

Convex polygons in Cartesian products

Citation for published version (APA):

De Carufel, J-L., Dumitrescu, A., Meulemans, W., Ophelders, T., Pennarun, C., Tóth, C. D., & Verdonshot, S. (2021). Convex polygons in Cartesian products. *Journal of Computational Geometry*, 11(2), 205-233. <https://doi.org/10.20382/jocg.v11i2a9>

Document license:

CC BY

DOI:

[10.20382/jocg.v11i2a9](https://doi.org/10.20382/jocg.v11i2a9)

Document status and date:

Published: 01/01/2021

Document Version:

Publisher's PDF, also known as Version of Record (includes final page, issue and volume numbers)

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.tue.nl/taverne

Take down policy

If you believe that this document breaches copyright please contact us at:

openaccess@tue.nl

providing details and we will investigate your claim.

CONVEX POLYGONS IN CARTESIAN PRODUCTS*

Jean-Lou De Carufel[†], Adrian Dumitrescu[‡], Wouter Meulemans[§], Tim Ophelders[§],
Claire Pennarun[¶], Csaba D. Tóth^{||}, and Sander Verdonschot^{**}

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 $\Omega(\log n)$ points in convex position 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 x - or y -monotone 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 a convex polygon with such unique coordinates can be efficiently approximated up to a factor of 2. Finally, we present exponential bounds on the maximum number of point sets in convex position in such 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 $n = 5$ (see 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 contains (resp., supports) a convex k -gon? What is the largest k such that *every* $n \times n$ grid supports a convex k -gon?

*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.* This work has been supported by NSERC.

[‡]adriandumitrescu.org, ad.dumitrescu@gmail.com

[§]*Department of Mathematics and Computer Science, TU Eindhoven, The Netherlands, {w.meulemans|t.a.e.ophelders@tue.nl}@tue.nl*

[¶]*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*

How many convex polygons does an $n \times n$ grid contain, or support? 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 contains (resp., supports) a convex polygon with $(1 - o(1)) \log n$ vertices¹; this bound is tight up to a constant factor: there are $n \times n$ grids that do not contain 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 $\prod_{i=1}^d X_i$, where $|X_i| = n$ 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 [19], 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 [17]. The current best upper bound, due to Suk [32], 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 [22] and Valtr [33] 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 [33]. 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 [23]. 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 [18].

There has been significant interest in *counting* the number of convex polygons in various point sets. Answering a question of Hammer, Erdős [16] proved that every set of

¹All logarithms in this paper are of base 2.

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 $\exp(O(n^{1/3}))$. Bárány and Vershik [4] generalized this bound to d -dimensions and showed that there are $\exp(O(n^{(d-1)/(d+1)}))$ convex polytopes in an $n \times \dots \times n$ 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 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 [26, 27, 28, 31], while the analogous statements for point 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 polyhedron; and it is in *general position* if no $d + 1$ points lie on a hyperplane. The *1-skeleton* of a polyhedron is the graph formed by its vertices and edges. In \mathbb{R}^d , we say that the x_d -axis is *vertical*, hyperplanes orthogonal to the x_d -axis are *horizontal*, and the above-below relationship is understood with respect to x_d coordinates.

We consider special types of convex polygons in the plane. Let P be a convex polygon with vertices $((x_1, y_1), \dots, (x_k, y_k))$ in clockwise order. We say that P is a *convex cap* if the x - or y -coordinates are strictly monotone, 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 $\{\curvearrowright, \curvearrowleft, \smile, \frown\}$. We have
 - $P \in \curvearrowright$ if and only if $(x_i)_{i=1}^k$ strictly increases;
 - $P \in \curvearrowleft$ if and only if $(y_i)_{i=1}^k$ strictly increases;
 - $P \in \smile$ if and only if $(x_i)_{i=1}^k$ strictly decreases;
 - $P \in \frown$ if and only if $(y_i)_{i=1}^k$ strictly decreases;
- *convex chains* come in four types $\{\nearrow, \searrow, \swarrow, \nwarrow\}$. We have
 - $\nearrow = \curvearrowleft \cap \curvearrowright$, $\searrow = \curvearrowright \cap \curvearrowleft$, $\swarrow = \smile \cap \curvearrowleft$, $\nwarrow = \frown \cap \smile$.

Initial observations. It is easy to see that for $n \in \{3, 4\}$, 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. Figure 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. Refer to Figure 2 for illustration. Let x_1, \dots, x_6 and y_1, \dots, y_6 denote the vertices of X and Y in increasing order. Consider first the 4×4 subgrid obtained by omitting the

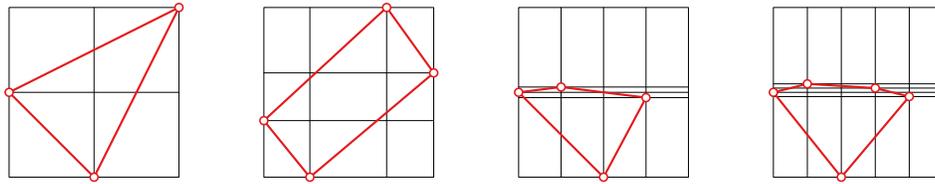


Figure 1: Maximum-size supported convex polygons of respective sizes 3, 4, 4, and 5 in $n \times n$ grids, where n is between 3 and 6.

extremal values, that is, $\{x_2, \dots, x_5\} \times \{y_2, \dots, y_5\}$. As the 2×2 grid $\{x_3, x_4\} \times \{y_3, y_4\}$ consists of four corners of a rectangle, at least two of these points are not collinear with (x_2, y_2) and (x_5, y_5) ; let (x', y') denote such a point and $x'' \in \{x_3, x_4\}$ and $y'' \in \{y_3, y_4\}$ denote the unused coordinates. Now, $P' = \{(x_2, y_2), (x', y'), (x_5, y_5)\}$ is a supported convex chain of the subgrid and thus of $X \times Y$. By construction, P' is either in \nearrow or \swarrow ; without loss of generality, assume the former.

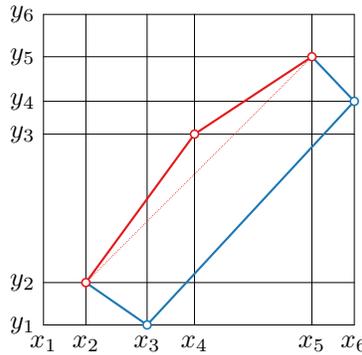


Figure 2: A 6×6 grid, in which the convex chain P' (red) has size 3, and it is extended to a convex pentagon P (red and blue) supported by the grid.

We now extend P' to a convex polygon P of size 5 by appending two vertices: (x_6, y'') and (x'', y_1) . By construction, all five vertices of P have distinct coordinates. Moreover, the points in P are in convex position, as each is an extreme point in P : (x_2, y_2) has the lowest x -coordinate; (x_5, y_5) has the highest y -coordinate; (x_6, y'') has the maximum x -coordinate; and (x'', y_1) has the minimum y -coordinate. Finally, (x', y') is an extreme point in P' , hence P' has a supporting halfplane h whose boundary contains (x', y') . Since $P' \subset h$, and we have $x_2 < x'' < x_4$ and $y_2 < y'' < y_4$, the halfplane h contains (x_6, y'') and (x'', y_1) , as well, (x', y') is an extreme point in P . Thus, we conclude that P is a supported convex polygon of size 5 on $X \times Y$. \square

Note that, although every pentagon supported by a grid lies on a 5×5 subgrid, the proof above does not work on a 5×5 grid. If we start with a 5×5 grid, and choose P' from an off-center 4×4 subgrid, then we might not be able to add two points from the remaining 2×2 subgrid. Specifically, the coordinates of the $4 \times t$ grid may force P' to be of a specific type: the example in Figure 2 prevents P' from being in \swarrow as all four central points

are above the dotted line. Starting with a 6×6 grid, we know there are two extra points on the outer boundary that can be added to P' , regardless of the type of the chain P' .

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–Székeres number cannot be used directly, since they 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 [15]; the maximum number of points in general position is conjectured to be $\frac{\pi}{\sqrt{3}}n(1+o(1))$ [20, 34]. 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 [25], but no similar statements are known in higher dimensions. We use a recent result from incidence geometry.

Lemma 2 (Payne and Wood [24]). *Every set of N points in the plane with at most ℓ collinear, where $\ell \leq O(\sqrt{N})$, 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 [32], this set of $\Omega(n/\sqrt{\log n})$ points contains a convex polygon with at least $(1 - o(1))(\log(n/\sqrt{\log n})) = (1 - o(1)) \log n$ vertices, which have distinct x - and y -coordinates by construction, as required. \square

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 3.

Lemma 4. *For every $n \in \mathbb{N}$, there exists an $n \times n$ grid that contains at most $\lceil \log n \rceil + 1$ points in a convex chain; consequently, 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, without loss of generality, that $n = 2^k$, and let $X = \{0, \dots, n - 1\}$. 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 $Y = \{\sum_{i=0}^{k-1} m_i (2n)^i : 0 \leq m \leq n - 1\}$ (see Figure 3). Both X and Y are symmetric: $X =$

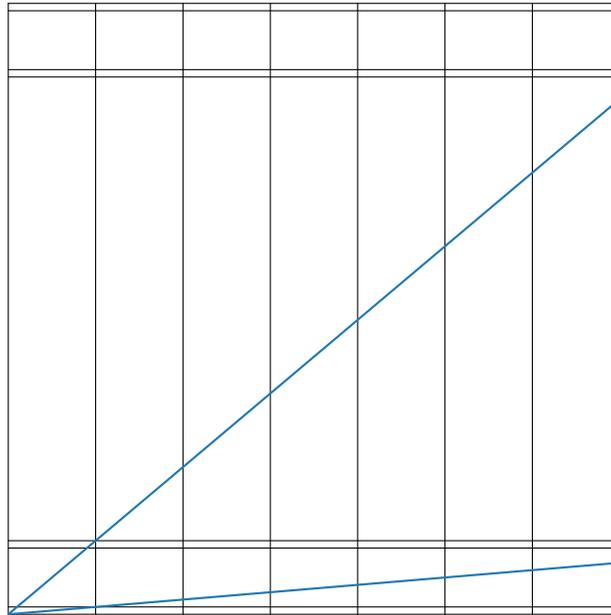


Figure 3: An 8×8 grid without convex chains of size greater than $4 = \log 8 + 1$, where $X = \{0, 1, \dots, 7\}$ and $Y = \{0, 1, 16, 17, 256, 257, 272, 273\}$. Two lines through pairs of grid points are drawn in blue.

$\{\max(X) - x : x \in X\}$ and $Y = \{\max(Y) - y : y \in Y\}$. Thus, it suffices to show that no convex chain $P \in \mathcal{C}$ 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' is $\text{slope}(p, p') = \sum_{i=0}^{k-1} (m'_i - m_i) (2n)^i / (x' - x)$. 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:

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

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

Hence, $\text{slope}(p, p') \in I_j$, where $I_j = (2 \cdot (2n)^{j-1}, 2 \cdot (2n)^j)$. Let us define the family of intervals I_0, I_1, \dots, I_{k-1} analogously, and note that these intervals are pairwise disjoint. Suppose that some convex chain $P \in \mathcal{C}$ 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 $\text{slope}(p, p')$ and $\text{slope}(p', p'')$ are in the same interval, say I_j . Assume that $y = \sum_{i=0}^{k-1} m_i (2n)^i$, $y' = \sum_{i=0}^{k-1} m'_i (2n)^i$, and $y'' = \sum_{i=0}^{k-1} m''_i (2n)^i$. 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 \curvearrowright of size greater than $k + 1$. Analogously, every convex chain in \curvearrowleft , \curvearrowright , or \curvearrowleft 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–Székeres numbers in d -dimensions for points in general position. Our 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 $X_i \subseteq \mathbb{R}$ for $i = 1, \dots, d$, such that the Cartesian product $X = \prod_{i=1}^d X_i$ contains at most $O(\log^{d-1} n)$ points in convex position.*

Before proving Lemma 5, we introduce some additional terminology and give a brief overview of key ideas. Let d be a positive integer. Let \mathbf{e}_d be the standard basis vector parallel to the x_d -axis in \mathbb{R}^d . For a point $a \in \mathbb{R}^d$, let a^{proj} denote the orthogonal projection of a to the horizontal hyperplane $x_d = 0$. Let $P \subset \mathbb{R}^d$ be a finite set in convex position; refer to Figure 4. The point set P is *full-dimensional* if no hyperplane contains P . The orthogonal projection of $\text{conv}(P)$ to the hyperplane $x_d = 0$ is a convex polytope in \mathbb{R}^{d-1} that we denote by $\text{conv}(P)^{\text{proj}}$. The *silhouette* of P is the subset of points in P whose orthogonal projection to $x_d = 0$ lies on the boundary of $\text{conv}(P)^{\text{proj}}$. Note that for every point $p \in P$ in the silhouette of P , the projection p^{proj} is either a vertex of $\text{conv}(P)^{\text{proj}}$ or lies in the relative interior of a face of $\text{conv}(P)^{\text{proj}}$. Since no three points in P are collinear, at most two points in P are projected to any point in the hyperplane $x_d = 0$. 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). For an integer $k \in \{1, \dots, d\}$, a k -dimensional flat (for short, *k-flat*) is *axis-aligned* if it is parallel to k coordinate axes; similarly, a k -dimensional polytope (*k-polytope*) is *axis-aligned* if its affine hull is an axis-aligned k -flat. A k -flat or k -polytope is *vertical* if its affine hull is parallel to the x_d -axis; or equivalently if it contains a nondegenerate line segment parallel to the x_d -axis.

The proof of Lemma 5 is constructive. For integers $0 \leq i \leq j$, we recursively define a $2^j \times \dots \times 2^j \times 2^i$ grid in \mathbb{R}^d , and denote it by $S_d(i, j)$. For $i \geq 1$, the grid $S_d(i, j)$ is the disjoint union of two translated copies of $S_d(i-1, j)$, one above the other. Then we prove that every d -dimensional subset $P \subset S_d(i, j)$ in convex position contains at most $2^{d(d+1)} \cdot i \cdot j^{d-2}$ upper (resp., lower) vertices. For $i = j = \log n$, this implies that $|P| \leq 2 \cdot 2^{d(d+1)} \cdot j^{d-1} = \Theta(\log^{d-1} n)$. The number of upper vertices in P can be bounded as follows (the case of lower vertices is analogous). We partition the upper vertices in P into three subsets: (1) Upper vertices $p \in P$ such that p^{proj} is a vertex of $\text{conv}(P)^{\text{proj}}$ in $(d-1)$ -dimensions. As p^{proj} may be an upper or a lower vertex in $S_{d-1}(j, j)$, by induction, there are at most $2 \cdot 2^{(d-1)d} \cdot j^{d-2}$ such points. (2) Upper vertices $p \in P$ such that p^{proj} lies in the interior of $\text{conv}(P)^{\text{proj}}$. We prove that one of the two translated copies of $S_d(i-1, j)$ contains these points as upper vertices, and by induction there are at most $2^{d(d+1)} \cdot (i-1) \cdot j^{d-2}$ such points. (3) Upper vertices $p \in P$ such that p^{proj} lies on the boundary of $\text{conv}(P)^{\text{proj}}$, but in the relative interior of some face of $\text{conv}(P)^{\text{proj}}$; see Figure 4 for examples. We show that only axis-aligned faces of $\text{conv}(P)^{\text{proj}}$

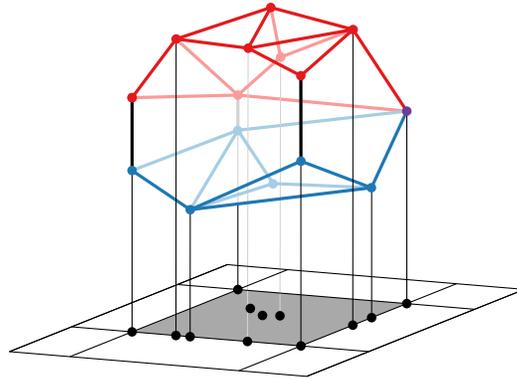


Figure 4: A convex polyhedron $\text{conv}(P)$ in \mathbb{R}^3 , whose projection $\text{conv}(P)^{\text{proj}}$ is a rectangle. Seven points in P are projected onto the four vertices of $\text{conv}(P)^{\text{proj}}$. Overall, the silhouette of P contains twelve points, five of which project into the relative interior of some edges of $\text{conv}(P)^{\text{proj}}$. Red (blue) vertices are upper (lower); the purple point is both upper and lower.

can contain interior grid points. We can control the number of such faces and such interior points by induction. Summation over axis-aligned faces of dimensions $1, \dots, d-1$ yields an upper bound of $(2^{d-1} - 2)2^{(d-1)d+3} \cdot j^{d-2}$ for such points. Summation over (1)–(3) yields the desired upper bound $2^{d(d+1)} \cdot i \cdot j^{d-2}$ for the number of upper (resp., lower) vertices.

Proof of Lemma 5. It is enough to prove the claim when $n = 2^j$ for integers $j \geq 0$. For integers $d \geq 3$ and $0 \leq i \leq j$, we recursively define the grid $S_d(i, j)$ as a Cartesian product of d sets of reals, where the first $d-1$ sets have 2^j elements and the last set has 2^i elements. In particular, $|S_d(i, j)| = 2^{i+(d-1)j}$. We then show that $S_d(i, j)$ does not contain the vertex set of any full-dimensional convex polyhedron with more than $2^{(d^2+d+1)} \cdot i \cdot j^{d-2}$ vertices.

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 $k \in \{1, \dots, d-2\}$, if a point $p \in S_{d-1}(j, j)$ lies in the relative interior of k -polytope F whose vertices are in $S_{d-1}(j, j)$, then F is axis-aligned.

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 $i = 1, \dots, j$, 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 a (large) scalar satisfying the following two conditions.

- (C1) For all $p \in A \cup B$ and $Q \subset A \cup B$, if p is in the relative interior of $\text{conv}(Q)$, then $\text{conv}(Q)$ is an axis-aligned polytope.
- (C2) For all $p_a \in A$, $p_b \in B$, and $Q \subset S_d(0, j) = S_{d-1}(j, j) \times \{0\}$, if both p_a^{proj} and p_b^{proj} are in the relative interior of $\text{conv}(Q)$, then the line $p_a p_b$ intersects the horizontal hyperplane $x_d = 0$ in the relative interior of $\text{conv}(Q)$.

We show that a scalar λ_d^i satisfying both (C1) and (C2) exists. First, we may assume that $\lambda_d^i > 0$ is sufficiently large so that A and B are separated by some horizontal hyperplane. Consider a set $Q \subset A \cup B$ such that Q contains points in both A and B . Let $F = \text{conv}(Q)$, $F_a = \text{conv}(Q \cap A)$, and $F_b = \text{conv}(Q \cap B)$. Since $\text{conv}(Q) = \text{conv}(F_a \cup F_b)$, we have $F \subset \text{conv}(F_a \cup B)$ and $F \subset \text{conv}(A \cup F_b)$. We set $\lambda_d^i > 0$ sufficiently large so that every point in $A \cap \text{conv}(F_a \cup B)$ is on a vertical line passing through a point in F_a , and every point in $B \cap \text{conv}(A \cup F_b)$ is on a vertical line passing through a point in F_b . Assume that this condition holds for all $Q \subset A \cup B$, where $Q \cap A \neq \emptyset$ and $Q \cap B \neq \emptyset$.

We show that (C1) holds if λ_d^i is chosen as specified above. Let $Q \subset A \cup B$ and let $p \in A \cup B$ be in the relative interior of $\text{conv}(Q)$. Let $F = \text{conv}(Q)$, $F_a = \text{conv}(Q \cap A)$, and $F_b = \text{conv}(Q \cap B)$, as above, and F is a k -dimensional polytope for some $k \in \{1, \dots, d\}$. We distinguish between several cases based on the location of p . (1) If p is in the relative interior of F_a (resp., F_b), then F_a (resp., F_b) is axis-aligned by induction, and it has the same dimension as F , hence F is also axis-aligned. (2) Otherwise p is contained in the relative interior of neither F_a nor F_b . By Carathéodory's theorem [12], p is in the relative interior of some simplex $F' = \text{conv}(Q')$ with vertex set $Q' \subset Q$ such that Q' has vertices in both A and B . By the choice of λ_d^i , point p is on the vertical line passing through a point in $Q' \cap A$ or $Q' \cap B$. This implies that p lies on a vertical line segment contained in F' , hence in F . Consequently, F is parallel to the x_d -axis, and F^{proj} is a $(k-1)$ -polytope. If $k=1$, this already proves that F is axis-aligned. Assume that $k > 1$ and note that $p^{\text{proj}} \in S_{d-1}(j, j)$ and it is in the relative interior of F^{proj} . By induction, F^{proj} is an axis-aligned $(k-1)$ -polytope. Overall, F is parallel to the x_d -axis and the $k-1$ coordinate axes parallel to F^{proj} , consequently F is axis-aligned, as required.

To satisfy condition (C2), notice that for any $p_a \in A$ and $p_b \in B$, the intersection of the line $p_a p_b$ and the hyperplane $x_d = 0$ can be arbitrarily close to p_a^{proj} if $\lambda_d^i > 0$ is sufficiently large. This completes the definition of $S_d(i, j)$ for $d \geq 3$ and $0 \leq i \leq j$.

Note, however, that in each iteration of the recursive contraction of $S_d(i, j)$, we only imposed lower bounds for the scalars λ_d^i , and different values of these parameters would yield different point sets. For all integers $d \geq 1$ and $0 \leq i \leq j$, let $\mathcal{S}_d(i, j)$ denote the collection of all $2^j \times \dots \times 2^j \times 2^i$ grids that can be obtained by the recursive construction above, in particular every grid in $\mathcal{S}_d(i, j)$ satisfies both (C1) and (C2) in each iteration. Our proof relies on the following relation between these constructions.

Claim 1. *Let $2 \leq \ell \leq d$, let $S_d(i, j) \in \mathcal{S}_d(i, j)$, and let W be an axis-aligned vertical ℓ -flat such that $S_d(j, j) \cap W \neq \emptyset$. Then $S_d(i, j) \cap W \in \mathcal{S}_\ell(i, j)$.*

To prove Claim 1, note first that $S_d(i, j) \cap W$ is a $2^j \times \dots \times 2^j \times 2^i$ grid. It can be obtained by the recursive construction above with suitable parameters λ_d^i in the recursive steps. Indeed, assume that $S_d(i, j) = A \cup B$, where $A = S_d(i-1, j)$ and $B = A + \lambda_d^i \mathbf{e}_d$. Then $S_d(i, j) \cap W = (A \cap W) \cup (B \cap W)$. As both (C1) and (C2) are formulated in terms of the relative interior of a set $Q \subset A \cup B$, both conditions hold for $Q \subset (A \cup B) \cap W$, as well. This completes the proof of Claim 1.

In the remainder of the proof, we establish the following claim by induction:

Claim 2. *Let $2 \leq d$ and $S_d(i, j) \in \mathcal{S}_d(i, j)$. If $P \subset S_d(i, j)$ is a d -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 $\text{conv}(P)$.*

We prove Claim 2 by induction on $d + i$. We focus on the number of upper vertices in P , the case of lower vertices follows analogously (e.g., by reflection in a horizontal hyperplane).

In the base case, we have $d = 2$ and $i = 0$. Then the set $S_d(0, j) = S_{d-1}(j, j) \times \{0\}$ lies in a horizontal hyperplane in \mathbb{R}^d , and so a subset $P \subset S_d(0, j)$ cannot be full-dimensional, hence the claim vacuously holds for all $P \subset S_d(0, j)$. We observe a key property for P (which will be used in the case $i = 1$). If $d = 2$ and $P \subset S_d(0, j)$ is a 1-dimensional set, then P has precisely two extreme points in \mathbb{R}^2 . If $d \geq 3$ and $P \subset S_d(0, j)$ is a $(d - 1)$ -dimensional set, then by induction it contains at most $2^{(d-1)d} \cdot j \cdot j^{d-3}$ upper (resp., lower) vertices in \mathbb{R}^{d-1} , hence it has at most $2 \cdot 2^{(d-1)d} \cdot j^{d-2}$ extreme points in \mathbb{R}^d .

Assume next that $i = 1$, and either $d = 2$ or the claim holds for $S_{d-1}(j, j)$. We prove the claim for $S_d(1, j)$. The set $S_d(1, j)$ is the disjoint union of $A = S_d(0, j)$ and $B = S_d(0, j) + \lambda_d^1 \mathbf{e}_d$. Let $P \subset S_d(1, j)$ be in convex position. Then every upper (resp., lower) vertex of P is an extreme vertex in $P \cap A$ or $P \cap B$, hence P contains at most $2 \cdot 2 \cdot 2^{(d-1)d} \cdot j^{d-2} = 2^{d^2-d+2} \cdot j^{d-2} \leq 2^{d(d+1)} \cdot 1 \cdot j^{d-2}$ upper (resp., lower) vertices.

In the general case, we assume that $2 \leq i \leq j$, the claim holds for $S_d(i - 1, j)$, and either $d = 2$ or the claim holds for $S_{d-1}(j, j)$, as well. 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 in convex position. We partition the upper vertices in P as follows. For every upper vertex p , we have $p^{\text{proj}} \in \text{conv}(P)^{\text{proj}}$, and p^{proj} is either a vertex of $\text{conv}(P)^{\text{proj}}$ or lies in the relative interior of a k -face of $\text{conv}(P)^{\text{proj}}$ for some $k \in \{1, \dots, d - 1\}$. Let P_0 be the set of upper vertices $p \in P$ such that p^{proj} is a vertex of $\text{conv}(P)^{\text{proj}}$. For all $k \in \{1, \dots, d - 1\}$, let $P_k \subset P$ be the set of upper vertices of P such that p^{proj} lies in the relative interior of a k -face of $\text{conv}(P)^{\text{proj}}$. Then $\bigcup_{k=0}^{d-1} P_k$ is the set of all upper vertices in P .

The orthogonal projection of $S_d(i, j)$ to the hyperplane $x_d = 0$ is $S_{d-1}(j, j)$, and the orthogonal projection of P_0 , denoted $P_0^{\text{proj}} \subset S_{d-1}(j, j)$, is the vertex set of a $(d - 1)$ -dimensional convex polyhedron. By induction, $|P_0| \leq 2 \cdot 2^{(d-1)d} \cdot j \cdot j^{d-3} = 2^{(d-1)d+1} j^{d-2}$. To derive an upper bound on $|P_k|$ for $k \in \{1, \dots, d - 1\}$, note that only axis-aligned faces of $\text{conv}(P)^{\text{proj}}$ can contain grid points of $S_{d-1}(j, j)$ in their relative interior by condition (C1). In order to complete the proof of Claim 2, we prove the following.

Claim 3. *For every axis-aligned face F of $\text{conv}(P)^{\text{proj}}$, either A or B contains all upper (resp., all lower) vertices that project to the relative interior of F .*

Let F be an axis-aligned k -face of $\text{conv}(P)^{\text{proj}}$ for $k \in \{1, 2, \dots, d - 1\}$. Let $P(F) \subset P$ be the set of upper vertices $p \in P$ such that p^{proj} lies in the relative interior of F , and let $V(F)$ be the vertex set of F . Since F is a face of the convex polytope $\text{conv}(P)^{\text{proj}}$ in the horizontal hyperplane $x_d = 0$, we have $V(F) \subset S_{d-1}(j, j)$ and $V(F)$ is in convex position. Consider the point set $P' = V(F) \cup P(F)$, and observe that if $P(F) \neq \emptyset$, then P' is the vertex set of the $(k + 1)$ -polytope $\text{conv}(P')$, in which F is one of the facets.

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 $P(F) \cap A$ with the maximum x_d -coordinate. By construction, every point in B has higher x_d -coordinate than any point in A (including p_a and all vertices of $V(F)$). The 1-skeleton of $\text{conv}(P')$ contains a x_d -monotonically increasing path from p_a to an x_d -maximal vertex in P' . Let p_b be the neighbor of p_a along such a path. By the maximality of p_a , we have $p_b \in B$. Then $p_a p_b$ is an edge of $\text{conv}(P')$, hence the line $p_a p_b$ is disjoint from the interior of $\text{conv}(P')$. However, by condition (C2), the line $p_a p_b$ intersects the relative interior of facet F of $\text{conv}(P')$. This contradiction completes the proof of Claim 3.

We can now finish the proof of Claim 2. We have seen that $|P_0| \leq 2^{(d-1)d+1} j^{d-2}$. It remains to bound $|P_k|$ for $k \in \{1, \dots, d-1\}$. Note that $\text{conv}(P)^{\text{proj}}$ is a $(d-1)$ -dimensional polytope. We can apply Claim 3 for $k = d-1$ and $F = \text{conv}(P)^{\text{proj}}$, and conclude that either A or B contains all points in P_{d-1} . The points in P_{d-1} are upper vertices of $\text{conv}(P_{d-1})$. Since both A and B are translates of $S_d(i-1, j)$, the induction hypothesis yields $|P_{d-1}| \leq 2^{d(d+1)} \cdot (i-1) \cdot j^{d-2}$.

For $k \in \{1, \dots, d-2\}$, the axis-aligned k -faces of $\text{conv}(P)^{\text{proj}}$ 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 $(d-k-1)$ -dimensional space (of the hyperplane $x_d = 0$), they become vertices of a $(d-k-1)$ -polytope. By induction, the number of (upper and lower) vertices of such a polytope is bounded above by $2 \cdot 2^{(d-k-1)(d-k)} j^{(d-k-1)-1} = 2^{(d-k-1)(d-k)+1} j^{d-k-2}$; this is an upper bound on the size of the equivalence class. Each axis-aligned k -face F of $\text{conv}(P)^{\text{proj}}$ is the orthogonal projection of some axis-aligned $(k+1)$ -face of $\text{conv}(P)$, that we denote by \overline{F} . Claim 1 with $\ell = k+1 \leq d-1$ implies that the induction hypothesis holds for \overline{F} , and it yields an upper bound of $2^{(k+1)(k+2)} \cdot j^k$ for the number of upper vertices in \overline{F} . Overall, we have

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

where we used that $(d-k-1)(d-k)+1+(k+1)(k+2) = (d-1)d+2k(k+2-d)+3 \leq (d-1)d+3$ for all $k \in \{1, \dots, d-2\}$ and $d \geq 3$. Altogether, the number of upper vertices is

$$\begin{aligned} \sum_{k=0}^{d-1} |P_k| &= |P_0| + \left(\sum_{k=1}^{d-2} |P_k| \right) + |P_{d-1}| \\ &\leq 2 \cdot 2^{(d-1)d} \cdot j^{d-2} + \left(\sum_{k=1}^{d-2} \binom{d-1}{k} 2^{(d-1)d+3} \cdot j^{d-2} \right) + 2^{d(d+1)} \cdot (i-1) \cdot j^{d-2} \\ &< 2^{d-1} \cdot 2^{(d-1)d+3} j^{d-2} + 2^{d(d+1)} \cdot (i-1) \cdot j^{d-2} \\ &< 2^{d(d+1)} j^{d-2} + 2^{d(d+1)} \cdot (i-1) \cdot j^{d-2} \\ &= 2^{d(d+1)} \cdot i \cdot j^{d-2}, \end{aligned}$$

as required, where we used the binomial theorem. This completes the proof of Claim 2.

Claim 2 immediately implies Lemma 5 by taking $n = 2^j$ and $X = S_d(j, j)$, since $2 \cdot 2^{d(d+1)} \cdot j^{d-1} = \Theta(\log^{d-1} n)$ for every fixed $d \geq 2$. \square

2.4 Lower bounds 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 $n \times \dots \times n$ 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 $\prod_{i=1}^d X_i$, where $|X_i| = n$ and d is fixed, contains $\Omega(\log^{d-1} n)$ points in convex position.*

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

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

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

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

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 $A = (a_1, \dots, a_n)$ is MD if and only if there exists a monotone (increasing or decreasing) convex function $f : \mathbb{R} \rightarrow \mathbb{R}$ such that $a_i = f(i)$ for all $i = 1, \dots, n$. MD sets have been studied in additive combinatorics [13, 21, 26, 30].

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 $(0, \dots, n-1)$ contains the subsequence $(i^2 : i = 0, \dots, \lfloor \sqrt{n-1} \rfloor)$. 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 $n \times \dots \times n$ 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* [29]) of size $\Omega(\log n)$ [6, Lemma 4.1]; see also [9, 14] 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 $X = (x_0, \dots, x_{n-1})$ be a strictly increasing sequence. Assume, without loss of generality, that $n = 2^\ell + 1$ for some $\ell \in \mathbb{N}$. We construct a sequence of nested intervals

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

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 $i = 1, \dots, \ell$.

We start with the interval $[a_0, b_0] = [x_0, x_{n-1}]$; and for every $i = 0, \dots, \ell - 1$, we divide $[a_i, b_i]$ into two intervals at the median, and recurse on the shorter interval; see Figure 5.

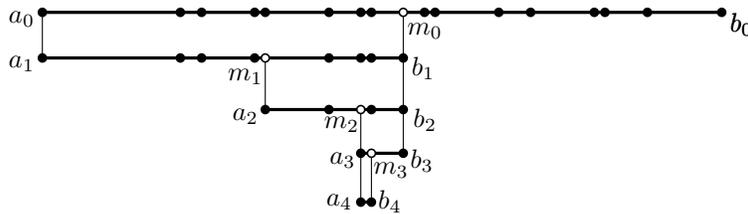


Figure 5: A sequence X of 17 elements and nested intervals $[a_0, b_0] \supset \dots \supset [a_4, b_4]$.

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 $i = 1, \dots, \ell$, 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 $A = (a_0, a_1, \dots, a_\ell)$ and $B = (b_\ell, b_{\ell-1}, \dots, b_0)$ both increase (not necessarily strictly), and at least one of them contains at least $1 + \ell/2$ distinct terms. Assume, without loss of generality, that A contains at least $1 + \ell/2$ distinct terms. Let $C = (c_0, \dots, c_k)$ 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 $i \in \{1, \dots, k - 1\}$. Assume that $c_i = a_j = \dots = a_{j'}$ for consecutive indices j, \dots, 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

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

For the upper bound, consider the $n \times n$ grid $\{0, \dots, n - 1\} \times Y$, defined in the proof of Lemma 4, for which every chain in \searrow or \swarrow supported by $\{0, \dots, n - 1\} \times Y$ has at most $\lceil \log n \rceil + 1$ vertices. We claim that the size of every MD subset of $Y \subset \mathbb{R}$ is at most $\lceil \log n \rceil + 1$. Let $\{b_0, \dots, b_{\ell-1}\} \subset Y$ be an MD subset such that $b_0 < \dots < b_{\ell-1}$. Then $\{(i, b_i) : i = 0, \dots, \ell - 1\} \subset X \times Y$ is in \searrow or \swarrow . Consequently, every MD subset of Y has at most $\lceil \log n \rceil + 1$ terms, as claimed. \square

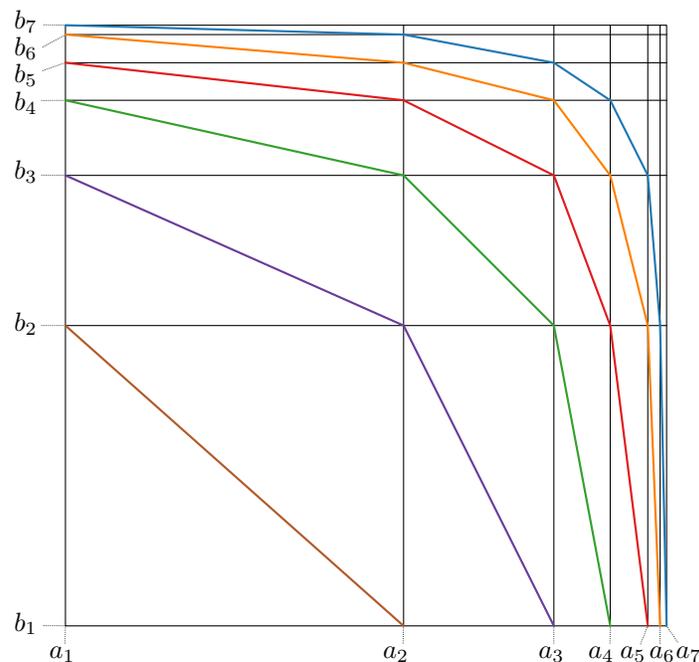


Figure 6: A 7×7 grid $\{a_1, \dots, a_7\} \times \{b_1, \dots, b_7\}$, where the differences between consecutive x -coordinates (resp., y -coordinates) decrease by factors of 2 or higher. The point sets $\{(0, 0)\} \cup \{(a_i, b_j) : i + j = k\}$, for $k = 2, \dots, 8$, form nested convex chains.

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 $A = \{a_1, \dots, a_n\}$ and $B = \{b_1, \dots, b_n\}$ be MD sets such that $a_i < a_{i+1}$ and $b_i < b_{i+1}$ for $i \in \{1, \dots, n-1\}$. 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 $i \in \{2, \dots, n-1\}$; see Figure 6.

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 $b_{n-i} - b_{n-i-1}$ are also positive and strictly decrease in i , the slopes $(b_{n-i} - b_{n-i-1}) / (a_i - a_{i-1})$ 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 $A = \{a_1, \dots, a_n\}$, $B = \{b_1, \dots, b_n\}$, and $C = \{c_1, \dots, c_n\}$ 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

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

for $i = 2, \dots, n - 1$. For $i, j, k \in \{1, \dots, n\}$, 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 $\text{conv}(P)$. We show that the remaining faces of $\text{conv}(P)$ 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 Figure 6. By construction, the union of these faces is homeomorphic to a sphere. It suffices to show that the dihedral angle between any two edge-adjacent triangles is convex. Without loss of generality, consider triangle $T'_{i,j,k}$, which shares an edge with (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 dihedral 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 $b_{i+1} - b_i < b_i - b_{i-1}$. \square

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 $d = 2$ and $d = 3$, Lemmas 8 and 9 prove the claim. Assume that $d \geq 4$, and the claim holds in lower dimensions. For every $i \in \{1, \dots, d\}$, let $A_i = \{a_{i,1}, \dots, a_{i,n}\} \subset \mathbb{R}$ be an MD set such that $a_{i,1} < \dots < a_{i,n}$. We may assume, without loss of generality, that the differences between consecutive elements in A_i strictly decrease for all $i \in \{1, \dots, d\}$.

For every vector $\mathbf{v} = (v_1, \dots, v_d) \in \{1, \dots, n\}^d$, let $p_{\mathbf{v}} = (a_{1,v_1}, a_{2,v_2}, \dots, a_{d,v_d}) \in \prod_{i=1}^d A_i$. 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 $P' = P \cup \{p_{1,\dots,1}\}$. We show that the points in P' are in convex position. We define a $(d - 1)$ -dimensional piecewise linear manifold M as the union of $(d - 1)$ -dimensional convex polyhedra (facets). We show that M is homeomorphic to the sphere \mathbb{S}^{d-1} . We also show that M is the boundary of a convex polytope by verifying that the dihedral angle between any two adjacent facets is convex. It follows that M is the boundary of $\text{conv}(P')$, consequently P' is in convex position.

For every $i \in \{1, \dots, d\}$, the points of P' lying in the hyperplane $x_i = a_{i,1}$ are in convex position by induction, hence they each span a $(d - 1)$ -dimensional facets of $\text{conv}(P')$. For every vector $\mathbf{v} \in \{1, \dots, n - 1\}^d$, consider the hyperrectangle $r_{\mathbf{v}} = \prod_{i=1}^d [a_{i,v_i}, a_{i,v_i+1}]$,

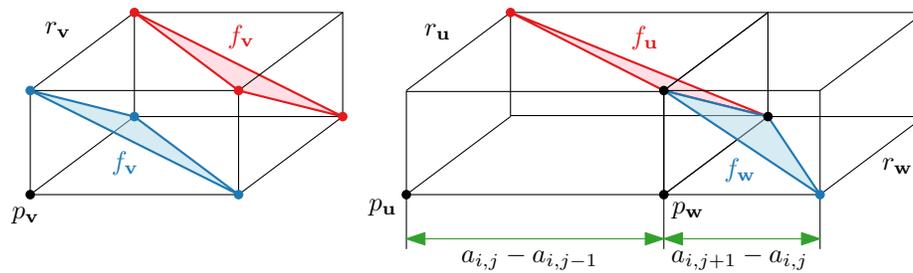


Figure 7: Left: A rectangle r_v in \mathbb{R}^3 , and two possible locations for f_v , depending on the parity of $v_1 + v_2 + v_3$. (In general, the position of f_v in r_v depends on $\sum_{i=1}^d v_i \pmod{d-1}$.) Right: Two adjacent facets $f_u, f_w \in \mathcal{F}$ in two adjacent hyperrectangles r_u and r_w .

where $[a_{i,v_i}, a_{i,v_i+1}]$ is the i -th extent of r_v . All vertices of each r_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_v , let $F_v = \text{conv}(P \cap r_v)$, that is, the convex hull of vertices of r_v that are in P ; see Figure 7(left). By construction, every F_v is at most $(d-1)$ -dimensional (possibly empty). Let \mathcal{F} be the set of $(d-1)$ -dimensional polyhedra F_v , where $v \in \{1, \dots, n-1\}^d$, 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 all $i \in \{1, \dots, d\}$, is homeomorphic to \mathbb{S}^{d-1} . Consider two facets $f_u, f_w \in \mathcal{F}$ that share a $(d-2)$ -dimensional face; see Figure 7(right). Then f_u and f_w lie in two adjacent hyperrectangles, say r_u and r_w , whose common boundary is $(d-1)$ -dimensional, and it is contained in a hyperplane $x_i = a_{i,j}$ for some $i \in \{1, \dots, d\}$ and $j \in \{2, \dots, n-1\}$. The facet f_u (resp., f_w) is parallel to the $(d-1)$ -simplex spanned by the d vertices of r_u (resp., r_w) adjacent to p_u (resp., p_w). Since A_i is MD, we have $a_{i,j+1} - a_{i,j} < a_{i,j} - a_{i,j-1}$. Note that $a_{i,j+1} - a_{i,j}$ and $a_{i,j} - a_{i,j-1}$ are the i -th extents of r_u and r_w , respectively; and r_u and r_w have the same extent in the remaining $d-1$ coordinates. Consequently, the dihedral angle between f_u and f_w is convex, as required. \square

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., 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. We briefly review these results as they provide an algorithm for the case of maximum chains supported by a grid, since the points in a convex chain have distinct x - and y -coordinates.

Given a set P of N points in the plane, Edelsbrunner and Guibas [11, Thm. 5.1.2] use the dual line arrangement and dynamic programming to find the maximum size of a convex

cap contained in \curvearrowright in $O(N^2)$ time and $O(N)$ space; the same bounds hold for \smile , ζ , and \succcurlyeq . A convex cap of maximum length can be also computed in $O(N^2 \log N)$ time and $O(N \log N)$ space. Their method sweeps the dual line arrangement, in which each vertex (intersection of two dual lines) corresponds to a line segment pq for $p, q \in P$, $p \neq q$. Specifically, for every pair of points $p, q \in P$, let $L(p, q)$ be the maximum length of a chain in \curvearrowright whose *last* two vertices are p and q , respectively. Similarly, let $R(p, q)$ be the maximum length of a chain in \curvearrowright whose *first* two vertices are p and q , respectively. The sweepline algorithm in [11] computes the values $L(p, q)$ (or all values $R(p, q)$) for all $p, q \in P$, $p \neq q$, in $O(N^2)$ time.

The following observation allows us to adapt the algorithm in [11] to find the maximum size of a convex cap in \curvearrowright (alternatively, to report a longest such chain) within the same time and space bounds. Analogous observations hold for \curvearrowleft , \curvearrowright , and \smile .

Observation 11. *Let P be a finite set of points in the plane, and let $p, q \in P$ such that $x(p) < x(q)$. If $\text{slope}(pq) > 0$, then $L(p, q)$ is the maximum length of a chain in \curvearrowright whose last two vertices are p and q .*

Consequently, the maximum length of a chain in \curvearrowright is $\max L(p, q)$, where the maximum is taken over all pairs $p, q \in P$ where $x(p) < x(q)$ and $y(p) < y(q)$. Since x - and y -coordinates do not repeat in a convex chain, we obtain the following result.

Theorem 12. *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.*

3.1 Convex caps

In order to compute the maximum size of a convex cap in \curvearrowright , we must be careful to use each y -coordinate at most once. We solve the more general problem of computing the maximum total size of two chains $A \in \curvearrowright$ and $B \in \curvearrowleft$ that use distinct y -coordinates. Note that in contrast to the chains constituting a convex cap, the chains A and B may cross, or can have overlapping x -projections.

We present below a dynamic program to solve this problem, essentially building up solutions from bottom to top. To define the appropriate subproblem, we need to be able to refer to the last segment on the upper side of a chain. Hence, we refer to the last two vertices of A as l_1 and l_2 , setting $l_2 := l_1$ if A is to consist of the single vertex l_1 . Similarly, we refer to the first two vertices of B as r_1 and r_2 , setting $r_2 := r_1$ if B is to consist of the single vertex r_1 . The subproblem $C(l_1, l_2, r_1, r_2)$ then expresses the maximum total size of two chains $A \in \curvearrowright$ and $B \in \curvearrowleft$ with the given end vertices for A and start vertices for B . We claim that $C(l_1, l_2, r_1, r_2)$ as defined below yields the desired quantity, or $-\infty$ if no such

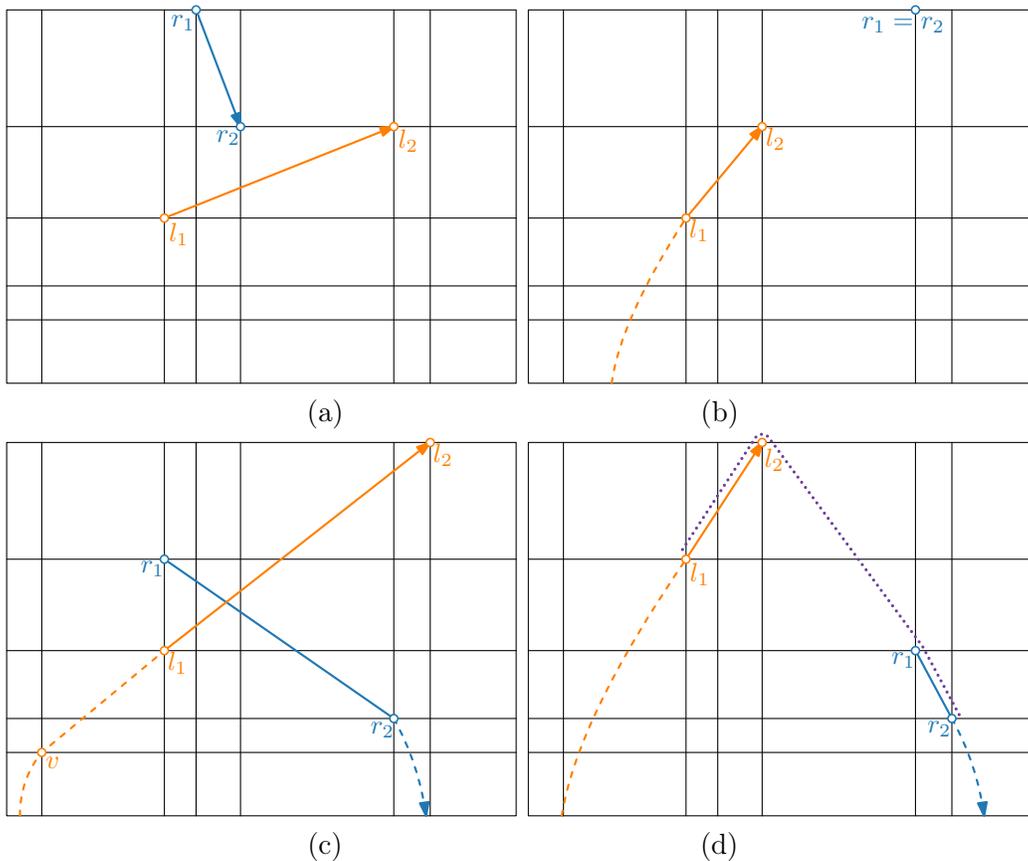


Figure 8: Illustration for the cases of $C(l_1, l_2, r_1, r_2)$. (a) Invalid configuration, as $y(r_2) = y(l_2)$. (b) $r_1 = r_2$ is a single point above l_2 , so we look for the longest chain ending in (l_1, l_2) . (c) Removing the topmost point (l_2 in this case), testing all valid possible v to find the longest chain. Note that the left and right chain may not complete to a cap – this is checked separately. (d) We need test only whether (l_1, l_2) and (r_1, r_2) combine to make a cap (purple dotted line) to check whether the entry (l_1, l_2, r_1, r_2) should be considered.

chains A and B exist.

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

Informally, the first clause of the definition of $C(l_1, l_2, r_1, r_2)$ gets rid of invalid arguments, for which no such chains $A \in \curvearrowright$ and $B \in \curvearrowleft$ exist; see also Figure 8. The

second clause is the base case where A and B each have one vertex. The third and fourth clause handle the case where the topmost chain consists of a single vertex—it can thus not interfere with the other (longer) chain anymore. Here, $L(p_1, p_2)$ (resp., $R(p_1, p_2)$) denotes 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 $p_1 = p_2$. We can use the dynamic programming algorithm of [11] to compute $L(p_1, p_2)$ and $R(p_1, p_2)$. The final two clauses deal with the case where the topmost chain consists of at least one edge, looking for the largest pair of chains that can be constructed by adding the topmost vertex to a ‘lower’ pair of chains. We prove the correctness of the formula in Lemma 15, using the following general observations.

Observation 13. *If a supported convex polygon P is in a set $\nearrow, \searrow, \swarrow, \nwarrow, \curvearrowright, \curvearrowleft, \curvearrowup, \curvearrowdown$, or \curvearrowright , then every subsequence of P is in the same set. That is, these classes are hereditary.*

Lemma 14. *Let $A = (a_1, \dots, a_k)$ and $B = (b_1, \dots, b_{k'})$ with $k, k' \geq 2$, $a_{k-1} = b_1$ and $a_k = b_2$. If A and B are both in the same set of $\searrow, \nearrow, \swarrow$, or \nwarrow , then their concatenation $D = (a_1, \dots, a_k, b_3, \dots, b_{k'})$ is in the same set.*

Proof. Consider the case that A and B are in \nearrow ; the other cases are analogous. This means that both x -coordinates and their y -coordinates strictly increase in A and B . As $a_k = b_2$, we know that $x(a_k) < x(b_i)$ and $y(a_k) < y(b_i)$ for all $i > 2$. Therefore, the concatenation $D = (a_1, \dots, a_k, b_3, \dots, b_{k'})$ has strictly increasing x - and y -coordinates. Except for its endpoints, each vertex of D is an interior vertex of A or B . Since A and B are convex, this readily implies that all interior vertices of D are convex and hence D is convex. Thus, D is in \nearrow as well. \square

Lemma 15. *Let k be the maximum total size of two chains $A \in \nearrow$ and $B \in \searrow$ that use distinct y -coordinates, subject to the constraints that A ends in a given edge (l_1, l_2) (or $A = (l_1)$ if $l_1 = l_2$) and B starts in a given edge (r_1, r_2) (or $B = (r_1)$ if $r_1 = r_2$), or let $k = -\infty$ if no such chains exists. Then $C(l_1, l_2, r_1, r_2) = k$.*

Proof. Suppose that $A \in \nearrow$ and $B \in \searrow$ realize the maximum total size over all chains with distinct y -coordinates under the given constraints. If the y -coordinates of $\{l_1, l_2\}$ and $\{r_1, r_2\}$ are not disjoint, then A and B do not use distinct y -coordinates, so $C(l_1, l_2, r_1, r_2)$ correctly returns $-\infty$. So suppose that the y -coordinates differ. If $l_1 = l_2$ and $r_1 = r_2$, then A and B both have size 1, and we correctly return 2.

If $l_1 \neq l_2$ and $A \in \nearrow$, then by Observation 13 we must have $(l_1, l_2) \in \nearrow$, so we correctly return $-\infty$ if $(l_1, l_2) \notin \nearrow$. Similarly, if $r_1 \neq r_2$ and $(r_1, r_2) \notin \searrow$, we correctly return $-\infty$.

For the remaining cases, the point of $\{l_1, l_2, r_1, r_2\}$ with maximum y -coordinate will be either l_2 or r_1 . If $y(l_2) < y(r_1)$ and $r_1 = r_2$, then B has size 1 and A has only y -coordinates at most that of l_2 . Since $y(l_2) < y(r_1)$, any chain of \nearrow that ends in (l_1, l_2) will have y -coordinates distinct from B , so A has the maximum size over all chains of \nearrow ending in (l_1, l_2) , as given by $L(l_1, l_2)$. Therefore we correctly return $L(l_1, l_2) + 1$. Symmetrically, if $y(r_1) < y(l_2)$ and $l_1 = r_2$, we correctly return $R(r_1, r_2) + 1$.

Two (symmetric) cases remain: either $y(l_2) < y(r_1)$ and $r_1 \neq r_2$, or $y(r_1) < y(l_2)$ and $l_1 \neq l_2$. We use induction on the maximum y -coordinate to prove that $C(l_1, l_2, r_1, r_2) = k$

in these cases. We show the case where $y(l_2) < y(r_1)$ and $r_1 \neq r_2$, the argument for the other case is symmetric. Removing the point r_1 from chain B results in a chain $B' \in \searrow$ of size one less than that of B . The chain B' starts in r_2 , and its y -coordinates are still disjoint from those of A . Since B had at least one edge, B' consists either of a single vertex (namely r_2), or it starts with an edge (r_2, v) for which $(r_1, r_2, v) \in \searrow$. Since these are exactly the terms we take the maximum over in the last clause of the definition of $C(l_1, l_2, r_1, r_2)$, we by induction return at least k .

It remains to show that, $C(l_1, l_2, r_1, r_2) \leq k$. For this, let $A'' \in \nearrow$ and $B'' \in \searrow$ be two chains that realize the maximum total size over all pairs of chains with distinct y -coordinates, where A'' ends in (l_1, l_2) and B'' consists either of the single vertex r_2 , or it starts with an edge (r_2, v) for which $(r_1, r_2, v) \in \searrow$ (if no such chains exist, we correctly return $-\infty + 1 = -\infty$). It suffices to show that adding r_1 in front of chain B'' yields a chain of \searrow with y -coordinates distinct from those of A'' . Indeed, since $y(r_1)$ is greater than all other y -coordinates, the y -coordinates remain distinct. Moreover, if B'' consists of the single vertex r_2 , adding r_1 in front of it yields (r_1, r_2) , which we have already verified to lie in \searrow (we would have returned $-\infty$ otherwise). If $B'' \in \searrow$ instead starts in an edge (r_2, v) for which $(r_1, r_2, v) \in \searrow$, then the chain with r_1 in front also lie in \searrow , by Lemma 14. Hence, $C(l_1, l_2, r_1, r_2)$ returns at most k , and hence exactly k . \square

Lemma 16. *The value of C for all inputs can be computed in $O(n^{10})$ time and $O(n^8)$ space.*

Proof. Deciding the applicable clause for an input takes constant time. Each clause takes at most $O(n^2)$ time to compute, assuming that the referenced terms C , L , and R are already computed. For the third and fourth clauses, we can precompute all $L(l_1, l_2)$ and $R(r_1, r_2)$ within the desired time and space. The last two clauses of the equation reference $O(n^2)$ inputs of C recursively, but with a smaller maximum y -coordinate, so this recurrence is well-defined. There are $O(n^8)$ possible inputs to C , each of which takes constant additional space and $O(n^2)$ time to compute, which gives the desired bounds. \square

Any cap (of size at least 2) can be split into two chains $A \in \nearrow$ and $B \in \searrow$ with distinct y -coordinates by removing the topmost edge. Since unless $n = 1$, the largest cap has size at least 2, we can compute the size of the largest cap using C by taking the maximum over all inputs for whose concatenation (omitting l_1 if $l_1 = l_2$ and r_2 if $r_1 = r_2$) lies in \curvearrowright , as shown in Lemma 17.

Lemma 17. *Suppose that $A = (l_1, \dots, l_a) \in \nearrow$ and $B = (r_1, \dots, r_b) \in \searrow$ use distinct y -coordinates, and let D be the concatenation of the last two vertices (or single vertex if A has size 1) of A and the first two vertices of B (or single vertex if B has size 1). If $D \in \curvearrowright$, then the concatenation $D' = (l_1, \dots, l_a, r_1, \dots, r_b)$ of A and B lies in \curvearrowright .*

Proof. Since $D \in \curvearrowright$, we have $x(l_a) < x(r_1)$, so the x -coordinates of D' are strictly increasing, and the y -coordinates are distinct. Except l_1 and r_b , every vertex of D' is a non-endpoint vertex of A , B or D . Since these are convex, the vertex is convex in D' as well; the entire concatenation is convex as a result. Thus, D' lies in \curvearrowright . \square

Testing if an input of C results in a cap takes constant time, so Theorem 18 follows.

Theorem 18. *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 $X = \{x_1, \dots, x_n\}$ and $Y = \{y_1, \dots, y_n\}$ with $x_i < x_{i+1}$ and $y_i < y_{i+1}$. To test whether there is a chain of size n in \curvearrowright , it suffices to test whether the chain $((x_1, y_1), \dots, (x_n, y_n))$ is in \curvearrowright , in $O(n)$ time.

To test whether there is a supported convex cap of size n in \curvearrowleft , we adapt the algorithm of Theorem 18. Suppose P is a cap of \curvearrowleft of size n , with A_P and B_P its maximal components in \curvearrowright and in \curvearrowleft respectively. Then P uses all coordinates of X , which restricts the types of chains A_P and B_P considerably. In particular $C(l_1, l_2, r_1, r_2)$ can be modified to consider only edges l and r that use consecutive x -coordinates.

For $1 < k < n$ consider the subchains $A_k \in \curvearrowright$ and $B_k \in \curvearrowleft$ 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 $\{y_1, \dots, y_k\}$. 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 $C(l_1, l_2, r_1, r_2)$. 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 $C(l_1, l_2, r_1, r_2)$ 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.

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 , \curvearrowright , and \curvearrowleft 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 $\Theta^*(\cdot)$ notation hides polynomial factors in n .

Theorem 19. *The following bounds hold:*

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

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 $\nearrow_P, \searrow_P, \swarrow_P, \nwarrow_P$, with only extreme vertices of P appearing in multiple chains. Let $\zeta_P = \nearrow_P \cup \searrow_P$ and $\upsilon_P = \swarrow_P \cup \nwarrow_P$. To encode P , we assign the following number to each of the $2n$ grid lines ℓ (see Figure 9 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 ζ_P if ℓ is horizontal, or on υ_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).

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

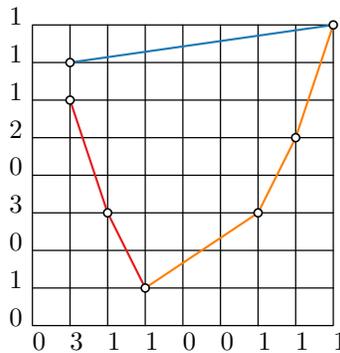


Figure 9: Encoding the grid lines.

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 $n = 2m + 4$, where $m \in \mathbb{N}$ satisfies suitable divisibility conditions, as needed. All four lower bounds use the same grid, constructed as follows (see Figure 10).

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

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 $1 \leq i < j < k \leq n$.

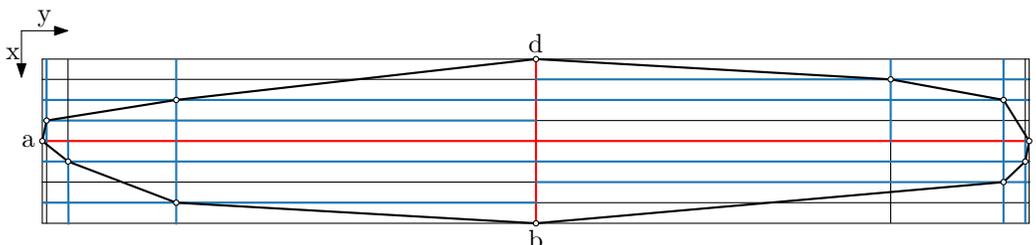


Figure 10: The $n \times n$ grid defined in Section 4.2, with $n = 9 = 2m + 3$ for $m = 3$, before doubling the median lines (red). The segments (parts of grid lines) incident to vertices are drawn in blue.

Then the slope of pq is strictly smaller than the slope of qr , since

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

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 \searrow if they are in the lower half and \nwarrow 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 lowest vertices are the intersections of the duplicate vertical 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

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

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 the number of vertices with the two different chains) to make up our subgrid and plug in the bound on $\bar{G}(n)$, which yields

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

Finally, for the bound on $F(n)$, we select $3m/4$ grid lines on each side of each median (balancing the number of 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

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

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 20. *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 $n \times \dots \times n$ 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 [22, 33] 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 $n \times \dots \times n$ 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 [16] 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.

Acknowledgments. We are grateful to the anonymous referees for their careful reading of the paper that helped clarify several subtle details in the inductive proof in Section 2.

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](https://doi.org/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](https://doi.org/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. doi:[10.1201/9781315119601](https://doi.org/10.1201/9781315119601).
- [6] Boris Bukh and Jiří Matoušek. Erdős–Szekeres-type statements: Ramsey function and decidability in dimension 1. *Duke. Math. J.*, 163(12):2243–2270, 2014. doi:[10.1215/00127094-2785915](https://doi.org/10.1215/00127094-2785915).
- [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](https://doi.org/10.1007/s00022-016-0323-5).
- [8] Václav Chvátal and Gheza Thomas 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](https://doi.org/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](https://doi.org/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](https://doi.org/10.1016/0022-0000(89)90038-X).
- [12] Harold Gordon Eggleston. Helly’s theorem and its applications. In *Convexity*, Cambridge Tracts in Mathematics, chapter 2, pages 33–44. Cambridge University Press, 1958. doi:[10.1017/CB09780511566172.003](https://doi.org/10.1017/CB09780511566172.003).
- [13] 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](https://doi.org/10.1006/jnth.1999.2386).
- [14] Marek Eliáš and Jiří Matoušek. Higher-order Erdős–Szekeres theorems. *Advances in Mathematics*, 244:1–15, 2013. doi:[10.1016/j.aim.2013.04.020](https://doi.org/10.1016/j.aim.2013.04.020).
- [15] 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](https://doi.org/10.1112/jlms/s1-26.3.198).
- [16] Paul Erdős. Some more problems on elementary geometry. *Gazette of the Australian Mathematical Society*, 5(2):52–54, 1978.

- [17] Paul Erdős and György Szekeres. On some extremum problems in elementary geometry. *Annales Universitatis Scientiarum Budapestinensis de Rolando Eötvös Nominatae Sectio Mathematica*, 3-4:53–62, 1960/1961.
- [18] 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](https://doi.org/10.1007/978-3-642-40450-4_46).
- [19] 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. doi:[10.1201/9781315119601](https://doi.org/10.1201/9781315119601).
- [20] Richard K. Guy and Patrick A. Kelly. The no-three-in-line-problem. *Canadian Mathematical Bulletin*, 11:527–531, 1968. doi:[10.4153/CMB-1968-062-3](https://doi.org/10.4153/CMB-1968-062-3).
- [21] Norbert Hegyvári. On consecutive sums in sequences. *Acta Mathematica Hungarica*, 48(1-2):193–200, 1986. doi:[10.1007/BF01949064](https://doi.org/10.1007/BF01949064).
- [22] 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](https://doi.org/10.1007/s00454-003-0009-4).
- [23] 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](https://doi.org/10.1016/0020-0190(95)00130-5).
- [24] 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](https://doi.org/10.1137/120897493).
- [25] 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](https://doi.org/10.1007/s00453-006-0158-9).
- [26] 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](https://doi.org/10.1016/j.comgeo.2016.12.002).
- [27] 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](https://doi.org/10.1353/ajm.2016.0033).
- [28] 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](https://doi.org/10.1215/00127094-3674103).
- [29] David Alan Rosenthal. *The classification of the order indiscernibles of real closed fields and other theories*. PhD thesis, University of Wisconsin–Madison, 1981.

- [30] Tomasz Schoen and Ilya D. Shkredov. On sumsets of convex sets. *Combinatorics, Probability and Computing*, 20(5):793–798, 2011. doi:10.1017/S0963548311000277.
- [31] 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.
- [32] Andrew Suk. On the Erdős-Szekeres convex polygon problem. *Journal of the American Mathematical Society*, 30:1047–1053, 2017. doi:10.1090/jams/869.
- [33] 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.
- [34] Eric W. Weisstein. No-three-in-a-line-problem. Wolfram MathWorld, 2005. URL: <http://mathworld.wolfram.com/No-Three-in-a-Line-Problem.html>.