily chosen) element if the first three are all contained in D_j. All the other elements are discarded. From the last block we keep the first two elements.

Let us consider the greedy partition of S^*. By induction on j, it is easy to see that for $j < m$, the jth block B_j^* starts with $a_{i_{j-1}}$, ends with a_{i_j}, and the sign of $D_j \cup \{a_{i_j+1}\}$ is different from σ_j, which is the sign of (all) $(k+1)$-tuples in B_j^*.

It follows that every string of $2k+5$ consecutive elements of S^* contains a full block B_j^* plus the next element a_{i_j+1}. The sign of the first $k+1$ elements of B_j^* is different from $\mathrm{sgn}(D_j \cup \{a_{i_j+1}\})$. □

A k-sequence $S = (a_1, \dots, a_n)$ is called a *flip k-sequence* if it has the property as in Lemma 3.1; that is, for every k-element $A \subset \{a_1, \dots, a_n\}$, the *sign sequence* of A

$$\big(\mathrm{sgn}(\{a_i\} \cup A) : i = 1, \dots, n,\ a_i \notin A\big) \tag{4.1}$$

has at most one sign change. The following result of combinatorial nature is the key step in the proof of Theorem 1.2.

Theorem 4.2. *For every $k \ge 1$ there is $c(k)$ such that the greedy partition of every flip k-sequence has at most $c(k)$ blocks.*

We prove this result in the next section. Now we show how it implies Theorem 1.2.

Proof of Theorem 1.2. We assume that $P = (p_1, \dots, p_n) \subset \mathbb{R}^d$ is in general position. Let $\pi = p_1 \cdots p_n$ be the corresponding polygonal path. Lemma 3.1 shows that $(p_1, \dots, p_n)$ with the sign of $(d+1)$-tuples given by their orientation is a flip d-sequence. Theorem 4.2 says that its greedy partition has at most $c(d)$ blocks. All $(d+1)$-tuples in B_j have the same sign, so $B_j = (p_{i_{j-1}}, \dots, p_{i_j})$ is order-type homogeneous, and thus the polygonal path $p_{i_{j-1}} \cdots p_{j_i}$ is convex. It follows that $M(d) \le c(d)$. □

5. Proof of Theorem 4.2

We proceed by induction on k.

The case $k = 1$. We will show that $c(1) = 3$ (instead of reading this part, the reader may perhaps prefer to find a simple proof of $c(1) \le 5$, say).

Let $S = (a_1, \dots, a_n)$ be a flip 1-sequence, and let $B_1, \dots, B_m$ be the blocks of its greedy partition. Each B_i has the form $(b_i, x_i, \dots, c_i)$ where $b_{i+1} = c_i$, and B_i contains an element d_i such that $\mathrm{sgn}(d_i, x_{i+1}) \ne \sigma_i$. Note that x_1 and d_m are undefined.

Observation. If B_i and B_{i+1} are two consecutive blocks, both positive, then d_i, $c_i = b_{i+1}$, and x_{i+1} are three distinct elements of S. Moreover, for every $a \in S$ preceding d_i we have (a, x_{i+1}) negative, and similarly, for every a following x_{i+1} we have (d_i, a) negative.

Only the last two statements need an explanation. Since (c_i, x_{i+1}) is positive and (d_i, x_{i+1}) is negative, (a, x_{i+1}) must be negative for a preceding d_i, for otherwise there are two sign changes in the sign sequence of $\{x_{i+1}\}$. The statement about (d_i, a) is proved in the same way.

The proof of $c(1) \le 3$ comes in five steps. We assume without loss of generality that (a_1, a_2) is positive.

Step 1. If all (a_i, a_{i+1}) are positive, then $m < 4$. Indeed, supposing B_4 exists, all blocks are positive, (d_1, x_2) is negative, and (d_1, d_3) is negative by the Observation above. Also, (d_3, x_4) is negative and there are two sign changes in the sign sequence of $\{d_3\}$, since (b_3, d_3) or (d_3, c_3) (or both) are positive.

Step 2. If j is the smallest index with (a_j, a_{j+1}) negative, then $a_j = c_i = b_{i+1}$, B_i is a positive block, and B_{i+1} is a negative one. Assume B_{i-1} exists. Then it is positive, (d_{i-1}, x_i) is negative, and thus (d_{i-1}, a_j) is negative by the Observation. But then there are two sign changes in the sign sequence of $\{a_j\}$: (d_{i-1}, a_j) and (a_j, x_{i+1}) are negative and (b_i, a_j) is positive. Thus B_{i-1} cannot exist, $i = 1$, and there is a single block before a_j.

Step 3. Thus B_1 is positive and B_2 negative. Assume B_3 negative; then (d_2, x_3) is positive and so is (b_2, x_3) by the Observation. Consequently, there are two sign changes in the sign sequence of $\{b_2\}$: (b_1, b_2) and (b_2, x_3) are positive but (b_2, x_2) is negative. We conclude that B_3 is a positive block.

Step 4. Assume B_4 exists and is positive. Then (d_3, x_4) is negative and so is (b_3, x_4) by the Observation. Then there are two sign changes in the sign sequence of $\{b_3\}$: (b_2, b_3) and (b_3, x_4) are negative and (b_3, c_3) is positive.

Step 5. We are left with the case when B_1, B_3 are positive and B_2, B_4 negative. If (b_2, b_4) is positive, then there are two sign changes in the sign sequence of $\{b_2\}$: (b_1, b_2) and (b_2, b_4) are positive and (b_2, c_2) negative. Similarly, if (b_2, b_4) is negative, then there are two sign changes in the sign sequence of $\{b_4\}$.
Consequently, B_4 does not exist: $m < 4$ and so $c(1) \le 3$.

The example $S = (a_1, a_2, a_3, a_4)$ with a_1, a_2 and a_3, a_4 positive and all other pairs negative shows that $c(1) = 3$.

The inductive step from $k - 1$ to k. Assuming that the greedy partition of each flip $(k-1)$-sequence has at most $c(k-1)$ blocks, we will show that the greedy partition of an arbitrary flip k-sequence $S = (a_1, \dots, a_n)$ has at most $c(k) := 1 + (4k + 10)c(k-1)/k$ blocks.

So we suppose on the contrary that S as above has $m > c(k)$ blocks. We can further assume that S is reduced in the sense of Lemma 4.1, for otherwise we can replace S by S^*. Since each B_i, $i < m$, has at least $k+1$ elements, and $|B_m| = 2$, the length of S is at least

$$n \ge (m-1)k + 2 > (4k + 10)c(k-1) + 2.$$

We consider the sequence $T = (a_1, \dots, a_{n-1})$ and regard it as a $(k-1)$-sequence by defining, for a k-element $A \subset \{a_1, \dots, a_{n-1}\}$, the sign $\operatorname{sgn} A := \operatorname{sgn}(A \cup \{a_n\})$. It is clear that T is a flip $(k-1)$-sequence, and so its greedy partition has at most $c(k-1)$ blocks. One of the blocks, which we call B, has at least $(n-1)/c(k-1) \ge 4k+10$ elements. We may assume without loss of generality that $\operatorname{sgn} A = +1$ for every k-element subset of B.

Since S is reduced, there is a positive $(k+1)$-tuple $(b_1, \dots, b_{k+1})$ among the first $2k+5$ elements of B, and a negative $(k+1)$-tuple $(b_{k+2}, \dots, b_{2k+2})$ among the last $2k+5$ elements of B. The sign of the $(k+1)$-tuple $(b_i, \dots, b_{i+k})$ changes from $+1$ to -1 as i moves through $1, \dots, k+2$, and so there is some j with $\operatorname{sgn}(b_j, \dots, b_{j+k+1}) = +1$ and $\operatorname{sgn}(b_{j+1}, \dots, b_{j+k+2}) = -1$.

We set $A = \{b_{j+1}, \ldots, b_{j+k+1}\}$. Then $\mathrm{sgn}(\{b_j\} \cup A) = +1$ and $\mathrm{sgn}(A \cup \{b_{j+k+2}\}) = -1$, while $\mathrm{sgn}(A \cup \{a_n\}) = +1$ by the choice of the block B. Hence the sign sequence of A has at least two sign changes, contradicting the assumption that S is a flip k-sequence. This contradiction finishes the proof of Theorem 4.2. □

6. From polygonal paths to curves: proof of Theorem 1.1

Here we show how Theorem 1.2 implies Theorem 1.1. We assume that $\gamma : I \to \mathbb{R}^d$ is a $(\leq d+1)$-crossing curve.

Let us say that an n-tuple $T = (t_1, \ldots, t_n)$, $t_1, \ldots, t_n \in I$, $t_1 < \cdots < t_n$, is an *ε-sample* if every subinterval of I of length ε contains some t_i. Let $\pi = \pi(\gamma, T) = \gamma(t_0)\gamma(t_1)\cdots\gamma(t_n)$ be the polygonal line determined by T.

First we observe that for every $\varepsilon > 0$, there is an ε-sample T with $\pi(\gamma, T)$ in general position. Indeed, having already placed k points of T so that their γ-images are in general position, we consider the finitely many hyperplanes spanned by d-tuples of these γ-images. Since γ is $(\leq d+1)$-crossing, each of these hyperplanes contains at most one extra point of γ, and so at every step of the construction, we have only finitely many excluded points of I. Thus, we can construct an ε-sample as desired.

Next, for every $\varepsilon > 0$, we fix an ε-sample $T = T(\varepsilon)$ with $\pi(\gamma, T(\varepsilon))$ in general position. Let $M = M(d)$ be as in Theorem 1.2; by that theorem, we can also fix a subdivision of I into M subintervals such that the restriction of $\pi(T(\varepsilon), \gamma)$ on each of them is convex. By compactness, these subdivisions have a cluster point for $\varepsilon \to 0$; we denote its intervals by $I_1, \ldots, I_M$.

It remains to show that γ restricted to each I_j is convex. This follows from the next lemma, applied with $I = I_j$ and $\gamma = \gamma_j$.

Lemma 6.1. *Let $\gamma : I \to \mathbb{R}^d$ be a $(\leq d+1)$-crossing curve, and suppose that for every $\varepsilon > 0$ there is an ε-sample $T(\varepsilon)$ such that the corresponding polygonal path $\pi(\gamma, T(\varepsilon))$ is in general position and convex. Then γ is convex as well.*

Proof. For contradiction, we suppose that there is a hyperplane h intersecting γ in at least $d+1$ points.

First we observe that these points can be assumed to span h: if their affine hull F had dimension smaller than $d-1$, then since $\gamma \not\subset F$, we could rotate h around F and thus get more than $d+1$ intersections.

Let us say that a point $\gamma(t) \in h$, $t \in I$, is a *generic intersection* with h if for an arbitrarily small neighborhood U of t, $\gamma(U)$ intersects both of the open halfspaces bounded by h (as usual, we count generic intersections with multiplicity, so the generic intersection is actually determined by t). We claim that there is a hyperplane h' with at least $d+1$ generic intersections.

For easier description, let us imagine h horizontal. An intersection that is not generic is either an endpoint of γ, or it is a point p where γ touches h, with a sufficiently small open neighborhood of p on γ lying all strictly above h or all strictly below it; let us call such intersections *top-touching* or *bottom-touching*.

Let $q_1, \ldots, q_k$ be the nongeneric intersections of γ with h. At least $k-1$ of these are affinely independent, say $q_1, \ldots, q_{k-1}$, and thus we can make an arbitrarily small movement of h so that a prescribed subset of $\{q_1, \ldots, q_{k-1}\}$ ends up below h and its complement above h. The previously generic intersections remain generic, provided that the movement was sufficiently small.

Now if q_i was bottom-touching and it lies above h after the move, then it yields (at least) two generic intersections with h, and similarly for top-touching. If q_i is an endpoint, then it yields at least one generic intersection, provided that h was moved in the right direction.

Hence by an appropriate move we can always get at least $d+1-k+2(k-3)+2 = d+k-3$ generic intersections, which is at least $d+1$ for $k \geq 4$. So it remains to discuss the cases $1 \leq k \leq 3$.

For $k \leq 2$, the nongeneric intersections are distinct and thus affinely independent, and so we can get k new generic intersections by a suitable move. For $k=3$, there are two affinely independent nongeneric intersections, at least one of them top-touching or bottom-touching, and hence we can also get three new generic intersections by a suitable move. Thus, we have obtained a hyperplane h' with at least $d+1$ generic intersections as required.

Let $t_1, \ldots, t_{d+1} \in I$, $t_1 < \cdots < t_{d+1}$, be the parameter values corresponding to these generic intersections with h'. To finish the proof of the lemma, we fix a sufficiently small $\varepsilon > 0$ and intervals $J_1^+, J_1^-, \ldots, J_{d+1}^+, J_{d+1}^- \subset I$, each of length at least ε, such that J_i^+ and J_i^- are in a small neighborhood of t_i (and thus they lie to the left of $J_{i+1}^+ \cup J_{i+1}^-$), and $\gamma(J_i^+)$ lies above h' and $\gamma(J_i^-)$ below it.

Suppose that J_1^+ precedes J_1^-, for example. We choose points $u_0, u_1, \ldots, u_{d+2} \in T(\varepsilon)$ with $u_0 \in J_1^+$, $u_1 \in J_1^-$, $u_2 \in J_2^+$, $u_3 \in J_3^-$, $u_4 \in J_4^+$, etc. Then the polygonal line $\pi(\gamma, T(\varepsilon))$ changes sides of h' at least $d+1$ times, and thus it has at least $d+1$ intersections with h'. Since the position of h' is generic, this shows that $\pi(\gamma, T(\varepsilon))$ is not convex—a contradiction proving the lemma, and also concluding the proof of Theorem 1.1. □

7. The lower bound for order-type homogeneous subsequences

Super-order type homogeneity. The following strengthening of order-type homogeneity was considered in [EMRS14]: a point sequence $P = (p_1, \ldots, p_n)$ in $\mathbb{R}^d$ is *super-order type homogeneous* if, for every $k = 1, \ldots, d$, the projection of P to the first k coordinates is order-type homogeneous (this includes the assumption that all of these projections are in general position—let us abbreviate this by saying that P is *in super-general position*).

It is easily seen, e.g. by Ramsey's theorem, that for every d and n there is N such that every N-point sequence in super-general position in $\mathbb{R}^d$ contains a super-order type homogeneous subsequence of length n. Let us denote the corresponding Ramsey function by $\mathrm{OT}_d^*(n)$.

It was shown in [EMRS14] that $\mathrm{OT}^*_d(n) \geq \mathrm{twr}_d(n-d)$. Thus, to prove Theorem 1.3, the lower bound for OT_d, and having Theorem 1.2 at our disposal, it suffices to verify the following.

Lemma 7.1. *For all $d \geq 2$, $\mathrm{OT}_d(n) \geq \mathrm{OT}^*_d(\Omega(n))$.*

Proof. Given n, let us set $N = \mathrm{OT}_d(n)$, and consider an N-point sequence in super-general position in $\mathbb{R}^d$. By definition, it contains an n-point order-type homogeneous subsequence P_1.

By Lemma 2.1, the polygonal path given by P_1 is convex, i.e., $(\leq d)$-crossing, and hence its projection onto the first $d-1$ coordinates is $(\leq d)$-crossing as well. So by the assumption, it can be subdivided into at most $M(d-1)$ polygonal paths that are $(\leq d-1)$-crossing. One of them corresponds, by Lemma 2.1 again, to a subsequence P_2 of P_1 of length at least $n/M(d-1)$ whose projection to the first $d-1$ coordinates is order-type homogeneous.

Analogously we construct $P_3, \ldots, P_d$, where $|P_i| \geq |P_{i-1}|/M(d-i+1)$ and the projections of P_i to the first k coordinates, for $k = d-i+1, d-i+2, \ldots, d$, are order-type homogeneous. In particular, P_d is the desired super-order type homogeneous subsequence of length $\Omega(n)$. □

Acknowledgments. We would like to thank Marek Eliáš, Edgardo Roldán-Pensado, Zuzana Safernová, and Pavel Valtr for inspiring discussion at initial stages of this research, R. A. Zalik for answering our question about Chebyshev systems, and F. Santos, B. Shapiro, O. R. Musin, and an anonymous referee for useful comments and references.

The first and the second author were supported by the ERC Advanced Research Grant no. 267165 (DISCONV), and the first author by the Hungarian National Research Grant K 83767 as well.

References

[Arn04] Arnold, V. I. (ed.): Arnold's Problems. Springer, Berlin (2004) Zbl 1051.00002 MR 2078115

[BM13] Bárány, I., Matoušek, J.: Curves in $\mathbb{R}^d$ intersecting every hyperplane at most d+1 times. arXiv:1309.1147v1 (2013)

[CPZ98] Carnicer, J. M., Peña, J. M., Zalik, R. A.: Strictly totally positive systems. J. Approx. Theory **92**, 411–441 (1998) Zbl 0903.41015 MR 1609198

[CFP+14] Conlon, D., Fox, J., Pach, J., Sudakov, B., Suk, A.: Ramsey-type results for semi-algebraic relations. Trans. Amer. Math. Soc. **366**, 5043–5065 (2014) Zbl 1306.14027 MR 3217709

[EM13] Eliáš, M., Matoušek, J.: Higher-order Erdős–Szekeres theorems. Adv. Math. **244**, 1–15 (2013) Zbl 1283.05175 MR 3077863

[EMRS14] Eliáš, M., Matoušek, J., Roldán-Pensado, E., Safernová, Z.: Lower bounds on geometric Ramsey functions. SIAM J. Discrete Math. **28**, 1960–1970 (2014) Zbl 1326.14136 MR 3281280

[ES35] Erdős, P., Szekeres, G.: A combinatorial problem in geometry. Compos. Math. **2**, 463–470 (1935) JFM 61.0651.04 MR 1556929

[KS66] Karlin, S., Studden, W. J.: Tchebycheff Systems: With Applications in Analysis and Statistics. Pure Appl. Math. 15, Interscience Publ. Wiley, New York (1966) Zbl 0153.38902 MR 0204922

[Mus98] Musin, O. R.: Chebyshev systems and zeros of a function on a convex curve. Trudy Mat. Inst. Steklova **221**, 247–256 (1998) (in Russian) Zbl 0971.52002 MR 1683698

[SS05] Sedykh, V., Shapiro, B.: On two conjectures concerning convex curves. Int. J. Math. **16**, 1157–1173 (2005) Zbl 1085.14028 MR 2182213

[SS00] Shapiro, B., Shapiro, M.: Projective convexity in $\mathbb{P}^3$ implies Grassmann convexity. Int. J. Math. **11**, 579–588 (2000) Zbl 1110.53301 MR 1768174

[Suk14] Suk, A.: A note on order-type homogeneous point sets. Mathematika **60**, 37–42 (2014) Zbl 1301.05353 MR 3164517

[Živ04] Živaljević, R. T.: Topological methods. In: J. E. Goodman and J. O'Rourke (eds.), Handbook of Discrete and Computational Geometry, 2nd ed., Chapman & Hall/CRC, Boca Raton, FL, 305–329 (2004) Zbl 1056.52001 MR 2082993