Convex Polygons in Cartesian Products*

Jean-Lou De Carufel[†] Adrian Dumitrescu[‡] Wouter Meulemans[§] Tim Ophelders[¶]
Claire Pennarun[‖] Csaba D. Tóth[**] Sander Verdonschot[††]

May 5, 2019

Abstract

We study several problems concerning convex polygons whose vertices lie in a Cartesian product of two sets of n real numbers (for short, *grid*). First, we prove that every such grid contains a convex polygon with $\Omega(\log n)$ vertices and that this bound is tight up to a constant factor. We generalize this result to d dimensions (for a fixed $d \in \mathbb{N}$), and obtain a tight lower bound of $\Omega(\log^{d-1} n)$ for the maximum number of points in convex position in a d-dimensional grid. Second, we present polynomial-time algorithms for computing the longest convex polygonal chain in a grid that contains no two points with the same x- or y-coordinate. We show that the maximum size of such a convex polygon can be efficiently approximated up to a factor of 2. Finally, we present exponential bounds on the maximum number of convex polygons in these grids, and for some restricted variants. These bounds are tight up to polynomial factors.

1 Introduction

Can a convex polygon P in the plane be reconstructed from the projections of its vertices to the coordinate axes? Assuming that no two vertices of P share the same x- or y-coordinate, we arrive at the following problem: given two sets, X and Y, each containing n real numbers, does the Cartesian product $X \times Y$ support a convex polygon with n vertices? We say that $X \times Y$ *contains* a polygon P if every vertex of P is in $X \times Y$; and $X \times Y$ *supports* P if it contains P and no two vertices of P share an x- or y-coordinate. For short, we call the Cartesian product $X \times Y$ an $n \times n$ *grid*.

Not every $n \times n$ grid supports a convex n-gon. This is the case already for $n = 5$ (Figure 1). Several interesting questions arise: can we decide efficiently whether an $n \times n$-grid supports a convex n-gon? How can we find the largest k such that it supports a convex k-gon? What is the largest k such that *every* $n \times n$ grid supports a convex k-gon? How many convex polygons does an $n \times n$

*This work was initiated at the 2017 Fields Workshop on Discrete and Computational Geometry (Carleton University, Ottawa, ON, July 31–August 4, 2017).

†School of Electrical Engineering and Computer Science, University of Ottawa, Canada, jdecaruf@uottawa.ca
‡Department of Computer Science, University of Wisconsin-Milwaukee, USA dumitres@uwm.edu
§Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands, w.meulemans@tue.nl
¶Department of Computational Mathematics, Science and Engineering, Michigan State University, USA, ophelder@egr.msu.edu
‖LIRMM, CNRS & Université de Montpellier, France, claire.pennarun@gmail.com
**Department of Mathematics, California State University Northridge, Los Angeles, CA; and Department of Computer Science, Tufts University, Medford, MA, USA, csaba.toth@csun.edu
††School of Computer Science, Carleton University, Canada, sander@cg.scs.carleton.ca

grid support, or contain? We initiate the study of these questions for convex polygons, and their higher dimensional variants for convex polyhedra.

Our Results. We first show that every $n \times n$ grid supports a convex polygon with $(1 - o(1)) \log n$ vertices[1]; this bound is tight up to a constant factor: there are $n \times n$ grids that do not support convex polygons with more than $4(\lceil \log n \rceil + 1)$ vertices. We generalize our upper and lower bounds to higher dimensions, and show that every d-dimensional Cartesian product $\prod_{i=1}^{d} Y_i$, where $|Y_i| = n$ and d is constant, contains $\Omega(\log^{d-1} n)$ points in convex position; this bound is also tight apart from constant factors (Section 2). Next, we present polynomial-time algorithms to find a maximum supported convex polygon that is x- or y-monotone. We show how to efficiently approximate the maximum size of a supported convex polygon up to a factor of two (Section 3). Finally, we present tight asymptotic bounds for the maximum number of convex polygons supported by an $n \times n$ grid (Section 4). We conclude with open problems (Section 5).

Related Work. Erdős and Szekeres proved, as one of the first Ramsey-type results in combinatorial geometry [17], that for every $k \in \mathbb{N}$, a sufficiently large point set in the plane in general position contains k points in convex position. The minimum cardinality of a point set that guarantees k points in convex position is known as the Erdős–Szekeres number, $f(k)$. They proved that $2^{k-2} + 1 \leq f(k) \leq \binom{2k-4}{k-2} + 1 = 4^{k(1-o(1))}$, and conjectured that the lower bound is tight [15]. The current best upper bound, due to Suk [30], is $f(k) \leq 2^{k(1+o(1))}$. In other words, every set of n points in general position in the plane contains $(1 - o(1)) \log n$ points in convex position, and this bound is tight up to lower-order terms.

In dimension $d \geq 3$, the asymptotic growth rate of the Erdős–Szekeres number is not known. By the Erdős–Szekeres theorem, every set of n points in general position in $\mathbb{R}^d$ contains $\Omega(\log n)$ points in convex position (it is enough to find points whose projections onto a generic plane are in convex position). For every constant $d \geq 2$, Károlyi and Valtr [20] and Valtr [31] constructed n-element sets in general position in $\mathbb{R}^d$ in which no more than $O(\log^{d-1} n)$ points are in convex position. Both constructions are recursive, and one of them is related to high-dimensional Horton sets [31]. These bounds are conjectured to be optimal apart from constant factors. Our results establish the same $O(\log^{d-1} n)$ upper bound for Cartesian products, for which it is tight apart from constant factors. However our results do not improve the bounds for points in general position.

Algorithmically, one can find a largest convex cap in a given set of n points in $\mathbb{R}^2$ in $O(n^2 \log n)$ time by dynamic programming [11], and a largest subset in convex position in $O(n^3)$ time [8, 11]. The same approach can be used for counting the number of convex polygons contained in a given point set [21]. While this approach applies to grids, it is unclear how to include the restriction that each coordinate is used at most once. On the negative side, finding a largest subset in convex position in a point set in $\mathbb{R}^d$ for dimensions $d \geq 3$ was recently shown to be NP-hard [16].

There has been significant interest in *counting* the number of convex polygons in various point sets. Answering a question of Hammer, Erdős [14] proved that every set of n points in general position in $\mathbb{R}^2$ contains $\exp(\Theta(\log^2 n))$ subsets in convex position, and this bound is the best possible. Bárány and Pach [3] showed that the number of convex polygons in an $n \times n$ section of the integer lattice is $\exp\left(O(n^{2/3})\right)$. Bárány and Vershik [4] generalized this bound to d-dimensions and showed that there are $\exp\left(O(n^{(d-1)/(d+1)})\right)$ convex polytopes in an $n \times \cdots \times n$ section of $\mathbb{Z}^d$. Note that the exponent is sublinear in n for every $d \geq 2$. We prove that an $n \times n$ Cartesian product can contain $\exp(\Theta(n))$ convex polygons, significantly more than integer grids, and our bounds are tight up to

[1]All logarithms in this paper are of base 2.

polynomial factors.

Motivated by integer programming and geometric number theory, lattice polytopes (whose vertices are in $\mathbb{Z}^d$) have been intensely studied; refer to [2, 5]. However, results for lattices do not extend to arbitrary Cartesian products. Recently, several deep results have been established for Cartesian products in incidence geometry and additive combinatorics [24, 25, 26, 29], while the analogous statements for points sets in general position remain elusive.

Definitions. A polygon P in $\mathbb{R}^2$ is *convex* if all of its internal angles are strictly smaller than π. A point set in $\mathbb{R}^2$ is in *convex position* if it is the vertex set of a convex polygon; and it is in *general position* if no three points are collinear. Similarly, a polyhedron P in $\mathbb{R}^d$ is *convex* if it is the convex hull of a finite set of points. A point set in $\mathbb{R}^d$ is in *convex position* if it is the vertex set of a convex polytope; and it is in *general position* if no $d+1$ points lie on a hyperplane. In $\mathbb{R}^d$, we say that the x_d-axis is *vertical*, hyperplanes orthogonal to x_d are *horizontal*, and understand the above-below relationship with respect to the x_d-axis. Let $\mathbf{e}_d$ be a standard basis vector parallel to the x_d-axis. A point set P in $\mathbb{R}^d$ is *full-dimensional* if no hyperplane contains P.

We consider special types of convex polygons. Let P be a convex polygon with vertices $((x_1, y_1), \ldots, (x_k, y_k))$ in clockwise order. We say that P is a *convex cap* if the x- or y-coordinates are strictly monotonic, and a *convex chain* if both the x- and y-coordinates are strictly monotonic. We distinguish four types of convex caps (resp., chains) based on the monotonicity of the coordinates as follows:

- *convex caps* come in four types {↷, ⤹, ⤻, ⤸}. We have
  $P \in$ ↷ if and only if $(x_i)_{i=1}^k$ strictly increases;
  $P \in$ ⤹ if and only if $(y_i)_{i=1}^k$ strictly increases;
  $P \in$ ⤻ if and only if $(x_i)_{i=1}^k$ strictly decreases;
  $P \in$ ⤸ if and only if $(y_i)_{i=1}^k$ strictly decreases;
- *convex chains* come in four types {↗, ↘, ↙, ↖}. We have
  ↗ = ⤹ $\cap$ ↷, ↘ = ↷ $\cap$ ⤸, ↖ = ⤻ $\cap$ ⤹, ↙ = ⤸ $\cap$ ⤻.

Initial observations. It is easy to see that for $n = 3, 4$, every $n \times n$ grid supports a convex n-gon. However, there exists a 5×5 grid that does not support any convex pentagon (cf. Fig. 1). Interestingly, every 6×6 grid supports a convex pentagon.

Lemma 1. *Every 6×6 grid $X \times Y$ supports a convex polygon of size at least 5.*

Proof. Let $X' = X \setminus \{\min(X), \max(X)\}$ and $Y' = Y \setminus \{\min(Y), \max(Y)\}$. The 4×4 grid $X' \times Y'$ supports a convex chain P' of size 3 between two opposite corners of $X' \times Y'$. Then one x-coordinate $x' \in X'$ and one y-coordinate $y' \in Y'$ are not used by P'. Without loss of generality, assume that $P' \in$ ↗. Then the convex polygon containing the points of P' and $(x', \min(Y))$ and $(\max(X), y')$ is a supported convex polygon of size 5 on $X \times Y$. $\square$

2 Extremal Bounds for Convex Polytopes in Cartesian Products

2.1 Lower Bounds in the Plane

In this section, we show that for every $n \geq 3$, every $n \times n$ grid supports a convex polygon with $\Omega(\log n)$ vertices. The results on the Erdős–Szekeres number cannot be used directly, since they

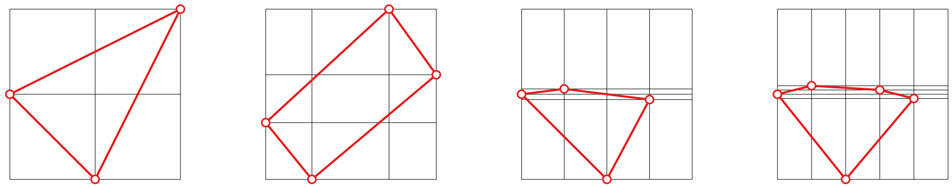

Figure 1: Maximum-size supported convex polygons of respective sizes 3, 4, 4, and 5 in $n \times n$ grids, where n is between 3 and 6.

crucially use the assumption that the given set of points is in general position. An $n \times n$ section of the integer lattice is known to contain $\Theta(n)$ points in general position [13], and this number is conjectured to be $\frac{\pi}{\sqrt{3}}n(1+o(1))$ [18, 32]. However, this result does not apply to arbitrary Cartesian products. It is worth noting that higher dimensional variants for the integer lattice are poorly understood: it is known that an $n \times n \times n$ section of $\mathbb{Z}^3$ contains $\Theta(n^2)$ points no three of which are collinear [23], but no similar statements are known in higher dimensions. We use a recent result from incidence geometry.

Lemma 2. *(Payne and Wood [22]) Every set of N points in the plane with at most ℓ collinear, $\ell \leq O(\sqrt{N})$, contains a set of $\Omega(\sqrt{N/\log \ell})$ points in general position.*

Lemma 3. *Every $n \times n$ grid supports a convex polygon of size $(1-o(1))\log n$.*

Proof. Every $n \times n$ grid contains a set of $\Omega(\sqrt{n^2/\log n}) = \Omega(n/\sqrt{\log n})$ points in general position by applying Lemma 2 with $N = n^2$ and $\ell = n$. Discarding points with the same x- or y-coordinate reduces the size by a factor at most $\frac{1}{4}$, so this asymptotic bound also holds when coordinates in X and Y are used at most once. By Suk's result [30], the grid supports a convex polygon with at least $(1-o(1))\big(\log(n/\sqrt{\log n})\big) = (1-o(1))\log n$ vertices. □

2.2 Upper Bounds in the Plane

For the upper bound, we construct $n \times n$ Cartesian products that do not support large convex chains. For $n = 8$, such a grid is depicted in Figure 2.

Lemma 4. *For every $n \in \mathbb{N}$, there exists an $n \times n$ grid that contains at most $4(\lceil \log n \rceil + 1)$ points in convex position.*

Proof. Let $g(n)$ be the maximum integer such that for all n-element sets $X, Y \subset \mathbb{R}$, the grid $X \times Y$ supports a convex polygon of size $g(n)$; clearly $g(n)$ is nondecreasing. Let k be the minimum integer such that $n \leq 2^k$; thus $\lceil \log n \rceil \leq k$ and $g(n) \leq g(2^k)$. We show that $g(2^k) \leq 4(k+1)$ and thereby establish that $g(n) \leq 4(k+1)$.

Assume w.l.o.g. that $n = 2^k$, and let $X = \{0, \ldots, n-1\}$. For a k-bit integer m, let m_i be the bit at its i-th position, such that $m = \sum_{i=0}^{k-1} m_i 2^i$. Let $Y = \{\sum_{i=0}^{k-1} m_i (2n)^i \mid 0 \leq m \leq n-1\}$ (see Fig. 2). Both X and Y are symmetric: $X = \{\max(X) - x \mid x \in X\}$ and $Y = \{\max(Y) - y \mid y \in Y\}$. Thus, it suffices to show that no convex chain $P \in \curvearrowright$ of size greater than $k+1$ exists.

Consider two points, $p = (x, y)$ and $p' = (x', y')$, in $X \times Y$ such that $x < x'$ and $y < y'$. Assume $y = \sum_{i=0}^{k-1} m_i (2n)^i$ and $y' = \sum_{i=0}^{k-1} m'_i (2n)^i$. The slope of the line spanned by p and p'

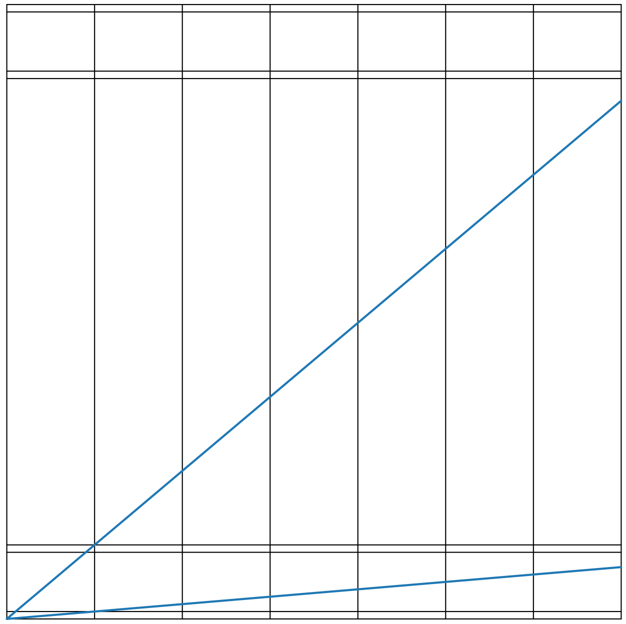

Figure 2: An 8×8 grid without convex chains of size greater than $4 = \log 8{+}1$, where $X = \{0, 1, \ldots, 7\}$ and $Y = \{0, 1, 16, 17, 256, 257, 272, 273\}$. Two lines through pairs of grid points are drawn in blue.

is $\mathrm{slope}(p, p') = \sum_{i=0}^{k-1}(m'_i - m_i)(2n)^i/(x' - x)$. Let j be the largest index such that $m_j \neq m'_j$. Then $y < y'$ implies $m_j < m'_j$, and we can bound the slope as follows:

$$\mathrm{slope}(p, p') \geq \frac{(2n)^j - \sum_{i=0}^{j-1}(2n)^i}{x' - x} > \frac{(2n)^j - 2(2n)^{j-1}}{n - 1} = 2 \cdot (2n)^{j-1},$$

$$\mathrm{slope}(p, p') \leq \frac{\sum_{i=0}^{j}(2n)^i}{x' - x} \leq \frac{\sum_{i=0}^{j}(2n)^i}{1} = \frac{(2n)^{j+1} - 1}{2n - 1} < 2 \cdot (2n)^j \,.$$

Hence, $\mathrm{slope}(p, p') \in I_j = (2 \cdot (2n)^{j-1}, 2 \cdot (2n)^j)$. Let us define the family of intervals $I_0, I_1, \ldots, I_{k-1}$ analogously, and note that these intervals are pairwise disjoint. Suppose that some convex chain $P \in \nearrow$ contains more than $k + 1$ points. Since the slopes of the first $k + 1$ edges of P decrease monotonically, by the pigeonhole principle, there must be three consecutive vertices $p = (x, y)$, $p' = (x', y')$, and $p'' = (x'', y'')$ of P such that both $\mathrm{slope}(p, p')$ and $\mathrm{slope}(p', p'')$ are in the same interval, say I_j. Assume that $y = \sum_{i=0}^{k-1} m_i(2n)^i$, $y' = \sum_{i=0}^{k-1} m'_i(2n)^{i+1}$, and $y'' = \sum_{i=0}^{k-1} m''_i(2n)^{i+1}$. Then j is the largest index such that $m_j \neq m'_j$, and also the largest index such that $m'_j \neq m''_j$. Because $m < m' < m''$, we have $m_j < m'_j < m''_j$, which is impossible since each of m_j, m'_j and m''_j is either 0 or 1.

Hence, $X \times Y$ does not contain any convex chain in $\nearrow$ of size greater than $k + 1$. Analogously, every convex chain in $\searrow$, $\swarrow$, or $\nwarrow$ has at most $k + 1$ vertices. Consequently, $X \times Y$ contains at most $4(k + 1)$ points in convex position. $\square$

2.3 Upper Bounds in Higher Dimensions

We construct Cartesian products in $\mathbb{R}^d$, for $d \geq 3$, that match the best known upper bound $O(\log^{d-1} n)$ for the Erdős–Szekeres numbers in d-dimensions for points in general position. Our

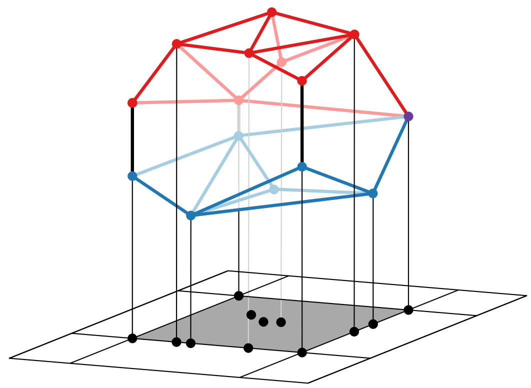

Figure 3: A polyhedron $\text{conv}(P)$ in $\mathbb{R}^3$, whose projection $\text{conv}(P)^{\text{proj}}$ is a rectangle. Seven points in P are projected onto the four vertices of $\text{conv}(P)^{\text{proj}}$. Overall the *silhouette* of P contains 12 points. Red (blue) vertices are upper (lower); the purple point is both upper and lower.

construction generalizes the ideas from the proof of Lemma 4 to d-space.

Lemma 5. *Let $d \geq 2$ be an integer. For every integer $n \geq 2$, there exist n-element sets $Y_i \subseteq \mathbb{R}$ for $i = 1, \ldots, d$, such that the Cartesian product $Y = \prod_{i=1}^{d} Y_i$ contains at most $O(\log^{d-1} n)$ points in convex position.*

Proof. We construct point sets recursively. For $d = 2$, the result follows from Lemma 4. For integers $d \geq 3$ and $0 \leq i \leq j$, we define $S_d(i,j)$ as a Cartesian product of d sets, where the first $d-1$ sets have 2^j elements and the last set has 2^i elements. We then show that $S_d(i,j)$ does not contain the vertex set pf any full-dimensional convex polyhedron with more than $2^{d^2-d+1} \cdot i \cdot j^{d-2}$ vertices (there is no restriction on lower-dimensional convex polyhedra it might contain).

To initialize the recursion, we define boundary values as follows: For every integer $j \geq 0$, let $S_2(j,j)$ be the $2^j \times 2^j$ grid defined in the proof of Lemma 4 that does not contain more than $4(j+1)$ points in convex position. Note that every line that contains 3 or more points from $S_2(j,j)$ is axis-parallel (this property was not needed in the proof of Lemma 4). Assume now that $d \geq 3$, and $S_{d-1}(j,j)$ has been defined for all $j \geq 0$; and for all $k = 1, \ldots, d$, every k-dimensional flat containing $2^k + 1$ or more points is parallel to at least one coordinate axis. Let j be a nonnegative integer. We now construct $S_d(i,j)$ for all integers $0 \leq i \leq j$ as follows.

Let $S_d(0,j) = S_{d-1}(j,j) \times \{0\}$. For $i = 1, \ldots, j$, we define $S_d(i,j)$ as the disjoint union of two translates of $S_d(i-1,j)$. Specifically, let $S_d(i,j) = A \cup B$, where $A = S_d(i-1,j)$ and $B = A + \lambda_d^i \mathbf{e}_d$, where $\lambda_d^i > 0$ is sufficiently large (as specified below) and algebraically independent from the coordinates of $S_d(i,j)$, such that for all $k = 1, \ldots, d$, every k-dimensional flat containing $2^k + 1$ or more points is parallel to at least one coordinate axis.

Let $P \subset S_d(i,j)$ be a full-dimensional set in convex position. The orthogonal projection of $\text{conv}(P)$ to the horizontal hyperplane $x_d = 0$ is a convex polytope in $\mathbb{R}^{d-1}$ that we denote by $\text{conv}(P)^{\text{proj}}$; refer to Fig. 3. The *silhouette* of P is the subset of vertices whose orthogonal projection to $x_d = 0$ lies on the boundary of $\text{conv}(P)^{\text{proj}}$. Since no three points in P are collinear, then at most two points in P are projected to a same point. A point $p \in P$ is an *upper* (resp., *lower*) vertex if P lies in the closed halfspace below (resp., above) some tangent hyperplane of $\text{conv}(P)$ at p (a point in p may be both upper and lower vertex).

We prove, by double induction on d and i, the following:

Claim 1. If $P \subset S_d(i,j)$ is a full-dimensional set in convex position, then P contains at most $2^{d(d-1)} \cdot i \cdot j^{d-2}$ upper (resp., lower) vertices of $\mathrm{conv}(P)$.

For $d = 2$ and $i = j$, this holds by definition (cf. Lemma 4). For $i = 0$, the set $S_d(0,j) = S_{d-1}(j,j) \times \{0\}$ lies in a horizontal hyperplane in $\mathbb{R}^d$, and so it is not full-dimensional, hence the claim vacuously holds. By induction, $S_{d-1}(j,j)$ contains at most $2^{(d-1)(d-2)} \cdot j \cdot j^{d-3}$ upper (resp., lower) vertices in $\mathbb{R}^{d-1}$, hence $S_d(0,j)$ has at most $2 \cdot 2^{(d-1)(d-2)} \cdot j \cdot j^{d-3}$ extreme points in $\mathbb{R}^d$. Assume that $d \geq 3$, $1 = i \leq j$, and the claim holds for $S_{d-1}(j,j)$. We prove the claim for $S_d(i,j)$. The set $S_d(1,j)$ is the disjoint union of $A = S_d(0,j)$ and $B = S_d(0,j) + \lambda_d^1 \mathbf{e}_d$. Every upper (resp., lower) vertex of $S_d(1,j)$ is an extreme vertex in A or B, hence $S_d(1,j)$ contains at most $4 \cdot 2^{(d-1)(d-2)} \cdot j^{d-2} = 2^{d^2-3d+4} \cdot j^{d-2} < 2^{d^2-d} \cdot 1 \cdot j^{d-2}$ upper (resp., lower) vertices, as required, where we used that $d \geq 3$.

In the general case, we assume that $d \geq 3$, $2 \leq i \leq j$, and the claim holds for $S_{d-1}(j,j)$ and $S_d(i-1,j)$. We prove the claim for $S_d(i,j)$. Recall that $S_d(i,j)$ is the disjoint union of two translates of $S_d(i-1,j)$, namely $A = S_d(i-1,j)$ and $B = S_d(i-1,j) + \lambda_d^i \mathbf{e}_d$. Let $P \subset S_d(i,j)$ be a full-dimensional set. We partition the upper vertices in P as follows. Let $P_0 \subset P$ be the set of upper vertices whose orthogonal projection to $x_d = 0$ is a vertex of $\mathrm{conv}(P)^{\mathrm{proj}}$. For $k = 1, \ldots, d-1$, let $P_k \subset P$ be the set of upper vertices whose orthogonal projection to $x_d = 0$ lies in the relative interior of a k-face of $\mathrm{conv}(P)^{\mathrm{proj}}$. By construction, only axis-aligned faces can contain interior points. For the $(d-1)$-dimensional polytope $\mathrm{conv}(P)^{\mathrm{proj}}$, the axis-aligned k-faces ($k = 1, \ldots, d-1$) can be partitioned into $\binom{d-1}{k}$ equivalence classes, based on the set of parallel coordinate axes.

The orthogonal projection of $S_d(i,j)$ to $x_d = 0$ is $S_{d-1}(j,j)$, and the orthogonal projection of P_0 is, $P_0^{\mathrm{proj}} \subset S_{d-1}(j,j)$, is the vertex set of a $(d-1)$-dimensional convex polyhedron. By induction, $|P_0| \leq 2 \cdot 2^{(d-1)(d-2)} \cdot j \cdot j^{d-3} = 2^{d^2-3d+3} j^{d-2}$. We show the following.

Claim 2. For every axis-aligned face F of $\mathrm{conv}(P)^{\mathrm{proj}}$, the set of upper vertices that project to the interior of F is contained in either A or B.

Let F be an axis-aligned k-face of $\mathrm{conv}(P)^{\mathrm{proj}}$ for $k \in \{1, 2, \ldots, d-1\}$. Let $P(F) \subset P$ be the set of upper vertices whose orthogonal projection lies in the interior of F, and let $P(\partial F)$ be the set of upper vertices whose orthogonal projection lies in the boundary of F. Let $P(\partial F)^{\mathrm{proj}}$ be the orthogonal projection of $P(\partial F)$ to the hyperplane $x_d = 0$. Consider the point set $P' = P(\partial F)^{\mathrm{proj}} \cup P(F)$, and observe that if $P(F) \neq \emptyset$, then it is a vertex set of a $(k+1)$-dimensional polytope in which all vertices are upper. It remains to show that $P(F) \subseteq A$ or $P(F) \subseteq B$. Suppose, for the sake of contradiction, that $P(F)$ contains points from both A and B. Let p_a be a vertex in $P(F)$ with the maximum x_d-coordinate. The 1-skeleton of $\mathrm{conv}(P')$ contains a x_d-monotonically decreasing path from p_a to an x_d-minimal vertex in P'. Let p_b be the neighbor of p_a along such a path. Then $p_b \in B$ by the choice of p_a. Every hyperplane containing p_a and p_b, in particular the tangent hyperplane of P' containing the edge $p_a p_b$, partitions $P(\partial F)^{\mathrm{proj}}$, which is a contradiction.

We can now finish the proof of Claim 1. We have seen that $|P_0| \leq 2 \cdot 2^{(d-1)(d-2)} \cdot j \cdot j^{d-3} = 2^{d^2-3d+3} j^{d-2}$. For $k = 1, \ldots, d-1$, the axis-aligned k-faces of $\mathrm{conv}(P)^{\mathrm{proj}}$ are partitioned into $\binom{d-1}{k}$ equivalence classes. In each equivalence class, the k-faces F are parallel to the same coordinate axes; in the projection to an orthogonal $(d-k)$-dimensional space, they become vertices of a $(d-k-1)$-dimensional polytope. By induction, the number of such vertices (hence the size of the equivalence class) is $2^{(d-k-1)(d-k-2)+1}j^{d-k-1}$. Overall, we have

$$\begin{aligned}|P_k| &\leq \binom{d-1}{k} \cdot 2^{(d-k-1)(d-k-2)+1}j^{d-k-1} \cdot 2^{k(k-1)} \cdot (i-1) \cdot j^{k-1} \\ &\leq \binom{d-1}{k} \cdot 2^{(d-1)(d-2)} \cdot (i-1) \cdot j^{d-2}.\end{aligned}$$

Altogether, the number of upper vertices is

$$\begin{aligned}\sum_{k=0}^{d-1} |P_k| &\leq 2^{d^2-3d+3}j^{d-2} + \sum_{k=1}^{d-1} \binom{d-1}{k} 2^{(d-1)(d-2)} \cdot (i-1) \cdot j^{d-2} \\ &< 2^{d(d-1)}j^{d-2} + 2^{d-1} \cdot 2^{(d-1)(d-2)} \cdot (i-1) \cdot j^{d-2} \\ &< 2^{d(d-1)} \cdot i \cdot j^{d-2},\end{aligned}$$

as required, where we used the binomial theorem and the inequality $3 \leq d$. □

2.4 Lower Bound in Higher Dimensions

The proof technique in Section 2.1 is insufficient for establishing a lower bound of $\Omega(\log^{d-1} n)$ for $d \geq 3$. Whereas a d-dimensional $n \times \ldots \times n$ grid contains $\Omega(n^\delta)$ points in general position for some $\delta = \delta(d) > 0$ [7], the current best lower bound on the number of points in convex position in any set of n points in general position in $\mathbb{R}^d$ is $\Omega(\log n)$; the conjectured value is $\Omega(\log^{d-1} n)$. Instead, we rely on the structure of Cartesian products and induction on d. Our main result in this section is the following.

Theorem 6. *Every d-dimensional Cartesian product $\prod_{i=1}^d Y_i$, where $|Y_i| = n$ and d is fixed, contains $\Omega(\log^{d-1} n)$ points in convex position.*

We say that a strictly increasing sequence of real numbers $A = (a_1, \ldots, a_n)$, has the *monotone differences* property (for short, A is *MD*), if

- $a_{i+1} - a_i > a_i - a_{i-1}$ for $i = 2, \ldots, n-1$, or
- $a_{i+1} - a_i < a_i - a_{i-1}$ for $i = 2, \ldots, n-1$.

Further, the sequence A is r-*MD* for some $r > 1$ if

- $a_{i+1} - a_i \geq r(a_i - a_{i-1})$ for $i = 2, \ldots, n-1$, or
- $a_{i+1} - a_i \leq (a_i - a_{i-1})/r$ for $i = 2, \ldots, n-1$.

A finite set $X \subseteq \mathbb{R}$ is *MD* (resp., r-*MD*) if its elements arranged in increasing order form an MD (resp., r-MD) sequence. These sequences are intimately related to convexity: a strictly increasing sequence $A = (a_1, \ldots, a_n)$ is MD if and only if there exists a monotone (increasing or decreasing) convex function $f : \mathbb{R} \to \mathbb{R}$ such that $a_i = f(i)$ for all $i = 1, \ldots, n$. MD sets have been studied in additive combinatorics [12, 19, 24, 28].

We first show that every n-element set $X \subseteq \mathbb{R}$ contains an MD subset of size $\Omega(\log n)$, and this bound is the best possible (Lemma 7). In contrast, every n-term arithmetic progression contains an MD subsequence of $\Theta(\sqrt{n})$ terms: for example $(0, \ldots, n-1)$ contains the subsequence

$(i^2 : i = 0, \ldots, \lfloor\sqrt{n-1}\rfloor)$. We then show that for constant $d \geq 2$, the d-dimensional Cartesian product of n-element MD sets contains $\Theta(n^{d-1})$ points in convex position. The combination of these results immediately implies that every $n \times \ldots \times n$ Cartesian product in $\mathbb{R}^d$ contains $\Omega(\log^{d-1} n)$ points in convex position.

The following lemma gives a lower bound for MD sequences. It is known that a monotone sequence of n reals contains a 2-MD sequence (satisfying the so-called *doubling differences condition* [27]) of size $\Omega(\log n)$ [6, Lemma 4.1]; see also [9] for related recent results.

Lemma 7. *Every set of n real numbers contains an MD subset of size $\lfloor(\log n)/2\rfloor + 1$. For every $n \in \mathbb{N}$, there exists a set of n real numbers in which the size of every MD subset is at most $\lceil\log n\rceil + 1$.*

Proof. Let $X = (x_0, \ldots, x_{n-1})$ be a strictly increasing sequence. Assume w.l.o.g. that $n = 2^\ell + 1$ for some $\ell \in \mathbb{N}$. We construct a sequence of nested intervals

$$[a_0, b_0] \supset [a_1, b_1] \supset \ldots \supset [a_\ell, b_\ell]$$

such that the endpoints of the intervals are in X and the lengths of the intervals decrease by factors of 2 or higher, that is, $b_i - a_i \leq (b_{i-1} - a_{i-1})/2$ for $i = 1, \ldots, \ell$.

We start with the interval $[a_0, b_0] = [x_0, x_{n-1}]$; and for every $i = 0, \ldots, \ell - 1$, we divide $[a_i, b_i]$ into two intervals at the median, and recurse on the shorter interval.

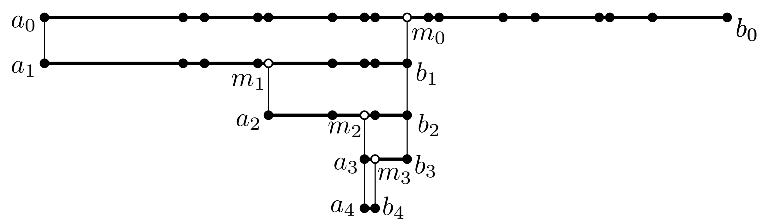

Figure 4: A sequence X of 17 elements and nested intervals $[a_0, b_0] \supset \ldots \supset [a_4, b_4]$.

By partitioning $[a_i, b_i]$ at the median, the algorithm maintains the invariant that $[a_i, b_i]$ contains $2^{\ell-i} + 1$ elements of X. Note that for every $i = 1, \ldots, \ell$, we have either ($a_{i-1} = a_i$ and $b_{i-1} < b_i$) or ($a_{i-1} = a_i$ and $b_i < b_{i-1}$). Consequently, the sequences $A = (a_0, a_1, \ldots, a_\ell)$ and $B = (b_\ell, b_{\ell-1}, \ldots, b_0)$ both increase (not necessarily strictly), and at least one of them contains at least $1 + \ell/2$ distinct terms. Assume w.l.o.g. that A contains at least $1 + \ell/2$ distinct terms. Let $C = (c_0, \ldots, c_k)$ be a maximal strictly increasing subsequence of A. Then $k \geq \ell/2 = \lfloor(\log n)/2\rfloor$.

We show that C is an MD sequence. Let $i \in \{1, \ldots, k-1\}$. Assume that $c_i = a_j = \ldots = a_{j'}$ for consecutive indices $j, \ldots, j'$. Then $c_{i-1} = a_{j-1}$, $c_i = a_j$, and $c_{i+1} = a_{j'+1}$. By construction, $c_i \in [a_{j-1}, b_{j-1}] = [c_{i-1}, b_j]$ such that $c_i - c_{i-1} \geq b_j - c_i$. Similarly, $c_{i+1} \in [a_{j'}, b_{j'}] = [c_i, b_{j'}]$ such that $c_{i+1} - c_i \geq b_{j'} - c_{i+1}$. However, $[a_j, b_j] \subset [a_{j-1}, b_{j-1}]$. As required, this yields

$$c_{i+1} - c_i = a_{j'+1} - a_j < b_j - a_j \leq \frac{b_{j-1} - a_{j-1}}{2} \leq a_j - a_{j-1} = c_i - c_{i-1}.$$

The upper bound construction is the point set Y defined in the proof of Lemma 4, for which every chain in $\nwarrow$ or $\swarrow$ supported by $\{0, \ldots, n-1\} \times Y$ has at most $\lceil\log n\rceil + 1$ vertices. Let $\{b_0, \ldots, b_{\ell-1}\} \subset Y$ be an MD subset such that $b_0 < \ldots < b_{\ell-1}$. Then $\{(i, b_i) : i = 0, \ldots, \ell - 1\} \subset X \times Y$ is in $\nwarrow$ or $\swarrow$. Consequently, every MD subset of Y has at most $\lceil\log n\rceil + 1$ terms, as claimed. $\square$

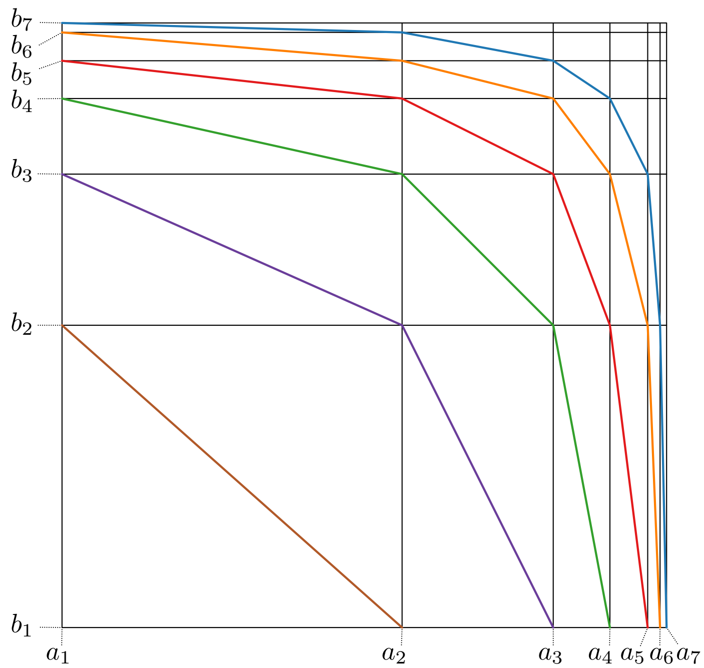

Figure 5: A 7×7 grid $\{a_1, \ldots, a_7\} \times \{b_1, \ldots, b_7\}$, where the differences between consecutive x-coordinates (resp., y-coordinates) decrease by factors of 2 or higher. The point sets $\{(0,0)\} \cup \{(a_i, b_j) : i + j = k\}$, for $k = 2, \ldots, 8$, form nested convex chains.

We show how to use Lemma 7 to establish a lower bound in the plane. While this approach yields worse constant coefficients than Lemma 3, its main advantage is that it generalizes to higher dimensions (see Lemma 10 below).

Lemma 8. *The Cartesian product of two MD sets, each of size n, supports n points in convex position.*

Proof. Let $A = \{a_1, \ldots, a_n\}$ and $B = \{b_1, \ldots, b_n\}$ be MD sets such that $a_i < a_{i+1}$ and $b_i < b_{i+1}$ for $i = 1, \ldots, n-1$. We may assume, by applying a reflection if necessary, that $a_{i+1} - a_i < a_i - a_{i-1}$ and $b_{i+1} - b_i < b_i - b_{i-1}$, for $i = 2, \ldots, n-1$ (see Fig. 5).

We define $P \subset A \times B$ as the set of n points (a_i, b_j) such that $i + j = n + 1$. By construction, every horizontal (vertical) line contains at most one point in P. Since the differences $a_i - a_{i-1}$ are positive and strictly decrease in i; and the differences $b_{\ell-i} - b_{\ell-i-1}$ are negative and their absolute values strictly increase in i, the slopes $(b_{\ell-i} - b_{\ell-i-1})/(a_i - a_{i-1})$ strictly decrease, which proves the convexity of P. $\square$

Lemma 9. *The Cartesian product of three MD sets, each of size n, contains $\binom{n+1}{2}$ points in convex position.*

Proof. Let $A = \{a_1, \ldots, a_n\}$, $B = \{b_1, \ldots, b_n\}$, and $C = \{c_1, \ldots, c_n\}$ be MD sets, where the elements are labeled in increasing order. We may assume, by applying a reflection in the x-, y-, or z-axis if necessary, that

$$a_{i+1} - a_i < a_i - a_{i-1}, \qquad b_{i+1} - b_i < b_i - b_{i-1}, \qquad c_{i+1} - c_i < c_i - c_{i-1},$$

for $i = 2, \ldots, n-1$. For $i, j, k \in \{1, \ldots, n\}$, let $p_{i,j,k} = (a_i, b_j, c_k) \in A \times B \times C$. We can now let $P = \{p_{i,j,k} : i + j + k = n + 2\}$. It is clear that $|P| = \sum_{i=1}^{n} i = \binom{n+1}{2}$. We let $P' = P \cup \{p_{1,1,1}\}$ and show that the points in P' are in convex position.

By Lemma 8, the points in P' lying in the planes $x = a_1$, $y = b_1$, and $z = c_1$ are each in convex position. These convex $(n+1)$-gons are faces of the convex hull of P, denoted $\mathrm{conv}(P)$. We show that the remaining faces of $\mathrm{conv}(P)$ are the triangles $T'_{i,j,k}$ spanned by $p_{i,j,k}$, $p_{i,j+1,k-1}$, and $p_{i+1,j,k-1}$; and the triangles $T''_{i,j,k}$ spanned by $p_{i,j,k}$, $p_{i,j-1,k+1}$, and $p_{i-1,j,k+1}$.

The projection of these triangles to an xy-plane is shown in Fig. 5. By construction, the union of these faces is homeomorphic to a sphere. It suffices to show that the dihedral angle between any two adjacent triangles is convex. Without loss of generality, consider triangle $T'_{i,j,k}$, which is adjacent to (up to) three other triangles: $T''_{i+1,j,k-1}$, $T''_{i,j+1,k-1}$, and $T''_{i+1,j+1,k-2}$. Consider first the triangles $T'_{i,j,k}$ and $T''_{i+1,j+1,k-2}$. They share the edge $p_{i+1,j-1,k+1}p_{i-1,j+1,k+1}$, which lies in the xy-plane $z = c_{k+1}$. The orthogonal projections of these triangles to an xy-plane are congruent, however their extents in the z-axis are $c_{i+1} - c_i$ and $c_i - c_{i-1}$, respectively. Since $c_{i+1} - c_i < c_i - c_{i-1}$, their diheral angle is convex. Similarly, the dihedral angles between $T'_{i,j,k}$ and $T''_{i+1,j,k-1}$ (resp., $T''_{i,j+1,k-1}$) is convex because $a_{i+1} - a_i < a_i - a_{i-1}$ and $b_{i+1} - b_i < b_i - b_{i-1}$. □

The proof technique of Lemma 9 generalizes to higher dimensions:

Lemma 10. *For every constant $d \geq 2$, the Cartesian product of d MD sets, each of size n, contains $\Omega(n^{d-1})$ points in convex position.*

Proof. We proceed by induction on d. For $d = 2, 3$, Lemmas 8-9 prove the claim. Assume that $d \geq 4$, and the claim holds in lower dimensions. For every $i = 1, \ldots, d$, let $A_i = \{a_{i,1} < \ldots < a_{i,n}\} \subset \mathbb{R}$ be an MD set such that the elements are labeled in increasing order. We may assume w.l.o.g. that the differences between consecutive elements in A_i strictly decrease for all $i = 1, \ldots, d$.

For every vector $\mathbf{v} = (v_1, \ldots, v_d) \in \{1, \ldots, n\}^d$, let $p_{\mathbf{v}} = (a_{1,v_1}, a_{2,v_2}, \ldots, a_{d,v_d}) \in \prod_{i=1}^{d} A_i$. Let $P = \{p_{\mathbf{v}} : \sum_{i=1}^{d} v_i = n + d - 1\}$. It is easy to see that $|P| = \Theta(n^{d-1})$. Let $P' = P \cup \{p_{1,\ldots,1}\}$. We show that the points in P' are in convex position. By induction, the points of P' lying in the hyperplanes $x_i = a_{i,1}$, for $i = 1, \ldots, d$, are each in convex position, hence they each define $(d-1)$-dimensional facets of $\mathrm{conv}(P')$. We define additional $(d-1)$-dimensional facets such that the union of these facets is homeomorphic to $\mathbb{S}^{d-1}$, and the union of their vertex sets is P'. We then verify that the dihedral angles between adjacent facets are convex, which implies that these are the facets of $\mathrm{conv}(P')$, and consequently P' is in convex position.

For every vector $\mathbf{v} \in \{1, \ldots, n-1\}^d$, consider the hyperrectangle $r_{\mathbf{v}} = \prod_{i=1}^{d} [a_{i,v_i}, a_{i,v_i+1}]$. Note that the vertices of each $r_{\mathbf{v}}$ are in $\prod_{i=1}^{d} A_i$, and the hyperrectangles jointly tile the bounding box $\prod_{i=1}^{d} [a_{i,1}, a_{i,n}]$. For every $r_{\mathbf{v}}$, let $F_{\mathbf{v}} = \mathrm{conv}(P \cap r_{\mathbf{v}})$, that is, the convex hull of vertices of $r_{\mathbf{v}}$ that are in P. By construction, every $F_{\mathbf{v}}$ is at most $(d-1)$-dimensional (possibly empty). Let $\mathcal{F}$ be the set of all $(d-1)$-dimensional sets $F_{\mathbf{v}}$; we call them *facets*. By construction, the union of the facets in $\mathcal{F}$, together with the facets in the hyperplanes $x_i = a_{i,1}$ for $i = 1, \ldots, d$, is homeomorphic to the sphere $\mathbb{S}^{d-1}$. Any two adjacent facets in $f_{\mathbf{u}}, f_{\mathbf{w}} \in \mathcal{F}$ lie in two adjacent hyperrectangles, $r_{\mathbf{u}}$ and $r_{\mathbf{w}}$, whose common boundary is $(d-1)$-dimensional, say, in the hyperplane $x_i = a_{i,j}$ for some $j = 2, \ldots, n-1$. The facet $f_{\mathbf{u}}$ (resp., $f_{\mathbf{w}}$) is parallel to the $(d-1)$-simplex spanned by the d vertices of $r_{\mathbf{u}}$ (resp., $r_{\mathbf{w}}$) adjacent to $p_{\mathbf{u}}$ (resp., $p_{\mathbf{w}}$). Since A_i is MD, we have $a_{i,j+1} - a_{i,j} < a_{i,j} - a_{i,j-1}$. These are the i-th extents of $r_{\mathbf{u}}$ and $r_{\mathbf{w}}$; their remaining $d-1$ extents are the same. Consequently, the dihedral angle between $f_{\mathbf{u}}$ and $f_{\mathbf{w}}$ is convex, as required. □

Now Theorem 6 follows from Lemma 7 and Lemma 10.

3 Algorithms

In this section, we describe polynomial-time algorithms for (i) finding convex chains and caps of maximum size; and (ii) approximating the maximum size of a convex polygon; where these structures are supported by a given grid. The main challenge is to ensure that the vertices of the convex polygon (resp., cap or chain) have distinct x- and y-coordinates. The coordinates of a point $p \in X \times Y$ are denoted by $x(p)$ and $y(p)$.

As noted in Section 1, efficient algorithms are available for finding a largest convex polygon or convex cap *contained* in a planar point set. Edelsbrunner and Guibas [11, Thm. 5.1.2] use the dual line arrangement of N points in the plane and dynamic programming to find the maximum size of a convex cap in $\curvearrowright$ in $O(N^2)$ time and $O(N)$ space; the same bounds hold for $\curvearrowleft$, ⤹, and ⤸. A longest convex cap can be also returned in $O(N^2 \log N)$ time and $O(N \log N)$ space. It is straightforward to adapt their algorithm to find the maximum size of a convex cap in $\nearrow$, and report a longest such chain within the same time and space bounds. Since x- and y-coordinates do not repeat in a convex chain, we obtain the following.

Theorem 11. *In a given $n \times n$ grid, the maximum size of a supported convex chain can be computed in $O(n^4)$ time and $O(n^2)$ space; and a supported convex chain of maximum size can be computed in $O(n^4 \log n)$ time and $O(n^2 \log n)$ space.*

We make use of the following general observation:

Observation 12. *If a supported convex polygon P is in a set $\nearrow$, $\searrow$, $\swarrow$, $\nwarrow$, $\curvearrowright$, ⤸, $\curvearrowleft$, or ⤹, then every subsequence of P is in the same set. That is, these classes are hereditary.*

3.1 Convex caps

For computing the maximum size of a convex cap in $\curvearrowright$, we need to be careful to use each y-coordinate at most once. We design an algorithm that finds the maximum size of *two* convex chains that use distinct y-coordinates by dynamic programming. Specifically, for two edges $l = (l_1, l_2)$ and $r = (r_1, r_2)$, we compute the maximum total size $C(l, r)$ of a pair of chains $A \in \nearrow$ and $B \in \searrow$ such that their vertices use distinct y-coordinates and such that the last two vertices of A are l_1 and l_2 (or $A = (l_1)$ if $l_1 = l_2$), and the first two vertices of B are r_1 and r_2 (or $B = (r_1)$ if $r_1 = r_2$). We use the dynamic programming algorithm of [11] to find $L(p_1, p_2)$ (resp., $R(p_1, p_2)$), the size of a largest convex chain P in $\nearrow$ (resp., $\searrow$), ending (resp., starting) with vertices p_1 and p_2, or $P = (p_1)$ if $p_1 = p_2$.

The desired quantity $C(l, r)$ can now be computed by dynamic programming. By Observation 12, we can always safely eliminate the highest vertex of the union of the two chains, to find a smaller subproblem, as this vertex cannot be (implicitly) part of the optimal solution to a subproblem. In particular, if l is a single vertex and it is highest, we can simply use the value of $R(r_1, r_2)$, incrementing it by one for the one vertex of l. Analogously, we handle the case if r is or both l and r are a single vertex. The interesting case is when both chains end in an edge. Here, we observe that we can easily check whether l and r use unique coordinates. If not, then this subproblem is invalid; otherwise, we may find a smaller subproblem by eliminating the highest vertex and comparing all possible subchains that could lead to it.

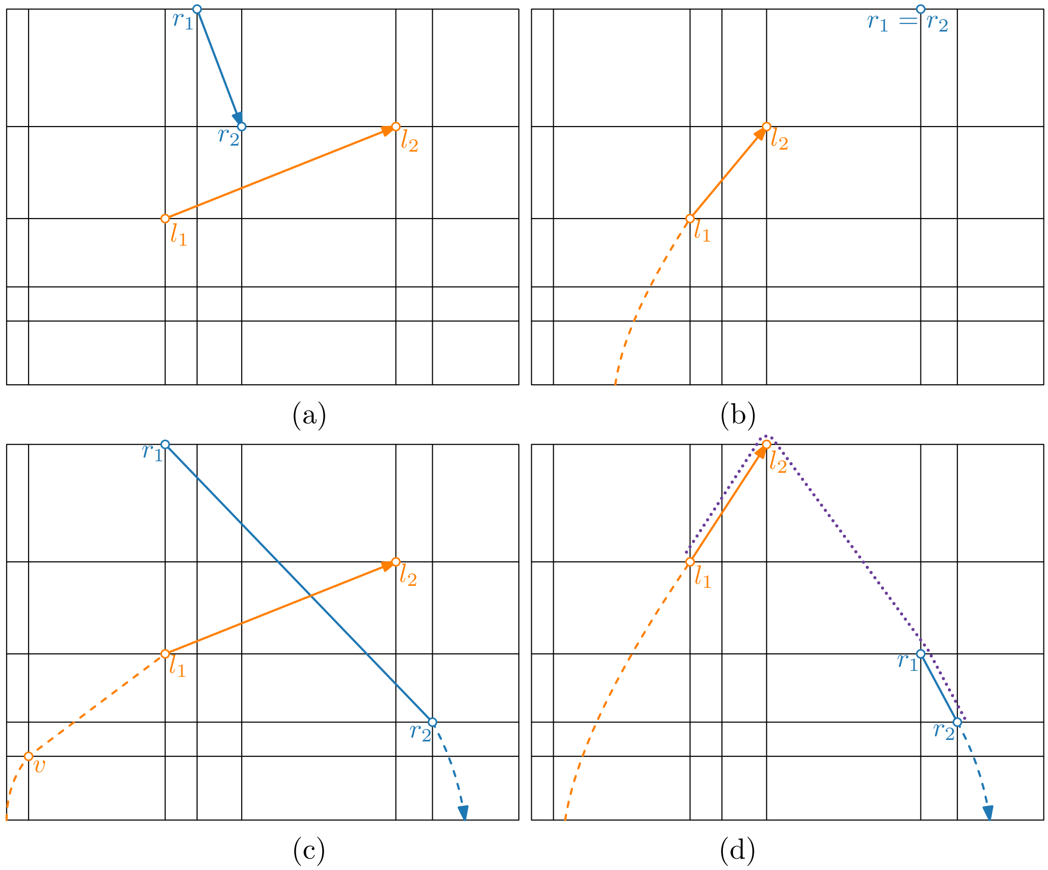

Figure 6: Illustration for the cases of $C(l, r)$. (a) Invalid configuration, as $r_2.y = l_2.y$. (b) r is a single point above $l_2.y$, so we look for the longest chain ending in l. (c) Removing the topmost point (from l in this case), testing all valid possible v to find the longest chain. Note that the left and right chain may not complete to a cap – this is checked separately. (d) We need to test only whether l and r make a cap (purple dotted line) to check whether the $C(l, r)$ entry should be considered.

With the reasoning above, we obtain the recurrence below; see Fig. 6(a–c) for illustration. The first case eliminates invalid edges, and edge pairs that use a y-coordinate more than once. In all remaining cases, we assume that $l \in \nearrow$, $r \in \searrow$, and l and r use distinct y-coordinates.

$$C(l,r) = \begin{cases} -\infty & \text{if } l_1 \neq l_2 \text{ and } l \notin \nearrow, \text{ or} \\ & \quad r_1 \neq r_2 \text{ and } r \notin \searrow, \text{ or} \\ & \quad \{l_1.y, l_2.y\} \cap \{r_1.y, r_2.y\} \neq \emptyset \\ 2 & \text{otherwise, if } l_1 = l_2 \text{ and } r_1 = r_2 \\ L(l_1, l_2) + 1 & \text{otherwise, if } r_1 = r_2 \text{ and } l_2.y < r_1.y \\ R(r_1, r_2) + 1 & \text{otherwise, if } l_1 = l_2 \text{ and } l_2.y > r_1.y \\ \max_{(v,l_1,l_2) \in \nearrow \text{ or } v = l_1} C((v, l_1), r) + 1 & \text{otherwise, if } l_2.y > r_1.y \\ \max_{(r_1,r_2,v) \in \searrow \text{ or } v = r_1} C(l, (r_2, v)) + 1 & \text{otherwise, } l_2.y < r_1.y. \end{cases}$$

Let $E = (X \times Y)^2$ denote the number of pairs (edges) in the grid, from which we take l and r. As $|E| = O(n^4)$, we can compute $C(l, r)$ for all l and r in $O(|E|^2 |X \times Y|) = O(n^{10})$ time and $O(|E|^2) = O(n^8)$ space. With $C(l, r)$, we can easily find the size of a maximum size cap P in $\curvearrowright$, using the observation below, and analogous observations for the special case $k = 1$ and/or $\ell = 1$ (see Fig. 6(d)).

Observation 13. *If $A = (a_1, \ldots, a_k) \in \nearrow$ and $B = (b_1, \ldots, b_\ell) \in \searrow$ with $k \geq 2$, $\ell \geq 2$ and $(a_{k-1}, a_k, b_1, b_2) \in \curvearrowright$ and A and B use distinct y-coordinates, then $(a_1, \ldots, a_k, b_1, \ldots, b_\ell)$ lies in $\curvearrowright$ and has size $k + \ell$.*

Note that the condition $(a_{k-1}, a_k, b_1, b_2) \in \curvearrowright$ implies that the x-coordinates are disjoint.

Theorem 14. *For a given $n \times n$ grid, a supported convex cap of maximum size can be computed in $O(n^{10})$ time and $O(n^8)$ space.*

3.2 Convex n-chains and n-caps

If we are solely interested in deciding whether the grid $X \times Y$, where $|X| = |Y| = n$, supports a convex chain or cap with precisely n vertices, we can improve upon the previous algorithms considerably. Let $X = \{x_1, \ldots, x_n\}$ and $Y = \{y_1, \ldots, y_n\}$ with $x_i < x_{i+1}$ and $y_i < y_{i+1}$. To test whether there is a chain of size n in $\nearrow$, it suffices to test whether the chain $((x_1, y_1), \ldots, (x_n, y_n))$ is in $\nearrow$, in $O(n)$ time.

To test whether there is a supported convex cap of size n in $\curvearrowright$, we adapt the algorithm of Theorem 14. Suppose P is a cap of $\curvearrowright$ of size n, with A_P and B_P its maximal components in $\nearrow$ and in $\searrow$ respectively. Then P uses all coordinates of X, which restricts the types of chains A_P and B_P considerably. In particular $C(l, r)$ can be modified to consider only edges l and r that use consecutive x-coordinates.

For $1 < k < n$ consider the subchains $A_k \in \nearrow$ and $B_k \in \searrow$ of A_P and B_P consisting only of vertices with y-coordinate at most y_k. These chains have length k in total and use all of the coordinates $\{y_1, \ldots, y_k\}$. Let (l_1, l_2) be the last edge of A_k and let (r_1, r_2) be the first edge of B_k. Then the coordinates y_{k-1} and y_k are used by l, or by r, or by l_2 and r_1. Moreover, since the total length of A_k and B_k is k, there are $n - k$ unused X-coordinates between l_2 and r_1, so if $l_2.x = x_i$, then $r_1.x = x_{i+n-k+1}$. So for a fixed value of k, we need only consider $O(|Y|^2|X|)$ inputs for $C(l, r)$. Moreover, the recursive calls in the last two cases need only consider $O(|Y|)$ values of v.

This implies that there are $O(|Y|^3|X|)$ possible inputs to $C(l, r)$ over all k. As an entry now depends on $O(|Y|)$ subproblems and each is evaluated in constant time, the corresponding values can be computed in $O(|Y|^4|X|) = O(n^5)$ time and $O(n^4)$ space. Similarly, we can test whether there is a cap of size n in $\curvearrowright$ within the same time and space bounds.

3.3 Approximations

Although computing the maximum size of a supported convex polygon remains elusive, we can easily devise a constant-factor approximation algorithm by eliminating duplicate coordinates as follows. Compute a maximum size convex polygon P (possibly with duplicate coordinates) in a given $n \times n$ grid in $O(n^6)$ time and $O(n^2)$ space [11, Thm. 5.1.3]. Define a *conflict graph* on the vertices of P, where two vertices are *in conflict* if they share an x- or y-coordinate. Since each conflict corresponds to a horizontal or vertical line, the conflict graph has maximum degree at most 2 and contains no odd cycles, hence it is bipartite. One of the two sets in the bipartition contains at least half of the vertices of P without duplicate coordinates, and so it determines a supported convex polygon.

Since P has $O(n)$ vertices, the conflict graph can be computed in $O(n)$ time. Overall, we obtain a $\frac{1}{2}$-approximation for the maximum supported convex polygon in $O(n^6)$ time and $O(n^2)$ space. The same strategy provides an $\frac{1}{2}$-approximation for the maximum supported polygon in $\curvearrowright$, $\curvearrowleft$, $\circlearrowleft$, and $\circlearrowright$ in $O(n^4)$ time and $O(n^2)$ space using [11, Thm. 5.1.2].

4 The Maximum Number of Convex Polygons

Let $F(n)$ be the maximum number of convex polygons that can be present in an $n \times n$ grid, with no restriction on the number of times each coordinate is used. Let $G(n)$ be this number where all $2n$ grid lines are used (i.e., each grid line contains at least one vertex of the polygon). Let $\bar{F}(n)$ and $\bar{G}(n)$ be the corresponding numbers where each grid line is used at most once (so $\bar{F}(n)$ counts the maximum number of supported convex polygons). By definition, we have $F(n) \geq G(n) \geq \bar{G}(n)$ and $F(n) \geq \bar{F}(n) \geq \bar{G}(n)$ for all $n \geq 2$. We prove the following theorem, in which the $\Theta^*(.)$ notation hides polynomial factors in n.

Theorem 15. *The following bounds hold:*

$$F(n) = \Theta^*(16^n), \qquad \bar{F}(n) = \Theta^*(9^n), \qquad G(n) = \Theta^*(9^n), \qquad \bar{G}(n) = \Theta^*(4^n).$$

4.1 Upper bounds

We first prove that $F(n) = O(n \cdot 16^n)$ by encoding each convex polygon in a unique way, so that the total number of convex polygons is bounded by the total number of encodings. Recall that a convex polygon P can be decomposed into four convex chains $\nearrow_P, \nwarrow_P, \swarrow_P, \searrow_P$, with only extreme vertices of P appearing in multiple chains. Let $\circlearrowleft_P = \nearrow_P \cup \nwarrow_P$ and $\curvearrowleft_P = \nwarrow_P \cup \swarrow_P$. To encode P, we assign the following number to each of the $2n$ grid lines ℓ (see Fig. 7 for an example): 0 if ℓ is not incident on any vertex of P, 3 if ℓ is incident on multiple vertices of P, 1 if ℓ is incident on one vertex of P and that vertex lies on $\circlearrowleft_P$ if ℓ is horizontal, or on $\curvearrowleft_P$ if ℓ is vertical, and 2 otherwise.

In addition, we record the index of the horizontal line containing the leftmost vertex of P (pick the topmost of these if there are multiple leftmost points).

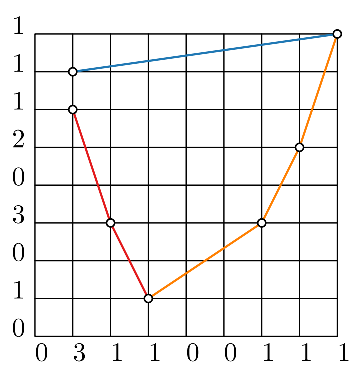

Figure 7: Encoding the grid lines.

Since each of the $2n$ grid lines is assigned one of 4 possible values, and there are n horizontal lines, the total number of encodings is $O(n \cdot 4^{2n}) = O(n \cdot 16^n)$. All that is left to show is that each encoding corresponds to at most one convex polygon.

First, observe that if P is a convex chain, say in $\nearrow$, then the set of grid lines containing a vertex of P uniquely defines P: since both coordinates change monotonically, the i-th vertex of P must be the intersection of the i-th horizontal and vertical lines. So all we need to do to reconstruct P is to identify the set of lines that make up each convex chain.

Since we know the location of the (topmost) leftmost vertex of P, we know where $\nearrow_P$ starts. Every horizontal line above this point labelled with a 1 or 3 must contain a vertex of $\nearrow_P$; let k be the number of such lines. Since the x-coordinates are monotonic as well, $\nearrow_P$ ends at the k-th vertical line labelled with a 2 or 3. The next chain, $\searrow_P$, starts either at the end of $\nearrow_P$, if the horizontal line is labelled with a 1, or at the intersection of this horizontal line with the next vertical line labelled with a 2 or 3, if this horizontal line is labelled with a 3. We can find the rest of the chains in a similar way. Thus, $F(n) = O(n \cdot 16^n)$.

The upper bounds for $\bar{F}(n)$, $G(n)$, and $\bar{G}(n)$ are analogous, except that certain labels are excluded. For the number of supported convex polygons $\bar{F}(n)$, each grid line is used at most once, which means that the label 3 cannot be used. Thus, $\bar{F}(n) = O(n \cdot 3^{2n}) = O(n \cdot 9^n)$. Similarly, for $G(n)$, all grid lines contain at least one vertex of the polygon, so the label 0 cannot be used. Therefore $G(n) = O(n \cdot 3^{2n}) = O(n \cdot 9^n)$. Finally, for $\bar{G}(n)$, every grid line contains exactly one vertex of the polygon, so neither 0 nor 3 can be used as labels. This gives $\bar{G}(n) = O(n \cdot 2^{2n}) = O(n \cdot 4^n)$ possibilities.

4.2 Lower bounds

Assume that $n = 2m + 3$, where $m \in \mathbb{N}$ satisfies suitable divisibility conditions, as needed. All four lower bounds use the same grid, constructed as follows (see Fig. 8).

$$\begin{aligned} X &= \{1, \ldots, n-1\} \qquad & Y^- &= \{y_1, \ldots, y_{m+2}\}, \text{ where } y_i = n^i \\ Y &= Y^- \cup Y^+ \qquad & Y^+ &= \{z_1, \ldots, z_{m+2}\}, \text{ where } z_i = 2 \cdot y_{m+2} - y_i \end{aligned}$$

Note that this results in an $(n-1) \times (n-1)$ grid, since $y_{m+2} = z_{m+2}$. To obtain an $n \times n$ grid, we duplicate the median grid lines in both directions and offset them by a sufficiently small distance $\varepsilon > 0$. The resulting grid has the property that any three points p, q, r in the lower half $X \times Y^-$ with $x(p) < x(q) < x(r)$ and $y(p) < y(q) < y(r)$ make a left turn at q. To see this, suppose that $y(p) = n^i$, $y(q) = n^j$, and $y(r) = n^k$, for some $1 \le i < j < k \le n$. Then the slope of pq is strictly smaller than the slope of qr, since

$$\mathrm{slope}(qr) = \frac{n^k - n^j}{x(r) - x(q)} \ge \frac{n^{j+1} - n^j}{n-1} = n^j > \frac{n^j - n^i}{1} \ge \frac{n^j - n^i}{x(q) - x(p)} = \mathrm{slope}(pq).$$

Thus, any sequence of points with increasing x- and y-coordinates in the lower half is in $\swarrow$. By symmetry, such a sequence in the upper half $X \times Y^+$ is in $\nearrow$. Analogously, points with increasing x-coordinates and decreasing y-coordinates are in $\nwarrow$ if they are in the lower half and $\searrow$ if they are in the upper half.

We first derive lower bounds on $\bar{G}(n)$ and $G(n)$ by constructing a large set of convex polygons that use each grid line at least once. Then we use these bounds to derive the bounds on $\bar{F}(n)$ and $F(n)$. The polygons we construct all share the same four extreme vertices, which lie on the intersections of the grid boundary with the duplicated median grid lines. Specifically, the leftmost and rightmost vertices are the intersections of the duplicate horizontal medians with the left and right boundary, and the highest and lowers vertices are the intersections of the duplicate vertical

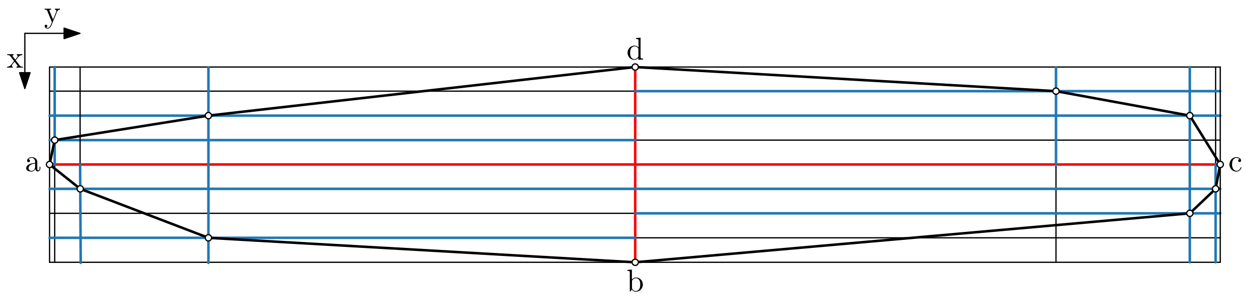

Figure 8: The $n \times n$ grid defined in Section 4.2, with $n = 9 = 2m + 3$ for $m = 3$, before doubling the median lines. The segments (parts of grid lines) incident to vertices are drawn in blue.

medians with the top and bottom boundary. Since each of these median lines now contain a vertex, we can choose additional vertices from the remaining $2m$ grid lines in each direction.

To construct each polygon, select $m/2$ vertical grid lines left of the median to participate in the bottom chain, and do the same right of the median. Likewise, select $m/2$ horizontal grid lines above and below the median, respectively, to participate in the left chain. The remaining grid lines participate in the other chain (top or right). This results in a polygon with $m/2$ vertices in each quadrant of the grid (excluding the extreme vertices). The convexity follows from our earlier observations. The total number of such polygons is

$$\binom{m}{\frac{m}{2}}^4 = \Theta\left(\left(m^{-\frac{1}{2}}2^m\right)^4\right) = \Theta(m^{-2}2^{4m}) = \Theta(n^{-2}2^{2n}) = \Theta(n^{-2}4^n) = \Theta^*(4^n).$$

The first step uses the following estimate, which can be derived from Stirling's formula for the factorial [10]. Let $0 < \alpha < 1$, then

$$\binom{n}{\alpha n} = \Theta(n^{-\frac{1}{2}}2^{H(\alpha)n}), \text{ where } H(\alpha) = -\alpha \log_2 \alpha - (1-\alpha)\log_2(1-\alpha).$$

For the lower bound on $G(n)$, the only difference is that we now allow grid lines to contain vertices in two chains. We obtain a maximum when we divide the grid lines evenly between the three groups (bottom chain, top chain, both chains). Thus, we select $m/3$ vertical grid lines left of the median to participate in the bottom chain, another $m/3$ to participate in the top chain and the remaining $m/3$ participate in both. We repeat this selection to the right of the median and on both sides of the median horizontal line. As before, this results in a convex polygon with the same number of vertices in each quadrant of the grid—exactly $2m/3$ this time. The number of such polygons is

$$\begin{aligned}\binom{m}{\frac{m}{3}}^4 \binom{\frac{2m}{3}}{\frac{m}{3}}^4 &= \Theta\left(\left(m^{-\frac{1}{2}}2^{H(\frac{1}{3})m} \cdot m^{-\frac{1}{2}}2^{H(\frac{1}{2})\frac{2m}{3}}\right)^4\right) \\ &= \Theta\left(m^{-4}2^{4m(\log_2 3 - \frac{2}{3} + \frac{2}{3})}\right) = \Theta\left(n^{-4}2^{2n\log_2 3}\right) = \Theta\left(n^{-4}9^n\right) = \Theta^*\left(9^n\right).\end{aligned}$$

To translate these bounds to bounds on $\bar{F}(n)$ and $F(n)$, where some grid lines may contain no vertices of the polygon, we observe that the arguments for the bounds above also work for a subgrid of $X \times Y$, provided that the subgrid includes the boundary and medians and has the same number of grid lines on each side of the median in both directions. For $\bar{F}(n)$, we select $2m/3$ grid lines on each side of each median (balancing no vertices with the two different chains) to make up our subgrid and plug in the bound on $\bar{G}(n)$, which yields

$$\binom{m}{\frac{2m}{3}}^4 \Omega^*(4^{\frac{2n}{3}}) = \Omega^*(2^{2n(H(\frac{2}{3})+\frac{2}{3})}) = \Omega^*(2^{2n\log_2 3}) = \Omega^*(9^n).$$

Finally, for the bound on $F(n)$, we select $3m/4$ grid lines on each side of each median (balancing no vertices with the three different options for a grid line in the proof of $G(n)$), to make up our subgrid and plug in the bound on $G(n)$, giving

$$\binom{m}{\frac{3m}{4}}^4 \Omega^*(9^{\frac{3n}{4}}) = \Omega^*(2^{2n(H(\frac{3}{4})+\frac{3}{4}\log_2 3)}) = \Omega^*(2^{4n}) = \Omega^*(16^n).$$

4.3 The maximum number of weakly convex polygons

Let $W(n)$ denote the maximum number of weakly convex polygons that contained in an $n \times n$ grid. A polygon P in $\mathbb{R}^2$ is *weakly convex* if all of its internal angles are less than or equal to π. Here we identify each polygon by its set of boundary vertices, so different polygons may have identical convex hulls. Since $W(n) \geq F(n)$, we have $W(n) = \Omega^*(16^n)$. In fact, a slightly better lower bound trivially holds even for the $n \times n$ section of the integer lattice $Z_0 = [n] \times [n]$. Consider all polygons whose vertices are the four extreme vertices of Z_0, and an arbitrary subset of the remaining $4n-8$ grid points in $\partial\text{conv}(Z_0)$. There are $2^{4n-8} = \Omega(16^n)$ such subsets, and the lower bound $W(n) = \Omega(16^n)$ follows.

To show that $W(n) = O^*(16^n)$, we modify the previous encoding used to show that $F(n) = O^*(16^n)$. While the four grid lines along the boundary of the bounding box of P can be incident to arbitrarily many vertices, we still use at most two vertices for each such line, namely at most two extreme vertices. For each weakly convex polygon P, record the at most 8 extreme vertices incident to ∂B together with a vector (sequence) of length $2n$: n elements corresponding to the horizontal lines (from the lowest to the highest), and n elements corresponding to the vertical lines (from left to right). As previously, we encode each grid line by an element of $\{0, 1, 2, 3\}$, where 3 stands for a line incident to at least two vertices. By (weak) convexity, a grid line can be incident to 3 or more vertices of P only if it is one of the four lines along the bounding box of P.

From the recorded information, we can reconstruct a weakly convex polygon in the $n \times n$ grid. Consequently, the number of convex polygons in the grid is bounded from above by the number of encodings, namely $W(n) = O(n^8 \cdot 4^{2n}) = O(n^8 \cdot 16^n) = O^*(16^n)$. We summarize the bounds we have obtained in the following.

Theorem 16. *Let $W(n)$ denote the maximum number of weakly convex polygons that can be present in an $n \times n$ grid. Then $W(n) = \Omega(16^n)$ and $W(n) = O^*(16^n)$.*

5 Conclusions

We studied combinatorial properties of convex polygons (resp., polytopes) in Cartesian products in d-space. Similar questions for point sets in general position or for lattice polygons (resp., polytopes) have been previously considered. We showed that every $n \times \ldots \times n$ Cartesian product in $\mathbb{R}^d$ contains $\Omega(\log^{d-1} n)$ points in convex position, and this bound is the best possible. Our upper bound matches previous bounds [20, 31] for points in general position, which are conjectured to be tight. Our lower bound, however, does not yield any improvement for points in general position. In contrast, an $n \times \ldots \times n$ section of the integer lattice $\mathbb{Z}^d$ contains significantly more, namely $\Theta(n^{d(d-1)/(d+1)})$, points in convex position [1].

The *maximum number* of convex polygons in an $n \times n$ Cartesian product is $F(n) = \Theta^*(16^n)$. This bound is tight up to polynomial factors, and is significantly larger than the corresponding bound in an $n \times n$ section of the integer lattice [3]. In contrast, n^2 points in convex (hence general) position trivially determine $2^{n^2} - 1$ convex polygons. Erdős [14] proved that the *minimum number* of convex polygons determined by n points in general position is $\exp(\Theta(\log^2 n))$. Determining (or estimating) the minimum number of convex polygons in an $n \times n$ Cartesian product and in higher dimensions remain as open problems.

Our motivating problem was the reconstruction of a convex polygon from the x- and y-projections of its vertices. We presented a $\frac{1}{2}$-approximation for computing the maximal size of a convex polygon supported by a grid $X \times Y$. Finding an efficient algorithm for the original problem, or proving its hardness, remains open. As our dynamic program does not directly extend to $d \geq 3$, approximation algorithms in higher dimensions are also of interest.

References

[1] George E. Andrews. A lower bound for the volume of strictly convex bodies with many boundary lattice points. *Transactions of the American Mathematical Society*, 106:270–279, 1963. doi:10.1090/S0002-9947-1963-0143105-7.

[2] Imre Bárány. Random points and lattice points in convex bodies. *Bulletin of the American Mathematical Society*, 45:339–365, 2008. doi:10.1090/S0273-0979-08-01210-X.

[3] Imre Bárány and János Pach. On the number of convex lattice polygons. *Combinatorics, Probability and Computing*, 1(4):295–302, 1992. doi:10.1017/S0963548300000341.

[4] Imre Bárány and Anatoly Moiseevich Vershik. On the number of convex lattice polytopes. *Geometric and Functional Analysis*, 2(4):381–393, 1992. doi:10.1007/BF01896660.

[5] Alexander Barvinok. Lattice points and lattice polytopes. In Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 7, pages 185–210. CRC Press, Boca Raton, FL, 3rd edition, 2017.

[6] Boris Bukh and Jiří Matoušek. Erdős–Szekeres-type statements: Ramsey function and decidability in dimension 1. Manuscript, 2012. URL: https://arxiv.org/abs/1207.0705.

[7] Jean Cardinal, Csaba D. Tóth, and David R. Wood. A note on independent hyperplanes and general position subsets in d-space. *Journal of Geometry*, 108(1):33–43, 2017. doi:10.1007/s00022-016-0323-5.

[8] Václav Chvátal and G. T. Klincsek. Finding largest convex subsets. *Congressus Numerantium*, 29:453–460, 1980.

[9] David Conlon, Jacob Fox, János Pach, Benny Sudakov, and Andrew Suk. Ramsey-type results for semi-algebraic relations. *Transactions of the American Mathematical Society*, 366:5043–5065, 2014. doi:10.1090/S0002-9947-2014-06179-5.

[10] Adrian Dumitrescu, André Schulz, Adam Sheffer, and Csaba D. Tóth. Bounds on the maximum multiplicity of some common geometric graphs. *SIAM Journal on Discrete Mathematics*, 27(2):802–826, 2013. doi:10.1137/110849407.

[11] Herbert Edelsbrunner and Leonidas J. Guibas. Topologically sweeping an arrangement. *Journal of Computer and System Sciences*, 38(1):165–194, 1989. doi:10.1016/0022-0000(89)90038-X.

[12] György Elekes, Melvyn B. Nathanson, and Imre Z. Ruzsa. Convexity and sumsets. *Journal of Number Theory*, 83(2):194–201, 2000. doi:10.1006/jnth.1999.2386.

[13] Paul Erdős. Appendix, in Klaus F. Roth, On a problem of Heilbronn. *Journal of the London Mathematical Society*, 26:198–204, 1951. doi:10.1112/jlms/s1-26.3.198.

[14] Paul Erdős. Some more problems on elementary geometry. *Gazette of the Australian Mathematical Society*, 5(2):52–54, 1978.

[15] Paul Erdős and György Szekeres. On some extremum problems in elementary geometry. *Annales Universitatis Scientiarium Budapestinensis de Rolando Eötvös Nominatae Sectio Mathematica*, 3-4:53–62, 1960/1961.

[16] Panos Giannopoulos, Christian Knauer, and Daniel Werner. On the computational complexity of Erdős-Szekeres and related problems in $\mathbb{R}^3$. In *Proc. 21st European Symposium on Algorithms*, volume 8125 of *LNCS*, pages 541–552. Springer, 2013. doi:10.1007/978-3-642-40450-4_46.

[17] Ronald L. Graham. Euclidean Ramsey theory. In Jacob E. Goodman, Joseph O'Rourke, and Csaba D. Tóth, editors, *Handbook of Discrete and Computational Geometry*, chapter 11, pages 281–297. CRC Press, Boca Raton, FL, 3rd edition, 2017.

[18] Richard K. Guy and Patrick A. Kelly. The no-three-in-line-problem. *Canadian Mathematical Bulletin*, 11:527–531, 1968.

[19] Norbert Hegyvári. On consecutive sums in sequences. *Acta Mathematica Hungarica*, 48(1-2):193–200, 1986. doi:10.1007/BF01949064.

[20] Gyula Károlyi and Pavel Valtr. Configurations in d-space without large subsets in convex position. *Discrete & Computational Geometry*, 30(2):277–286, 2003. doi:10.1007/s00454-003-0009-4.

[21] Joseph S.B. Mitchell, Günter Rote, Gopalakrishnan Sundaram, and Gerhard Woeginger. Counting convex polygons in planar point sets. *Information Processing Letters*, 56(1):45–49, 1995. doi:10.1016/0020-0190(95)00130-5.

[22] Michael S. Payne and David R. Wood. On the general position subset selection problem. *SIAM Journal on Discrete Mathematics*, 27(4):1727–1733, 2013. doi:10.1137/120897493.

[23] Attila Pór and David R. Wood. No-three-in-line-in-3D. *Algorithmica*, 47(4):481–488, 2007. doi:10.1007/s00453-006-0158-9.

[24] Orit E. Raz, Micha Sharir, and Ilya D. Shkredov. On the number of unit-area triangles spanned by convex grids in the plane. *Computational Geometry: Theory and Applications*, 62:25–33, 2017. doi:10.1016/j.comgeo.2016.12.002.

[25] Orit E. Raz, Micha Sharir, and József Solymosi. Polynomials vanishing on grids: The Elekes–Rónyai problem revisited. *American Journal of Mathematics*, 138(4):1029–1065, 2016. `doi:10.1353/ajm.2016.0033`.

[26] Orit E. Raz, Micha Sharir, and Frank De Zeeuw. Polynomials vanishing on Cartesian products: The Elekes–Szabó theorem revisited. *Duke Mathematical Journal*, 165(18):3517–3566, 2016. `doi:doi:10.1215/00127094-3674103`.

[27] David Alan Rosenthal. *The classification of the order indiscernibles of real closed fields and other theories*. PhD thesis, University of Wisconsin–Madison, 1981.

[28] Tomasz Schoen and Ilya D. Shkredov. On sumsets of convex sets. *Combinatorics, Probability and Computing*, 20(5):793–798, 2011. `doi:10.1017/S0963548311000277`.

[29] Ryan Schwartz, József Solymosi, and Frank de Zeeuw. Extensions of a result of Elekes and Rónyai. *Journal of Combinatorial Theory, Series A*, 120(7):1695–1713, 2013. `doi:10.1016/j.jcta.2013.06.004`.

[30] Andrew Suk. On the Erdős-Szekeres convex polygon problem. *Journal of the American Mathematical Society*, 30:1047–1053, 2017. `doi:10.1090/jams/869`.

[31] Pavel Valtr. Sets in $\mathbb{R}^d$ with no large empty convex subsets. *Discrete Mathematics*, 108(1–3):115–124, 1992. `doi:10.1016/0012-365X(92)90665-3`.

[32] Eric W. Weisstein. No-three-in-a-line-problem. online, 2005. URL: `http://mathworld.wolfram.com/No-Three-in-a-Line-Problem.html`.